ABSTRACT

Title of Dissertation: ESSAYS ON EMPIRICAL ASSET PRICING

Jinming Xue
Doctor of Philosophy, 2020

Dissertation Directed by: Professor Russ Wermers
Department of Finance

and

Professor Albert "Pete" Kyle
Department of Finance

This dissertation contains three essays on empirical asset pricing.

Chapter 1 measures the time-varying provision of liquidity by buy-side customers (e.g., mutual funds and pension funds), relative to bond dealers, in corporate bond markets using a structural vector autoregression (SVAR) model. As indicated by my simple theory model, shocks to the relative willingness of customers and bond dealers to provide liquidity affect, in opposite directions, the choice of bond dealers between market-making (principal) and matchmaking (riskless principal) transactions. Motivated by this model, my SVAR empirically disentangles these shocks to customers versus bond dealers. My SVAR-derived patterns of these structural shocks provide fundamental insights into the mechanics in corporate bond markets following recent events, such as exposing the increased role of buy-side customers for liquidity provision after the many regulatory changes following the 2008 financial crisis. Furthermore, my empirical approach generates "factors" that provide

an improved time-series asset-pricing model for yield spreads of corporate bonds of different credit ratings.

In Chapter 2, with Gurdip Baskhi and Xiaohui Bakshi, we consider an approach in which the conditional expectation of asset return quantities (under the real-world probability measure) can be synthesized from the prices of the risk-free bond, the asset, and options on the asset. The method is free of distributional assumptions, and we use it to study empirical questions related to (i) conditional probability of a disaster and return upside and (ii) spanning hypothesis in the Treasury market. We examine empirical consistency and show that our theoretical treatment is relevant.

In Chapter 3, with Haibo Jiang and Georgios Skoulakis, we study the importance of dissecting oil price changes in predicting global equity returns. Based on data until the mid-2000s, oil price changes were shown to predict international equity index returns with a negative predictive slope. Extending the sample to 2015, we document that this relationship has been reversed over the last ten years and therefore has not been stable over time. We then posit that oil price changes are still useful for forecasting equity returns once complemented with relevant information about oil supply and global economic activity. Using a structural VAR approach, we decompose oil price changes into oil supply shocks, global demand shocks, and oil-specific demand shocks. The hypothesis that oil supply shocks and oil-specific demand shocks (global demand shocks) predict equity returns with a negative (positive) slope is supported by the empirical evidence over the 1986-2015 period. The results are statistically and economically significant and do not appear to be consistent with time-varying risk premia.

Essays on Empirical Asset Pricing

by

Jinming Xue

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2020

Advisory Committee:
Professor Russ Wermers, Chair
Professor Albert "Pete" Kyle, Co-Chair
Professor Gurdip S. Bakshi
Professor Alberto Rossi
Professor John Chao (Dean's Representative)

Acknowledgments

I would like to express my sincere gratitude to all the people without whom I would not have made it through, and because of whom my graduate experience has been one that I will cherish forever.

First and foremost I thank my advisor, Professor Russ Wermers, for encouraging me and supporting me to pursue the degree and giving me invaluable opportunities to work on various challenging and extremely exciting research projects. We have known each other for almost ten years, and I have benefited tremendously from him both academically and personally. This dissertation wouldn't be possible without countless discussions with Russ and with many reputable scholars that he introduced to me. During my most difficult times, Russ was indeed a great friend who encouraged me to make it through.

I would also like to thank my co-advisor, Professor Albert "Pete" Kyle, who is a great scholar and an excellent doctoral mentor. Pete brought me into the market microstructure literature. He gives freely of his time to me, as well as other students, which makes the Ph.D. program at Smith fantastic.

It is my pleasure to work with Professor Gurdip Bakshi, to whom I am tremendously grateful. Gurdip is an exceptional coauthor and mentor who helps me cultivate my ability in conducting research and thinking critically. Talking with Gurdip has always been, and will be, fun and inspiring.

Thanks are due to Professor Alberto Ross and John Chao for agreeing to serve on my thesis committee. Alberto has given me invaluable advice in the past six years. John is one of the best teachers I have ever met, and I enjoyed his econometric courses.

I owe the deepest thanks to my mother and father, who have always been supportive through my career. Words cannot express the gratitude I owe them. I would also like to thank my dear colleagues Professor Xiaohui Gao Bakshi, Professor Haibo Jiang, and Professor Georgios Skoulakis.

Most importantly, a thousand special thanks go to our favorite, excellent, and beloved Justina Blanco—she entirely takes care of all the Ph.D. students and let me focus on my research.

Taiyi Yan is the best housemate and gym buddy. I'd like to express my gratitude to my peers Ziwei Cao, Wen Chen, Bo Hu, Xiaoyuan Hu, Shuaiqi Li, Liyue Yan, Zhang Zhang, and Wei Zhou for their friendship and support.

It is impossible to remember all, and I apologize to those I've inadvertently left out.

Lastly, thank you all!

Table of Contents

List of Tables

List of Figures

Chapter 1: Measuring Liquidity Provision by Customers in Corporate Bond Markets

1.1 Introduction

US corporate bond markets have long been a focal point of both academics and regulators in their search to identify who provides liquidity in this (largely) over-the-counter (OTC) market, the size of which has grown to over \$9.2tn in 2018, compared to \$5.5tn in 2008.[1] Furthermore, the structure of bond market trading has been substantially affected by regulatory and technology developments since the financial crisis, including the Volcker Rule, the Basel banking accords, and the emergence of electronic trading platforms ([1]). These developments raise the fundamental question of who provides liquidity in this market sector and how such liquidity provision changes during market stress events. Liquidity provision by customers has attracted considerable attention, as bond dealers appear to have become increasingly reluctant to provide liquidity due to regulatory changes ([2])).

Although various approaches have been used in the literature to measure aggregate liquidity provision in corporate bond markets,[2] it remains unclear how,

[1] See `https://www.sifma.org/resources/research/fixed-income-chart/`.

[2] See, for example, [3], [4], [5], [6], [7], [8], and [9].

exactly, to measure the liquidity provision by buy-side customers. In response to this question, my study provides a structural vector autoregression (SVAR) – with predictions supported by my simple theory model – to measure the willingness of buy-side customers to provide liquidity in corporate bond markets. My SVAR measures the willingness of customers to provide liquidity using the amount of riskless principal transactions (RPTs), relative to the total transacting volume, conducted by bond dealers.

In RPTs, dealers buy from customers and immediately sell to other customers who are the ultimate liquidity providers,[3] whereas in inventory transactions, dealers provide liquidity and hold the resulting inventory. Thus, the relative amount of RPTs is a natural candidate for measuring customer liquidity provision. This amount, however, reflects the competing willingness of customers versus dealers to provide liquidity. Either a decrease in the risk-taking willingness of bond dealers or an increase in the risk-taking willingness of customers tends to lower the observed relatively amount of RPTs.

A separation of these two potential influences relies on an auxiliary measurement—the incremental trading costs of inventory transactions relative to that of RPTs—that reflects the marginal willingness for bond dealers to provide liquidity. A decrease in the risk-taking willingness of bond dealers will increase the incremental costs, whereas an increase in the risk-taking willingness of customers will lower this incremental costs because increased customer liquidity provision willingness leads

[3]Dealers help customers search for potential customer liquidity providers and get compensated for their efforts in the searching process.

to less dealer inventory.

This is a crucial differentiating intuition of my SVAR model in identifying shocks to customers versus bond dealers. At the core of the SVAR are the following three time series: (i) the bid-ask spreads implied by inventory transactions, (ii) the bid-ask spreads implied by RPTs, and (iii) the fraction of RPT volume over total trading volume. The SVAR is governed by (negative) shocks to three structural variables, which are the risk-taking willingness of (i) customers who demand liquidity (hereafter "hedgers"), (ii) customer liquidity providers, and (iii) bond dealers. Structural shocks to the risk-taking willingness of dealers and customer liquidity providers are taken to be my new measures of liquidity.

What is the rationale to include these variables in the SVAR, and what are the economic connections between the structural shocks and the observed trading activities? First, the market-making willingness of bank-affiliated broker dealers might have been compromised by the above-noted regulatory changes, resulting in lowered capital commitments ([10]); if so, this is a structural shock lowering their risk-taking willingness. Another example is the run on the repo market in 2008 and ensuing deterioration of the securitized banking system ([11]). When the required return on capital is high due to capital constraints, bond dealers tend to conduct RPTs and charge higher bid-ask spreads for inventory transactions.

Second, the risk-taking willingness of customer liquidity providers is my measure of customer liquidity. The empirical literature has focused on post-crisis regulatory shocks to bank-affiliated bond dealers and the resulting high demand for liquidity provision by buy-side customers ([2]). Distinct from this perspective, my measure captures the supply side dynamics. One important example of supply side dynamics is capital inflows into fixed income mutual funds, which may lead to an increased willingness to buy distressed bonds at favorable prices. In such cases, more RPTs will be conducted.

Third, the risk-taking willingness of hedgers approximates their eagerness to liquidate the bond. For instance, unexpected losses due to natural disasters force insurance companies to liquidate their positions ([12]). Shocks that lower the risk-taking willingness of hedgers lower their reservation prices and, thus, increase the bid-ask spreads due to the market power of bond dealers.

I separately estimate the SVAR for investment-grade bond markets and high-yield bond markets. The SVAR-derived patterns of the structural shocks are sensible and instructive about important historic events. Bond dealers aggressively built inventory before the 2008 crisis when "too big to fail" guarantees incentivized dealers toward excessive risk-taking activities ([1]). Conversely, bond dealers were much less willing to provide liquidity during the collapse of Lehman Brothers and the post-regulation periods.

In sharp contrast, customer liquidity provision increases for investment-grade bond markets during the post-regulation period, consistent with the view that recently launched alternative trading systems (ATSs), a new but growing mechanism that is parallel to the classic OTC market-making mechanism, manage to locate more customer liquidity providers. The assets traded in ATSs tend to be liquid bonds with investment-grade ratings ([13]). My measure also uncovers an increase in customer liquidity provision during the 2008 financial crisis, especially for investment-grade bonds. Evidence suggests that customer liquidity providers crowded into corporate bond markets during the crisis. This is a "flight-to-safety" phenomenon, as documented by [14].

Measuring customer liquidity provision has potentially significant welfare implications. Theoretically, bank-affiliated dealers who face stringent regulations tend to increase investments on matchmaking technology ([15]). Improved matchmaking technology could reach more customer liquidity providers such that bond dealers can fulfill more orders from customers who have hedging demands but otherwise would not trade. Complementary to this theory, my contribution is to provide empirical evidence that developments in matchmaking technology, indeed, attract more customer liquidity providers.

Customer liquidity provision could potentially supplement aggregate liquidity when the willingness of bond dealers to provide liquidity contracts ([16]). In contrast, aggregate market liquidity could be largely suppressed if the willingness of both dealers and customers to provide liquidity decreases. Therefore, measuring the supply of liquidity provision by customers is important from the perspective of regulators and investors. Overall, my proposed SVAR approach can be utilized as a tool in understanding and evaluating market conditions and the associated impact on the tendency of dealers versus customers to provide liquidity.

Furthermore, I study the asset-pricing implications of my derived liquidity measures. The question asked is whether my measures are related to time-series changes in yield spreads of corporate bonds. My analysis reveals that changes in yield spreads for bonds of different credit ratings respond differentially to my measures of liquidity. High quality (above A-rated) corporate bonds are more exposed to my measure of customer liquidity, whereas lower quality bonds (below A-rated) are mostly subject to my measure of dealer liquidity.

The differential impact of shocks to dealers versus customer liquidity providers also translates into different exposures of yield spreads in response to recent regulations (decrease in liquidity provision by dealers) and technology developments (increase in liquidity provision by buy-side customers). Raising concerns for regulators, the mixed consequence of stringent regulations and technology improvements could point to an outcome that safer bonds become more liquid, whereas markets are even thinner for bonds that are riskier and, likely, already illiquid. For example, [9] document that illiquidity of downgraded bonds is even more severe after the implementation of the Volcker Rule. [17] show that prices of investment-grade bonds increase, while prices of speculative-grade bonds fall during illiquidity periods. My paper attributes such effects to the conflicting exposures of corporate bonds to my measures of customer liquidity providers and bond dealers. [13] find that, for downgraded corporate bonds, trading activities shift from electronic platforms to voice trading.

Empirically, the identification of a SVAR requires certain restrictions that need to be justified by economic theory. To motivate my identification strategy, I present a theoretical model of segmented markets. In each market, there is a market maker and two types of customers (hedgers and customer liquidity providers). Customers can trade *only* with the local market maker, and all market makers strategically trade in a centralized system. A market maker may buy from hedgers, sell a fraction to customer liquidity providers (RPTs), sell in the interdealer market, and keep the remaining as inventory.

This model sheds light on the three identification restrictions required in the SVAR. The risk-taking willingness of neither dealers nor customer liquidity providers will affect the bid-ask spreads for RPTs. The intuition is that dealers' own capital is not involved. In addition, an increase in customer liquidity provision improves both the bid and the ask prices in parallel, but not the bid-ask spreads. The third restriction is that the fraction of RPT volume responds only to the relative changes in the risk-taking willingness among customers and dealers. In other words, this fraction remains stagnant when the risk-taking willingness of all market participants increases in proportion.

In addition, I exploit an alternative empirical method to identify the SVAR, which imposes sign restrictions on the impulse response functions ([18]), called the "agnostic method." The method focuses on the shock of interest (either hedgers, dealers, or customer liquidity providers) but remains agnostic about other shocks; only partial identification is achieved under this method. The advantage is that the imposed restrictions exploit relatively weak prior beliefs, and these restrictions are independent from the three zero restrictions. Compared with the full identification approach, this agnostic method, as a safeguard, yields similar outcomes.

This paper contributes to the literature that investigates the impact of post-crisis regulations on bond dealers and customer liquidity providers.[4] [2] find increased RPT volume and show that bid-ask spreads for inventory transactions in-

[4] See, for example, [19], [20], [10], and [21].

deed become larger during the post-regulation periods. My paper is among the few that employ a structural VAR model to study liquidity of various OTC-traded assets. [22] and [23] use SVAR with sign restrictions to investigate the liquidity supply by primary dealers in Treasury and corporate bond markets. I depart from these papers in that my paper is the first to decompose demand-side versus supply-side shocks that jointly drive prices and quantities of principal transactions versus RPTs, and my focus is on the liquidity supplied by customers versus bond dealers.

My decomposition is extremely useful in understanding both the mechanism of liquidity provision as well as how this mechanism varies across bond sectors of differential credit qualities. Without this decomposition, deriving such a measure is challenging. Whether an investor, or a sector of investors, is a liquidity demander or a liquidity provider may change over time and across holdings, depending on the state of financial markets and the original position of the investor. For example, hedge funds provide liquidity to corporate bond markets during normal times, but were compelled to unwind their pre-crisis "basis trades" positions in the months following Lehman's collapse ([24]). [16] find that in the corporate bond market some mutual funds provide liquidity, whereas some mutuals demand liquidity. My SVAR approach produces the desired aggregate market level measures by using prices and quantities from transaction data, without relying on the identities of the customers that are usually unavailable to researchers.

1.2 An SVAR approach

This section discusses the SVAR approach. This approach can help explain the source of shocks that affect the observed market trading patterns. My design of the SVAR separately yields measures of shocks to customer liquidity providers and shocks to bond dealers.

The SVAR incorporates three market-wide aggregate time series: (i) bid-ask spreads from riskless principal transactions (denoted by $\text{Spread}_{\text{t}}^{\text{RPT}}$), (ii) bid-ask spreads from inventory transactions (denoted by $\text{Spread}_{\text{t}}^{\text{IVT}}$), and (iii) fraction of customer-dealer-customer riskless principal transaction volume over total customer-dealer volume (denoted by %RPT).

Let the vector $\mathbf{y_t}$ contain the three series:

$$\mathbf{y_t} = \left[\log\left(\text{Spread}_{\text{t}}^{\text{RPT}}\right),\ \log\left(\text{Spread}_{\text{t}}^{\text{IVT}}\right),\ \%\text{RPT}_{\text{t}}\right]^T. \tag{1.1}$$

The SVAR specification[5] reads as

$$\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i}\mathbf{y_{t-i}} + \underbrace{\begin{bmatrix} b_1 & b_{1,2} & b_{1,3} \\ b_{2,1} & b_2 & b_{2,3} \\ b_{3,1} & b_{3,2} & b_3 \end{bmatrix}}_{\mathbf{B}} \underbrace{\begin{bmatrix} \epsilon_t^H \\ \epsilon_t^L \\ \epsilon_t^d \end{bmatrix}}_{\boldsymbol{\epsilon_t}}, \tag{1.2}$$

where $\boldsymbol{\epsilon_t} = \left[\epsilon_t^H,\ \epsilon_t^L,\ \epsilon_t^d\right]^T$ is a vector of unit variance orthogonal shocks to hedgers

[5]Logarithmic transformation is applied in order to be consistent with the conclusions drawn from the model in Section 1.5. Empirical results are economically similar without the logarithmic transformation.

(ϵ_t^H), customer liquidity providers (ϵ_t^L), and dealers (ϵ_t^d). A positive innovation in any element in $\epsilon_\mathbf{t}$ means that the corresponding group of participants in aggregation becomes more risk averse and less willing to provide liquidity.

I assume the orthogonality conditions in my main analysis. Appendix 1.7.4 presents a case in which the orthogonality conditions are violated. In that case, the recovered structural shocks can be interpreted as shocks to hedgers, shocks to dealers that are orthogonal to the shocks to the hedgers, and shocks to customer liquidity providers that are orthogonal to the other two types of shocks.

The SVAR does not incorporate proxies for the dynamics of information asymmetry, which could be a potential concern for market makers to charge bid-ask spreads. Empirical literature, however, shows that adverse selection is not a significant component for the effective bid-ask spreads in corporate bond markets.[6]

1.2.1 What is the expected impact of the structural shocks?

The specification in equation (1.2) links the dynamics of unobserved structural shocks to various market participants ($\epsilon_\mathbf{t}$) to the observable market trading patterns of interest ($\mathbf{y_t}$) to academics and regulators. The matrix $\mathbf{B}$ governs these links and determines the direction and the magnitude of the impact of each shock. My

[6] See [25], [26], and [27].

hypotheses on the directions of the elements in $\mathbf{B}$ are summarized below.

$$\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i y_{t-i}} + \underbrace{\begin{bmatrix} + & 0 & 0 \\ + & ? & + \\ ? & - & + \end{bmatrix}}_{\mathbf{B}} \underbrace{\begin{bmatrix} \epsilon_t^H \\ \epsilon_t^L \\ \epsilon_t^d \end{bmatrix}}_{\epsilon_\mathbf{t}} . \tag{1.3}$$

There are five priors (marked as "+" or "−") and two zero restrictions (marked as "0"); the associated economic intuitions are discussed below. The directions of the remaining two elements (marked as "?") are undetermined. I solicit answers based on the empirical data.

ϵ_t^H captures the eagerness of hedgers to liquidate a bond position due to certain exogenous liquidity shocks or changes in the desired inventory level. Positive shocks enlarge the gap of the reservation prices between hedgers and (both dealer and customer) liquidity providers.[7] Thus, profit-optimizing bond dealers charge higher $\text{Spread}_\text{t}^\text{RPT}$, when conducting offsetting transactions between two customers; that is, $b_1 > 0$. Besides, this widened gap entices bond dealers to charge higher $\text{Spread}_\text{t}^\text{IVT}$ because risk-averse dealers take more inventory and require higher compensation: $b_{2,1} > 0$.

A positive ϵ_t^d captures increased cost of market-making during the 2008 financial crisis and the post-regulation period. Facing such shocks, bond dealers conduct more riskless principal transactions ($\%\text{RPT}_\text{t}$) and request more compensations for market-making ($\text{Spread}_\text{t}^\text{IVT}$). Therefore, I hypothesize $b_3 > 0$ and $b_{2,3} > 0$. [2]

[7] A reservation price is the private value of the bond from the perspective of a participant; that is the marginal price at which a participant would be optimal if she does not trade.

show economically insignificant difference between $\mathrm{Spread}_{\mathrm{t}}^{\mathrm{RPT}}$ in the pre-crisis period (when market-making cost is low) and that in the post-regulation period. The cost of market-making (ϵ_t^d), thus, does not directly relate to $\mathrm{Spread}_{\mathrm{t}}^{\mathrm{RPT}}$, namely, $b_{1,3} = 0$.

The shock ϵ_t^L is considered as a parallel shift of the supply curve of aggregate customer liquidity due to various reasons. For instance, cash flows shift between equity markets and fixed income markets. Another reason is that bond dealers invest in matchmaking technology so that they can easily find a customer liquidity provider and conduct more riskless principal transactions. Higher ϵ_t^L, or less customer liquidity provision, is expected to lower $\%\mathrm{RPT}_{\mathrm{t}}$: $b_{3,2} < 0$. Note that such an impact is distinct from a shock to dealers who choose the optimal quantities along the supply curve of customer liquidity.

The shock ϵ_t^L does not affect $\mathrm{Spread}_{\mathrm{t}}^{\mathrm{RPT}}$, specifically $b_{1,2} = 0$. Customer liquidity providers trade (indirectly) with hedgers because of their differential reservation prices. ϵ_t^L affects neither of the two reservation prices. Hence, it is optimal for a dealer to charge the same $\mathrm{Spread}_{\mathrm{t}}^{\mathrm{RPT}}$, despite the shifts of both the bid and the ask prices.

The two zero restrictions are important in estimating the SVAR. The other priors are not used in the estimation procedure, but they serve the purpose of testing the soundness of the SVAR, as predicted by economic intuitions and my theoretical model.

1.2.2 Estimation procedure for the SVAR

The following three restrictions on $\mathbf{B}$ are imposed to identify the SVAR:

$$b_{1,2} = 0, \qquad b_{1,3} = 0, \qquad \text{and} \qquad b_{3,1} + b_{3,2} + b_3 = 0. \tag{1.4}$$

I restrict that $b_{1,2} = b_{1,3} = 0$ because ϵ_t^d and ϵ_t^L have no impact on $\mathrm{Spread}_\mathrm{t}^{\mathrm{RPT}}$. The third restriction ($b_{3,1} + b_{3,2} + b_3 = 0$) obeys the intuition that $\%\mathrm{RPT}_\mathrm{t}$ depends on the relative change in shocks between dealers and customers. If the risk-taking willingness of hedgers, dealers, and customer liquidity providers were all halved, the equilibrium prices would improve, but all volumes would remain the same. When every participant becomes proportionally less willing to hold this bond,[8] hedgers are more eager to sell, and dealers and customer liquidity providers become reluctant to hold this bond. Dealers and customer liquidity providers will still buy the same units of the bond from hedgers since the change is proportional. In Section 1.5, I show that $\frac{\partial \%\mathrm{RPT}_\mathrm{t}}{\partial \log(\delta_L)} + \frac{\partial \%\mathrm{RPT}_\mathrm{t}}{\partial \log(\delta_H)} + \frac{\partial \%\mathrm{RPT}_\mathrm{t}}{\partial \log(\delta_d)} = 0$, where δ_H, δ_L, and δ_d are the inverse of risk-taking willingness of the corresponding market participants, and logarithms of the δs represent $\epsilon_\mathbf{t}$ in the SVAR.

Assuming the restrictions stated in equation (1.4), the SVAR is estimated in two steps [28, p. 372]. In the first step, loadings on the lagged $\mathbf{y_t}$ terms, $\hat{\mathbf{A}} = \left[\hat{\mathbf{A}}_\mathbf{0} \dots \hat{\mathbf{A}}_\mathbf{I}\right]$, are estimated via ordinary least squares (OLS). The optimal number of lags I is determined by the Schwarz-Bayesian information criterion (SBIC). Next,

[8] Assume a positive net supply of the asset.

loadings on the structural shocks, $\hat{\mathbf{B}}$, are estimated by maximizing the log-likelihood,

$$\log \mathcal{L}_c(\mathbf{B}) = \text{constant} - \frac{T}{2}\log|\mathbf{B}| - \frac{T}{2}\text{trace}\left(\mathbf{B}'^{-1}\mathbf{B}^{-1}\tilde{\Sigma}_u\right), \tag{1.5}$$

where $\tilde{\Sigma}_u$ is the covariance matrix estimated from the residuals in the first step.

In closing, I present the design of a SVAR in analyzing trading activities in corporate bond markets. In doing so, I outline my priors for coefficients in matrix **B** and the procedure for estimation.

1.3 Data

This section introduces the academia TRACE data and the steps to clean the database. I describe the method that I use to classify each customer-dealer transaction in TRACE as a riskless principal transaction or an inventory transaction. My focus is on aggregating the transaction level data into (i) the bid-ask spreads from riskless principal transactions ($\text{Spread}_\text{t}^\text{RPT}$), (ii) the bid-ask spreads from inventory transactions ($\text{Spread}_\text{t}^\text{IVT}$), and (iii) the ratio of riskless principal transaction volume over total customer-dealer volume ($\%\text{RPT}_\text{t}$). These time-series are inputs into the SVAR.

1.3.1 Academia TRACE

The Trade Reporting and Compliance Engine (TRACE) is a FINRA-developed vehicle for mandatory reporting in OTC markets. It collects all the secondary market transactions conducted by FINRA registered broker-dealers in the U.S. corporate bond markets. The standard TRACE contains a limited set of information for each transaction, including capped transaction quantity, transaction price, execution date and time, indicator for dealer/client buy and sell, among others.

In addition to that information, the academia TRACE provides uncapped transaction quantity and, importantly, masked identifications of the dealers that facilitate the transactions. Those fields allow me to identify riskless principal transactions conducted by the same dealers within a short time period.

The academia TRACE provides two fields: (a) **Reporting Market Participant Identifier** is the dealer who reports the transaction and likely the dealer who executes the transaction. (b) **Contra Identifier** could be either the contra-party dealer, if the field is populated by a masked dealer ID, or a customer denoted by a letter C. According to these two fields I can find riskless principal transactions, provided that the buyer in one transaction is the seller in the other transaction and that other transactional information matches.

I take care of the exceptions that the reporting party is not the party that executes the transaction, that is, "locked in" and "give up" transactions. In those cases, the additional fields **Reporting Side Give Up Participant Identifier** and/or **Contra Give Up Identifier** are populated.[9]

Table 1.1 lists the main steps that I implement for data cleaning. The database contains more than 159 million transactions for 106,873 distinct issues facilitated by

[9]One example is that Dealer A hires Dealer B to report transactions conducted by Dealer A. In the example, Reporting Side Give Up Participant Identifier is Dealer A, and Reporting Market Participant Identifier is Dealer B. Another example is that two dealers transact in an alternative trading system (ATS). The ATS, on behalf of two dealers, reports the transaction to TRACE. Overall, the dealers in the Give Up fields should be treated as the actual buyer or seller in the transactions whenever the Give Up fields are populated.

3,815 unique dealers over the entire sample period from 2002 to 2015. I apply the procedure from [29] to clean duplicated records in TRACE when a dealer attempts to cancel or modify a reported transaction, resulting in a total of 113 million transaction records.

Please insert Table 1.1 here

I focus on non-puttable US Corporate Debentures and US Corporate Bank Notes as identified by FISD bond type (CDEB or USBN) after the end of phase-in period (Dec 31, 2005). Here I adopt two exclusionary criteria. First, I exclude transactions that occur within 30 days of the offering date. Second, I eliminate affiliated transactions in which a FINRA member dealer transfers its inventory to an affiliated party for bookkeeping purposes. Affiliated transactions are identified as those offsetting transactions conducted by a dealer at the exact same prices for both legs, and the dealer makes zero profits.[10]

The cleaned sample consists of 24,754 unique Cusips and approximately 54 million transactions conducted by 3,389 unique dealers. The cleaned TRACE data is merged with Mergent Fixed Income Securities Database (FISD) for information of bond characteristics and credit ratings.

[10] The detailed identification procedure is described in Appendix A of [2]. TRACE started to require reporting dealers to indicate affiliated transactions on 11/02/2015. This affiliated transaction indicator only covers the last two months of the data, so that I opted not to use it.

1.3.2 Classifying riskless principal transactions

I follow the "last in, first out" (LIFO) method in [2] to classify riskless principal transactions. In general, a set of transactions on a single Cusip are considered as riskless principal transactions if the following conditions are satisfied:

- The transactions are conducted by the same dealer.
- The dealer buys and sells within a small time interval.
- The total volumes of buys and sells are approximately equal.

Appendix 1.7.5 lists the detailed steps in identifying riskless principal transactions.

Based on the outcome of this algorithm, every customer-dealer transaction is classified based on the identity of the counterparty in the offsetting transaction. I focus on two types of transactions: (i) customer-dealer-customer riskless principal transactions (RPTs) and (ii) customer-dealer-customer inventory transactions (IVTs).

In a pair of RPTs, a customer sells to a dealer and the dealer sells to another customer within 15 minutes. In IVTs, a customer sells to a dealer and this dealer takes no selling action in the next 15 minutes. Transactions in which a dealer buys from a customer and sells to another dealer immediately are not considered. If a transaction partially follows the categories above, the classification is based on the largest portion. For instance, if a customer sells to a dealer 1,000 shares, and the dealer sells 800 shares to another customer immediately, it is considered an RPT.

1.3.3 Computation of bid-ask spreads and trading volume

This subsection describes my method of computing the bid-ask spreads and presents the summary statistics of the monthly average bid-ask spreads and trading volume for RPTs and IVTs. Bid-ask spreads are computed for transactions with par value greater than (inclusive) $100K. The full bid-ask spread is twice the difference between the traded price and the reference price, which is the volume-weighted average interdealer prices for the same bond on the same day.

I focus on the largest dealers, which tend to be bank-affiliated dealers, selected based on a K-mean clustering method. The number of selected dealers per month is between 10 and 15 during the sample period of 2006 to 2015. Figure 1.11 in Appendix demonstrates the resulted clusters. These numbers are consistent with the top 70% sample in [10].

Please insert Figure 1.11 here

K-mean clustering is an algorithm to partition observations into K clusters so that the total distance between each observation and the mean of the belonged cluster is minimized. The advantage of the K-mean clustering method is to jointly consider multidimensional activities of each dealer, including the customer-dealer trading volume, the interdealer trading volume, and the amount of offsetting transactions. It is better than a fixed threshold because there are large banks that enter and exit the market, and the aggregate trading volume is time-varying. The method excludes interdealer brokers that arrange agency trades only for dealers.

Table 1.2 presents the summary statistics for the trading activities. I first average the bid-ask spreads and volumes for all transactions over a given month and then report the mean, standard deviation and percentiles for these monthly aggregated quantities.

Please insert Table 1.2 here

In investment-grade (IG) bond markets, the mean of $\text{Spread}_t^{\text{IVT}}$ is 47.46bps, while the average of $\text{Spread}_t^{\text{RPT}}$ is 22.12bps. In an average month, there are \$16.8 billion par value and 1,588 distinct Cusips transacted for IVTs and \$1.3 billion par value and 252 distinct Cusips transacted for RPTs. The average number of transactions per month with par value above (inclusive) \$1MM is 2,674 for IVTs and 173 for RPTs.

The number of transactions used to compute $\text{Spread}_t^{\text{RPT}}$ is limited. Note that the level of reference price, which is hard to accurately pin down, is not important in the computation of $\text{Spread}_t^{\text{RPT}}$. It does not necessarily require a large number of transactions to infer an aggregate $\text{Spread}_t^{\text{RPT}}$. To further reduce noise, especially for $\text{Spread}_t^{\text{IVT}}$, I also include more transactions, that is, those with par value between \$100K and \$1MM. Transactions with a size smaller than \$100K tend to incur significantly higher bid-ask spreads, so they are excluded from my aggregations.

Similar patterns are observed for high-yield (HY) bonds. HY bonds are more likely to be transacted in RPT and dealers tend to charge higher bid-ask spreads for HY bonds.

Please insert Figure 1.1 here

Figure 1.1 plots the time series used for the SVAR. Presented are the monthly average $\text{Spread}_t^{\text{RPT}}$ and $\text{Spread}_t^{\text{IVT}}$ for IG and HY bonds, based on Moody's Corporation. Both spreads reach the peak during the 2008 financial crisis. $\text{Spread}_t^{\text{IVT}}$ are always greater than $\text{Spread}_t^{\text{RPT}}$. The differences between $\text{Spread}_t^{\text{IVT}}$ and $\text{Spread}_t^{\text{RPT}}$ are widened during the post-regulation period, relative to the pre-crisis period. Also presented is the ratio of riskless principal transaction volume over total customer-dealer volume ($\%\text{RPT}_t$). $\%\text{RPT}_t$ is higher during the post-regulation periods. These patterns are consistent with the observations in [2].

1.4 Customer liquidity provision: Empirical results

To streamline my discussion, recall that my notation for the structural shocks is the following: ϵ_t^H is the shock to hedgers, ϵ_t^L is the shock to customer liquidity providers, and ϵ_t^D is the shock to dealers. A positive shock represents a decrease in the risk-taking willingness of the corresponding party.

This section implements the SVAR, introduces the estimates of $\hat{\mathbf{B}}$, and provides economic interpretations. Then, I construct my measures of liquidity from the SVAR, which allows me to address the following questions:

- How do my measures of liquidity, especially the measure of customer liquidity provision, behave during significant historic events?
- What are the economic connections between my measures of liquidity and the

prices of corporate bonds?

These questions are asked in light of the fact that dealers are subject to post-crisis regulations and the demand for liquidity provided by customers is heightened. As I show, my approach is supportive of the view that dealers have started to improve their matchmaking technology in the hope of attracting more customer liquidity providers (actions that improve social welfare). Finally, I use my empirical setting to explore the asset pricing implications of my new dealer and customer liquidity measures.

1.4.1 Estimates of $\hat{\mathbf{B}}$ and economic interpretation

The SVAR in equation (1.2) is estimated using the three series presented in Figure 1.1 for IG bonds and HY bonds, respectively. Selected based on the SBIC, the optimal number of lagged terms is one for IG bonds and two for HY bonds. Table 1.3 reports the estimated $\hat{\mathbf{B}}$s. The sign of each estimated coefficient is consistent with my hypotheses in Section 1.2.1.

Please insert Table 1.3 here

First, $\text{Spread}_{\text{t}}^{\text{RPT}}$ are only affected by ϵ_t^H via b_1. The estimates of $\hat{b}_1$ are 0.23 and 0.14 for IG and HY bonds, suggesting that a positive ϵ_t^H increases $\text{Spread}_{\text{t}}^{\text{RPT}}$. They are both statistically significant at the 1% level. Note that $\text{Spread}_{\text{t}}^{\text{RPT}}$ enters into the SVAR in the logarithm format. The magnitudes of 23% and 14% movements of $\text{Spread}_{\text{t}}^{\text{RPT}}$ for a standard deviation change in ϵ_t^H are economically relevant.

Second, the impacts of ϵ_t^H, ϵ_t^L, and ϵ_t^D on $\mathrm{Spread}_\mathrm{t}^\mathrm{IVT}$ are not restricted in the SVAR. The actual impacts are reflected in the second row ($\left[\hat{b}_{2,1},\ \hat{b}_2,\ \hat{b}_{2,3}\right]$) of $\hat{\mathbf{B}}$. For IG bonds, the corresponding estimated coefficients are [0.08, 0.01, 0.12]. $\hat{b}_{2,1}$ and $\hat{b}_{2,3}$ are statistically significant at the $p = 1\%$ level, whereas $\hat{b}_2$ is insignificant. One standard deviation of a positive ϵ_t^H (ϵ_t^D) will increase $\mathrm{Spread}_\mathrm{t}^\mathrm{IVT}$ by 8% (12%).

In contrast, for HY bonds, the corresponding estimated coefficients are [0.05, 0.05, 0.09]. All the coefficients are statistically significant at the $p = 1\%$ level. Interpreting the estimates in the HY market, a positive shock to the risk-taking willingness for any of the three market participants will increase $\mathrm{Spread}_\mathrm{t}^\mathrm{IVT}$. Among the three structural shocks, ϵ_t^d exhibits the highest impact. A one standard deviation of increase in ϵ_t^d leads to 9% increase in $\mathrm{Spread}_\mathrm{t}^\mathrm{IVT}$, while ϵ_t^H and ϵ_t^l have a marginal impact of 5% on $\mathrm{Spread}_\mathrm{t}^\mathrm{IVT}$.

Finally, the third rows in $\hat{\mathbf{B}}$s demonstrate the impact of structural shocks on the fraction of riskless principal trading volume (i.e., $\%\mathrm{RPT}_\mathrm{t}$). For the IG bond market, the estimates are [0.31, −1.26, 0.95]. My estimates imply that ϵ_t^L and ϵ_t^d are the important drivers of $\%\mathrm{RPT}_\mathrm{t}$.

Building further on the economic insights, when customer liquidity providers receive a one standard deviation positive shock, they are expected to provide less liquidity; that is, $\%\mathrm{RPT}_\mathrm{t}$ decreases by -1.26%. The impact is sizable, considering the average of $\%\mathrm{RPT}_\mathrm{t}$ is 8.67% over the sample period. Providing complementary evidence, when dealers receive a one standard deviation positive shock, $\%\mathrm{RPT}_\mathrm{t}$ increases by 0.95%.

This finding is consistent with the literature that shows that dealers conduct more riskless principal transactions both during the financial crisis and over the post-regulation periods (e.g., [2]). Similar intuition can be garnered when considering the impact of structural shocks on $\%\mathrm{RPT_t}$ for the HY market. Overall, my estimates of $\hat{\mathbf{B}}$s are aligned with theory and economic intuition.

1.4.2 What drives $\mathbf{Spread_t^{IVT}}$?

Recall that the second row in $\hat{\mathbf{B}}$ reflects the relationship between the three structural shocks and $\mathrm{Spread_t^{IVT}}$. $\mathrm{Spread_t^{IVT}}$ is an important object because it approximates the cost for immediacy in an OTC market, i.e., the marginal cost that hedgers pay, compared with the interdealer price. Since the structural shocks are orthogonal to each other, the fluctuation in $\mathrm{Spread_t^{IVT}}$ can be attributed to the variation of each shock. The impact of customer liquidity provisions on $\mathrm{Spread_t^{IVT}}$ can be determined in this setting.

In IG markets, my results suggest that most of the fluctuation (about 67%)[11] in $\mathrm{Spread_t^{IVT}}$ is driven by ϵ_t^d, while the remaining 33% results from ϵ_t^H. In contrast, the total fluctuation of $\mathrm{Spread_t^{IVT}}$ in HY markets can be decomposed into shocks to the three market participants such that 20% of the fluctuation is due to ϵ_t^H, 21% results from ϵ_t^L, and the rest 59% is explained by ϵ_t^d.

The markets of IG and HY bonds exhibit some commonalities. The majority of the fluctuations of $\mathrm{Spread_t^{IVT}}$ for both markets is driven by ϵ_t^d. The direction of such an impact is consistent with the view that decreased dealers' risk-taking

[11] Under the assumption that structural shocks are orthogonal, one can attribute the total variation of the unexpected change in $\mathbf{y_t}$ into three components. For this specific case, 67% is computed by $\frac{\hat{b}_{2,3}^2}{\hat{b}_{2,1}^2+\hat{b}_2^2+\hat{b}_{2,3}^2}$

willingness leads to higher $\text{Spread}_t^{\text{IVT}}$ ([19]). I quantify the fraction of fluctuations in $\text{Spread}_t^{\text{IVT}}$ with regard to this type of shock.

The two markets are distinct, when considering the impact of ϵ_t^L on $\text{Spread}_t^{\text{IVT}}$. Only in HY markets, a positive ϵ_t^L significantly increases $\text{Spread}_t^{\text{IVT}}$, and ϵ_t^L explains a meaningful amount of fluctuations in $\text{Spread}_t^{\text{IVT}}$. The results answer the question that shocks to customer liquidity provision affect the $\text{Spread}_t^{\text{IVT}}$. However, it suggests that the magnitude of impact differs between IG and HY bond sectors.

In Section 1.5, my theoretical model provides an explanation to this divergence in results in the two different bond sectors. In the model, I show that the marginal impact of ϵ_t^L on $\text{Spread}_t^{\text{IVT}}$ is nonlinear. It depends on the relative amount of liquidity provision between dealers and customer liquidity providers. Change in customer liquidity provision exhibits a stronger marginal impact on $\text{Spread}_t^{\text{IVT}}$, when there is relatively more customer liquidity provision on the market.

The intuition is the following. Dealers charge $\text{Spread}_t^{\text{IVT}}$ for compensation for taking inventory risk. The $\text{Spread}_t^{\text{IVT}}$ charged by a dealer is linear in the ultimate absolute inventory level, which is nonlinear in the amount of customer liquidity provision. Interpreting this nonlinearity in inventory level is the key to understanding the explanation of the nonlinear impact of customer liquidity provision on $\text{Spread}_t^{\text{IVT}}$.

If the aggregate level of customer liquidity provision is high, dealers' relative inventory level is low. When a shock hits customer liquidity providers, dealers quickly step in to provide liquidity. Dealers' inventory level moves significantly, as does $\text{Spread}_t^{\text{IVT}}$. In contrast, if dealers are responsible for most of the liquidity provision, incremental changes in customer liquidity provision result in small changes in dealers' inventory level and $\text{Spread}_t^{\text{IVT}}$. Thus, the relative level of liquidity provision, between dealers and customers, plays a large role in the relative impact of liquidity shocks to dealers versus customers in a certain bond market sector.

Empirically, the marginal impact of shocks to customer liquidity on $\text{Spread}_t^{\text{IVT}}$ is much stronger in HY markets. In other words, dealers rely on customer liquidity provision to a much greater degree in HY markets. As shown in Figure 1.1, the percentage of riskless principal transaction volume ($\%\text{RPT}_t$) is consistently higher in the HY market (15%–33%) than that in IG markets (3%–17%). Evidence supporting the empirical fact is documented in [30], who find that the tendency of overnight holdings for dealers are lower for HY bonds.

1.4.3 Evolution of historical structural shocks

Analyzing realized structural shocks yields insights on the conditions of corporate bond markets. The SVAR yields two separated liquidity measures for dealers and customer liquidity providers. To demonstrate their usefulness, this subsection discusses the linkages between the structural shocks and the most known historic events that affect corporate bond markets.

Figure 1.2 presents the backward accumulated structural shocks using a rolling window of twelve months.[12] Facilitating detected patterns in Figure 1.2, Table 1.4 presents statistical evidence, in which inferences are drawn based on a bootstrap

[12]For example, the data point labeled 2009 captures the cumulative sum of $\{\hat{\epsilon}_t\}$ from Jan 2008 to Dec 2008.

method described in Appendix 1.7.7.

Please insert Table 1.4 here

Please insert Figure 1.2 here

I divide the sample into the following subperiods: pre-crisis period (Jan 2006 to Jun 2007), the 2008 financial crisis (Jul 2007 to Apr 2009), the post crisis period (May 2009 to May 2012), the Basel 2.5 period (Jun 2012 to Jun 2013), the Basel III period (Jul 2013 to Mar 2014), and the Volcker Rule period (post Apr 2014). The selected dates follow [10] and [9]. I also include a period of the collapse of Lehman Brothers (Aug 2008 to Dec 2008) which is nested in the 2008 financial crisis period.

A. During the pre-crisis period, negative ϵ_t^ds are observed (i.e., expansion of dealer liquidity provision). It reflects the aggressive market-making strategies for bond dealers. Primary dealers accumulated up to \$225B inventory in corporate assets at the end of Aug 2007, compared with a net position of \$155B in Dec 2005.[13] [1] characterize this period with implicit "too big to fail" guarantees and low costs for market-making, which may not be socially optimal.

B. In the 2008 financial crisis, no significant ϵ_t^ds are observed during the overall period of Aug 2007 to Dec 2008. The evidence shows that dealers were not reluctant to provide liquidity, at least during the initial stage of the 2008 financial crisis.

Instead, positive ϵ_t^Hs and negative ϵ_t^Ls are observed. The evidence suggests

[13] Data of the >1 year holdings are from the Federal Reserve Bank of New York: `https://www.newyorkfed.org/markets/gsds/search`.

that some investors are eager to sell, whereas other customers are willing to provide liquidity, possibly due to capitals from other markets that are riskier than corporate bond markets. These two shocks drive the increase of bid-ask spreads and the amount of riskless principal transactions during late 2007 and early 2008.

The findings of increased customer liquidity provision during the crisis are called "flight-to-safety." [7] document this phenomenon for IG bonds only. To interpret the distinction of customer liquidity provision between IG and HY markets, it requires further investigation into the months following the collapse of Lehman Brothers.

C. It was not until the collapse of Lehman Brothers that dealers become extremely cautious in liquidity provision. The nested dummy of Lehman Brothers is significantly positive for dealers. The aggregate net position for primary dealers on corporate securities with time-to-maturity over one year dropped from $155B in May 2008 to $73B in December 2008. One explanation is that increased haircut in repo transactions and the resulting run on repo limited the bond dealers' ability to market the market.

At the peak of the crisis, ϵ_t^L for HY turns significantly positive, while ϵ_t^L for IG does not. The evidence shows that although customer liquidity providers tend to participate in the riskier HY markets at the beginning of the crisis, they quickly flee away and shift toward the higher quality IG markets.

Overall, the results suggest that customers facilitate liquidity provision during the 2008 financial crisis, complementing the findings in [31] that dealers are unwilling to expand their inventory, especially for those bonds that clients were selling the most at that moment.

D. In the periods of Basel Accords, positive ϵ_t^ds for HY bonds are detected in Figure 1.2 and Table 1.4. The effect is concentrated in the period of the Basel 2.5. There are two plausible explanations. First, banks start to comply with Basel III before the actual implementation date. Second, this finding agrees with the survey evidence in [32]: *For the corporate bond... ("Basel 2.5") were seen to have had the largest impact ... the Basel III requirements, in turn, was expected to have only a minor impact.*

Note that significant impacts are only detected for HY bonds. Basel 2.5 and Basel III aim at mitigating the risk of banks' portfolios, and thus lead to a stronger effect in the riskier HY markets. For example, Basel 2.5 introduces an incremental risk capital charge, which accounts for default and migration risk ([33]). Basel III requires banks to maintain a minimum liquidity coverage ratio, which differentially treats IG and HY bonds.[14]

E. The Volcker Rule, which went into effect on April 1, 2014, requires banks to report inventory turnover as well as other statistics to ensure that banks do not engage in proprietary trading. The initial purpose of the Volcker Rule is to prevent bank holding companies from risky activities that take regulatory advantages, such as the FDIC insurance and "too big to fail" guarantees. However, [19] observes that the nature of market-making activities is a form of proprietary trading and predicts that the implementation of the Volcker Rule could potentially disincentivize liquidity

[14] The liquidity coverage ratio is calculated by dividing the high quality liquid asset amount (HQLA) by the total net cash flows of the bank over a 30-day stress period. All HY bonds are not considered HQLA, while IG corporate debt securities issued by non-financial sector corporations are considered as Level 2 HQLA. See [34] for more details.

provision by bond dealers. Consistent with this prediction, it is shown that ϵ_t^ds in my SVAR are positive; that is, dealers receive positive shocks and become less willing to make the market.

The aftermath of the stringent regulations is also a period when dealers invest more in matchmaking technology for more customer liquidity providers. Many electronic trading systems emerge in corporate bond markets, mostly using a method called "requests for quotations" (RFQs). Those platforms aim to attract more customer liquidity providers into corporate bond markets. My contribution is to explicitly display this increase in the supply of customer liquidity provision.

In IG markets, the measure ϵ_t^L exhibits a sharp decrease since July 2014. Note that the plot presents a cumulative backward 12-month shock, and that data point captures the raising of trading volume in the electronic platform in the second half of 2013 and early 2014. Electronic trading of IG corporate bonds constitutes 8% of total market volume in 2013 and 16% in 2014 [35, p. 5]. As the market share of electronic trading platforms grows, ϵ_t^L keeps at a negative level (i.e., plentiful customer liquidity provision) until the end of my sample.

In contrast, movements of ϵ_t^L in HY markets are mild and insignificant. In Panel C of Figure 1.2, ϵ_t^L in HY markets does not increase until late 2015. This evidence is consistent with the documentation in [36] that RFQs in MarketAxess are more suitable for large size IG bonds.[15]

In short, this subsection discusses the movements of the derived measures of

[15] MarketAxess reports a 85% market share in electronically facilitated corporate bond trading ([1]. MarketAxess is a dominating electronic platform that accounts for 20% (12%) of total TRACE trading volume for IG bonds in Dec 2015 (Jan 2014).

liquidity for both customers and dealers over the last decade. The behavior of these generated series are consistent with the expected directions in the literature. My SVAR could be used to measure future market conditions for regulators when data is available. For example, what is the consequence of the undergoing back push that proposes to loosen the restrictions around the Volcker Rule?

1.4.4 How do my measures relate to stress indicators?

It has been well documented that illiquidity of corporate bonds comoves with the aggregate market conditions. In this subsection, I investigate the relation between the stress indicators and my measures of liquidity. I show that ϵ_t^H and ϵ_t^d are positively correlated with the stress indicators, whereas ϵ_t^L is not, due to other driving forces such as the "flight-to-quality" phenomenon and the improvement of matchmaking technology.

[6] approximate the aggregate market condition using the Chicago Board Options Exchange Volatility Index (VIX).[16] In addition to the VIX, I relate my measures to a volatility index for 30-year US Treasury bonds, called TVIX. Compared to the VIX, the TVIX is more specific to the corporate bond markets. The volatility of the Treasury bond prices has a larger exposure to the volatility of interest rates ([37]), which is also an important factor in corporate bond markets. Furthermore, most primary dealers in the Treasury bond market also make markets in corporate bond markets. The risk-taking attitudes for dealers in the two markets are tied closely.

[16] Also see [31], [37], and [26].

Specifically, I consider the price of a volatility contract whose payoff is the squared return of holding a fully collateralized 30-year Treasury bond futures for one period (approximately a month) under the risk-neutral measure,

$$\mathrm{TVIX}_t^2 = R_{f,t+1}^{-1}\mathbb{E}_t^Q\left(\log\left\{\frac{F_{t+1}}{F_t}\right\}^2\right), \tag{1.6}$$

where F_t is the 30-year Treasury bond futures price at period t. The derivation of the TVIX exploits the spanning engine in [38] and [39]. A detailed description is available in Appendix 1.7.6.

Please insert Figure 1.3 here

Figure 1.3 plots the time series of my TVIX; also plotted is the VIX at the corresponding dates. Although the correlation between the TVIX and the VIX is high (0.81), the two indexes reflect different stress events. For instance, on Aug 24, 2015, world stock markets plunged due to the drop in commodity prices. This global equity market disaster, which caused a spike in the VIX, did not propagate any fear in fixed income markets. In contrast, the TVIX increased from 6.6 in Jun 2014 to 14.3 in Feb 2015, while the VIX did not change much in the same period. This period corresponds to Federal Reserve's decision to halt the quantitative easing (QE) purchases.[17] Hence, it is necessary to investigate the relation between the shocks from SVAR and the two indexes, respectively.

[17] The amount of outright holdings of Treasury bills and mortgage-backed securities approached the peak of $4.2 trillion around October 2014 and remained stagnant afterwards. See data from `https://fred.stlouisfed.org/graph/?g=qHw`.

I investigate the relation between $\hat{\epsilon}_{\mathbf{t}}$ and the volatility indexes. Table 1.5 presents the results from the following regression:

$$\mathrm{y}_{t+1} = \text{Constant} + \beta_1 \sum_{s=t-l-1}^{t} \hat{\epsilon}^{H}_{s-1\to s} + \beta_2 \sum_{s=t-l-1}^{t} \hat{\epsilon}^{L}_{t-1\to t} + \beta_3 \sum_{s=t-l-1}^{t} \hat{\epsilon}^{d}_{t-1\to t} + u_t, \quad (1.7)$$

where y_{t+1} is $\mathrm{VIX_{t+1}}$ or $\mathrm{TVIX_{t+1}}$ and $l = 1$, 3, 6, or 12 is the length of the backward rolling window.[18]

Please insert Table 1.5 here

Overall, evidence suggests that ϵ^d_t and ϵ^H_t are significantly correlated with VIX, whereas ϵ^L_t is not. My measure of customer liquidity provision is orthogonal to the commonly used stress indicators. It adds incremental value for the purpose of monitoring market conditions.

Panel A in Table 1.5 exhibits the results for IG bonds when the regressand is the VIX. Jointly, the three ϵ_ts are able to explain a significant portion of the variation in VIX. The R^2_{adj} is 0.31 (0.40) for a rolling window of three (six) months. When $l = 3$, one standard deviation increase in $\sum \hat{\epsilon}^H$ ($\sum \hat{\epsilon}^d$) is related to an increase of 3.17 (3.61) in VIX. In contrast, $\sum \hat{\epsilon}^L$ always fails to provide incremental explanatory power in the fluctuations of the VIX.

Evidence from Panel B in Table 1.5, where the regressand is the TVIX, reinforces these findings. The proposed measure of customer liquidity provision is not correlated with the TVIX as well. One distinction is that the TVIX exhibits a

[18]Time subscripts are properly aligned so that y_{t+1} is the next available $\mathrm{VIX_{t+1}}$ or $\mathrm{TVIX_{t+1}}$ following the $\{\hat{\epsilon}_{\mathbf{t}}\}$ up to time t. For instance, when $l = 3$, I use $\{\hat{\epsilon}_{\mathbf{t}}\}$ for the month of May, June, and July 2015 to predict the VIX on August 31, 2015 and the TVIX on August 24, 2015.

stronger relation to $\sum \hat{\epsilon}^d$. When $l = 6$, the coefficients of $\sum \hat{\epsilon}^d$ and $\sum \hat{\epsilon}^H$ for the TVIX (VIX) are 2.43 (3.91) and 1.68 (4.62). One explanation is that government policies impact dealers' market-making behavior in both corporate bond markets and Treasury markets. Hence, the stress indicator in the Treasury market is more relevant to my measure of shocks to dealers.

Panels C and D provide results for the HY market. Patterns are mostly consistent with those in the IG market. Closing this section, evidence suggests that the derived shocks to dealers and hedgers are relevant to the stress indicators. My measure of customer liquidity provision, however, is orthogonal to the stress indicators. In the next section, I explore the usefulness of the information contained in my measure of customer liquidity provision.

1.4.5 Liquidities and corporate bond yields

The willingness of market participants to provide liquidity affects not only market microstructure measures, but also equilibrium asset prices ([40], [41], and [26]). In this subsection, I examine the asset pricing implications of my liquidity measures of both customer liquidity providers and dealers. The variable of interest is the yield spread of corporate bonds, denoted as $\mathrm{YS}_{i,t}$, which is the difference between the yield of a corporate bond i and the yield of a corresponding Treasury bond of the same maturity. My prediction is that a decrease in risk-taking willingness (or an increase in my measure of structural shocks) could potentially decrease equilibrium asset prices and increase yield spreads.

To compute $\text{YS}_{i,t}$, I use all transactions of a Cusip i that are executed during the second half of the month t.[19] To relieve market microstructure noises, I focus on transactions in which the par value is greater than or equal to $100K and the maturity is greater than one month. The computation of yields is extremely sensitive to small change in prices when a bond is close to maturity.[20]

My analysis exploits the following panel regression:

$$\Delta \text{YS}_{i,t} = \alpha_i + \beta_1 \text{Customer Liquidity}_t + \beta_2 \text{Dealer}_t + \beta_3 \text{Hedger}_t + \beta_{\mathbf{4}} \mathbf{\Delta X_{i,t}} + \beta_{\mathbf{5}} \mathbf{\Delta X_t} + u_{i,t}, \tag{1.8}$$

where $\text{Customer Liquidity}_\text{t}$, Dealer_t, and Hedger_t are my SVAR-based measures of the willingness of liquidity provision for customer liquidity providers, dealers and hedgers. In my analysis, I use the three-month rolling sum of structural shocks and standardize the three-month measures to zero mean and unit variance for better interpretation of the economic consequences. $\Delta \text{YS}_{i,t}$ is measured in basis points (bps). α_i is the Cusip fixed effect. $\mathbf{\Delta X_{i,t}}$ and $\mathbf{\Delta X_t}$ are systematic and firm level factors that explain yield spread changes, following [37]. $\mathbf{\Delta X_{i,t}}$ includes the market return of the issuers' common equity.[21] $\mathbf{\Delta X_t}$ includes the change in risk-free rate (10-year Treasure), the squared change in risk-free rate, the change in the slope of the yield curve (10-year Treasure minus 2-year Treasure), the change in the VIX,

[19] The results are similar if using all transactions during the entire month, or the last ten transactions in the month.

[20] Also excluded are erroneous transactions in which the reported yield is 200% higher or 80% lower than both the preceding and subsequent transactions for this Cusip. For instance, the reported transaction yields (in percentages) are 5, 30, 5.1 for the first, second and third transactions. In this case, the second transaction will be excluded.

[21] I use the change in the leverage of the issuers produces similar results. Most of the fluctuations in the change in leverage are due to the change in the market value of equities.

the S&P 500 return, and the jump factor that captures the tail risks.

The aim is to evaluate the impact of my SVAR-based measures of dealers and customer liquidity providers on changes in yield spreads. Table 1.6 reports the results for bonds in different rating categories. Statistical inferences for the results are double-clustered at Cusip and year-month level. Clustering at year-month level is crucial because it takes care of the plausible comovements of the residuals between different bonds in the same month; that is, $cov(u_{i,t}, u_{j,t}) \neq 0$ for bonds i and j. The literature has shown that a large fraction of the change in yield spreads of corporate bonds is driven by a single component ([26]).

Please insert Table 1.6 here

Consistent with my prediction, the coefficients in front of the structural shocks are all positive whenever they are statistically significant. However, the changes in yield spreads for bonds of different credit ratings differentially respond to the measures of liquidity, emphasizing the economic importance of customer liquidity provision, especially for high quality corporate bonds.

For high quality bonds, I find that the coefficient in front of the customer liquidity measure is 2.30 (2.15), with a p-value of 0.055 (0.065) for Aaa-rated (Aa-rated) bonds. One standard deviation increase in customer liquidity provision decreases the yield spreads by approximately 2.3bps. The economic significance is meaningful, considering that the median of the absolute change in yield spread for both Aaa and Aa-rated bonds is 11bps.

For lower quality bonds (Baa-rated and below), the impact of customer liquidity measure is still positive but lacks statistical power. In contrast, for A-rated bonds, my measures of both dealers and customer liquidity providers are able to explain the movements of yield spreads. When considering A-rated bonds and high-yield bonds, I find that both the measure of dealers and the shock to hedgers are able to explain yield spread changes.

The coefficients in front of the control variables are all consistent with [26]. A negative stock return ($\mathrm{R}_{i,t}$) leads to a higher probability of default and increases the yield spreads, especially for lower quality bonds. Increases in the risk-free rate ($\Delta\mathrm{RF}_t$) also lower the yield spreads. The static effect of increased spot rate is to increase the drift term of the stochastic process for the firm value and thus decreases the probability of default ([42]). Other variables (RM_t, $\Delta\mathrm{VIX}_t$, $\Delta\mathrm{Jump}_t$) are not statistically significant, thanks to the robust standard errors.

Table 1.7 shows that the observed patterns are the same for bonds of different maturity groups. The explanatory power of the model in equation (1.8) is stronger for bonds with longer maturities. For Aa-rated bonds, the Adjusted-R^2 is 54% for bonds with maturity over eight years and 18% for bonds with maturity less than eight years. Yield spreads for shorter maturity bonds are more sensitive to market microstructure noises, so that the changes in yield spreads are harder to explain.

Please insert Table 1.7 here

One standard deviation of my measure of customer liquidity will impact the yield spreads of long-term IG bonds by 2.5bps. Considering a bond with a duration of 10 years, a change of 2.5bps implies a price movement of 25bps, which is sizable for a transaction over $100K. The magnitude is similar across bonds of lower ratings. In contrast, one standard deviation of my measure of dealers will impact the price of a 10-year Aaa-rated (Baa-rated) bond by 14bps (64bps). Note that the focus is the monthly changes in yield spreads. Events that introduce shocks over several months will lead to additive effects on the prices of corporate bonds.

Overall, the messages conveyed in the results are that high quality (Aaa and Aa) bonds are more exposed to my measure of customer liquidity, whereas lower quality bonds (below Aa) are mostly subject to liquidity provision of dealers. The results are consistent with the findings in [14]. They find that flight-to-quality is confined to AAA-rated bonds during the subprime crisis. I explicitly attribute this effect to customer liquidity providers, rather than to dealers.

This differential impact of shocks to dealers versus customer liquidity providers translates into different exposures of yield spreads in response to recent regulations (i.e., decrease in liquidity provision from dealers) and technological developments (i.e., increase in liquidity provision from buy-side customers). Regulations imposed on bank-afflicted dealers and recent launches of ATSs could yield disproportional impact on bonds of different qualities. ATSs tend to include more frequently traded bonds (potentially high quality), and further enhance the liquidity of assets. In contrast, dealers who are restricted from market-making activities become even more reluctant to trade illiquid assets.

Hence, the impact of recent regulations on the corporate bond market could be misunderstood if one only examines the aggregate conditions of corporate markets because data on available transactions are biased even more toward liquid assets. The mixed consequence of stringent regulations and technological improvements could point to an outcome that safe bonds become more liquid, whereas the markets are even thinner for those bonds that are riskier and, likely, already illiquid.

1.4.6 An alternative method: sign restrictions

The full identification of the SVAR relies on the equality restrictions specified in equation (1.4). The restrictions are motivated by economic intuitions and my theoretical model, but they cannot be empirically tested.

To examine the robustness of my full identification strategy, I consider an alternative approach that imposes only sign restrictions, using the "pure-sign-restriction approach" in [18]. The advantage is that theories and economic intuitions can provide clearer guidance for sign restrictions, compared with zero restrictions. For example, when considering the impact of shocks to customer liquidity providers, imposed restrictions are that an increase in ϵ_t^L (i) decreases $\%\mathrm{RPT}_\mathrm{t}$ (i.e., $b_{3,2} \leq 0$) and (ii) weakly increases $\log\left(\mathrm{Spread}_\mathrm{t}^\mathrm{IVT}\right)$ (i.e., $b_2 \geq 0$).

The limitation is that the procedure only achieves partial identification for the shock of interest and remains agnostic about the impact of other structural shocks. In the rest of this subsection, I briefly discuss this "agnostic" method,[22] and show that the obtained shocks are similar to my main results.

Recall that the object of interest in the SVAR, as stated in equation (1.2), is

[22] See [18] and [43] for a more detailed description.

$\mathbf{B}$, which satisfies $\mathbf{B}\mathbb{E}\left[\epsilon_{\mathbf{t}}\epsilon_{\mathbf{t}}'\right]\mathbf{B}' = \Sigma_u$, where Σ_u is the covariance matrix from the associated reduced form VAR. The idea is to find an impulse vector b, which is a column element in $\mathbf{B}$, such that the resulting impulse response functions, up to certain horizon, obey all prespecified inequality restrictions. Specifically, the method can be summarized in the following steps.

1. Specify restrictions on the impulse response functions (for example, my shock of interest has positive impact on y_1 and negative impact on y_2 up to horizon 3).
2. Run the SVAR using a Cholesky decomposition and denote the result as $\tilde{\mathbf{B}}$, such that $\tilde{\mathbf{B}}\tilde{\mathbf{B}}' = \Sigma_u$. This Cholesky decomposition has no economic meaning but is used as elementary building blocks.
3. Randomly generate an impulse vector b, defined as $b = \tilde{\mathbf{B}}\beta$, where β is a vector randomly drawn from the unit sphere. Use this b to generate the impulse response functions.
4. Check whether the impulse response functions satisfy the sign restrictions in step (1). If yes, keep this b, drop otherwise.
5. Repeat steps (3) and (4) until a desired number of satisfying impulse vectors is obtained. Use each impulse vector to compute the impulse response functions and infer the structural shocks.
6. Report the median and other percentiles of the posterior distribution of the impulse response functions and structural shocks.

Bringing this method to my SVAR, I separately apply two inequality restrictions for each of the three structural shocks to be recovered. To recover the shocks to hedgers, I impose the restrictions that an increase in ϵ_t^H increases $\log\left(\text{Spread}_t^{\text{RPT}}\right)$ (i.e., $b_1 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,1} \geq 0$). To obtain the shocks to dealers, I apply the restrictions that an increase in ϵ_t^d increases $\%\text{RPT}_t$ (i.e., $b_3 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,3} \geq 0$). Finally, to derive the shocks to customer liquidity providers, I use the restrictions that an increase in ϵ_t^L decreases $\%\text{RPT}_t$ (i.e., $b_{3,2} \leq 0$) and increases $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_2 \geq 0$).

Please insert Figure 1.4 here

Figure 1.4 presents the obtained impulse response functions with sign restrictions for IG bond markets. Figure 1.5 exhibits the histograms of the initial impulse response functions ($t = 0$). The directions of these initial impulse response functions all agree with the estimates based on my main results using equality restrictions (i.e., Table 1.3), providing robust evidence that the restrictions I impose in equation (1.4) are effective.

Please insert Figure 1.5 here

ϵ_t^d and ϵ_t^L are positively associated with $\text{Spread}_t^{\text{RPT}}$ for the median estimates, although the 2.5th and 97.5th percentiles straddle zero. One explanation is that these two shocks are positively correlated with ϵ_t^H, which will affect $\text{Spread}_t^{\text{RPT}}$. Appendix 1.7.4 shows that my main SVAR is valid. The difference is that my main results recover the part of shocks that are orthogonal to ϵ_t^H, whereas the agnostic method generates full shocks. Another explanation is that, facing shocks to dealers or customer liquidity providers, bond dealers endogenously adjust their matchmaking costs, which is a determinant of $\mathrm{Spread}_t^{\mathrm{RPT}}$, for example, they put more effects to exploit their network to search for potential buyers.[23]

Please insert Figure 1.6 here

The left panel of Figure 1.6 compares the obtained monthly structural shocks from the full identification strategy versus the same shocks from the agnostic approach in IG markets. All the spots in the scatter plots are close to the diagonal line $y = x$, suggesting that the shocks from the two methods are quantitatively the same. The correlation coefficients of the shocks derived from the two methods are 0.86, 0.86, and 0.82 for ϵ_t^H, ϵ_t^L and ϵ_t^d, respectively.

Plots in the right panel exhibit the 12-month rolling sum time series for each shock. All the described features over historic events are preserved for shocks recovered based on a SVAR using inequality restrictions.

Therefore, I show that my liquidity measures of customer liquidity providers and dealers are robust under the agnostic method with the least controversial restrictions. Using the agnostic method, the analysis for HY markets in Figures 1.12 and 1.14 also yields similar results, compared to the full identification method.

Please insert Figure 1.12 here

[23] I thank Charles Calomiris for suggesting this point.

Please insert Figure 1.14 here

1.5 A model of dealer and customer liquidity supply

Frictions and searching in OTC markets render dealers market power to charge bid-ask spreads. Empirical literature also documents the importance of network structure and relationship.[24] In this section, I present a simple setting, which sheds light on my empirical treatment in this paper.

My model relates to the setup in [49], who model the behavior of a monopoly market maker who deals with customer hedgers and customer liquidity providers. [50] develop a decentralized exchange model in which privately informed institutional investors strategically trade with each other under a generic network structure. One can trade in certain clubs that consist of a subset of investors.

1.5.1 My setup

Consider K segmented markets and one tradable risky asset whose payoff is normally distributed $\tilde{V} \sim N(\bar{V}, \sigma_v^2)$. All participants have the same information regarding the probability distribution of payoff for the tradable asset. They trade the risky asset with *accessible* counterparties in the first period and the payoff realizes in the second period.

In each segmented market, there are three types of participants: a continuum of identical customer liquidity providers with mass N_L, a continuum of identical customers (called hedgers) with N_H who suffer an exogenous shock and demand

[24] See [44], [45], [46], [31], [47], [48], among others.

liquidity, and $N_M = 1$ designated market maker. Customer liquidity providers and hedgers can only trade with this market maker. This assumption attempts to capture the real-world customer-dealer relationship in which a customer may only obtain favorable prices from a set of dealers with whom she has a strong business tie, especially in the short run. For simplicity, it is assumed that customers can trade only with one market maker.

Dealers and customer liquidity providers are endowed with zero unit of the asset before trading occurs. Assume CARA utility functions for all participants and the risk aversions for market makers, hedgers, and customer liquidity providers are δ_d , δ_H and δ_L, respectively. Figure 1.7 exhibits an example of the market structure with three market makers.

Hedgers in the k-th market are subject to a liquidity shock by receiving X_k units of another non-tradable risky asset,[25] which has a per-unit payoff of $\tilde{L} \sim N\left(0, \sigma_l^2\right)$, and the payoff has a covariance of $\sigma_{lv} > 0$ (w.l.o.g) with the tradable asset.

All K market makers can trade in a centralized system by submitting their demand schedules. Hence, hedgers can indirectly hedge across markets via the local market maker and the interdealer network. The segmented dealer-customer markets and the interdealer market are cleared simultaneously.

The k-th market maker observes atomistic customer demand schedule, denoted as $\Theta_{L,k}\left(P_k^L\right)$ and $\Theta_{H,k}\left(X_k, P_k^H\right)$, and sets the optimal quantities (prices). These

[25]The underlying data-generating process for $\{X_k\}_{k=1}^K$ is not crucial. Only the actual realizations of $\{X_k\}_{k=1}^K$ matter. It is equivalent if hedgers are assumed to receive heterogeneous shocks and the market maker cannot price discriminate these atomistic hedgers.

demand schedules imply that market makers can trade with customers at different prices P_k^L and P_k^H. Interpreting this setup, the market maker knows the identity of customers or she can infer the type of customers by observing their submitted demand schedules. In practice, dealers can differentiate customers by whether the customer or the dealer initiates the transaction.[26]

The market maker sets P_k^L and P_k^H to clear the k-th market, resulting in an inventory level $\beta_k - \alpha_k$, where $\alpha_k\left(P_k^L, P_k^H\right) = \sum_{i=H,L} N_i \Theta_{i,k}^+$ and $\beta_k\left(P_k^L, P_k^H\right) = \sum_{i=H,L} N_i \Theta_{i,k}^-$, denoting the aggregate amount, which the market maker transacts with the hedgers and the liquidity providers in this market. If both β_k and α_k are strictly positive, the market maker conducts a typical customer-dealer-customer riskless principal transaction for the volume $\min\left(\beta_k, \alpha_k\right)$, generating a bid-ask spread for riskless principal transactions.

The market maker can hold a fraction of the resulted inventory for one period and trade the rest in the interdealer market to exchange excessive inventory or provide liquidity to other market makers. Assume symmetricity and conjecture the following demand curve for each market maker:

$$M_k\left(X_k, P^D\right) = \nu - \eta X_k - \gamma P^D, \tag{1.9}$$

where P^D is the price in the interdealer market. The demand for the k-th market

[26]Since the total number of outstanding issues is large, uninformed investors who do not hold the particular issue pay little attention to this issue until being contacted by their dealers. For instance, [51] consider this type of customer liquidity as latent: *"if a buy order comes in to a dealer ... the dealer could 'work the order' by contacting customers to ... convince someone to sell her the bond."*

maker is negatively related to the interdealer price P^D and the local hedgers' shock X_k.

Apply the market clearing condition in the interdealer market $\sum_{k=1}^{K} M_k = 0$, or

$$M_k \left(X_k, P^D\right) + \sum_{l \neq k} \left(\nu - \eta X_l - \gamma P^D\right) = 0. \tag{1.10}$$

The residual supply schedule is $P^D (m_k) = \frac{\nu}{\gamma} - \frac{\eta}{\gamma} \bar{X}_{-k} + \frac{1}{(K-1)\gamma} m_k$, where $\bar{X}_{-k} = \frac{1}{K-1} \sum_{l \neq k} X_l$ is the average liquidity shock from other segmented markets.

In each market, denote the reservation prices for CARA hedgers and customer liquidity providers as

$$P^R_{H,k} = \bar{V} - \delta_H \sigma_{lv} X_k \qquad \text{and} \qquad P^R_{L,k} = P^R_L = \bar{V}. \tag{1.11}$$

Individual customers maximize their utilities as prices (P^H_k, P^L_k) are given. Their demand schedules are denoted as

$$\Theta_{H,k} = \frac{1}{\delta_H \sigma_v^2} \left(P^{R,H}_k - P^H_k\right) \qquad \text{and} \qquad \Theta_{L,k} = \frac{1}{\delta_L \sigma_v^2} \left(P^{R,L} - P^L_k\right). \tag{1.12}$$

The market maker's problem is

$$\max_{P^L_k, P^H_k, m_k} E\left[-e^{-\delta_d \tilde{w}_k}\right], \tag{1.13}$$

subject to

$$
\begin{aligned}
\tilde{w}_k &= \underbrace{N_H \Theta_{H,k} P_k^H + N_L \Theta_{L,k} P_k^L}_{\text{Proceeds from sales to customers}} \underbrace{- m_k \left(\frac{\nu}{\gamma} - \frac{\eta}{\gamma} X_{-k} + \frac{1}{(K-1)\gamma} m_k \right)}_{\text{Payments for purchases from other dealers}} \\
&\quad + \underbrace{\left(-N_H \Theta_{H,k} - N_L \Theta_{L,k} + m_k\right) \tilde{V}}_{\text{Realized portfolio value}}. \qquad (1.14)
\end{aligned}
$$

1.5.2 Equilibrium

Denote the quantity $\mathcal{D}^{-1} = \left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)$. $\mathcal{D}^{-1}$ captures the total risk-bearing capacity in the economy. The equilibrium transaction prices are as follows:

$$
\begin{aligned}
P_k^L &= \bar{V} - \frac{1}{2}\mathcal{D}\sigma_{lv} N_H \left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K} X_k \right), & (1.15) \\
P_k^H &= \bar{V} - \frac{1}{2}\mathcal{D}\sigma_{lv} N_H \left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K} X_k \right) - \frac{\sigma_{lv}\delta_H}{2} X_k, \quad \text{and} & (1.16) \\
P^D &= \bar{V} - \mathcal{D}\sigma_{lv} N_H \bar{X}, & (1.17)
\end{aligned}
$$

where $\bar{X} = \frac{1}{K}\sum X_k$ is the average liquidity shock for all K markets. The equilibrium trading quantities are as follows:

$$
\begin{aligned}
\Theta_{L,k} &= \frac{1}{2}\frac{\mathcal{D}}{\delta_L}\frac{\sigma_{lv} N_H}{\sigma_v^2} \left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K} X_k \right), & (1.18) \\
\Theta_{H,k} &= \frac{1}{2}\frac{\mathcal{D}}{\delta_H}\frac{\sigma_{lv} N_H}{\sigma_v^2} \left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K} X_k \right) - \frac{\sigma_{lv}}{2\sigma_v^2} X_k, \quad \text{and} & (1.19) \\
m_k &= \frac{1}{2}\frac{K-2}{K}\frac{\sigma_{lv} N_H}{\sigma_v^2} \left(\bar{X}_{-k} - X_k \right). & (1.20)
\end{aligned}
$$

The inventory level of the k-th dealer is

$$-\underbrace{N_H \Theta_{H,k}}_{\substack{\text{Units to}\\ \text{hedgers}}} - \underbrace{N_I \Theta_{I,k}}_{\substack{\text{Units to}\\ \text{liquidity providers}}} + \underbrace{m_k}_{\substack{\text{Units from}\\ \text{interdealer market}}} = \frac{\mathcal{D}}{\delta_d} \frac{\sigma_{lv} N_H}{\sigma_v^2} \left(\frac{K-2}{K} \bar{X}_{-k} + \frac{2}{K} X_k \right) \tag{1.21}$$

Appendix 1.7.1 provides the proof.

Equation (1.15) presents the equilibrium price that a market maker charges the customer liquidity providers. This price is equal to the reservation value of customer liquidity providers $\bar{V}$ discounted by a quantity proportional to the size of liquidity shocks $\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right)$. Since interdealer transactions propagate shocks across markets, the amount of discount relates to the local shock X_k and the average of shocks from other markets $\bar{X}_{-k}$. Price impact of dealers in the interdealer market prevents them from perfect risk sharing, resulting in an extra weight on X_k.

The amount of discount depends on the size of covariance σ_{lv}, the mass of hedgers who receive the shocks, and the total risk-bearing capacity in the market. If any of the participants in the market are risk-neutral, $\mathcal{D}$ is zero, and there is no discount in P_k^L.

The price that the market maker charges hedgers P_k^H is shown in equation (1.16). The difference between P_k^H and P_k^L is $\frac{\sigma_{lv}\delta_H}{2}X_k$. The market maker maximizes profits by setting the bid-ask spread as half of the gap in the reservation prices of hedgers and customer liquidity providers. Equations (1.15) and (1.16) show that bid-ask spreads from customer-dealer-customer riskless principal transactions are mainly driven by δ_H, but neither δ_d nor δ_L. The result supports my empirical identification strategy in the SVAR.

The interdealer price in equation (1.17) is equal to reservation price of market makers, discounted by the average of shocks $\bar{X}$ multiplied by $\mathcal{D}\sigma_{lv}N_H$. Equations (1.18) to (1.21) show the equilibrium trading quantities and the inventory level of dealers. Equation (1.18) shows that customer liquidity providers are buyers when the weighted average shock $\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right)$ is positive. In such cases, the local market is hit by a positive shock X_k, and/or other segmented markets receive an average positive shock $\bar{X}_{-k}$.

The amount that hedgers trade are shown in equation (1.19). As indirect participants in the global markets, hedgers play an additional role in absorbing global shocks, constituting the first term in equation (1.19), which resembles the term for customer liquidity providers. The second term in equation (1.19) shows that, as the actual receivers of the local shock, competitive hedgers liquidate half of the shocks.

Equation (1.20) shows that market makers will buy in the interdealer market when the local shock X_k is less than the average global shocks $\bar{X}_{-k}$ (and vice versa). Dealers absorb the market average shocks but are biased toward the local market.

In the following subsections, I demonstrate the intuition of the model by a numerical example and provide comparative statics under the parameters chosen as below, unless otherwise stated. The choice of parameters is harmless for most of the implications.

Parameter	K	N_H	N_L	$\bar{V}$	σ_v^2	σ_{lv}	δ_d	δ_H	δ_L
Value	3	1	1	100	1	1	2	2	2

1.5.3 A numerical example

Figure 1.7 presents an OTC market incorporated by three segmented markets. Hedgers in the market of Dealer A receive a positive shock $X_A = 7$, which is greater than the shocks $X_B = 4$ and $X_C = -9$. For hedging purposes, the hedgers in the markets of Dealers A and B would like to sell and the hedgers in the market of Dealer C would like to buy. Customer liquidity providers transact against their local hedgers.

All prices and volumes shown in Figure 1.7 are equilibrium quantities computed from equations (1.15) to (1.20). Dealer A buys three units from hedgers at a price of 92 and then sells 0.5 (1.6) units at a price of 99 (99.7) to customer liquidity providers in the local market (other dealers in the interdealer market), resulting in a net inventory of 0.9 units.

The market for Dealer A experiences a relatively severer positive shock so that the dealer has to rely on both her own customer liquidity pool, where she works the orders with potential customer buyers without liquidity shocks and trades with other dealers in the interdealer market. In this example, Dealer A conducts 0.5 units of customer-dealer-customer riskless principal transactions, and the spread for this transaction is 7.[27]

Hedgers with Dealer B receive a moderate shock. Therefore, they transact fewer quantities at a better price. Hedgers with Dealer C receive a negative shock,

[27]For illustration purposes, the absolute level of spread 7 in the example is large, relative to the expected payoff of the asset $\bar{V} = 100$. Ceteris paribus, increase in $\bar{V}$ would not impact the equilibrium spreads and volumes. For instance, let $\bar{V} = 10,000$ and the spread of 7 could be interpreted as approximately 7bps, within a reasonable ballpark from empirical observations.

and they buy 4.3 units from Dealer C at the price of 110.8, whereas customer liquidity providers partially provide 0.6 units liquidity to Dealer C at the price of 101.1. Dealer C purchases 2.5 units from the interdealer market, leaving a net balance of negative 1.2 units on her own book.

The aggregate global shock is positive, whereas perfect risk-sharing is not achieved. Market makers have monopoly power in local markets and face price impact in the interdealer market. Dealers and customer liquidity providers, considered as a whole, provide liquidity to hedgers and hold positive inventory. However, dealers and customer liquidity providers in market C remain with negative inventory.

Due to the positive aggregate shock, the interdealer price is 99.7, providing a positive expected return. Since markets are segmented, difference transaction prices are observed from different counterparties. Next, I discuss how to infer bid-ask spreads for riskless principal transactions and inventory transactions from the model. The interdealer price is regarded as the fundamental value of the asset at the first period. Note that the expected payoff of the asset $\bar{V} = 100$ is not a good candidate of fair price in this model.

1.5.4 Customer-dealer-customer riskless principal transactions

RPTs are empirically defined as a set of transactions in which a dealer buys from some customers and sells to other customers for appropriately equal quantities within a short period. Bringing the context into my model, I define RPT volume as the smaller one of the two absolute amounts that a dealer transacts with the two types of customers as follows:

$$\min(N_H|\Theta_{H,k}|, N_L|\Theta_{L,k}|)1_{\Theta_{H,k}\times\Theta_{L,k}<0}. \tag{1.22}$$

The amount in excess of the RPT volume is considered inventory transactions. In the numerical example, Dealer A conducts 2.5 units[28] of inventory transactions (IVT) with hedgers. In both RPT and IVT, hedgers sell at the price of 92.

This outcome, at first glance, seems counterintuitive. In fact, RPTs are traded at lower spreads because dealers opt to offer favorable prices to customer liquidity providers. In this example, the distance between interdealer price and liquidity provider price is smaller than the distance between interdealer price and hedger price. The consequence is a wider spread of 15.4 for IVT trades,[29] compared with a RPT spread of 7. In the RPT, the dealer offers the customer liquidity provider a negative half spread of -0.7.

The framework agrees with the intuition in [2]: *... dealer might find a non-dealer (C2) who is willing to provide liquidity for a fee (i.e., buying at a lower price than the fundamental value of the bond) ... C2 pays an even smaller spread (a negative spread in this scenario)....* Evidence in their Table 2 supports my formulation of RPTs. They document that the leg of a customer buy in pairs of RPTs is more likely to cross the interdealer price. Customer buyers in corporate bond markets are usually viewed as liquidity providers.

This framework differs from [52], in which hedgers can either trade immediately

[28] $3.0 - 0.5 = 2.5$

[29] $(99.7 - 92) \times 2 = 15.4$

with the dealer for an inventory transaction at a higher spread or wait until the dealer finds another potential buyer and trades as an RPT at a lower spread. Alternatively, [52] benchmark the transaction price with the expected payoff. The delay of trading is not modeled in my paper. Equivalently, if the hedgers decide to wait, my model implies that hedgers sell less to the dealer. My model emphasizes that market makers do not price discriminate hedgers.

1.5.5 Comparative statics

Figure 1.8 provides comparative statics for the expectation of bid-ask spreads and transaction volumes using a numerical method, in which liquidity shocks $\{X_k\}_{k=1}^K$ are randomly and independently drawn from identical and independent normal distributions. Computed are the bid-ask spread for RPTs, the bid-ask spread for IVTs, the ratio of RPT volume over total customer-dealer volume, and the average dealer inventory level. Appendix 1.7.2 describes the details in computing these quantities.[30]

Please insert Figure 1.10 here

Panel A of Figure 1.8 exhibits comparative statics when δ_d increases. It mimics the scenarios both during the 2008 financial crisis and the post-regulation periods, when dealer capital is constrained. The left plot in Panel A of Figure 1.8 shows that

[30] I exclude the realizations of $\{X_k\}_{k=1}^K$ in which local hedgers and customer liquidity providers trade in the same direction. It occurs when the absolute level of X_k is small and hedgers effectively become liquidity providers. This exclusion is sensible, especially considering large dealers, who conduct much more dealer-customer transactions than interdealer transactions. Figures 1.10 in the Appendix show that the main results are the same without this exclusion.

the bid-ask spreads for RPTs remain flat when δ_d increases, whereas the spreads for IVTs increase. The increase of δ_d leads to a higher risk aversion for the economy, that is, $\mathcal{D}$ in equations (1.15) and (1.16). Thus, customer-dealer prices P_k^Hs for different markets become more diverged. Since hedgers in high X_k markets (e.g., P_A^H in Figure 1.7) conduct IVT sell, and hedgers in low X_k markets conduct IVT buy (e.g., P_C^H in Figure 1.7), the IVT spread (e.g., gap between P_A^H and P_C^H) becomes wider.

Please insert Figure 1.7 here

Dealers conduct more RPTs and hold less inventory when they become more risk averse, as shown in the right plot in Panel A of Figure 1.8. It explains the findings in [2] that the average bid-ask spreads do not increase due to higher RPT volume during the post-regulation period.

Panel B of Figure 1.8 presents the comparative statics when δ_H increases and hedgers are more eager to sell, for instance, the 2008 financial crisis and the Eurozone crisis. The left plot in Panel B of Figure 1.8 demonstrates that δ_H is an important determinant for both types of bid-ask spreads.

One may not conclude that δ_H drives most of the fluctuations in the spreads for IVTs, because the variations of δ_H and δ_d are of different scales. Dealers sometimes could be extremely risk averse. For instance, [53] let δ_d be infinite in their model.

The message in Panel C of Figure 1.8 is that δ_L does not impact the RPT spread. It may impact the IVT spread, but the magnitude largely depends on selected parameters. The right plot in Panel C of Figure 1.8 shows that fewer customer liquidity providers, or, equivalently, an increase of δ_L, result in less riskless principal transactions and higher dealer inventory.

Please insert Figure 1.8 here

Overall, the predictions from my model are consistent with the empirical findings in the SVAR, as stated in equation 1.3.

1.6 Conclusion

This paper provides a measure of customer liquidity provision in corporate bond markets. This measure incorporates the information contained in the trading activities when bond dealers conduct inventory transactions and make markets versus arrange riskless principal transactions and match customer liquidity demanders and providers.

To infer this measure, I exploit a SVAR approach and decompose the changes in bid-ask spreads in corporate bond market into structural shocks to the risk-taking willingness of various market participants. This decomposition could be useful in monitoring conditions of different market participants in the corporate bond market and understanding the actual impact from regulatory events or market crises.

The obtained measures contain important information about the aggregate market conditions. Customer liquidity provisions are shown to increase during the post-regulation period, as suggested by [15], because dealers improve their technology in matchmaking.

I employ my new measures of both customer liquidity providers and dealers to explain the change in yield spreads for corporate bonds of different credit ratings. I show that yield spreads of high quality bonds respond more to the measure of customer liquidity provision, whereas yield spreads of riskier bonds are more exposed to the measure of dealers. It suggests that the decrease of liquidity due to stringent regulation and the increase of liquidity due to matchmaking technology might not be simply offset.

1.7 Appendix

1.7.1 Solve the equilibrium

The problem is equivalent to:

$$\max_{P_k^H, P_k^L, m_k} N_H \Theta_{H,k} P_k^H + N_L \Theta_{L,k} P_k^L - m_k \left(\frac{\nu}{\gamma} - \frac{\eta}{\gamma} X_{-k} + \frac{1}{(K-1)\gamma} m_k \right)$$
$$+ \left(-N_H \Theta_{H,k} - N_L \Theta_{L,k} + m_k\right) \bar{V} - \frac{1}{2} \delta_d \sigma_v^2 \left(-N_H \Theta_{H,k} - N_L \Theta_{L,k} + m_k\right)^2 . \tag{1.23}$$

The optimal demand for the customers H and L are:

$$\Theta_{H,k} = \frac{P_k^{R,H} - P_k^H}{\delta_H \sigma_v^2} \quad \text{and} \tag{1.24}$$

$$\Theta_{L,k} = \frac{P^{R,L} - P_k^L}{\delta_L \sigma_v^2}. \tag{1.25}$$

F.O.C. wrt P_k^H is

$$N_H \Theta_{H,k} - N_H \frac{P_k^H}{\delta_H \sigma_v^2} + \frac{N_H \bar{V}}{\delta_H \sigma_v^2} - \frac{\delta_d}{\delta_H} N_H \left(-N_H \Theta_{H,k} - N_L \Theta_{L,k} + m_k\right) = 0. \tag{1.26}$$

It reduces to

$$\left(2\delta_H + \delta_d N_H\right) \Theta_{H,k} + \delta_d \left(N_L \Theta_{L,k} - m_k\right) = -\frac{\delta_H \sigma_{lv} X_k}{\sigma_v^2}, \tag{1.27}$$

F.O.C. wrt P_k^L is

$$N_L \Theta_{L,k} - N_L \frac{P_k^L}{\delta_L \sigma_v^2} + \frac{N_L \bar{V}}{\delta_L \sigma_v^2} - \frac{\delta_d}{\delta_L} N_L \left(-N_H \Theta_{H,k} - N_L \Theta_{L,k} + m_k \right) = 0. \tag{1.28}$$

It reduces to

$$\left(2\delta_L + \delta_d N_L \right) \Theta_{L,k} + \delta_d \left(N_H \Theta_{H,k} - m_k \right) = 0. \tag{1.29}$$

F.O.C. wrt m_k is

$$-\left(\frac{\nu}{\gamma} - \frac{\eta}{\gamma} X_{-k} + \frac{1}{(K-1)\gamma} m_k \right) - \frac{1}{(K-1)\gamma} m_k + \bar{V} - \delta_d \sigma_v^2 \left(-N_H \Theta_{H,k} - N_L \Theta_{L,k} + m_k \right) = 0 \tag{1.30}$$

It reduces to:

$$P^{R,L} - P^D - \frac{1}{(K-1)\gamma} m_k + \delta_d \sigma_v^2 \left(N_H \Theta_{H,k} + N_L \Theta_{L,k} \right) - \delta_d \sigma_v^2 m_k = 0. \tag{1.31}$$

The solution to the system of three FOCs is

$$\begin{aligned}
\Theta_{H,k} &= -\frac{2\delta_d \delta_L \gamma \sigma_v^2 (K-1) \left(P^D(m_k) - P_k^{R,H} \right) + \left(2\delta_H \delta_L + \delta_d \delta_H N_L \right) \sigma_{lv} X_k}{2\sigma_v^2 \left(\delta_d \delta_L N_H + \delta_d \delta_H N_L + 2\delta_H \delta_L \left(1 + \delta_d \gamma (K-1) \sigma_v^2 \right) \right)} \\
\Theta_{L,k} &= \frac{\delta_d \delta_H \left(-2\gamma \sigma_v^2 (K-1) \left(P^D(m_k) - P^{R,L} \right) + N_H \sigma_{lv} X_k \right)}{2\sigma_v^2 \left(\delta_d \delta_L N_H + \delta_d \delta_H N_L + 2\delta_H \delta_L \left(1 + \delta_d \gamma (K-1) \sigma_v^2 \right) \right)} \\
m_k &= \frac{\gamma (K-1) \left(\left(P^{R,L} - P^D(m_k) \right) \left(2\delta_H \delta_L + \delta_d \delta_L N_H + \delta_d \delta_H N_L \right) - \delta_d \delta_H \delta_L N_H \sigma_{lv} X_k \right)}{\delta_d \delta_L N_H + \delta_d \delta_H N_L + 2\delta_H \delta_L \left(1 + \delta_d \gamma (K-1) \sigma_v^2 \right)}.
\end{aligned} \tag{1.32}$$

Next, I solve for ν, γ and η by matching parameters from equation (1.9) and equation (1.32). Rearrange equation (1.32):

$$\begin{aligned} m_k = & \frac{\gamma\left(K-1\right)\left(2\delta_H\delta_L+\delta_d\delta_L N_H+\delta_d\delta_H N_L\right)P^{R,L}}{\delta_d\delta_L N_H+\delta_d\delta_H N_L+2\delta_H\delta_L\left(1+\delta_d\gamma\left(K-1\right)\sigma_v^2\right)} \\ & -\frac{\gamma\left(K-1\right)\left(2\delta_H\delta_L+\delta_d\delta_L N_H+\delta_d\delta_H N_L\right)}{\delta_d\delta_L N_H+\delta_d\delta_H N_L+2\delta_H\delta_L\left(1+\delta_d\gamma\left(K-1\right)\sigma_v^2\right)}P^D\left(m_k\right) \\ & -\frac{\gamma\left(K-1\right)\delta_d\delta_H\delta_L N_H\sigma_{lv}}{\delta_d\delta_L N_H+\delta_d\delta_H N_L+2\delta_H\delta_L\left(1+\delta_d\gamma\left(K-1\right)\sigma_v^2\right)}X_k. \end{aligned}$$

Match the coefficients with the conjectured demand curve in equation (1.9) and get the following system:

$$\nu = \frac{\gamma\left(K-1\right)P_U^R\left(2\delta_c+\delta\left(N_I+N_U\right)\right)}{\delta\left(N_U+N_I\right)+2\delta_c+2\delta_c\delta\gamma\left(K-1\right)\sigma_v^2} \tag{1.33}$$

$$\gamma = \frac{\gamma\left(K-1\right)\left(2\delta_c+\delta\left(N_I+N_U\right)\right)}{\delta\left(N_U+N_I\right)+2\delta_c+2\delta_c\delta\gamma\left(K-1\right)\sigma_v^2} \tag{1.34}$$

$$\eta = \frac{\gamma\left(K-1\right)\delta\delta_c N_I\sigma_{lv}}{\delta\left(N_U+N_I\right)+2\delta_c+2\delta_c\delta\gamma\left(K-1\right)\sigma_v^2}. \tag{1.35}$$

The solution is:

$$\gamma = \frac{\left(K-2\right)\left(2\delta_H\delta_L+\delta_d\delta_L N_H+\delta_d\delta_H N_L\right)}{2\delta_d\delta_L\delta_H\left(K-1\right)\sigma_v^2} \tag{1.36}$$

$$\nu = \frac{\left(K-2\right)\left(2\delta_H\delta_L+\delta_d\delta_L N_H+\delta_d\delta_H N_L\right)}{2\delta_d\delta_H\left(K-1\right)\sigma_v^2}P^{R,L} \tag{1.37}$$

$$\eta = \frac{\left(K-2\right)}{2\left(K-1\right)\sigma_v^2}N_H\sigma_{lv}. \tag{1.38}$$

The optimal m_k^* is

$$m_k^* = \gamma \left(P^{R,L} - P^D\right) - \frac{(K-2)}{2\,(K-1)\,\sigma_v^2} N_H \sigma_{lv} X_k. \tag{1.39}$$

The market clearing condition in equation (1.10) implies the interdealer price is

$$\sum_{k=1}^{K} \gamma \left(P^{R,L} - P^D\right) - \frac{(K-2)\,N_H \sigma_{lv} X_k}{2\,(K-1)\,\sigma_v^2} = 0$$

and

$$P^{R,L} - P^D = \frac{N_H \sigma_{lv} \bar{X}}{\left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)}, \tag{1.40}$$

where $\bar{X}$ is defined as $\frac{1}{K}\sum_{k=1}^{K} X_k$. And thus we have

$$m_k = \frac{(K-2)\,N_H \sigma_{lv}}{2\,(K-1)\,\sigma_v^2} \left(\bar{X} - X_k\right). \tag{1.41}$$

The optimal trading quantities and prices are:

$$\Theta_{L,k} = \frac{K-2}{2K} \frac{\delta_d \delta_H}{(\delta_d \delta_L N_H + \delta_d \delta_H N_L + 2\delta_H \delta_L)} \frac{\sigma_{lv} N_H}{\sigma_v^2} \bar{X}_{-k} + \frac{\delta_d \delta_H}{K\,(\delta_d \delta_L N_H + \delta_d \delta_H N_L + 2\delta_H \delta_L)} \frac{\sigma_{lv} N_H}{\sigma_v^2} X_k \tag{1.42}$$

$$\Theta_{H,k} = \frac{K-2}{2K} \frac{\delta_d \delta_L}{(\delta_d \delta_L N_H + \delta_d \delta_H N_L + 2\delta_H \delta_L)} \frac{\sigma_{lv} N_H}{\sigma_v^2} \bar{X}_{-k} + \frac{\delta_d \delta_L}{K\,(\delta_d \delta_L N_H + \delta_d \delta_H N_L + 2\delta_H \delta_L)} \frac{\sigma_{lv} N_H}{\sigma_v^2} X_k - \frac{\sigma_{lv}}{2\sigma_v^2} X_k \tag{1.43}$$

$$P_k^L = \bar{V} - \frac{K-2}{2K} \frac{1}{\left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)} \sigma_{lv} N_H \bar{X}_{-k} - \frac{1}{K\left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)} \sigma_{lv} N_H X_k \tag{1.44}$$

$$P_k^H = \bar{V} - \frac{K-2}{2K} \frac{1}{\left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)} \sigma_{lv} N_H \bar{X}_{-k} - \frac{1}{K\left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)} \sigma_{lv} N_H X_k - \frac{\sigma_{lv} \delta_H}{2} X_k \tag{1.45}$$

$$P^D = \bar{V} - \frac{1}{\left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)} N_H \sigma_{lv} \bar{X}. \tag{1.46}$$

The ultimate inventory level is

$$-N_U\Theta_{U,k} - N_I\Theta_{I,k} + m_k = \frac{\delta_H\delta_L}{\sigma_v^2\left(\delta_d\delta_L N_H + \delta_d\delta_H N_L + 2\delta_H\delta_L\right)}\left(N_H\sigma_{lv}\bar{X}\right). \quad (1.47)$$

Denote the quantity $\mathcal{D}^{-1} = \left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)$, $\mathcal{D}^{-1}$ captures the total risk bearing capacity in the economy. The optimal traded quantities and prices reduce to

$$P_k^L = \bar{V} - \frac{1}{2}\mathcal{D}\sigma_{lv}N_H\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right) \quad (1.48)$$

$$P_k^H = \bar{V} - \frac{1}{2}\mathcal{D}\sigma_{lv}N_H\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right) - \frac{\sigma_{lv}\delta_H}{2}X_k \quad (1.49)$$

$$P^D = \bar{V} - \mathcal{D}\sigma_{lv}N_H\bar{X} \quad (1.50)$$

$$\Theta_{L,k} = \frac{1}{2}\frac{\mathcal{D}}{\delta_L}\frac{\sigma_{lv}N_H}{\sigma_v^2}\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right) \quad (1.51)$$

$$\Theta_{H,k} = \frac{1}{2}\frac{\mathcal{D}}{\delta_H}\frac{\sigma_{lv}N_H}{\sigma_v^2}\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right) - \frac{\sigma_{lv}}{2\sigma_v^2}X_k \quad (1.52)$$

$$-N_U\Theta_{U,k} - N_I\Theta_{I,k} + m_k = \frac{\mathcal{D}}{\delta_d}\frac{\sigma_{lv}N_H}{\sigma_v^2}\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right). \quad (1.53)$$

1.7.2 Numerical computation of spreads and volumes

First, transaction prices and volumes are computed based on equations (1.15) to (1.21). Equation 1.21 directly yields the average inventory level.

First, compute the RPT volume as described in equation (1.22) in each local market, denoted as

$$\Theta_k^{CDC} = \min(N_H|\Theta_{H,k}|, N_L|\Theta_{L,k}|)1_{\Theta_{H,k}\times\Theta_{L,k}<0}. \tag{1.54}$$

It follows that the ratio of customer-dealer-customer riskless principal transaction volume over total customer-dealer volume is given by:

$$\frac{\sum_{k=1}^{K}\Theta_k^{CDC}}{\frac{1}{2}\sum_{k=1}^{K}\left(N_H|\Theta_{H,k}| + N_L|\Theta_{L,k}|\right)} \tag{1.55}$$

The weighted average RPT spread is the difference between:

$$P_{Buy}^{CDC} = \frac{1}{\sum_{k=1}^{K} 1_{\Theta_k^{CDC}>0}} \sum_{k=1}^{K} 1_{\Theta_k^{CDC}>0}\left(1_{\Theta_{H,k}>0}P_k^H + 1_{\Theta_{L,k}>0}P_k^L\right) \tag{1.56}$$

and

$$P_{Sell}^{CDC} = \frac{1}{\sum_{k=1}^{K} 1_{\Theta_k^{CDC}>0}} \sum_{k=1}^{K} 1_{\Theta_k^{CDC}>0}\left(1_{\Theta_{H,k}<0}P_k^H + 1_{\Theta_{L,k}<0}P_k^L\right). \tag{1.57}$$

Then, compute the spread for inventory transactions, define the trading volume for inventory transaction as

$$\Theta_{H,k}^{IVT} = \min(N_H|\Theta_{H,k}| - \Theta_k^{CDC}, 0) \tag{1.58}$$

$$\Theta_{L,k}^{IVT} = \min(N_L|\Theta_{L,k}| - \Theta_k^{CDC}, 0) \tag{1.59}$$

The volume weighted average IVT spread is the difference between:

$$P^{IVT}_{Buy} = \frac{\sum_{k=1}^{K}\left(1_{\Theta^{IVT}_{H,k}>0}1_{\Theta_{H,k}>0}P^H_k + 1_{\Theta^{IVT}_{L,k}>0}1_{\Theta_{L,k}>0}P^L_k\right)}{\sum_{k=1}^{K}\left(1_{\Theta^{IVT}_{H,k}>0}1_{N_H\Theta_{H,k}>0} + 1_{\Theta^{IVT}_{L,k}>0}1_{\Theta_{L,k}>0}\right)} \tag{1.60}$$

and

$$P^{IVT}_{Sell} = \frac{\sum_{k=1}^{K}\left(1_{\Theta^{IVT}_{H,k}>0}1_{\Theta_{H,k}<0}P^H_k + 1_{\Theta^{IVT}_{L,k}>0}1_{\Theta_{L,k}<0}P^L_k\right)}{\sum_{k=1}^{K}\left(1_{\Theta^{IVT}_{H,k}>0}1_{N_H\Theta_{H,k}<0} + 1_{\Theta^{IVT}_{L,k}>0}1_{\Theta_{L,k}<0}\right)}. \tag{1.61}$$

In extreme cases that realizations of shocks are similar for all markets, all liquidity provider only provides liquidity to local hedgers. All the customer buy transactions in this realization are RPT. For instance, suppose each market receives the exact same size of positive shock, there will be no interdealer transactions and every dealer buys from hedgers and sell a fraction to local liquidity providers. A consequence is that P^{IVT}_{Buy} is not well defined. Without loss of generality, P^{IVT}_{Buy} is assumed be to P^{RPT}_{Buy} in those situations. This assumption is conservative because $P^{IVT}_{Buy} \geq P^{RPT}_{Buy}$. In addition, it follows the intuition that marginal increments of inventory buys would be traded at least at the current RPT transaction customer buy price.

Please insert Figure 1.9 here

Risk aversions of market makers and customer liquidity providers may slightly impact the spread for riskless principal transactions if hedgers and customer liquidity providers transact in the same direction with the market maker in a given local market, i.e. $sign(\Theta_{I,k}) = sign(\Theta_{U,k})$. As shown in Panel A of Figure1.9, it occurs when the absolute level of local shocks are too small, relative to the aggregate global shock. In those scenarios, there is no riskless principal transaction in this particular market. It is not a desirable feature since hedgers are effectively the same as customer liquidity providers because their main role is to provide liquidity to the global market, whereas the risk aversions of them are different. Panel B of Figure1.9 illustrate the simulated realized distribution of iid shocks of X_k drawn from $N(0, 5)$. Comparative statics in Figure 1.8 excludes those realizations when at least one local dealer trades with hedgers and liquidity providers in the same direction. Results including those unsatisfying trails are reported in Figure 1.10, which verifies that the impact is minimal. Panel C of Figure1.9 shows that those unsatisfying trails can be perfectly avoided if the data generating process for shocks are restricted at some distance away from zero. For example, shocks are uniformly drawn from $[-2, -1] \cup [1, 2]$.

1.7.3 The 3^{rd} restriction in SVAR

The optimal trading quantities $\Theta_{L,k}(\delta_L, \delta_H, \delta_d)$ and $\Theta_{H,k}(\delta_L, \delta_H, \delta_d)$:

$$\begin{aligned}
\Theta_{L,k} &= \frac{1}{2}\frac{\mathcal{D}}{\delta_L}\frac{\sigma_{lv} N_H}{\sigma_v^2}\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right) \\
\Theta_{H,k} &= \frac{1}{2}\frac{\mathcal{D}}{\delta_H}\frac{\sigma_{lv} N_H}{\sigma_v^2}\left(\frac{K-2}{K}\bar{X}_{-k} + \frac{2}{K}X_k\right) - \frac{\sigma_{lv}}{2\sigma_v^2}X_k
\end{aligned}$$

are homogeneous functions of degree zero, where $\mathcal{D}^{-1} = \left(\frac{2}{\delta_d} + \frac{N_H}{\delta_H} + \frac{N_L}{\delta_L}\right)$. If hedgers, dealers and customer liquidity providers all become twice as risk averse as they were, volumes remains the same.

Formally, consider a function $f\left(g_H, g_L\right)$, where $g_H\left(\delta_L, \delta_H, \delta_d\right) = \frac{\mathcal{D}(\delta_L,\delta_H,\delta_d)}{\delta_H}$ and $g_L\left(\delta_L, \delta_H, \delta_d\right) = \frac{\mathcal{D}(\delta_L,\delta_H,\delta_d)}{\delta_L}$. We can rewrite:

$$g_H = g'_H = N_L + exp\left(\log\left(\delta_L\right) + \log\left(N_H\right) - \log\left(\delta_H\right)\right) + 2exp\left(\log\left(\delta_L\right) - \log\left(\delta_d\right)\right)$$

$$g_L = g'_L = N_H + exp\left(\log\left(\delta_H\right) + \log\left(N_L\right) - \log\left(\delta_L\right)\right) + 2exp\left(\log\left(\delta_H\right) - \log\left(\delta_d\right)\right).$$

Apply the chain rule:

$$\begin{aligned}
\frac{\partial f}{\partial \log\left(\delta_L\right)} &= f_1 \frac{\partial g'}{\partial \log\left(\delta_L\right)} + f_2 \frac{\partial h'}{\partial \log\left(\delta_L\right)} \qquad (1.62)\\
&= f_1 exp\left(\log\left(\delta_L\right) + \log\left(N_H\right) - \log\left(\delta_H\right)\right) + \\
&\quad 2f_1 exp\left(\log\left(\delta_L\right) - \log\left(\delta_d\right)\right) - f_2 exp\left(\log\left(\delta_H\right) + \log\left(N_L\right) - \log\left(\delta_L\right)\right)\\
\frac{\partial f}{\partial \log\left(\delta_H\right)} &= f_1 \frac{\partial g'}{\partial \log\left(\delta_H\right)} + f_2 \frac{\partial h'}{\partial \log\left(\delta_H\right)} \qquad (1.63)\\
&= -f_1 exp\left(\log\left(\delta_L\right) + \log\left(N_H\right) - \log\left(\delta_H\right)\right) + \\
&\quad f_2 exp\left(\log\left(\delta_L\right) + \log\left(N_H\right) - \log\left(\delta_H\right)\right) + 2f_2 exp\left(\log\left(\delta_H\right) - \log\left(\delta_d\right)\right)\\
\frac{\partial f}{\partial \log\left(\delta_d\right)} &= f_1 \frac{\partial g'}{\partial \log\left(\delta_d\right)} + f_2 \frac{\partial h'}{\partial \log\left(\delta_d\right)} \qquad (1.64)\\
&= -2f_1 exp\left(\log\left(\delta_L\right) - \log\left(\delta_d\right)\right) - 2f_2 exp\left(\log\left(\delta_H\right) - \log\left(\delta_d\right)\right). \qquad (1.65)
\end{aligned}$$

Hence,

$$\frac{\partial f}{\partial \log\left(\delta_L\right)} + \frac{\partial f}{\partial \log\left(\delta_H\right)} + \frac{\partial f}{\partial \log\left(\delta_d\right)} = 0 \qquad (1.66)$$

So is % RPT, $\mathbb{E}\left(\sum_{k=1}^{K} \min(N_I|\Theta_{I,k}|, N_U|\Theta_{U,k}|) 1_{\Theta_{I,k}\times\Theta_{U,k}<0}\right)$:

$$\frac{\partial \% RPT}{\partial \log\left(\delta_L\right)} + \frac{\partial \% RPT}{\partial \log\left(\delta_H\right)} + \frac{\partial \% RPT}{\partial \log\left(\delta_d\right)} = 0.$$

1.7.4 What if shocks are correlated?

I show that my SVAR method is still valid when the orthogonality condition does not hold. Let

$$\begin{bmatrix} \epsilon_t^H \\ \epsilon_t^L \\ \epsilon_t^d \end{bmatrix} = \begin{bmatrix} \upsilon_t^H \\ \rho_2 \upsilon_t^H + \rho_3 \upsilon_t^d + \upsilon_t^L \\ \rho_1 \upsilon_t^H + \upsilon_t^d \end{bmatrix}, \tag{1.67}$$

such that $\upsilon_{\mathbf{t}} = \begin{bmatrix} \upsilon_t^H \\ \upsilon_t^L \\ \upsilon_t^d \end{bmatrix}$ are orthogonal shocks, and ρ_1, ρ_2, and ρ_3 govern the covariance among shocks to market participants ($\epsilon_{\mathbf{t}}$).

υ_t^d is the shocks to dealers that are orthogonal to the shocks to the hedgers, whereas υ_t^L is the shocks to customer liquidity providers that are orthogonal to the other two types of shocks. Thus, the SVAR in equation (1.2) can be written as

$$\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i y_{t-i}} + \begin{bmatrix} b_1 & 0 & 0 \\ b_{2,1} & b_2 & b_{2,3} \\ b_{3,1} & b_{3,2} & b_3 \end{bmatrix} \begin{bmatrix} \upsilon_t^H \\ \rho_2 \upsilon_t^H + \rho_3 \upsilon_t^d + \upsilon_t^L \\ \rho_1 \upsilon_t^H + \upsilon_t^d \end{bmatrix}. \tag{1.68}$$

We can rewrite the SVAR:

$$
\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i}\mathbf{y_{t-i}} + \begin{bmatrix} b_1 & 0 & 0 \\ b_{2,1} + b_2\rho_2 + b_{2,3}\rho_1 & b_2 & b_2\rho_3 + b_{2,3} \\ b_{3,1} + b_{3,2}\rho_2 + b_3\rho_1 & b_{3,2} & b_{3,2}\rho_3 + b_3 \end{bmatrix} \underbrace{\begin{bmatrix} \upsilon_t^H \\ \upsilon_t^L \\ \upsilon_t^d \end{bmatrix}}_{\boldsymbol{\upsilon}_\mathbf{t}}. \quad (1.69)
$$

Denote $\frac{b_{3,1}+b_{3,2}\rho_2+b_3\rho_1}{b_{3,1}} = \Xi_H$ and $\frac{b_{3,2}\rho_3+b_3}{b_3} = \Xi_d$, the system reads as

$$
\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i}\mathbf{y_{t-i}} + \underbrace{\begin{bmatrix} \Xi_H^{-1} b_1 & 0 & 0 \\ \Xi_H^{-1}\left(b_{2,1} + b_2\rho_2 + b_{2,3}\rho_1\right) & b_2 & \Xi_d^{-1}\left(b_2\rho_3 + b_{2,3}\right) \\ b_{3,1} & b_{3,2} & b_3 \end{bmatrix}}_{\mathbf{B}_\Xi} \begin{bmatrix} \Xi_H \upsilon_t^H \\ \upsilon_t^L \\ \Xi_d \upsilon_t^d \end{bmatrix}. \quad (1.70)
$$

Applying the same restrictions, as stated in equation (1.4), I can recover $\left[\Xi_H \upsilon_t^H,\ \upsilon_t^L,\ \Xi_d \upsilon_t^d\right]^T$, which are the desired series $\boldsymbol{\upsilon}_\mathbf{t}$, up to the scale factors Ξ_H and Ξ_d. The scale factors are harmless as long as they are positive. The positivity can be verified by (i) the movements of $\Xi_H \upsilon_t^H$ and $\Xi_d \upsilon_t^d$ over economic cycles and (ii) the signs in the estimates of the coefficients in $\mathbf{B}_\Xi$.

1.7.5 Identify riskless principal transactions

The identification method for riskless principal transactions follows a Last In, First Out (LIFO) strategy for each dealer-Cusip-date sub-dataset of transactions records, following [2]. Offsetting transactions are defined as two subsequent transactions where the dealer buys in one transaction and sell in another one within 15 minutes. Note that it is not necessary that the dealer buys in the earlier transaction and sells in the later one. It is also not necessary that the amounts are matched in the two transactions.

The implementation of the LIFO strategy is a recursive procedure described below. For every dealer-Cusip-date sub-dataset, first sorted by execution time. Denote this sub-dataset as $\mathbb{C}$, the collection of all transactions conducted by a dealer on a specific Cusip during a day. Denote $\mathbb{C}^{RPT}$ as the collection of riskless principal transactions. Denote $\mathbb{C}^{IVT}$ as the collection of inventory transactions. $\mathbb{C}^{RPT}$ and $\mathbb{C}^{IVT}$ are empty sets when initialized. The following steps are iterated until no new transactions are identified as riskless principal transaction.

1. Denote a dummy $RPT = 0$. Let n start from 1. Let N be the total number of transactions in $\mathbb{C}$. Focus on the n-th transaction. Break the loop as soon as $RPT = 1$ or $n = N$
2. Check if the dealer trades the same direction in the next transaction $n + 1$:

   if yes, $n = n + 1$; return to step 1;
3. Check if time difference between the n-th trade and the $n+1$-th trade is more than 15 minutes:

   if yes, $n = n + 1$; return to step 1;
4. Since both steps 2 and 3 are satisfied, the n-th and $n + 1$-th transactions are identified (partially) as riskless principal transactions. The volume of riskless principal transactions is the minimum of the trading volume of n-th and $n+1$-

th transactions.

5. If the volumes of n-th and $n+1$-th transactions are the same, move both transactions from $\mathbb{C}$ to $\mathbb{C}^{RPT}$;

   If the volume of n-th transaction is greater than that of the $n+1$-th transaction, move the $n+1$-th from $\mathbb{C}$ to $\mathbb{C}^{RPT}$, move the n-th from $\mathbb{C}$ to $\mathbb{C}^{RPT}$ but only with the riskless principal amount in step 4, and modify the volume of remaining part of the n-th in $\mathbb{C}$ as the difference between its original amount and riskless principal amount.

   if the volume of n-th transaction is smaller than that of the $n+1$-th transaction, do the opposite.

6. Let $RPT = 1$.

If the procedure (1) to (6) breaks because $RPT = 1$, applied this procedure (1) to (6) on the resulted $\mathbb{C}$. If the procedure (1) to (6) breaks because $n = N$, move all the remaining transactions to $\mathbb{C}^{IVT}$.

Considering the type of counterparties, each transaction can be classified as riskless principal transaction to customers, riskless principal transaction to dealers, and inventory transaction. If a transaction is attributed to multiple types, the ultimate classification is type of the highest volume.

This algorithm will capture possible riskless principal transactions where there are multiple transactions in buy and sell legs. For example, considering the following example trading pattern $\mathbb{C}$:

	$\mathbb{C}$		
ID	Buy/Sell	amt	time
1	S	120	9:31
2	B	100	9:32
3	B	200	10:25
4	S	100	10:30
5	S	120	10:35
6	B	20	10:35
7	S	200	14:25

$\mathbb{C}$ is applied on the described algorithm. First, transactions 1 and 2 are identified as (partial) riskless principal transactions.

	$\mathbb{C}$					$\mathbb{C}^{RPT}$		
ID	Buy/Sell	amt	time	ID	Buy/Sell	amt	time	counter ID
1	S	20	9:31	1	S	100	9:31	2
				2	B	100	9:32	1
3	B	200	10:25					
4	S	100	10:30					
5	S	120	10:35					
6	B	20	10:35					
7	S	200	14:25					

In the 2^{nd} iteration, transactions 3 and 4 are identified as (partial) riskless principal transactions.

	$\mathbb{C}$				$\mathbb{C}^{RPT}$			
ID	Buy/Sell	amt	time	ID	Buy/Sell	amt	time	counter ID
1	S	20	9:31	1	S	100	9:31	2
				2	B	100	9:32	1
3	B	100	10:25	3	B	100	10:25	4
				4	S	100	10:30	3
5	S	120	10:35					
6	B	20	10:35					
7	S	200	14:25					

In the 3^{nd} iteration, transactions 3 and 5 are identified as (partial) riskless principal transactions.

	$\mathbb{C}$				$\mathbb{C}^{RPT}$			
ID	Buy/Sell	amt	time	ID	Buy/Sell	amt	time	counter ID
1	S	20	9:31	1	S	100	9:31	2
				2	B	100	9:32	1
				3	B	100	10:25	4
				4	S	100	10:30	3
				3	B	100	10:25	5
5	S	20	10:35	5	S	100	10:35	3
6	B	20	10:35					
7	S	200	14:25					

In the 4^{nd} iteration, transactions 5 and 6 are identified as (partial) riskless principal transactions.

$\mathbb{C}$				$\mathbb{C}^{RPT}$				
ID	Buy/Sell	amt	time	ID	Buy/Sell	amt	time	counter ID
1	S	20	9:31	1	S	100	9:31	2
				2	B	100	9:32	1
				3	B	100	10:25	4
				4	S	100	10:30	3
				3	B	100	10:25	5
				5	S	100	10:35	3
				5	S	20	10:35	6
				6	B	20	10:35	5
7	S	200	14:25					

In the last iteration, $\mathbb{C}$ contains $N = 2$ records and the loop breaks because $n = N$. So all the remaining transactions are moved to $\mathbb{C}^{IVT}$:

$\mathbb{C}^{IVT}$				$\mathbb{C}^{RPT}$				
ID	Buy/Sell	amt	time	ID	Buy/Sell	amt	time	counter ID
1	S	20	9:31	1	S	100	9:31	2
				2	B	100	9:32	1
				3	B	100	10:25	4
				4	S	100	10:30	3
				3	B	100	10:25	5
				5	S	100	10:35	3
				5	S	20	10:35	6
				6	B	20	10:35	5
7	S	200	14:25					

Transaction 1 is considered as riskless principal transaction because 83% of the volume is offset within 15 minutes. Transactions 3 to 6 are also riskless principal transactions. Despite that they all have different trading volume, the aggregated buy and sell volume matches and they all occur within 15 minutes.

1.7.6 TVIX

Consider the price of a volatility contract payoff in the 30-year Treasury bond market:

$$TVIX_t^2 = R_{f,t+1}^{-1}\mathbb{E}_t^Q\left(\log\left\{\frac{F_{t+1}}{F_t}\right\}^2\right), \tag{1.71}$$

where F_t is the price of a future contact at month t which expires in month $t+1$ and the underlying is a 30-year Treasury bond.

Define $G\left(F_{t+1}\right) = \log\left\{\frac{F_{t+1}}{F_t}\right\}^2$ as the payoff of the squared contract. Note $G\left(F_{t+1}\right)$ is a twice-continuously differentiable function. Following the spanning engine in [38] and [39], the payoff $G\left(F_{t+1}\right)$ can be spanned by a continuum of out of the money (OTM) call options and put options. That is,

$$G\left(F_{t+1}\right) = \quad G\left(\bar{F}\right) + \left(F_{t+1} - \bar{F}\right) G_{F_{t+1}}\left(\bar{F}\right) + \int_{\bar{F}}^{\infty} G_{F_{t+1}F_{t+1}}\left(K\right)\left(F_{t+1} - K\right)^{+} dK \tag{1.72}$$

$$+ \int_{0}^{\bar{F}} G_{F_{t+1}F_{t+1}}\left(K\right)\left(K - F_{t+1}\right)^{+} dK, \tag{1.73}$$

where $\bar{F}$ represents a possible outcome of F_{t+1}, and $G_{F_{t+1}}\left(\bar{F}\right)$ is the first-order derivative for the payoff $G\left(F_{t+1}\right)$ with respect to F_{t+1} evaluated at $\bar{F}$. Similarly, $G_{F_{t+1}F_{t+1}}\left(K\right)$ stands for the second-order derivative for $G\left(F_{t+1}\right)$ with respect to F_{t+1} evaluated at a strike price K. $\left(F_{t+1} - K\right)^{+}$ and $\left(K - F_{t+1}\right)^{+}$ are payoffs for call options and put options, where $(x)^{+} = \max\left(0, x\right)$.

Thus, the price of the volatility contract payoff can be evaluated at $\bar{F} = F_t$

and be expressed as:

$$
\begin{aligned}
TVIX_t^2 &= R_{f,t+1}^{-1}\mathbb{E}_t^Q\left(\log\left\{\frac{F_{t+1}}{F_t}\right\}^2\right) && (1.74)\\
&= R_{f,t+1}^{-1}\mathbb{E}_t^Q\left(\int_{F_t}^{\infty} G_{F_{t+1}F_{t+1}}(K)\left(F_{t+1}-K\right)^+ dK\right) && (1.75)\\
&+ R_{f,t+1}^{-1}\mathbb{E}_t^Q\left(\int_{0}^{F_t} G_{F_{t+1}F_{t+1}}(K)\left(K-F_{t+1}\right)^+ dK\right) && (1.76)\\
&= \int_{F_t}^{\infty} 2\frac{1-\log\frac{K}{F_t}}{K^2}\mathbb{E}_t^Q\left\{R_{f,t+1}^{-1}\left(F_{t+1}-K\right)^+\right\} dK && (1.77)\\
&+ \int_{0}^{F_t} 2\frac{1-\log\frac{K}{F_t}}{K^2}\mathbb{E}_t^Q\left\{R_{f,t+1}^{-1}\left(K-F_{t+1}\right)^+\right\} dK && (1.78)\\
&= 2\int_{F_t}^{\infty}\frac{1}{K^2}\left(1-\log\frac{K}{F_t}\right)C_t(K)\,dK + 2\int_{0}^{F_t}\frac{1}{K^2}\left(1-\log\frac{K}{F_t}\right)P_t(K)\,dK, && (1.79)
\end{aligned}
$$

where $C_t(K)$ $(P_t(K))$ represent a call (put) option with the strike price K.

Options obtained from the CME group. over monthly expiration cycles on 30-year U.S. Treasury futures are employed to compute TVIX. The number of options used in computing the TVIX in each expiration cycle is sizable. The median is 40 options and the minimum is 15 options.

1.7.7 Bootstrap used to compute p-values

To derive robust statistical inferences from time series of a limited size, I perform the following bootstrap approach. Considering a regression $y_t = \beta X_{1,t} + \gamma X_{2,t} + u_t$, where $X_{1,t}$ and $X_{2,t}$ are two sets of dummy variables. For $t = 1, \ldots, T$, the interest is the null hypothesis that $\gamma = 0$. The following steps are implemented to compute the bootstrapped P-values:

1. Run OLS regression to compute $\hat{\beta}$ and $\hat{\gamma}$ and derive the residuals $\hat{u}_t$s. Label $\hat{u}_t$s into groups of periods (per-crisis, crisis, regulation periods, etc.).
2. For each trail, randomly draw u_s^b within each sub-period with replacement. For example, the period of Lehman Brothers is defined as 09/2008 - 12/2008. The set contains four months of $\hat{u}_t$s in this period. I randomly draw four observations from this set with replacement to construct the bootstrapped sample.
3. The resulted bootstrapped dataset contains the exact same number of data points as the original dataset in each sub-periods. Reconstruct y_s^b under the null hypothesis such that $y_s^b = \beta X_{1,s}^b + u_s^b$.
4. Run the regression $y_s^b = \beta X_{1,s}^b + \gamma X_{2,t}^b + \epsilon_s^b$ and obtain the OLS t-stats for $\hat{\gamma}$, denoted as t_γ^b.
5. Run steps (2) to (4) $10,000$ times and obtain the bootstrapped distribution of $\{t_\gamma^b\}_{b=1}^{100,000}$.
6. Compare the original t-stats t_γ with the bootstrapped distribution. Count the number of instances that $|t_\gamma|$ ¿ $|t_\gamma|^b$. The associated P-value is this total counts divided by the total number of bootstrapped, i.e., $100,000$.

The sub-period bootstrap is to ensure that the bootstrap sample has the same (small) size of observations when each dummy is one. For example, there are only four observations during the period of Lehman Brothers bankruptcy and the bootstrapped sample always have four observations when the Lehman Brothers dummy is one. Due to the small size, the procedure will generate a bootstrapped heavier-tail distribution of t_γ under the null for the dummy Lehman Brothers.

	Issues #	Trades #	Dealers #
Corporate bonds in TRACE (2002–2015)	106,873	159,307,893	3,815
Remove cancellations and correct corrections based on Dick-Nielsen (2014)	104,857	112,303,179	3,611
Keep only straight bonds (FISD bond type=CDEB or USBN)	30,755	77,146,341	3,441
Exclude transactions in which execution date is before 12/31/2005	24,999	62,552,200	3,441
Exclude transactions in which execution date is within 30 days of offering date	24,796	60,006,571	3,389
Exclude transactions in which dealers trade with affiliated counterparties	24,754	54,731,492	3,389

Table 1.1: **Main steps in data cleaning**
This table documents the steps in cleaning and filtering the academia TRACE data. Reported are the number of unique issues (Cusips), the number of total transactions, and the number of distinct dealers after each step.

Ratings	Variables	Trade type	mean	std	5%	50%	95%
Investment-grade	Average bid-ask spread (bps)	IVT	47.46	24.24	22.54	40.83	100.36
		RPT	22.12	11.30	10.61	19.69	44.78
	Volume ($M)	IVT	16,857	4,986	8,981	17,794	23,506
		RPT	1,304	574	531	1,311	2,308
	Total number of issues	IVT	1,588	475	842	1,695	2,191
		RPT	252	159	77	215	584
	Total number of trades:						
	Volume ¿= $1M	IVT	2,674	867	1,250	2,865	3,903
		RPT	173	87	63	166	322
	$100K ¡= Volume ¡ $1M	IVT	2,401	1,027	764	2,712	3,887
		RPT	115	150	11	54	490
	Volume ¡ $100K	IVT	5,907	2,722	1,905	6,281	9,702
		RPT	302	367	33	162	1,182
High-yield	Average bid-ask spread (bps)	IVT	58.36	11.44	44.29	56.26	80.50
		RPT	28.97	8.71	18.67	27.54	48.44
	Volume ($M)	IVT	9,036	2,637	4,936	9,110	13,347
		RPT	2,664	1,183	1,266	2,376	4,903
	Total number of issues	IVT	752	220	408	736	1112
		RPT	283	104	158	253	463
	Total number of trades:						
	Volume ¿= $1M	IVT	2,405	709	1,363	2,376	3,570
		RPT	553	226	280	528	976
	$100K ¡= Volume ¡ $1M	IVT	1,187	663	451	961	2,324
		RPT	57	37	23	47	129
	Volume ¡ $100K	IVT	2,482	1,696	675	1,628	5,250
		RPT	121	96	31	71	294

Table 1.2: **Summary statistics for monthly average trading activities**

This table reports the summary statistics for inventory transactions (IVT) and riskless principal transactions (RPT) for investment-grade bonds and high-yield bonds. In each month, I compute average bid-ask spreads, total trading volume, total number of unique traded Cusips, and total number of transactions. Reported are mean, standard deviation, and percentiles of these monthly quantities. Only customer-dealer transactions are considered. Average spreads are computed using transactions in which the par value exceeds $100K. The sample period is January 2006 to December 2015.

Panel A: Investment-Grade

$b_{ij}(b_i)$	$j=1$	2	3
$i=1$	.23***		
	(.02)		
2	.08***	.01	.12***
	(.01)	(.01)	(.01)
3	.31**	−1.26***	.95***
	(.15)	(.10)	(.11)

Panel B: High-Yield

$b_{ij}(b_i)$	$j=1$	2	3
$i=1$	.14***		
	(.01)		
2	.05***	.05***	.09***
	(.01)	(.01)	(.01)
3	.40**	−1.61***	1.20***
	(.19)	(.12)	(.14)

Table 1.3: **Estimates of the B matrix**
This table presents estimates for the $\mathbf{B}$ matrix in the system $\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i y_{t-i}} + \mathbf{B}\epsilon_\mathbf{t}$ for investment-grade bonds and high-yield bonds, respectively, where $\mathbf{y_t}$ is a vector that contains three variables: $\log\left(\text{Spread}_t^{\text{RPT}}\right)$, $\log\left(\text{Spread}_t^{\text{IVT}}\right)$, and $\%\text{RPT}_t$. $\log\left(\text{Spread}_t^{\text{RPT}}\right)$ is the monthly average bid-ask spread computed based on riskless principal transactions. $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ is the monthly average bid-ask spread computed based on inventory transactions. $\%\text{RPT}_t$ is the ratio of customer-dealer-customer riskless principal transaction volume over total customer-dealer volume. I impose the following three restrictions: $b_{1,2} = 0$, $b_{1,3} = 0$, and $b_{3,1} + b_{3,2} + b_3 = 0$. Data from 2006 to 2015 are used to estimate the system. b_{ij} is expressed as b_i when $i = j$. Associated standard errors are reported in parentheses. Asterisks denote significance levels (*** = 1%, ** = 5%, and * = 10%).

	Panel A: IG			Panel B: HY		
	$\hat{\epsilon}_t^H$	$\hat{\epsilon}_t^L$	$\hat{\epsilon}_t^D$	$\hat{\epsilon}_t^H$	$\hat{\epsilon}_t^L$	$\hat{\epsilon}_t^D$
Pre-Crisis	0.17	-0.21	-0.77***	-0.20	-0.15	-1.04*
Jan 2006–Jun 2007	[0.47]	[0.18]	[0.01]	[0.41]	[0.53]	[0.07]
2008 Crisis	0.87*	-0.93*	0.12	0.79***	-0.76***	-0.22
Jul 2007–Apr 2009	[0.08]	[0.09]	[0.77]	[0.01]	[0.01]	[0.30]
Lehman Brothers	0.74	0.76	1.90**	0.94***	1.48***	2.03***
Aug 2008–Dec 2008	[0.13]	[0.15]	[0.03]	[0.01]	[0.00]	[0.00]
Post-Crisis	0.05	0.49**	-0.31*	-0.26	0.45	0.15
May 2009–May 2012	[0.84]	[0.01]	[0.09]	[0.18]	[0.16]	[0.45]
Basel Accords 2.5	-0.11	0.50*	-0.17	-0.19	0.31	0.46*
Jun 2012–Jun 2013	[0.76]	[0.06]	[0.37]	[0.38]	[0.49]	[0.08]
Basel Accords III	-0.34**	0.47	-0.13	-0.60*	-0.33	-0.04
Jul 2013–Mar 2014	[0.04]	[0.11]	[0.60]	[0.09]	[0.27]	[0.90]
Volcker Rule	-1.08***	-0.44***	0.65***	-0.10	-0.31	-0.14
Apr 2014–Dec 2015	[0.00]	[0.01]	[0.00]	[0.72]	[0.39]	[0.58]
N	116	116	116	115	115	115
R^2_{adj}	0.46	0.29	0.36	0.28	0.22	0.31

Table 1.4: **Historical evolution of structural shocks**
This table presents the analysis for movements of structural shocks during the stressed periods. I regress each series of structural shocks $\hat{\epsilon}_t^H$, $\hat{\epsilon}_t^L$, and $\hat{\epsilon}_t^D$ (accumulated over three months) on a set of dummies variables indicating the pre-crisis period, the 2008 financial crisis, a nested period of Lehman Brothers bankruptcy during the financial crisis, the post-crisis period, the implementation of Basel Accords 2.5 and III, and the Volcker Rule. Reported are the coefficients in front of the dummies, total number of observations, and adjusted-R^2s. Constant term is omitted in these regressions. In squared brackets are bootstrapped P-values. Asterisks denote significance levels ($^{***} = 1\%$, $^{**} = 5\%$, and $^{*} = 10\%$).

Panel A: IG & VIX

	$l=1$	$l=3$	$l=6$	$l=12$
Const	20.81***	20.93***	21.13***	21.65***
	(2.03)	(1.78)	(1.52)	(1.35)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{H}_{s-1\to s}$	2.32***	3.17***	4.62***	5.11***
	(0.79)	(0.99)	(1.20)	(1.35)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{L}_{s-1\to s}$	0.72	1.05	1.61	0.29
	(0.45)	(0.79)	(0.99)	(1.38)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{D}_{s-1\to s}$	2.77***	3.61***	3.91***	2.66**
	(0.84)	(0.89)	(1.04)	(1.26)
N	118	116	113	107
R^2_{adj}	0.16	0.31	0.40	0.37

Panel B: IG & TVIX

	$l=1$	$l=3$	$l=6$	$l=12$
Const	11.51***	11.58***	11.72***	11.98***
	(1.01)	(0.88)	(0.73)	(0.66)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{H}_{s-1\to s}$	0.43	0.96**	1.68***	2.08***
	(0.28)	(0.48)	(0.56)	(0.60)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{L}_{s-1\to s}$	0.12	0.46	0.73*	0.20
	(0.26)	(0.35)	(0.39)	(0.74)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{D}_{s-1\to s}$	0.95***	1.85***	2.43***	1.90***
	(0.25)	(0.35)	(0.54)	(0.62)
N	117	115	112	106
R^2_{adj}	0.06	0.25	0.42	0.38

Panel C: HY & VIX

	$l=1$	$l=3$	$l=6$	$l=12$
Const	20.89***	21.00***	21.21***	21.71***
	(2.02)	(1.60)	(1.29)	(1.11)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{H}_{s-1\to s}$	2.46*	4.02**	4.89***	6.11***
	(1.41)	(1.65)	(1.57)	(1.34)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{L}_{s-1\to s}$	-0.63	1.10	2.74***	2.20***
	(0.47)	(0.75)	(1.03)	(0.69)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{D}_{s-1\to s}$	1.88*	2.83**	1.89**	-0.41
	(1.03)	(1.26)	(0.85)	(0.92)
N	117	115	112	106
R^2_{adj}	0.12	0.30	0.41	0.37

Panel D: HY & TVIX

	$l=1$	$l=3$	$l=6$	$l=12$
Const	11.54***	11.62***	11.76***	12.03***
	(0.98)	(0.75)	(0.61)	(0.62)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{H}_{s-1\to s}$	0.87	2.08***	2.40***	2.55***
	(0.53)	(0.69)	(0.68)	(0.56)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{L}_{s-1\to s}$	-0.25	0.37	0.78*	0.70
	(0.21)	(0.34)	(0.46)	(0.53)
$\sum_{t-l+1}^{t} \hat{\epsilon}^{D}_{s-1\to s}$	0.65	1.34**	1.31***	0.49
	(0.49)	(0.58)	(0.51)	(0.62)
N	116	114	111	105
R^2_{adj}	0.06	0.32	0.47	0.35

Table 1.5: **How do structural shocks relate to market conditions?**
This table reports the results from the following regressions:

$$\mathrm{y}_{t+1} = \text{Constant} + \beta_1 \sum_{s=t-l-1}^{t} \hat{\epsilon}^{H}_{s-1\to s} + \beta_2 \sum_{s=t-l-1}^{t} \hat{\epsilon}^{L}_{t-1\to t} + \beta_3 \sum_{s=t-l-1}^{t} \hat{\epsilon}^{D}_{t-1\to t} + u_t,$$

where y_{t+1} is $\mathrm{VIX_{t+1}}$ and $\mathrm{TVIX_{t+1}}$. Regressors are the rolling sums of estimated structural shocks obtained from the system $\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i y_{t-i}} + \mathbf{B}\epsilon_\mathbf{t}$. The backward rolling window is $l = 1$, 3, 6, and 12 months. The three regressors are normalized to zero mean and unit standard deviation. Sample period is from 2006 to 2015. Standard errors are based on [54] with 12 lags. Asterisks denote significance levels ($^{***} = 1\%$, $^{**} = 5\%$, and $^{*} = 10\%$).

Dep Var: Δ Yield Spread (bps)	(1) Aaa	(2) Aaa	(3) Aa	(4) Aa	(5) A	(6) A	(7) Baa	(8) Baa	(9) Ba	(10) Ba	(11) B	(12) B
Hegers_t		-1.10		-0.70		1.38		4.62**		11.26*		13.06*
		(1.32)		(1.21)		(1.61)		(2.32)		(5.74)		(7.46)
$\text{Customer Liquidity}_t$		2.30*		1.98*		2.33*		1.70		4.31		5.02
		(1.16)		(1.06)		(1.38)		(1.93)		(4.66)		(6.17)
Dealers_t		1.71		1.87		4.62**		8.42***		14.93**		24.44***
		(1.34)		(1.34)		(1.84)		(2.72)		(6.20)		(8.61)
$\text{R}_{i,t}$	0.05	0.08	0.09	0.09	-0.08	-0.08	-0.62***	-0.60***	-1.56***	-1.50***	-2.41***	-2.34***
	(0.17)	(0.17)	(0.13)	(0.13)	(0.07)	(0.06)	(0.13)	(0.14)	(0.33)	(0.30)	(0.37)	(0.35)
ΔRF_t	-0.58***	-0.56***	-0.51***	-0.48***	-0.59***	-0.50***	-0.80***	-0.61***	-1.54***	-1.17***	-1.63***	-1.09**
	(0.09)	(0.09)	(0.08)	(0.08)	(0.13)	(0.11)	(0.18)	(0.14)	(0.44)	(0.35)	(0.51)	(0.42)
ΔRF_t^2	-4.79***	-5.27***	-4.32**	-4.85***	-4.84*	-5.85**	-3.83	-5.15	-1.67	-4.61	7.16	2.65
	(1.76)	(1.75)	(1.71)	(1.74)	(2.73)	(2.59)	(4.19)	(3.74)	(10.06)	(9.52)	(12.88)	(11.45)
ΔSLOPE_t	0.98***	0.95***	0.91***	0.87***	0.99***	0.87***	1.15***	0.94***	1.47***	1.04***	1.70***	1.04**
	(0.11)	(0.11)	(0.10)	(0.10)	(0.15)	(0.13)	(0.23)	(0.17)	(0.40)	(0.31)	(0.55)	(0.42)
RM_t	-0.40	-0.54	-0.88	-0.97	-1.03	-0.94	-1.67	-1.16	-1.57	-0.87	-6.39	-5.17
	(0.58)	(0.59)	(0.70)	(0.69)	(1.00)	(0.91)	(1.42)	(1.23)	(2.91)	(2.64)	(4.10)	(3.68)
ΔVIX_t	0.44	0.33	0.49	0.37	0.78	0.68	0.43	0.48	1.55	1.34	-0.87	-1.04
	(0.39)	(0.37)	(0.44)	(0.43)	(0.52)	(0.49)	(0.55)	(0.53)	(0.95)	(1.14)	(1.58)	(1.93)
ΔJump_t	-0.53	-0.52	-0.20	-0.23	0.16	-0.09	1.49	0.90	4.42	3.29	6.86	5.09
	(0.61)	(0.59)	(0.75)	(0.72)	(1.03)	(0.93)	(1.51)	(1.27)	(3.03)	(2.53)	(4.41)	(3.74)
Observations	2086	2086	8748	8748	47176	47176	67633	67633	23078	23078	20294	20294
Adjusted R^2	0.269	0.278	0.296	0.303	0.288	0.306	0.251	0.279	0.233	0.259	0.230	0.251
Cusip FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 1.6: **Yield spread changes and measures of liquidity**
This table reports the results of the regression: $\Delta\text{YS}_{i,t} = \alpha_i + \beta_1\text{Customer Liquidity}_t + \beta_2\text{Dealer}_t + \beta_3\text{Hedger}_t + \beta_4\mathbf{\Delta X_{i,t}} + \beta_5\mathbf{\Delta X_t} + u_{i,t}$, where $\text{Customer Liquidity}_\text{t}$, Dealer_t, and Hedger_t are measures capturing the willingness of liquidity provision for customer liquidity providers, dealers, and hedgers. $\mathbf{\Delta X_{i,t}}$ includes the change in the leverage of the issuers and the market return of the issuers' common equity. $\mathbf{\Delta X_t}$ includes the change in risk-free rate, the squared change in risk-free rate, the change in the slope of the yield curve, the change in the VIX, and the S&P 500 return. Reported in parentheses are standard errors double-clustered at Cusip and year-month level. The sample starts in 2006 and ends in 2015. Asterisks denote significance levels (*** = 1%, ** = 5%, and * = 10%).

Panel A: Maturity: Long (over eight years)

Dep Var: Δ Yield Spread (bps)	(1) Aaa	(2) Aa	(3) A	(4) Baa	(5) Ba	(6) B
Hedgers_t	-1.07	-0.90	0.91	3.15	6.88	8.33
	(1.68)	(1.48)	(1.83)	(2.25)	(4.32)	(5.37)
Customer Liquidity_t	2.49*	2.31**	2.57*	2.27	5.05	4.67
	(1.21)	(1.14)	(1.37)	(1.75)	(3.38)	(4.01)
Dealers_t	1.44	1.55	3.41*	6.37**	9.03**	14.06**
	(1.58)	(1.56)	(2.04)	(2.46)	(4.06)	(5.69)
Observations	1144	3533	20388	27108	6530	4342
Adjusted R^2	0.519	0.542	0.465	0.393	0.327	0.324
Cusip FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Maturity: Short (less than eight years)

Dep Var: Δ Yield Spread (bps)	(1) Aaa	(2) Aa	(3) A	(4) Baa	(5) Ba	(6) B
Hedgers_t	-1.31	-0.69	1.42	5.65**	14.01**	15.49*
	(1.47)	(1.19)	(1.63)	(2.55)	(6.59)	(8.31)
Customer Liquidity_t	1.87	1.94*	2.29	1.59	4.75	5.93
	(1.30)	(1.15)	(1.51)	(2.13)	(5.32)	(6.77)
Dealers_t	1.65	1.99	5.30***	9.71***	16.55**	26.15***
	(1.60)	(1.30)	(1.85)	(2.97)	(7.03)	(9.55)
Observations	933	5160	26584	40259	16512	15924
Adjusted R^2	0.092	0.184	0.218	0.229	0.257	0.253
Cusip FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Table 1.7: **Yield spread changes and measures of liquidity: Different maturity**

This table reports the results of the regression: $\Delta \text{YS}_{i,t} = \alpha_i + \beta_1 \text{Customer Liquidity}_t + \beta_2 \text{Dealer}_t + \beta_3 \text{Hedger}_t + \beta_\mathbf{4} \mathbf{\Delta X_{i,t}} + \beta_\mathbf{5} \mathbf{\Delta X_t} + u_{i,t}$, where $\text{Customer Liquidity}_\text{t}$, Dealer_t, and Hedger_t are measures capturing the willingness of liquidity provision for customer liquidity providers, dealers, and hedgers. $\mathbf{\Delta X_{i,t}}$ includes the change in the leverage of the issuers and the market return of the issuers' common equity. $\mathbf{\Delta X_t}$ includes the change in risk-free rate, the squared change in risk-free rate, the change in the slope of the yield curve, the change in the VIX, and the S&P 500 return. Reported in parentheses are standard errors double-clustered at Cusip and year-month level. The sample starts in 2006 and ends in 2015. Asterisks denote significance levels (*** = 1%, ** = 5%, and * = 10%).

Panel A: Investment-grade

	$l = 1$	$l = 3$	$l = 6$	$l = 12$
$\sum_{t-l+1}^{t} \hat{\epsilon}^{H}_{s-1\rightarrow s}$	0.185	0.433	0.472	0.569
$\sum_{t-l+1}^{t} \hat{\epsilon}^{L}_{s-1\rightarrow s}$	0.025	0.074	0.020	-0.274
$\sum_{t-l+1}^{t} \hat{\epsilon}^{D}_{s-1\rightarrow s}$	0.182	0.411	0.469	0.303

Panel B: High-yield

	$l = 1$	$l = 3$	$l = 6$	$l = 12$
$\sum_{t-l+1}^{t} \hat{\epsilon}^{H}_{s-1\rightarrow s}$	0.239	0.433	0.528	0.617
$\sum_{t-l+1}^{t} \hat{\epsilon}^{L}_{s-1\rightarrow s}$	-0.006	-0.017	0.190	-0.012
$\sum_{t-l+1}^{t} \hat{\epsilon}^{D}_{s-1\rightarrow s}$	0.235	0.375	0.489	0.231

Table 1.8: **Correlation between the structural shocks and the VIX**
This table presents the relationship between the structural shocks and the CBOE VIX index for investment-grade bonds and high-yield bonds, respectively. Computed are the correlation coefficients between the VIX and cumulative shocks to hedgers ($\sum_{t-l+1}^{t} \hat{\epsilon}^{H}_{s-1\rightarrow s}$), to liquidity providers ($\sum_{t-l+1}^{t} \hat{\epsilon}^{L}_{s-1\rightarrow s}$), or to dealers $\sum_{t-l+1}^{t} \hat{\epsilon}^{D}_{s-1\rightarrow s}$, over a backward rolling window of $l = 1$, 3, 6, and 12 months. Estimated structural shocks are obtained from the system $\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i y_{t-i}} + \mathbf{B}\epsilon_\mathbf{t}$. Sample period is from 2006 to 2015.

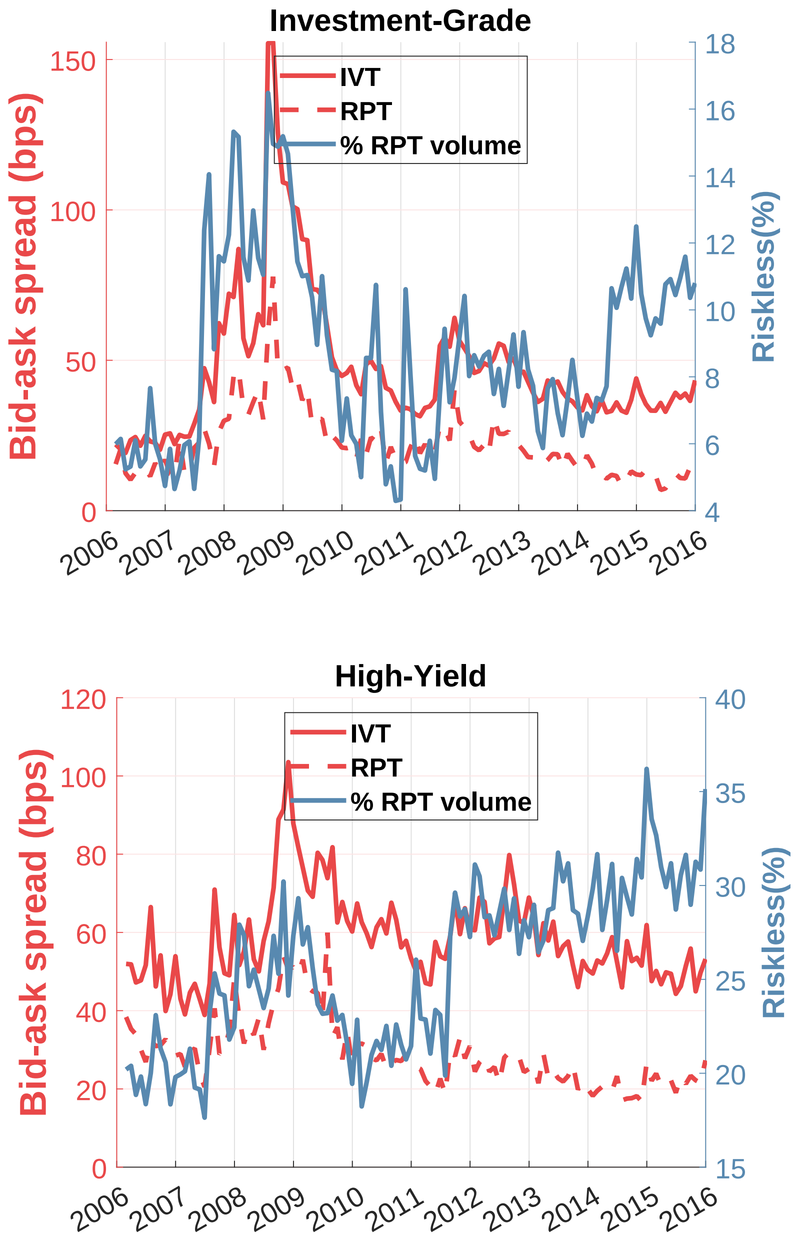

Figure 1.1: **Monthly average bid-ask spreads and fractions of riskless principal trading volumes**
These figures plot the bid-ask spreads for inventory transactions, the bid-ask spreads for riskless principal transactions, and the ratio of customer-dealer-customer riskless principal transaction volume over total customer-dealer volume for investment-grade bonds and high-yield bonds, respectively. The detailed algorithm to classify riskless principal transactions is documented in Appendix 1.7.5. The bid-ask spread charged for every customer-dealer transaction is twice the difference between the dealer-customer price and the benchmark interdealer price. Bid-ask spreads of each type are computed by averaging the spreads for all transactions of each type with par value above (inclusive) $100K.

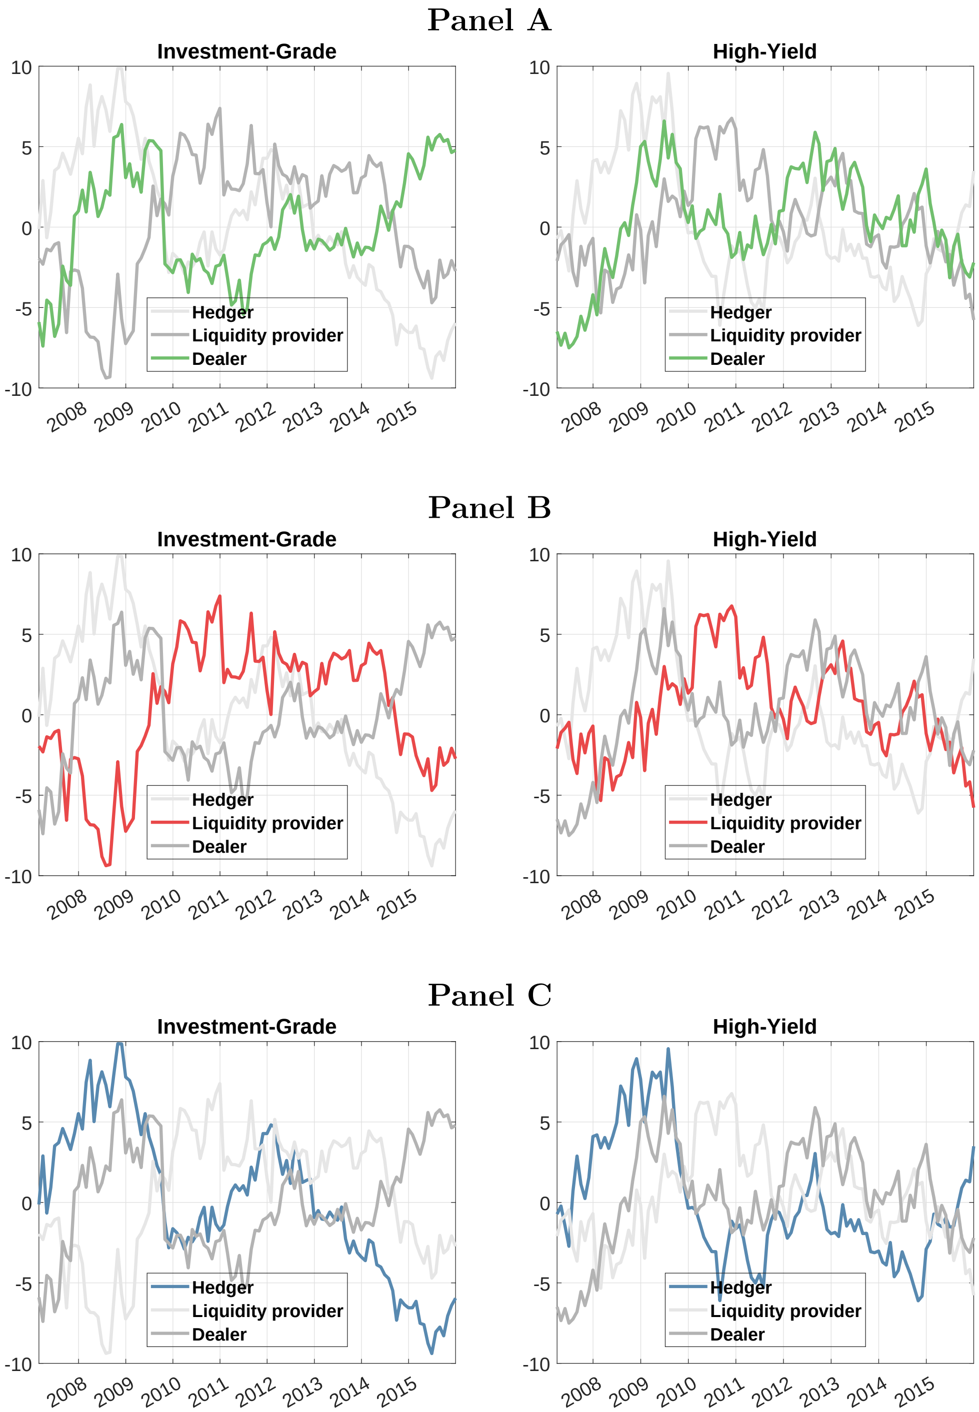

Figure 1.2: **Historical evolution of structural shocks**
These figures present the forward cumulated structural shocks using a rolling window of 12 months $\sum_{t}^{t+11} \hat{\epsilon}_{\mathbf{t}}$. $\epsilon_{\mathbf{t}} = \left[\epsilon_t^H \ \epsilon_t^L \ \epsilon_t^d\right]^T$ is a vector of unit variance orthogonal shocks to hedgers (ϵ_t^H), customer liquidity providers (ϵ_t^L), and dealers (ϵ_t^d). Each panel highlights the shocks to one of the three market participants. They are estimated from the system $\mathbf{y_t} = \mathbf{A_0} + \sum_i \mathbf{A_i y_{t-i}} + \mathbf{B}\epsilon_{\mathbf{t}}$ for investment-grade bonds and high-yield bonds, respectively. $\mathbf{y_t}$ is a vector that contains three variables: $\log\left(\text{Spread}_t^{\text{RPT}}\right)$, $\log\left(\text{Spread}_t^{\text{IVT}}\right)$, and $\%\text{RPT}_t$. $\log\left(\text{Spread}_t^{\text{RPT}}\right)$ is the monthly average bid-ask spread computed based on riskless principal transactions. $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ is the monthly average bid-ask spread computed based on inventory transactions. $\%\text{RPT}_t$ is the ratio of customer-dealer-customer riskless principal transaction volume over total customer-dealer volume. I impose the following three restrictions: $b_{1,2} = 0$, $b_{1,3} = 0$, and $b_{3,1} + b_{3,2} + b_3 = 0$. Data from 2006 to 2015 are used to estimate the system.

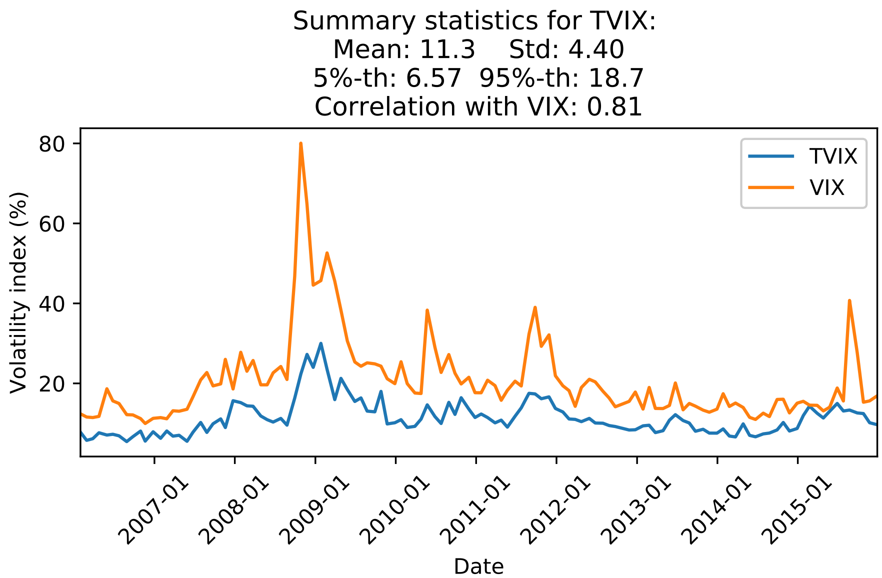

Figure 1.3: **Volatility index for equity and 30-year Treasury bond futures**
This table presents the time series of CBOE VIX and the volatility index for the 30-year Treasury bond futures (TVIX). VIX is obtained from yahoo.com. TVIX is computed based on options on the 30-year Treasury bond futures. Options data are obtained from the CME group. Sample begins in 2006 and ends in 2015.

Panel A: Hedgers | Panel B: Customer liquidity providers | Panel C: Dealers

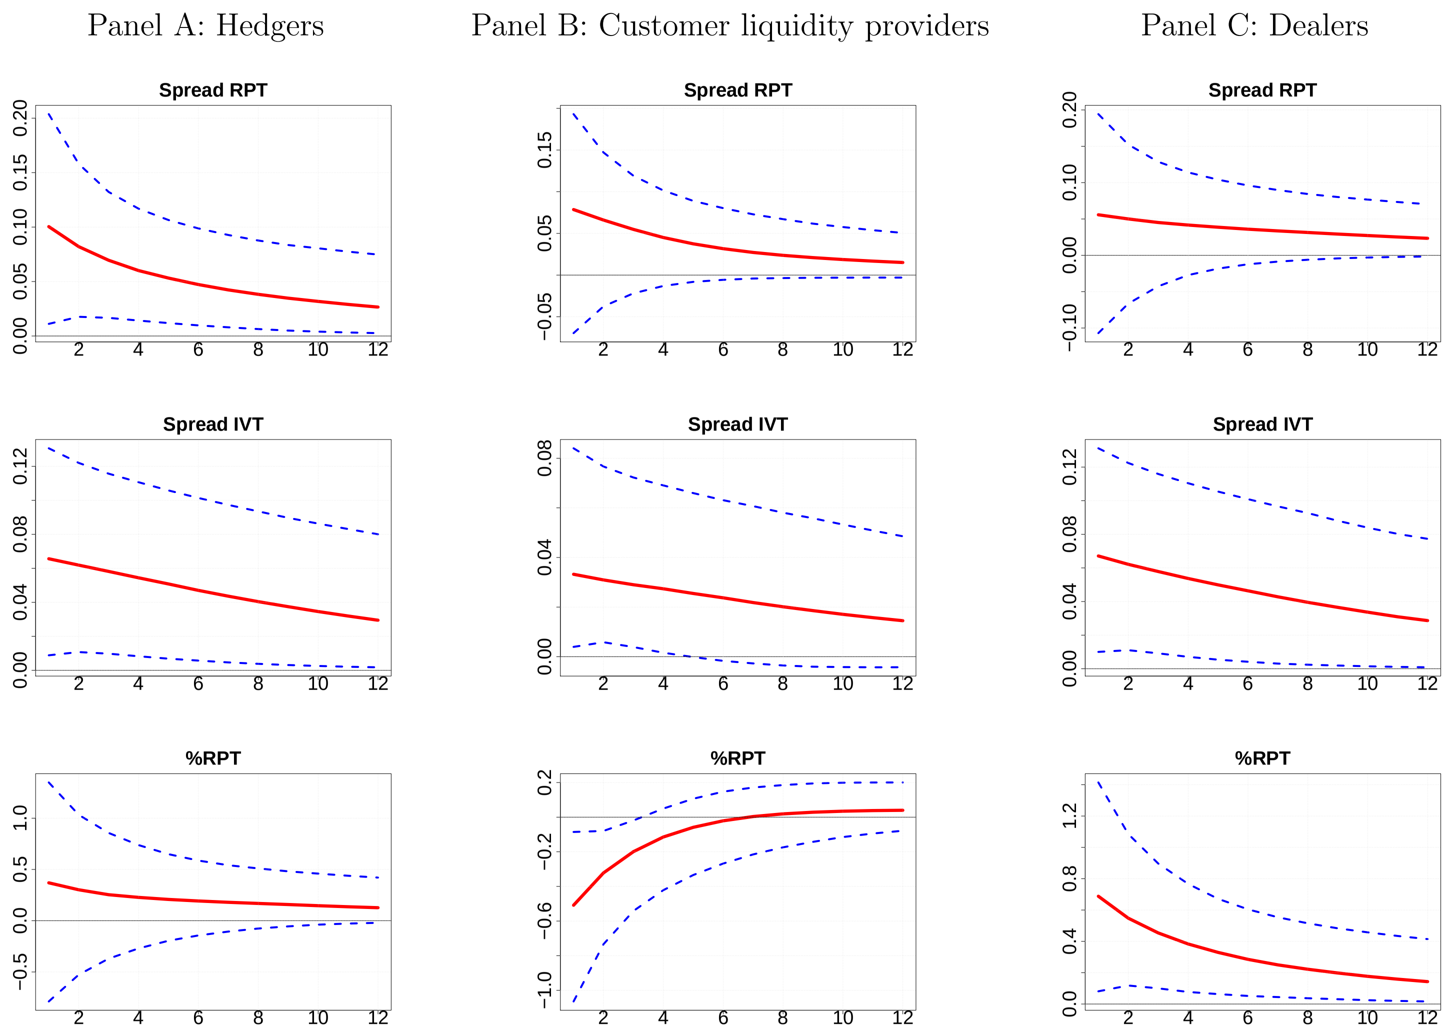

Figure 1.4: **Impulse response functions based on the agnostic method for investment-grade bond markets**
These figures present the impulse response functions of the SVAR model in equation (1.2) for investment-grade bond markets, using the agnostic method. The total number of lags selected is one. For each of the shocks to be recovered, I separately apply two inequality restrictions, up to a horizon of three periods, on candidate impulse response functions. To recover the shocks to hedgers, I impose the restrictions that an increase in ϵ_t^H increases $\log\left(\text{Spread}_t^{\text{RPT}}\right)$ (i.e., $b_1 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,1} \geq 0$). To obtain the shocks to dealers, I apply the restrictions that an increase in ϵ_t^d increases $\%\text{RPT}_t$ (i.e., $b_3 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,3} \geq 0$). Finally, to derive the shocks to customer liquidity providers, I use the restrictions that an increase in ϵ_t^d decreases $\%\text{RPT}_t$ (i.e., $b_2 \leq 0$) and increases $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{3,2} \geq 0$). Plotted are the median (in red) and 5-th and 95-th percentiles (in blue) of all feasible impulse response functions that satisfy the inequality constraints.

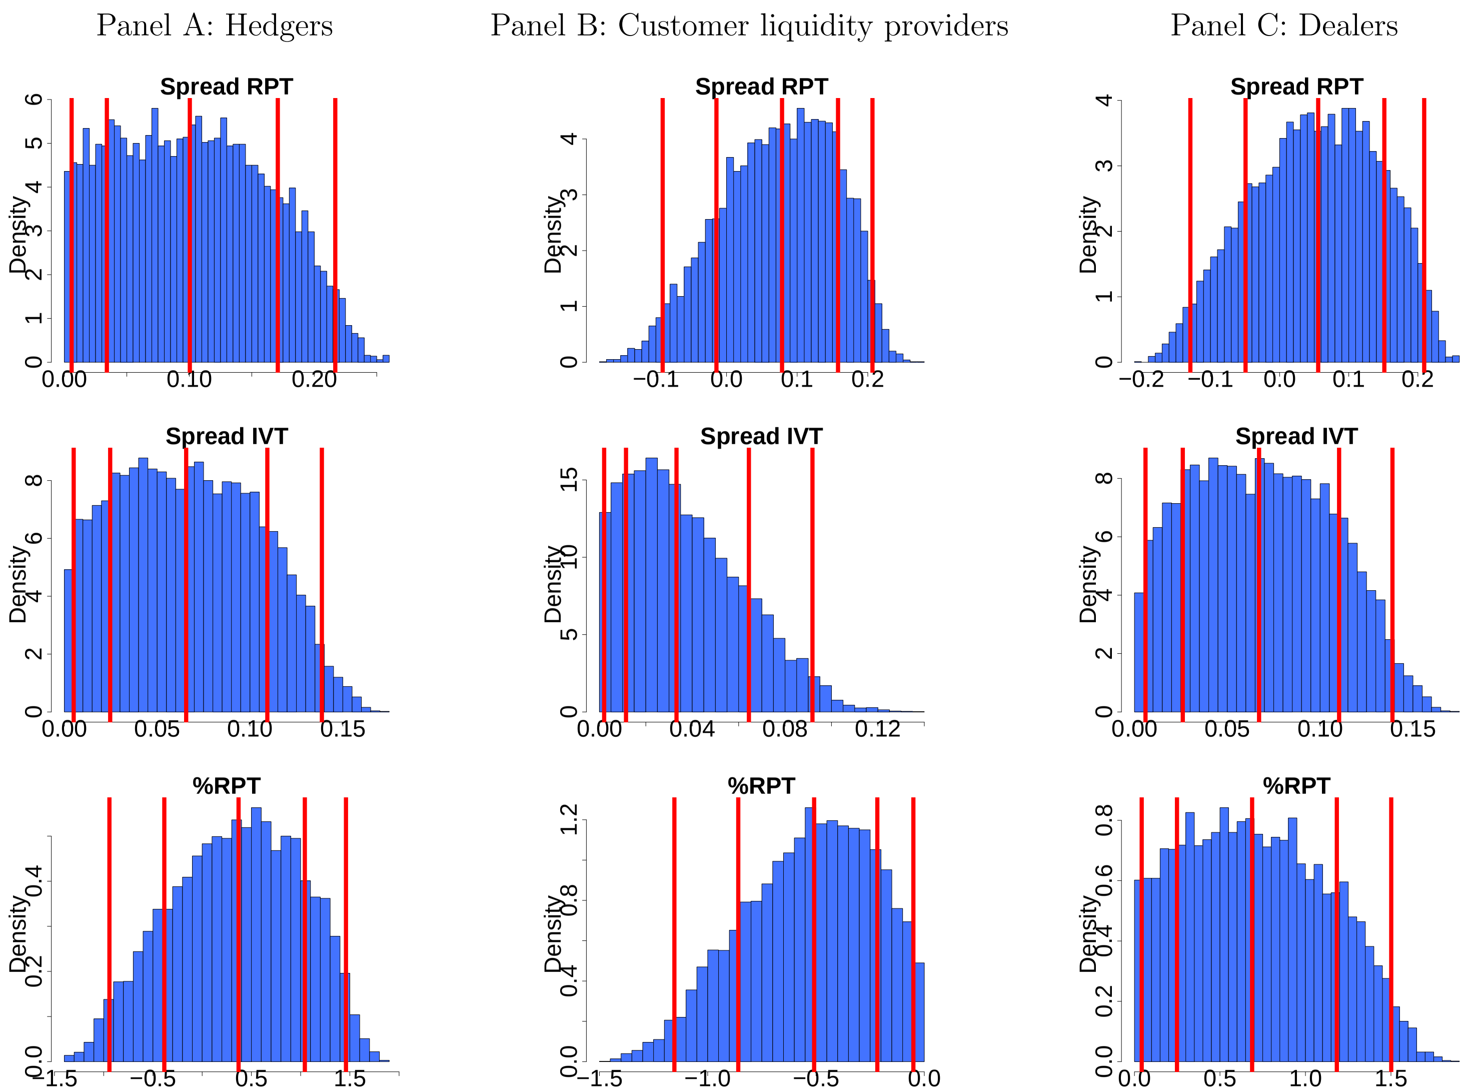

Figure 1.5: **Histograms of the initial impulse responses**
These figures plots the posterior distribution of the initial impulse responses for investment-grade bond markets, using the agnostic method. The total number of lags selected is one. For each of the shocks to be recovered, I separately apply two inequality restrictions, up to a horizon of three periods, on candidate impulse response functions. To recover the shocks to hedgers, I impose the restrictions that an increase in ϵ_t^H increases $\log\left(\text{Spread}_t^{\text{RPT}}\right)$ (i.e., $b_1 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,1} \geq 0$). To obtain the shocks to dealers, I apply the restrictions that an increase in ϵ_t^d increases $\%\text{RPT}_t$ (i.e., $b_3 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,3} \geq 0$). Finally, to derive the shocks to customer liquidity providers, I use the restrictions that an increase in ϵ_t^d decreases $\%\text{RPT}_t$ (i.e., $b_2 \leq 0$) and increases $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{3,2} \geq 0$). Red lines in the plots indicate the 2.5th, 16th, 50th, 84th, and 97.5th percentiles.

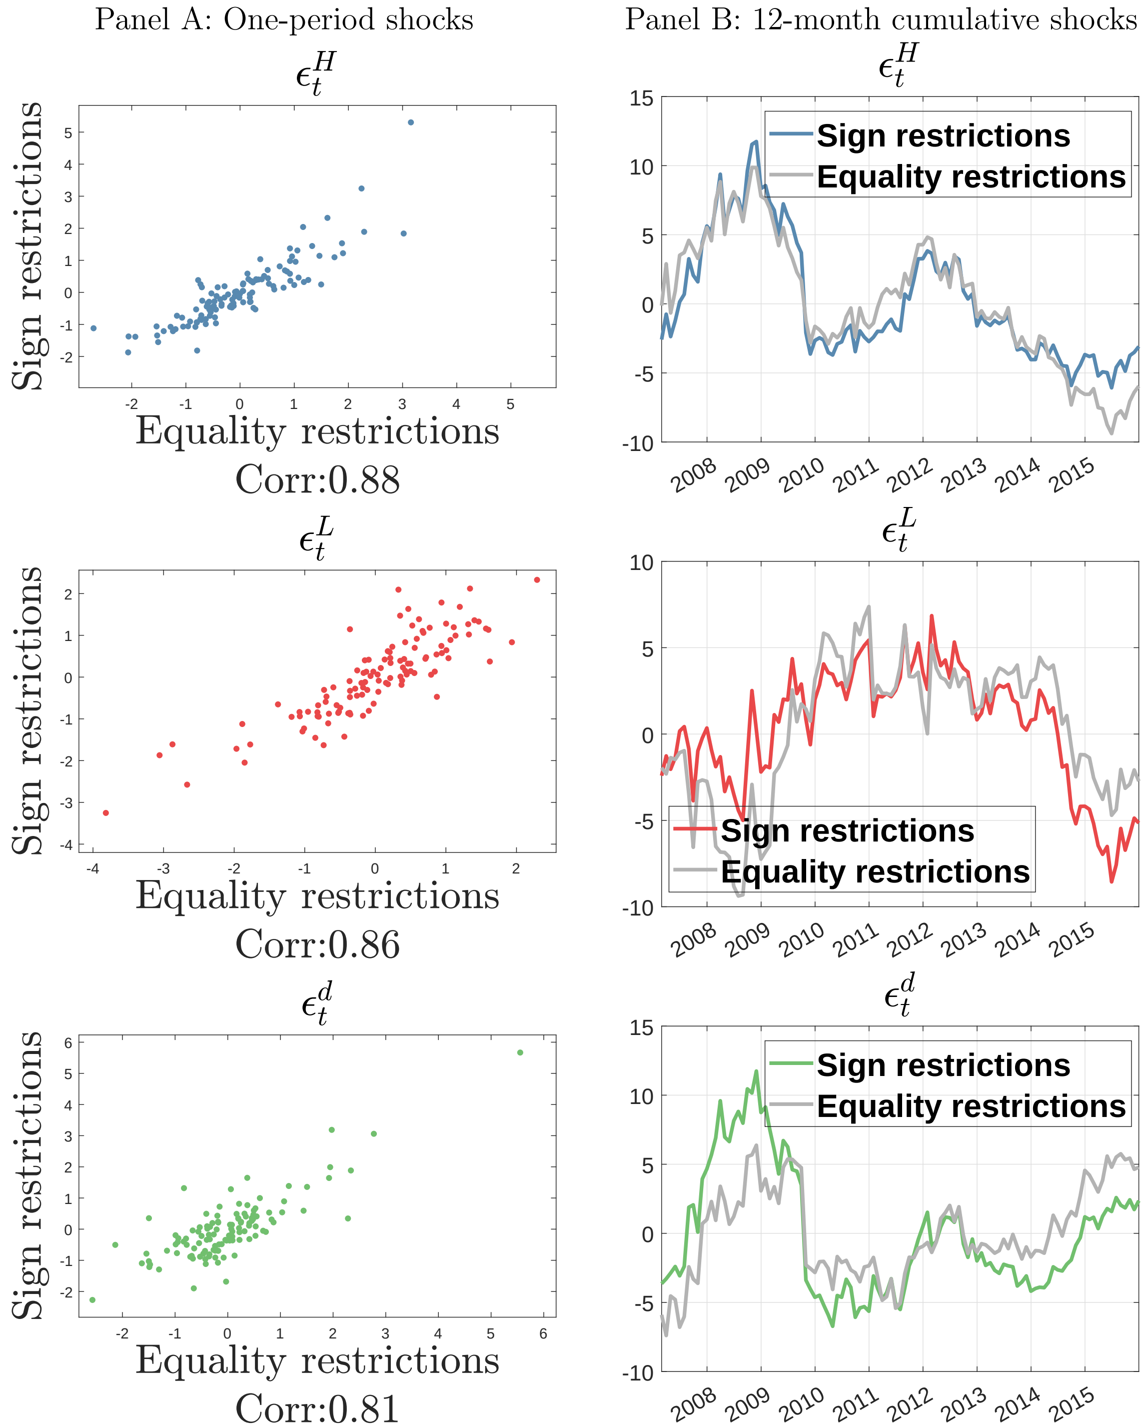

Figure 1.6: **Structural shocks from the agnostic method for investment-grade markets** These plots compare the obtained structural shocks based on the full identification method versus the agnostic method. Panel A presents the scatter plots where each spot represents a pair of the shocks based on the two methods in the same month. Also reported are the correlation coefficients of the shocks generated by the two methods. Panel B exhibits the time series of 12-month rolling sum of the structural shocks. Colored and solid lines are based on the sign restrictions, while gray and dashed lines are based on the equality restrictions.

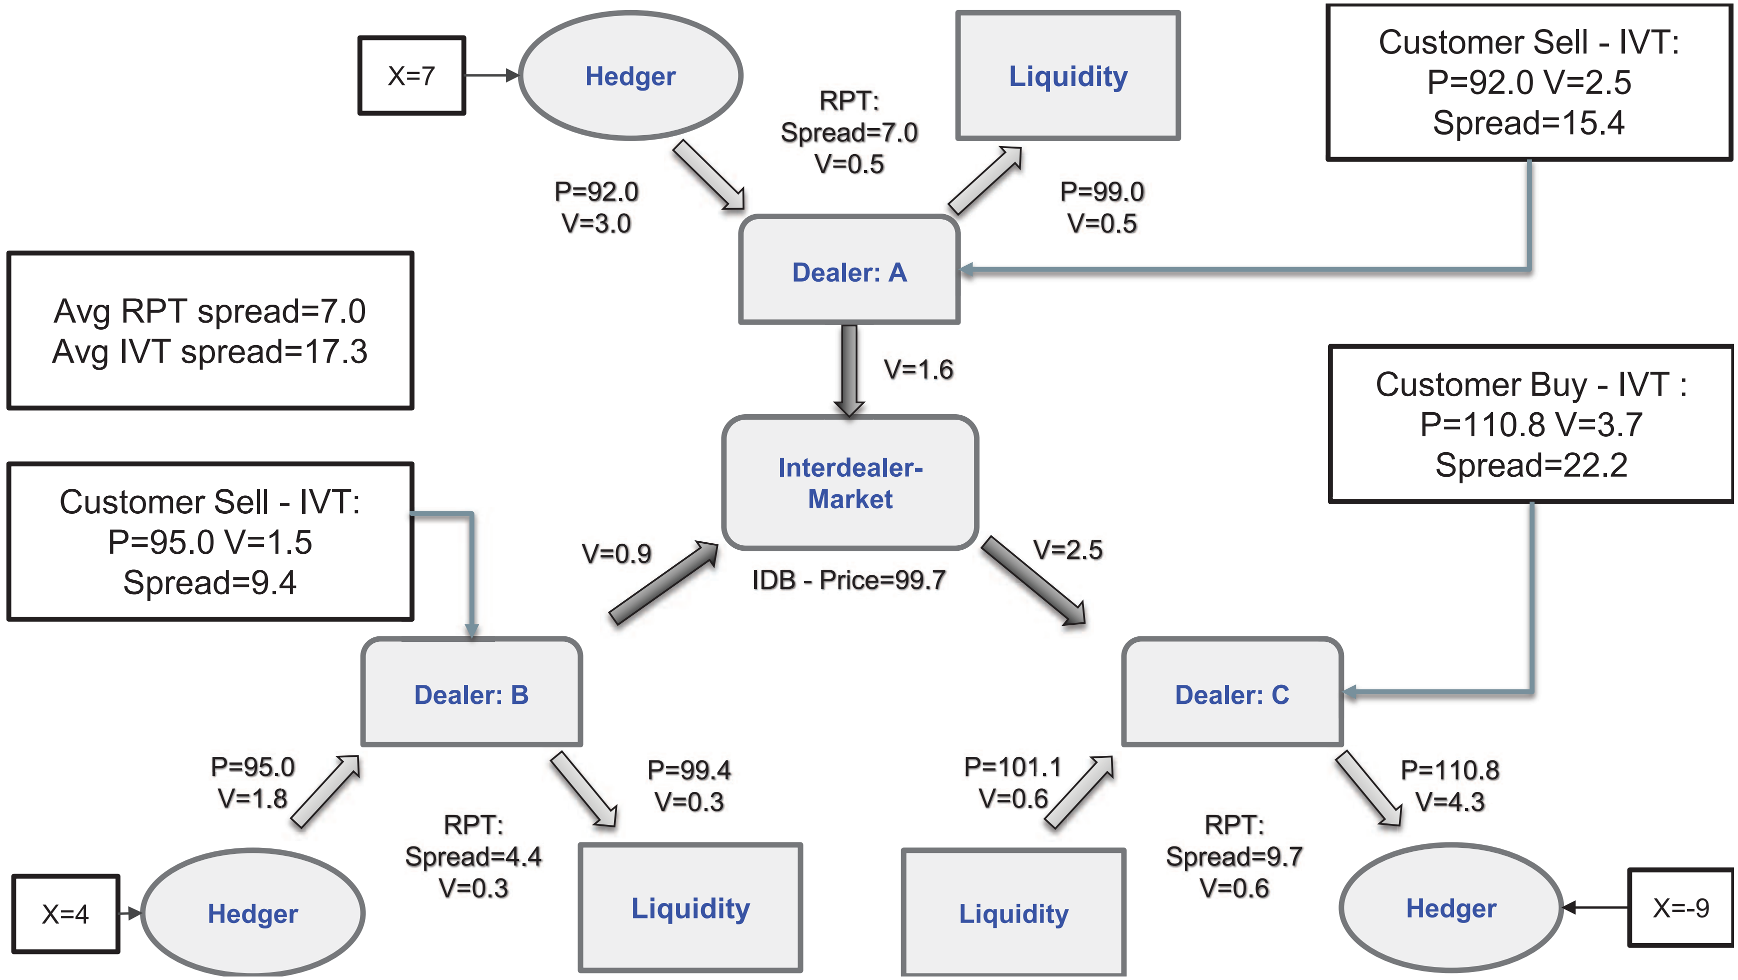

Figure 1.7: **A numerical illustration of the market for three local dealers**
This figure illustrates the modeled market structure. In this example, there are three local dealers. Each dealer has her own customers. Dealers can also trade in the interdealer market. Each dealer has two types of customers: hedgers who receive a liquidity shock from a correlated non-tradable asset of size X. The rectangles named Liquidity stand for customers who do not suffer from liquidity shocks, that is, customer liquidity providers. The arrows indicate the direction of transaction of the asset in equilibrium. For example, Dealer A buys from hedgers, sells to customer liquidity providers, and sells in the interdealer market. Numbers of equilibrium prices (P) and volumes (V) are presented on the side of the arrows, which represents the direction of transactions. Also reported are the total volume and average spreads for inventory trades and riskless principal transactions.

Panel A:

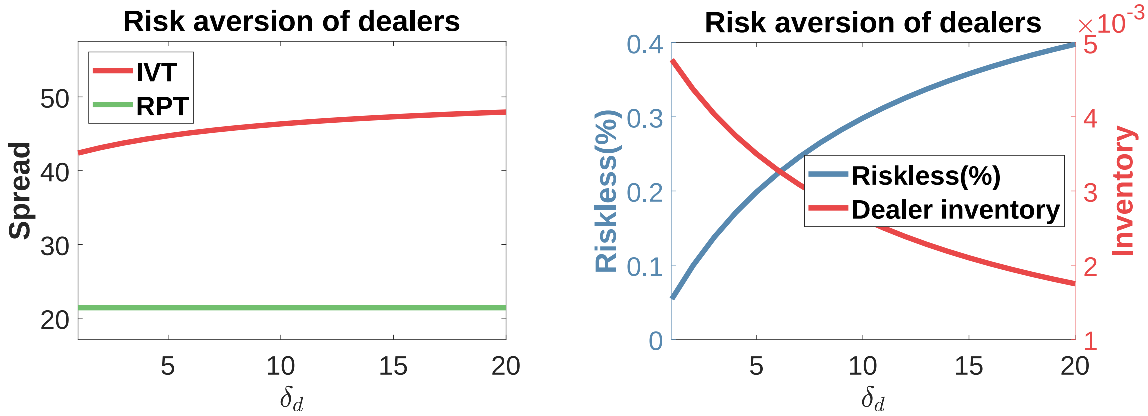

Panel B:

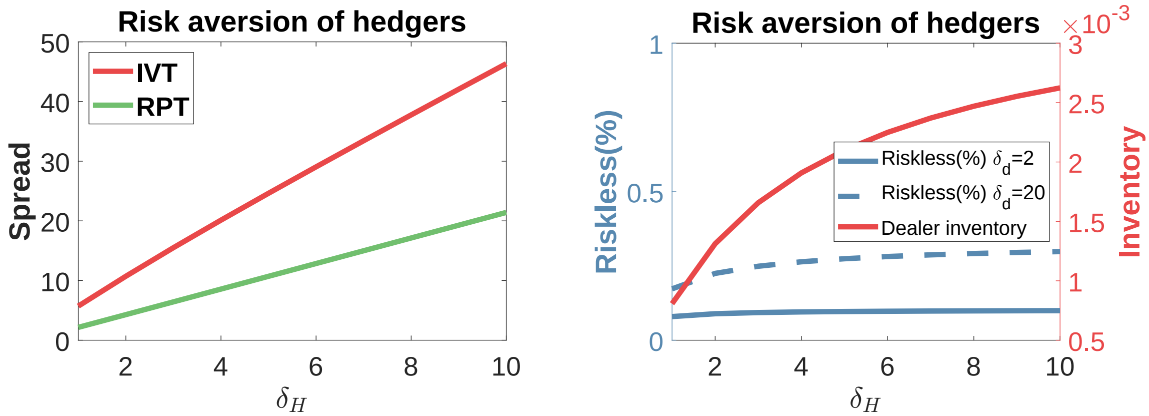

Panel C:

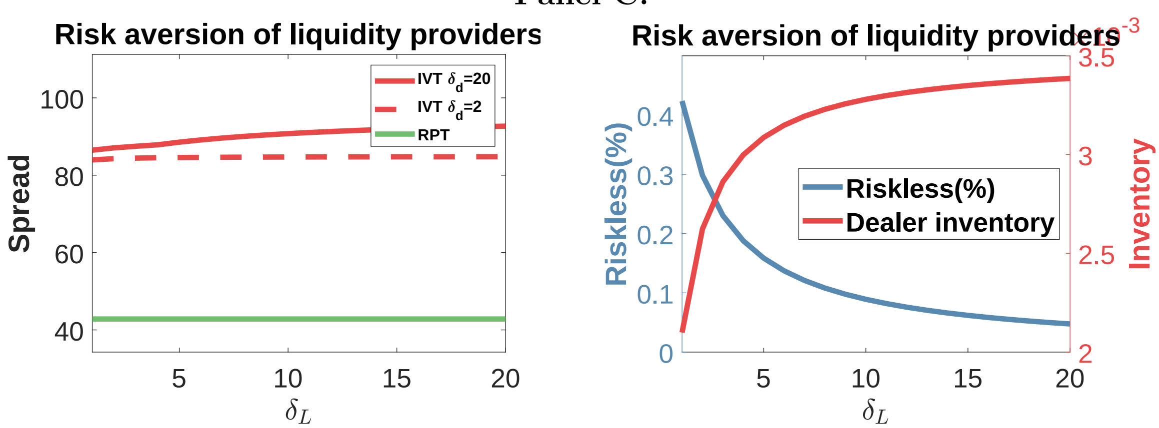

Figure 1.8: **Comparative statics**
These figures report the comparative statics of equilibrium spreads for inventory transactions, spreads for riskless principal transactions, fraction of riskless principal trading volume, and average dealer inventory level. Results are simulated by liquidity shocks $\{X_k\}_{k=1}^{K}$ that are randomly and independently drawn 100,000 times from identical normal distributions $N(0, 5^2)$. In each trail, I compute the corresponding quantities as described in Appendix 1.7.2. Reported are the average of each quantity over the 100,000 draws.

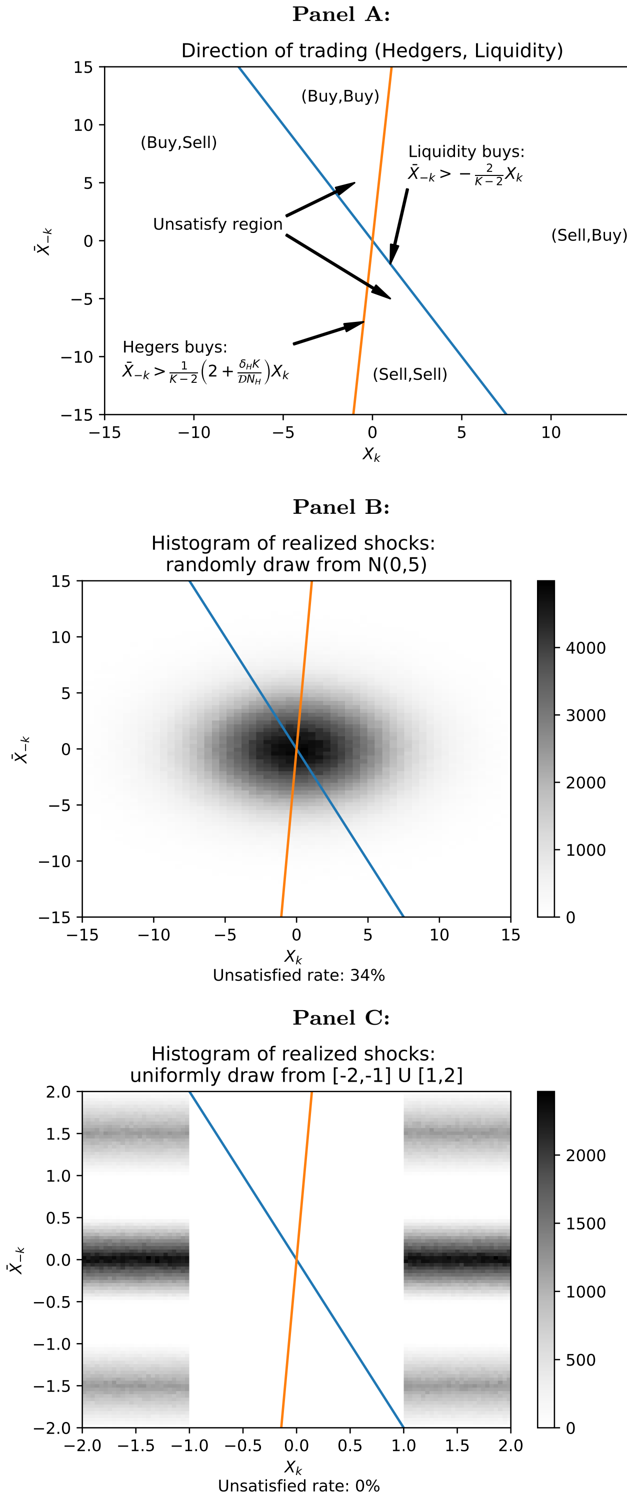

Figure 1.9: **Do hedgers and liquidity providers buy or sell?**
These figures describe the direction of trading for customers. x-axis represents the realization of local shock X_k whereas y-axis represents the average of shocks to other markets $\bar{X}_{-k}$. Hedgers buy if the local shock is small or negative, i.e. above the orange line. Liquidity providers buy if the local shock is large, i.e. above the blue line. Panel A presents the four regions indicating the four type of distinct trading decisions for pairs of local hedgers and liquidity providers. Panel B presents the realization of shocks uniformly drawn from $[-2,-1] \cup [1,2]$, while Panel C presents the realization of shocks from $N(0,5)$. Unsatisfied rate is the fraction of draws where at least one local dealer trades with hedgers and liquidity providers in the same direction.

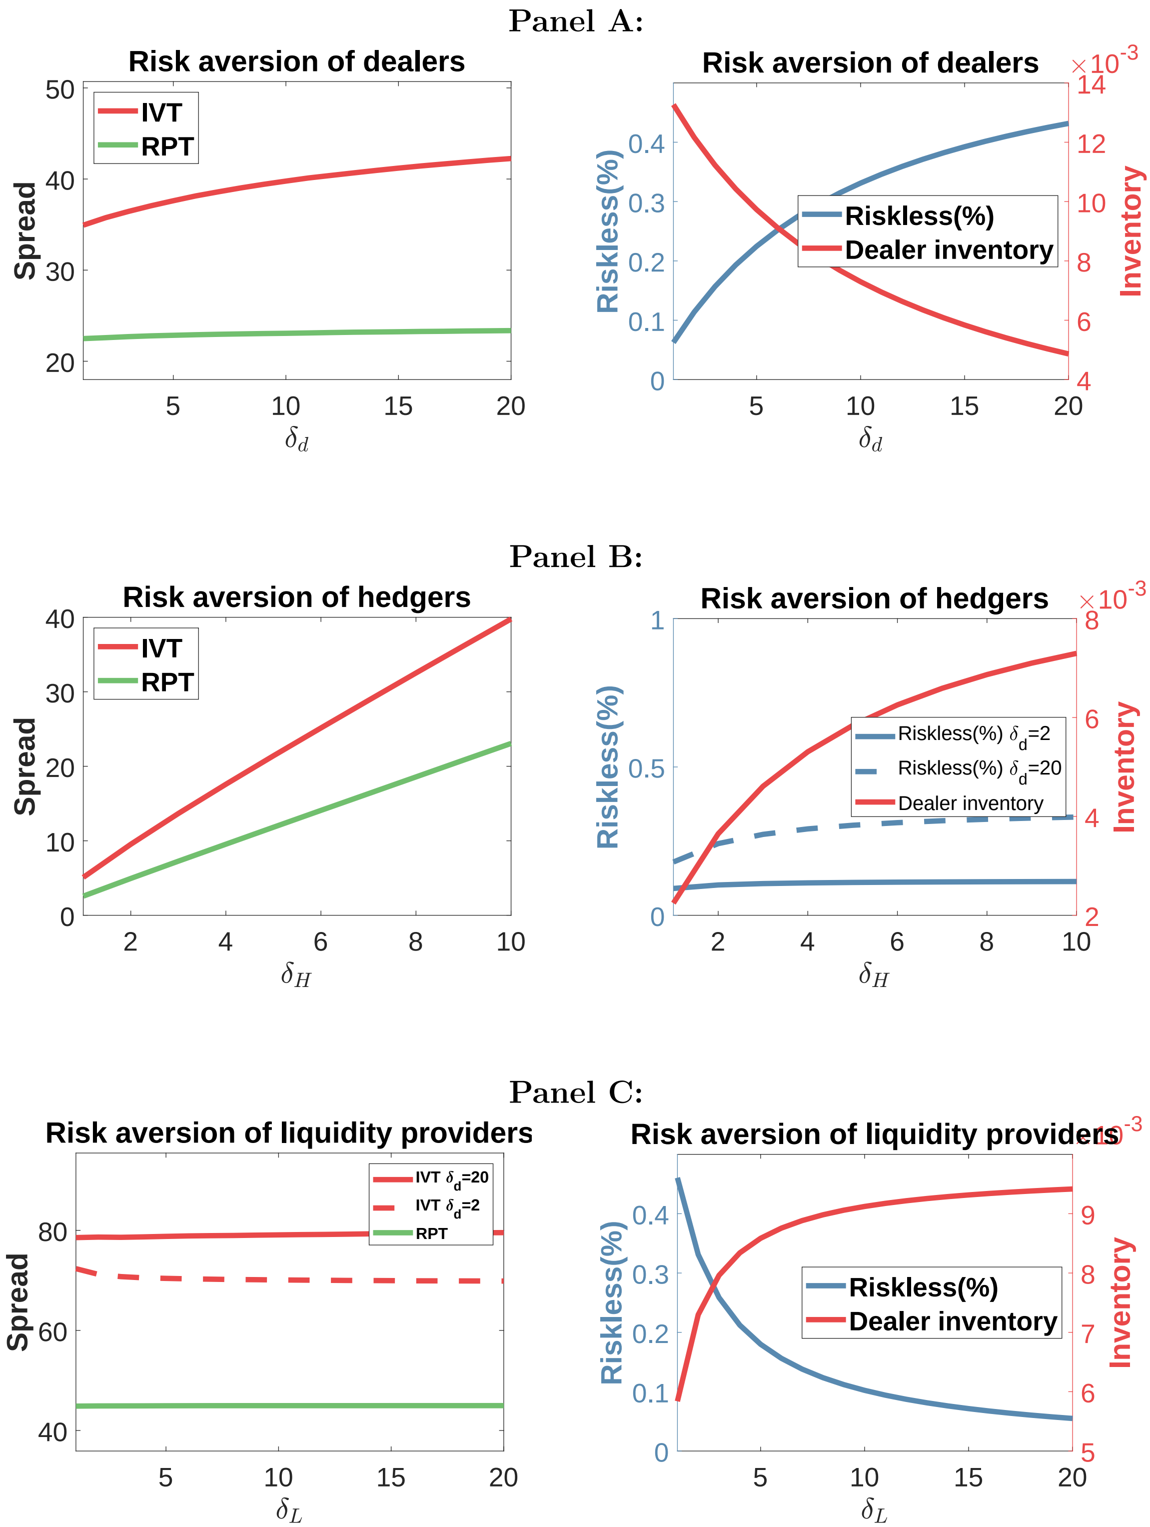

Figure 1.10: **Comparative statics: with trivial liquidity shocks to hedgers**
These figures report the comparative statics of equilibrium spreads for inventory transactions, spreads for riskless principal transactions, fraction of riskless principal trading volume, and average dealer inventory level. Results are simulated by liquidity shocks $\{X_k\}_{k=1}^{K}$ that are randomly and independently drawn 100,000 times from identical normal distributions $N(0, 5^2)$. In each trail, I compute the corresponding quantities as described in Appendix 1.7.2. Reported are the average of each quantity over the 100,000 draws.

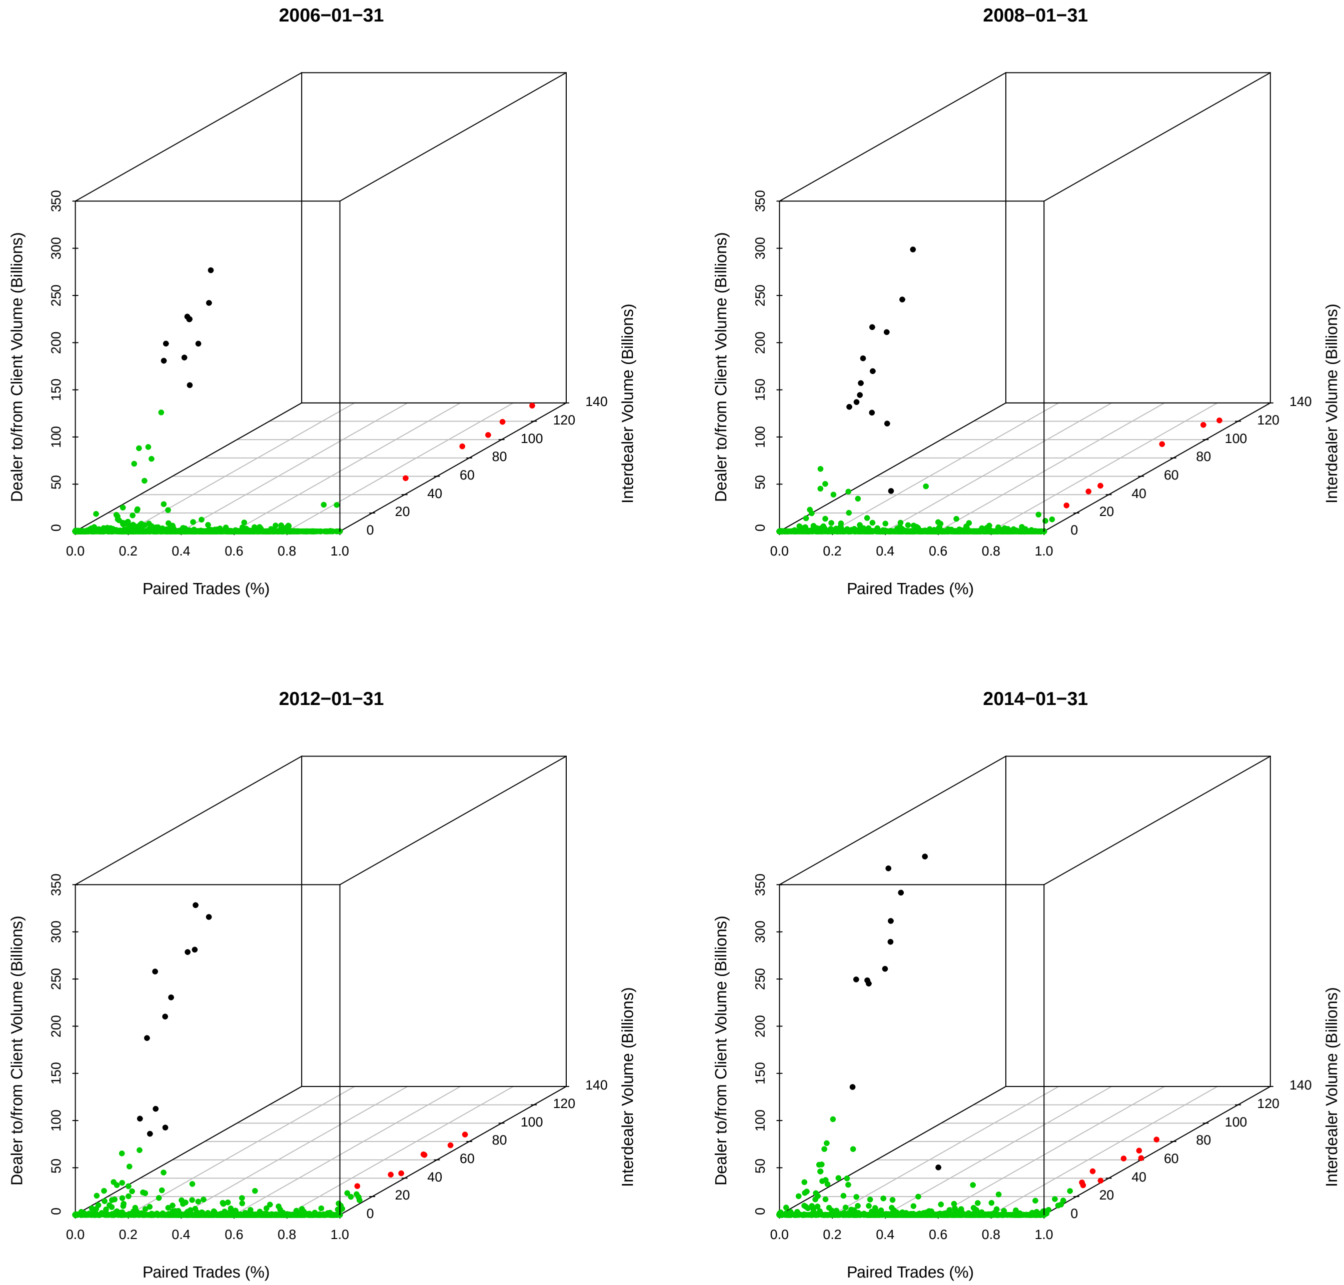

Figure 1.11: **Trading characteristics of dealers in corporate bond market.**
These 3D plots depict the characteristics for all dealers trading corporate bonds based on a rolling 12-month data ending on the stated date. Each dot in the plot represents a dealer. Z-axis is the total trading (the sum of buy and sell) volume between this dealer and its clients, and y-axis is the total trading volume between this dealer and other dealers. X-axis approximates the proportion of riskless principal transactions conducted by this dealer, i.e. the volume of paired trades over the total volume. For each dealer i and bond c on a given day t, consider the total volume that dealer i buys $V_{c,t}^{B,i}$ and sells $V_{c,t}^{S,i}$, and the paired volume is the smaller one of $(V_{c,t}^{B,i}, V_{c,t}^{S,i})$. Then, the paired volumes are summed over all trading days and all traded bonds, and finally I normalize this quantity by the average of buy and sell volume, i.e. $\frac{2\sum_{c,t}\min(V_{c,t}^{B,i},V_{c,t}^{S,i})}{\sum_{c,t}(V_{c,t}^{B,i}+V_{c,t}^{S,i}))}$.

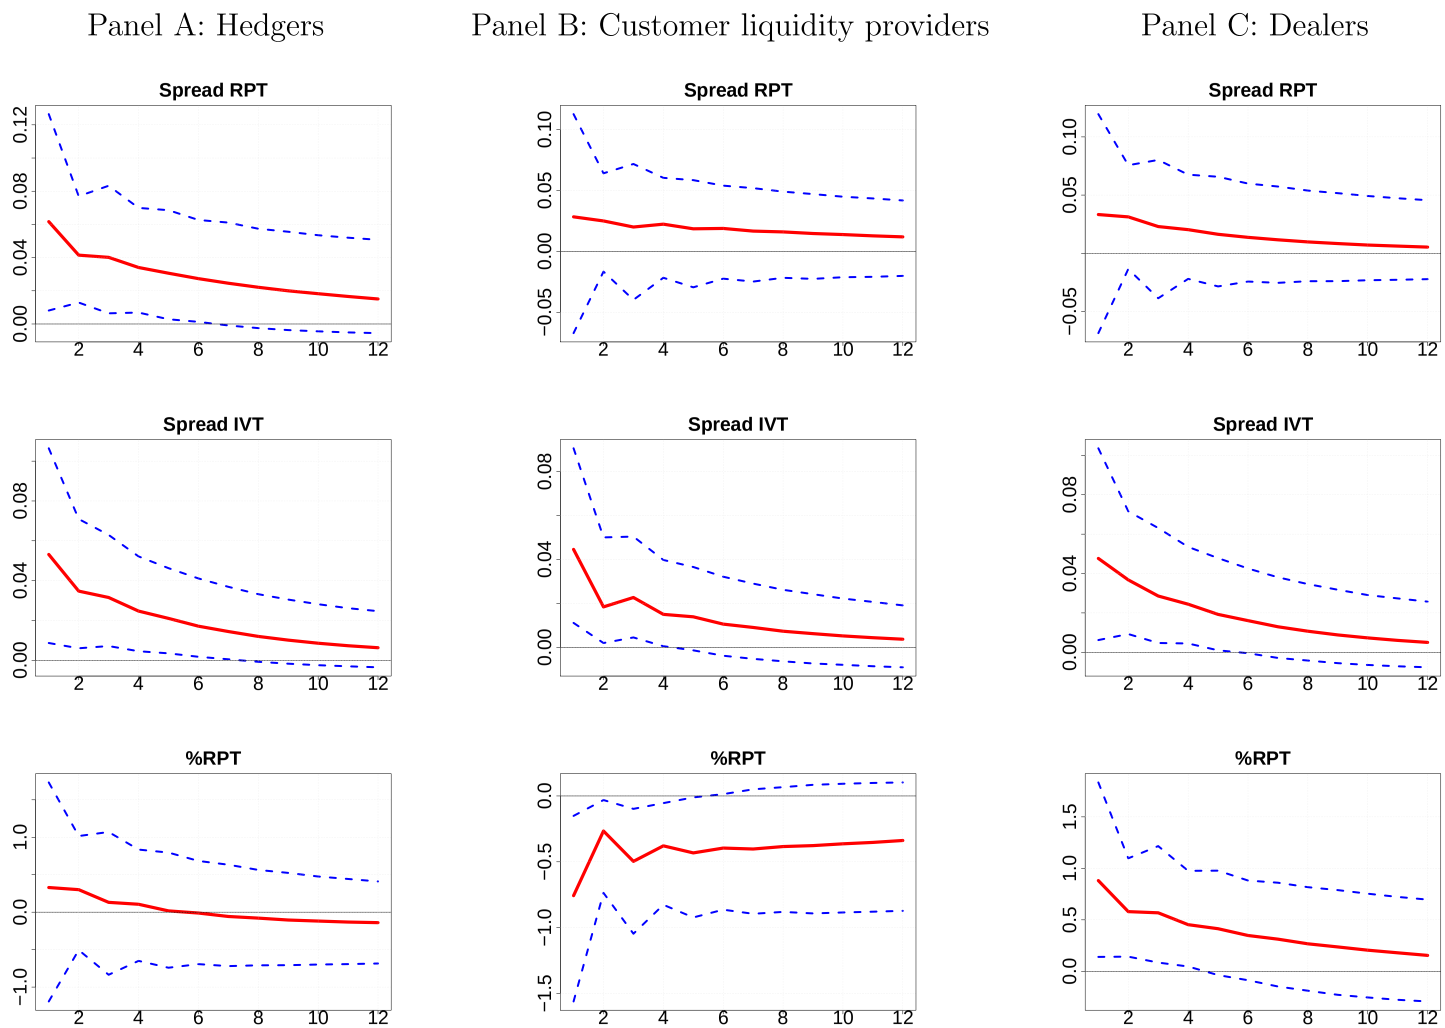

Figure 1.12: **Impulse response functions based on the agnostic method for high-yield bond markets**
These figures present the impulse response functions of the SVAR model in equation (1.2) for investment-grade bond markets, using the agnostic method. The total number of lags selected is one. For each of the shocks to be recovered, I separately apply two inequality restrictions, up to a horizon of three periods, on candidate impulse response functions. To recover the shocks to hedgers, I impose the restrictions that an increase in ϵ_t^H increases $\log\left(\text{Spread}_t^{\text{RPT}}\right)$ (i.e., $b_1 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,1} \geq 0$). To obtain the shocks to dealers, I apply the restrictions that an increase in ϵ_t^d increases $\%\text{RPT}_t$ (i.e., $b_3 \geq 0$) and $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{2,3} \geq 0$). Finally, to derive the shocks to customer liquidity providers, I use the restrictions that an increase in ϵ_t^d decreases $\%\text{RPT}_t$ (i.e., $b_2 \leq 0$) and increases $\log\left(\text{Spread}_t^{\text{IVT}}\right)$ (i.e., $b_{3,2} \geq 0$). Plotted are the median (in red) and 5-th and 95-th percentiles (in blue) of all feasible impulse response functions that satisfy the inequality constraints.

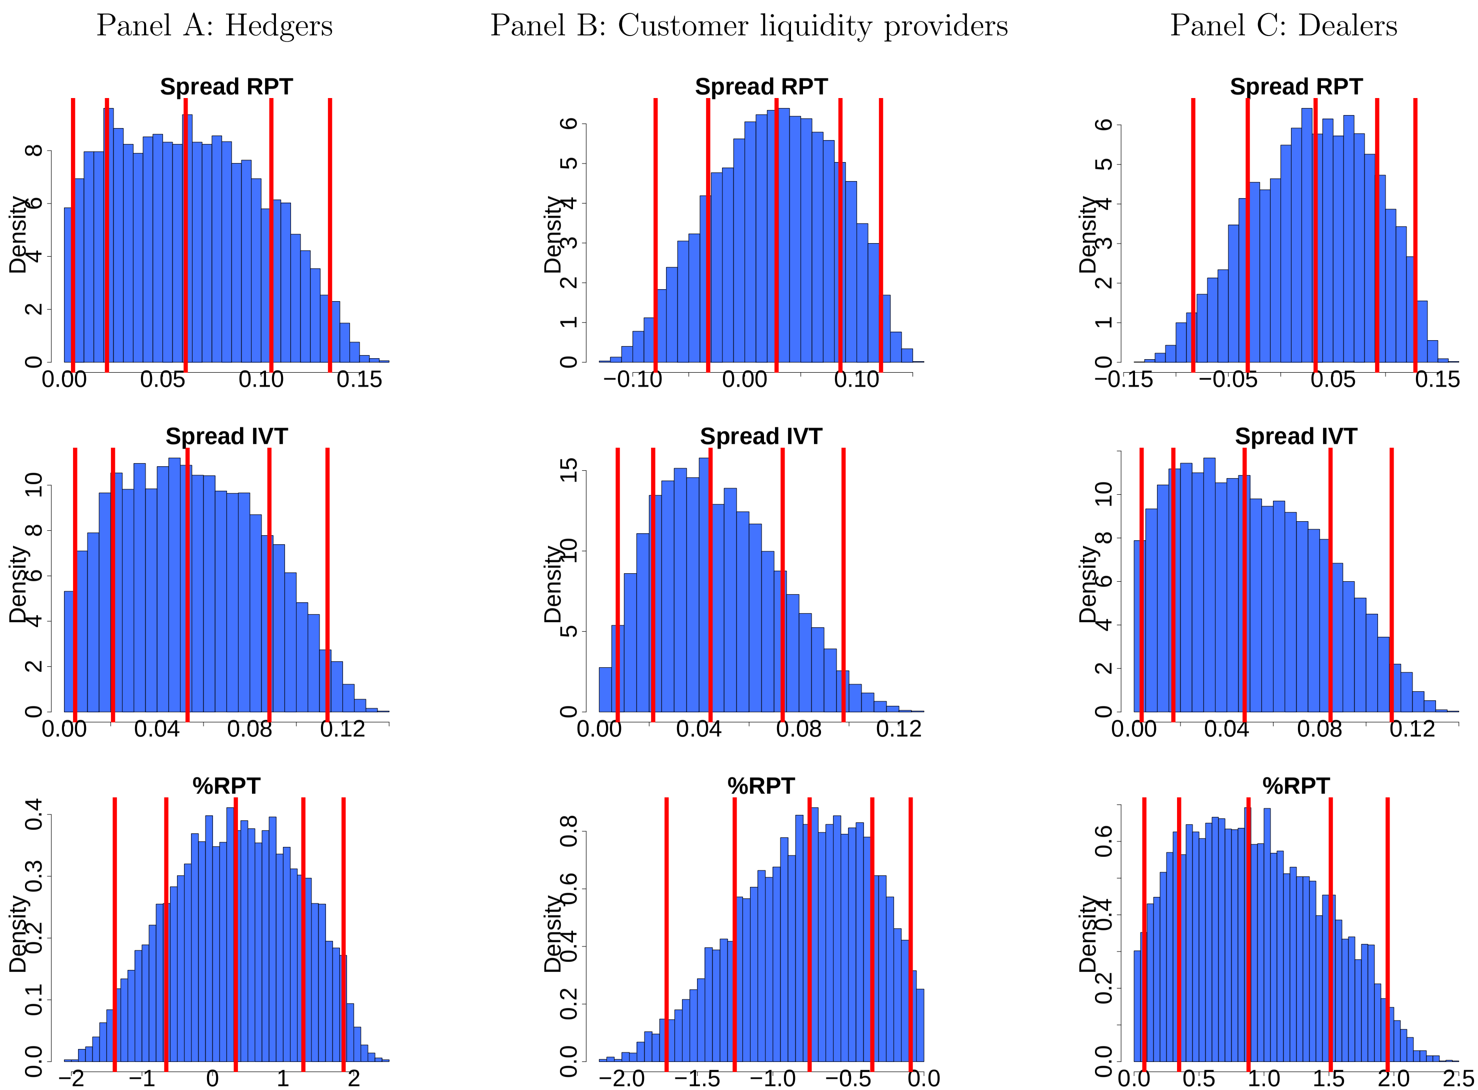

Figure 1.13: **Histograms of the initial impulse responses**
These figures plots the posterior distribution of the initial impulse responses for high-yield bond markets, using the agnostic method. The total number of lags selected is two. For each of the shocks to be recovered, I separately apply two inequality restrictions, up to a horizon of three periods, on candidate impulse response functions. To recover the shocks to hedgers, I impose the restrictions that an increase in ϵ_t^H increases $\log\left(\mathrm{Spread}_t^{\mathrm{RPT}}\right)$ (i.e., $b_1 \geq 0$) and $\log\left(\mathrm{Spread}_t^{\mathrm{IVT}}\right)$ (i.e., $b_{2,1} \geq 0$). To obtain the shocks to dealers, I apply the restrictions that an increase in ϵ_t^d increases $\%\mathrm{RPT}_t$ (i.e., $b_3 \geq 0$) and $\log\left(\mathrm{Spread}_t^{\mathrm{IVT}}\right)$ (i.e., $b_{2,3} \geq 0$). Finally, to derive the shocks to customer liquidity providers, I use the restrictions that an increase in ϵ_t^d decreases $\%\mathrm{RPT}_t$ (i.e., $b_2 \leq 0$) and increases $\log\left(\mathrm{Spread}_t^{\mathrm{IVT}}\right)$ (i.e., $b_{3,2} \geq 0$). Red lines in the plots indicate the 2.5th, 16th, 50th, 84th, and 97.5th percentiles.

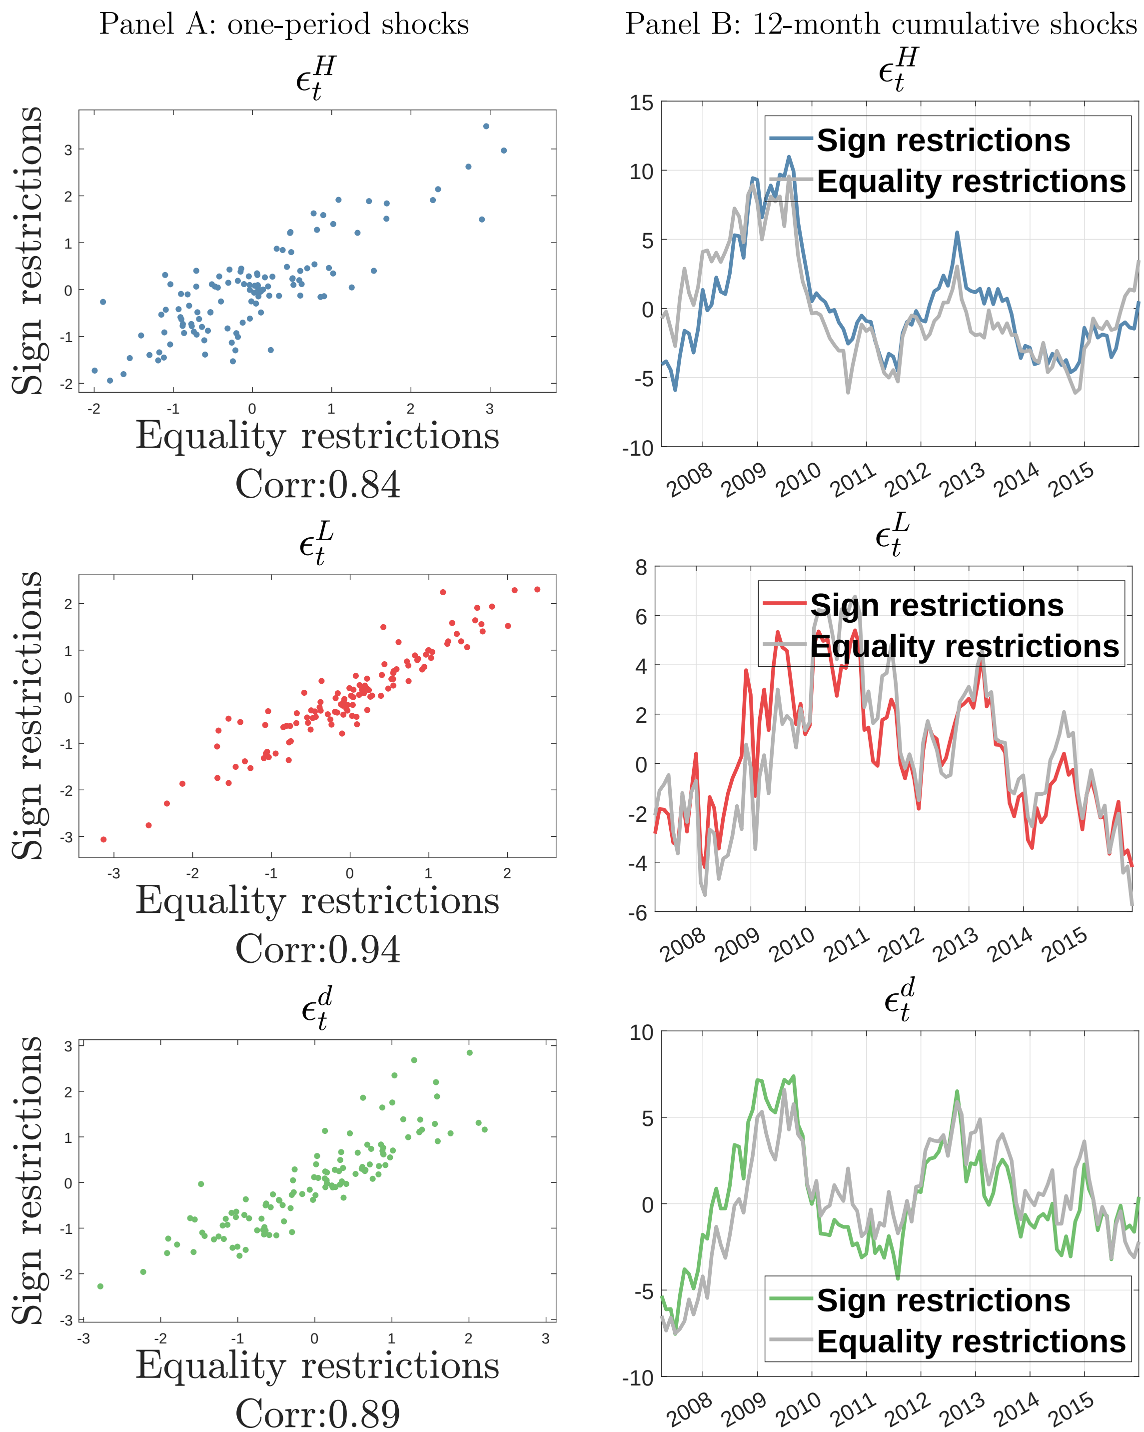

Figure 1.14: **Structural shocks from the agnostic method for high-yield markets**
These plots compare the obtained structural shocks based on the full identification method versus the agnostic method. Panel A presents the scatter plots where each spot represents a pair of the shocks based on the two methods in the same month. Also reported are the correlation coefficients of the shocks generated by the two methods. Panel B exhibits the time series of 12-month rolling sum of the structural shocks. Colored and solid lines are based on the sign restrictions, while gray and dashed lines are based on the equality restrictions.

Chapter 2: Recovery

2.1 Introduction

Financial market participants often have a forward-looking objective, which can be uncovering expectations of the asset return, volatility, or movements to the downside or the upside. This paper formalizes a framework that allows one to synthesize the conditional expectation of return quantities under the physical (real-world) probability measure from the prices of a set of spanning securities (e.g., the risk-free bond, the asset price, and options on the asset). Our treatment is formulaic in that we can express the expectation of functions of an asset return under the real-world probability measure using observed prices. We are motivated by a view that the literature has yet to articulate a coherent way of computing the conditional expectation of return quantities under the real-world probability measure.

The idea is to estimate the stochastic discount factor (SDF) projected onto returns to transform the risk-neutral probability measure — deduced from options on the asset — to compute expectations under the real-world probability measure. Our method is analytically tractable and free of distributional assumptions; in addition, it can be applied to return functions that satisfy certain smoothness properties.

We utilize the notion that the risk-neutral expectation of a broad class of return functions can be inferred from option prices. The challenge to compute the same expectation under the real-world probability measure is the unobservable SDF that maps the real-world measure to the risk-neutral measure. By forming an estimate of the projection of the SDF, we move between the risk-neutral and real-world probability measures to compute the real-world expectation of return functions.

Our work complements and connects to the work on recovery by [55], [56], [57], [58], [59], [60], [61], [62], [63], and [64]. The central difference from these papers is that we present formulas for the conditional expectation of return functions under the real-world probability measure. One featured theoretical result is that the real-world probability of disasters can be inferred from a positioning in a binary put and spread positions in out-of-the-money (OTM) put options. We also feature a theoretical result that the probability of return upside can be obtained from the price of a binary call and a portfolio of OTM call options. The insights of our theoretical work are the subject of our empirical investigation that uses options on equity market index and Treasury bond futures.

The manner in which we use projected SDFs facilitates mathematical analysis and estimation and offers a way to build conditional expectations that are disciplined by asset-pricing theory. Exploiting this framework, we construct the estimates of conditional expected excess returns of Treasury bond futures and use them to examine the spanning hypothesis — as formalized in [65] and [66] — with an augmented variable. We show that monthly nonoverlapping excess returns of futures on the 10- and 30-year Treasury bonds are not exclusively predictable from yield curve variables but also from conditional expected excess returns recovered from options on

Treasury bond futures. Thus, our approach can help to address a set of important, but not fully resolved, questions.

2.2 The approach

We suppose that the time-t value of an asset (e.g., equity market index or futures on the Treasury bond) is denoted by A_t. Let A_{t+1} represent the uncertainty about asset value at date $t+1$. Define $\Omega \equiv \{A_{t+1} > 0\}$ and let the physical density of A_{t+1} be given by $p[A_{t+1}]$.

Suppose one is interested in computing the conditional expectation of some function of A_{t+1} under the physical probability measure represented by $\mathbb{E}_t^{\mathbb{P}}\left(f[A_{t+1}]\right)$, with $\left|\mathbb{E}_t^{\mathbb{P}}\left(f[A_{t+1}]\right)\right| < \infty$.

Additionally, suppose $m[A_{t+1}]$ is the *projection* of the SDF onto the space generated by A_{t+1} at date $t+1$. We restrict our attention to the set of $m[A_{t+1}]$'s such that $m[A_{t+1}] > 0$ with probability 1. Assume that $m[A_{t+1}] \in \mathcal{C}^2$, the space of twice-continuously differentiable functions. The gross interest rate on a risk-free bond is $R_{\mathbf{f},t+1} = \frac{1}{\mathbb{E}_t^{\mathbb{P}}(m[A_{t+1}])}$.

Let $\mathbb{E}_t^{\mathbb{Q}}(.)$ denote conditional expectation under the risk-neutral probability measure. We assume that the measure $\mathbb{Q}$, corresponding to the density $q[A_{t+1}]$, is absolutely continuous with respect to measure $\mathbb{P}$. When the measure $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$, all $\mathbb{P}$-zero sets are also $\mathbb{Q}$-zero sets, and the two measures are equivalent. Then, by the Radon-Nikodym theorem, we can write the change-of-measure density of $\mathbb{Q}$ with respect to $\mathbb{P}$ as $\frac{d\mathbb{Q}}{d\mathbb{P}}$ or the change-of-measure density of $\mathbb{P}$ with respect to $\mathbb{Q}$ as $\frac{d\mathbb{P}}{d\mathbb{Q}} \equiv \mathcal{L}[A_{t+1}]$. Assume $\mathbb{E}_t^{\mathbb{P}}(\{\frac{d\mathbb{Q}}{d\mathbb{P}}\}^2) < \infty$.Hence, the change-of-measure density, or the likelihood function, can be represented as

$$\mathcal{L}[A_{t+1}] = \frac{p[A_{t+1}]}{q[A_{t+1}]}. \tag{2.1}$$

Next, rearranging the expression for the risk-neutral density function $q[A_{t+1}] = \frac{m[A_{t+1}]\,p[A_{t+1}]}{\int_\Omega m[A_{t+1}]\,p[A_{t+1}]\,dA_{t+1}}$, which ensures that $q[A_{t+1}]$ is a proper density function satisfying $1 = \int_\Omega q[A_{t+1}]\,dA_{t+1}$, we can write $\mathcal{L}[A_{t+1}] = \frac{\int_\Omega m[A_{t+1}]\,p[A_{t+1}]\,dA_{t+1}}{m[A_{t+1}]}$. With $\mathcal{L}[A_{t+1}] > 0$, it follows that

$$\mathbb{E}_t^{\mathbb{P}}\left(f[A_{t+1}]\right) = \int_\Omega f[A_{t+1}]\,\overbrace{\mathcal{L}[A_{t+1}]\,q[A_{t+1}]}^{p[A_{t+1}]}\,dA_{t+1} \tag{2.2}$$

$$= \mathbb{E}_t^{\mathbb{Q}}\left(f[A_{t+1}]\,\mathcal{L}[A_{t+1}]\right) \quad (\text{next using } \mathcal{L}[A_{t+1}] = \frac{\int_\Omega m[A_{t+1}]\,p[A_{t+1}]\,dA_{t+1}}{m[A_{t+1}]}) \tag{2.3}$$

$$= \underbrace{\int_\Omega m[A_{t+1}]\,p[A_{t+1}]\,dA_{t+1}}_{R_{\mathbf{f},t+1}^{-1}}\,\mathbb{E}_t^{\mathbb{Q}}\left(\frac{f[A_{t+1}]}{m[A_{t+1}]}\right) \tag{2.4}$$

$$= R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}\left(g[A_{t+1}]\right), \qquad \text{where} \;\; g[A_{t+1}] \;\equiv\; \frac{f[A_{t+1}]}{m[A_{t+1}]}. \tag{2.5}$$

The formula in equation (2.3) is associated with reweighting, whereby the likelihood function $\mathcal{L}[A_{t+1}]$ transforms the measure $\mathbb{P}$ into $\mathbb{Q}$. The implication is that the deduced expected value of $f[A_{t+1}]$ under $\mathbb{P}$ is unaltered because $\mathcal{L}[A_{t+1}]$ is included in the expectation with respect to measure $\mathbb{Q}$.

Suppose that options on A_{t+1} of all strike prices are traded. The function $g[A_{t+1}]$ can be synthesized with a portfolio of options, and the market for claims on

A_{t+1} is complete. Let $P_t[K]$ ($C_t[K]$) be the time-t price of the European put (call) option, with strike price K, that expire at date $t+1$. We assume that $\left|\mathbb{E}_t^{\mathbb{Q}}(g[A_{t+1}])\right| < \infty$, so $\mathbb{E}_t^{\mathbb{P}}(f[A_{t+1}])$ deduced in equation (2.5) is well-defined. Our constructions do not require distributional assumptions about A_{t+1}.

The Girsanov theorem involves reweighting the original diffusion process to obtain a different process that modifies the drift, but not the diffusion coefficient. In contrast, the counterpart for semimartingales potentially alters all moments of the $\mathbb{Q}$ measure when moving from $\mathbb{P}$ to $\mathbb{Q}$ (e.g., [67, Section 4.1 and Theorem 1.30]). Instead, we operate in a discrete-time environment and propose going in the reverse, from $\mathbb{Q}$ to $\mathbb{P}$, and using option prices to identify $\mathbb{Q}$. This, in turn, can help to analytically characterize $\mathbb{P}$-measure conditional expectations without taking a stand on how to model A_{t+1}. We call the step in equation (2.5) "inverting the Girsanov theorem." The next step makes our approach complete.

Definition 1 (Recovery)**.** *Suppose $\frac{f[A_{t+1}]}{m[A_{t+1}]}$ is in the span of options. Recovery implies that $\mathbb{E}_t^{\mathbb{P}}\left(f[A_{t+1}]\right)$ can be synthesized from the prices of a set (or subset) of spanning securities, namely, the risk-free bond, the underlying asset, and options on the asset.*

What is the broadest class of $f[A_{t+1}]$ for which one could deduce $\mathbb{E}_t^{\mathbb{P}}(f[A_{t+1}])$, given $m[A_{t+1}] \in \mathcal{C}^2$? Our answer is that the framework is tractable when $f[A_{t+1}]$ is of the class $\mathcal{C}^{-1}$, where the space $\mathcal{C}^{-1}$ is the set of functions for which $\int_\Omega f[A_{t+1}]\, dA_{t+1}$ is continuous (e.g., [68, Table 1]). Our analysis reveals the mechanism by which theory assigns weights to the spanning securities (as we show for different choices of $f[A_{t+1}]$ and $m[A_{t+1}]$).

Indicating the generality of our approach, $\mathbb{E}_t^{\mathbb{P}}(f[A_{t+1}])$ is amenable to characterization when the function $g[A_{t+1}] = \frac{f[A_{t+1}]}{m[A_{t+1}]}$ is of the class L^1 (i.e., $\int_\Omega \left|g[A_{t+1}]\right| dA_{t+1} < \infty$). This follows from [69, Appendix A.4, equation A.12], who show how to construct $g[A_{t+1}] \in L^1$ using the standard Gaussian density building block. The standard Gaussian density is of class $\mathcal{C}^2$ and in the span of options.

Does the lack of state dependencies in $m[A_{t+1}]$ detract from generality? [70] show that when there are claims on A_{t+1} for which there are data, one can model change-of-measure densities that are conditioned only on A_{t+1}. It suffices to work with the projection of the SDF onto the space generated by A_{t+1} (e.g., [71]). Elaborating on this attribute of our analysis in terms of inverting the Girsanov theorem for deducing $\mathbb{P}$ measure expectations and then achieving recovery, we note that the value at time t of a claim, denoted by $\mathfrak{p}_t$, paying a single cash flow $\mathbf{c}[A_{t+1}]$ at date $t+1$ is

$$\begin{aligned}
\mathfrak{p}_t &= R_{\mathbf{f},t+1}^{-1} \overbrace{\mathbb{E}_t^{\mathbb{P}}\left(\frac{d\mathbb{Q}}{d\mathbb{P}}\,\mathbf{c}[A_{t+1}]\right)}^{\mathbb{E}_t^{\mathbb{Q}}(\mathbf{c}[A_{t+1}])} && (2.6)\\
&= R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{P}}\left(\mathbb{E}^{\mathbb{P}}\left(\frac{d\mathbb{Q}}{d\mathbb{P}}\,\mathbf{c}[A_{t+1}]\middle| A_{t+1}\right)\right) && (2.7)\\
&= R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{P}}\left(\mathbb{E}^{\mathbb{P}}\left(\frac{d\mathbb{Q}}{d\mathbb{P}}\middle| A_{t+1}\right)\mathbf{c}[A_{t+1}]\right) && (2.8)\\
&= R_{\mathbf{f},t+1}^{-1}\mathbb{E}_t^{\mathbb{P}}\left(m[A_{t+1}]\,\mathbf{c}[A_{t+1}]\right), && (2.9)
\end{aligned}$$

where $\frac{d\mathbb{Q}}{d\mathbb{P}}$ represents the change of probability, and $m[A_{t+1}] \equiv \mathbb{E}^{\mathbb{P}}(\frac{d\mathbb{Q}}{d\mathbb{P}} \mid A_{t+1})$. Thus, in the spirit of [55], [56], [57], [58], [59], [60], and [61], we need only model $\mathbb{E}^{\mathbb{P}}(\frac{d\mathbb{Q}}{d\mathbb{P}} \mid A_{t+1})$ without using state-dependent specifications.

It may be useful to differentiate our approach in equation (2.5) from that of

[70]. The crux of their approach is to estimate the physical return density from past returns on each date and then exploit S&P 500 index option prices to deduce the

parameter vector representing the projected SDF on a given day. Instead, our interest centers on computing distributional quantities $\mathbb{E}_t^{\mathbb{P}}(f[A_{t+1}])$ from option prices, and we accomplish this *without* estimating the return density function under the $\mathbb{P}$ measure.

Under our method, one can recover the Fourier transform of the $\mathbb{P}$ measure density analytically, but this is not the case for $\mathbb{P}$ measure density. To see this, consider the complex-valued payoff $f[A_{t+1}] = e^{i\,u\,\log(\frac{A_{t+1}}{A_t})} = \cos(u\log(\frac{A_{t+1}}{A_t})) + i\,\sin(u\log(\frac{A_{t+1}}{A_t}))$ for $i = \sqrt{-1}$ and some constant parameter $u \in \mathbb{R}$. Then, we can compute $\mathbb{E}_t^{\mathbb{P}}(e^{i\,u\,\log(\frac{A_{t+1}}{A_t})}) = R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\frac{e^{i\,u\,\log(\frac{A_{t+1}}{A_t})}}{m[A_{t+1}]})$, which is algebraic in the prices of spanning securities. Yet, the reverse step in going from the Fourier transform of the return density to imputing the density does not lend to closed-form tractability. Specifically, $p[A_{t+1}] = \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{Re}[e^{-i\,u\,\log(\frac{A_{t+1}}{A_t})}\{\mathbb{E}_t^{\mathbb{P}}(e^{i\,u\,\log(\frac{A_{t+1}}{A_t})})\}]\,du$, where Re[.] is the real part of the expression (e.g., [72, equation 4.5]).

While the true $\mathbb{E}^{\mathbb{P}}(\frac{d\mathbb{Q}}{d\mathbb{P}} \mid A_{t+1}) = m[A_{t+1}]$ is not directly identifiable, any choice must be disciplined through economic theory in conjunction with empirical performance. Considering this possibility, we confine attention to a set of $\mathbb{E}^{\mathbb{P}}(\frac{d\mathbb{Q}}{d\mathbb{P}} \mid A_{t+1})$ that are tractable, preserve positivity, correctly price a set of returns, and are parsimonious from the perspective of estimation.

2.3 Developing the expectation of a disaster

In this section, we take the asset to be the equity market index. Let the level of the equity market index be denoted by S_t and $r_{t+1} \equiv \log(\frac{S_{t+1}}{S_t})$ be the logarithmic equity return. Consider the following conditional expectation of a (wealth) disaster under the *physical* probability measure, corresponding to some breakpoint K^d:

$$\mathrm{Prob}_t \left(S_{t+1} < K^d\right) \equiv \mathbb{E}_t^{\mathbb{P}} \left(\theta[K^d \ - \ S_{t+1}]\right) = \mathbb{E}_t^{\mathbb{P}} \left(\theta[r^d \ - \ r_{t+1}]\right), \tag{2.10}$$

where $\theta[x] \in \mathcal{C}^{-1}$ is the Heaviside theta function, defined as ([73, Chapter 1]),

$$\theta[x] \ \equiv \ \frac{d}{dx} \max(x, 0) \ = \ \begin{cases} 0 & \text{for } x < 0, \\ 1 & \text{for } x > 0. \end{cases} \tag{2.11}$$

Whereas some authors define $\theta[x]$ as taking a value of $\frac{1}{2}$ at $x = 0$, this omission is inconsequential to our calculations, given that the probability of a single event $S_{t+1} = K^d$ is zero. We define the threshold $r^d \equiv \log(\frac{K^d}{S_t}) < 0$, and the set $\{r_{t+1} < r^d\}$ captures a (wealth) disaster over $[t, t+1]$.[1]

2.3.1 Recovery of real-world (physical) probabilities

To illustrate our approach to obtain $\mathbb{E}_t^{\mathbb{P}} \left(f[S_{t+1}]\right)$ from option prices, consider first the disaster event $\{r_{t+1} < r^d\}$ over $[t, t+1]$. Following equation (2.5), we

[1] For example, if the level of the S&P 500 equity index at the start of the options expiration cycle on 11/23/2015 is 2086.59, then choosing $r^d = -0.07257$ implies the breakpoint $K^d = e^{r^d} S_t = 1940.53$, and $\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$ measures the conditional expectation of a wealth disaster of size greater than 7% (i.e., $\frac{K^d}{S_t} - 1 = -0.07$) by the end of the expiration cycle on 12/19/2015.

express $\mathbb{E}_t^{\mathbb{P}}(f[S_{t+1}]) = R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}\,(g[S_{t+1}])$, where

$$f[S_{t+1}] = \theta[K^d - S_{t+1}] \quad \text{and} \quad g[S_{t+1}] \equiv \theta[K^d - S_{t+1}]\frac{1}{m[S_{t+1}]} \geq 0. \tag{2.12}$$

The function $g[S_{t+1}]$ inherits some of its smoothness properties from $f[S_{t+1}]$.

Result 1 (Disaster probability)**.** *The probability of disasters can be analytically computed from put option prices, as follows:*

$$\underbrace{\mathbb{E}_t^{\mathbb{P}}(\theta[K^d - S_{t+1}])}_{\text{Disaster probability}} = g[K^d]\,\underbrace{\mathbb{B}_t^{\text{put}}[K^d]}_{\text{Binary put}} - g'[K^d]\,P_t[K^d] + \underbrace{\int_0^{K^d} g''[K]\,P_t[K]\,dK}_{\text{Collection of puts with strike } K<K^d}, \tag{2.13}$$

where $\mathbb{B}_t^{\text{put}}[K^d] \equiv R_{\mathbf{f},t+1}^{-1}\mathbb{E}_t^{\mathbb{Q}}(\theta[K^d - S_{t+1}]) \approx \frac{P_t[K^d+\Delta K]-P_t[K^d-\Delta K]}{2\,\Delta\,K}$ *is the time-t price of the binary put (for small strike price increment* ΔK*). Additionally,* $g[K^d] \equiv g[S_{t+1}]\big|_{S_{t+1}=K^d}$, $g'[K^d] \equiv \frac{dg[S_{t+1}]}{dS_{t+1}}\big|_{S_{t+1}=K^d}$, *and* $g''[K] \equiv \frac{d^2 g[S_{t+1}]}{dS_{t+1}^2}\big|_{S_{t+1}=K<K^d}$.

Proof: See Appendix 2.6.1. ■

Equation (2.13) shows that the disaster probability over $[t, t+1]$ can be inferred from OTM put prices. This calculation involves spread positions that combine buying and selling a particular set of put options of various OTM strike prices. The analyticity of the formulation can be traced to the feature that $\theta[K^d - S_{t+1}]\frac{1}{m[S_{t+1}]}$ can be synthesized from elementary payoffs of the type $\theta[K - S_{t+1}]$ and $\max(K - S_{t+1}, 0)$.

The expected return of the claim with payoff $f[r_{t+1}] = \theta[r^d - r_{t+1}]$ is $\mu_t[r^d] \equiv \frac{\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])}{R_{\mathbf{f},t+1}^{-1}\mathbb{E}_t^{\mathbb{Q}}(\theta[r^d - r_{t+1}])}$

1. Accordingly, we can define the disaster risk premium (in log form) as

$$\mathrm{drp}_t[r^d] \;\equiv\; \log(1+\mu_t[r^d]) \;-\; \log(R_{\mathbf{f},t+1}) = \log\left(\frac{\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])}{\mathbb{E}_t^{\mathbb{Q}}(\theta[r^d - r_{t+1}])}\right), \tag{2.14}$$

where $\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$ can be inferred based on equation (2.13), and $\mathbb{E}_t^{\mathbb{Q}}(\theta[r^d - r_{t+1}]) = R_{\mathbf{f},t+1}\,\mathbb{B}_t^{\mathrm{put}}[K^d]$.

Equation (2.13) quantifies the disaster probability under the $\mathbb{P}$ measure, and can be viewed as an analog to [74, equation (29)]. Their idea is to measure the realized jumps with high-frequency data and translate that information into forward-looking $\mathbb{P}$ measure quantities using the Extreme Value Theory. Our innovation (in equations (2.13) and (2.14)) is to characterize the physical probability of disasters and disaster risk premiums by relying on options data without modeling wealth.

Our companion result for upside probability below is also new and examined in our empirical work.

Result 2 (Upside probability). *Let K^u be some threshold satisfying $\frac{S_t}{K^u} - 1 < 0$ and $g[S_{t+1}] = \theta[S_{t+1} - K^u]\frac{1}{m[S_{t+1}]}$. Then*

$$\underbrace{\mathrm{Prob}[S_{t+1} > K^u]}_{\text{Upside probability}} \;=\; g[K^u]\,\underbrace{\mathbb{B}_t^{\mathrm{call}}[K^u]}_{\text{Binary call}} + g'[K^u]\,C_t[K^u] + \underbrace{\int_{K^u}^{\infty} g''[K]\,C_t[K]\;dK}_{\text{Collection of calls with strike } K>K^u}, \tag{2.15}$$

where $\mathbb{B}_t^{\mathrm{call}}[K^u] \equiv R_{\mathbf{f},t+1}^{-1}\mathbb{E}_t^{\mathbb{Q}}(\theta[S_{t+1} - K^u]) \approx \frac{C_t[K^u - \Delta K] - C_t[K^u + \Delta K]}{2\,\Delta\,K}$ is the time-t price of the binary call. Here, $g[K^u] \equiv g[S_{t+1}]\big|_{S_{t+1}=K^u}$, $g'[K^u] \equiv \frac{dg[S_{t+1}]}{dS_{t+1}}\big|_{S_{t+1}=K^u}$, and $g''[K] \equiv \frac{d^2 g[S_{t+1}]}{dS_{t+1}^2}\big|_{S_{t+1}=K>K^u}$.

Proof: See Appendix 2.6.2. ■

Now $\theta[S_{t+1} - K^u]\frac{1}{m[S_{t+1}]}$ can be synthesized from elementary payoffs of the type $\theta[S_{t+1} - K]$ and $\max(S_{t+1} - K, 0)$. The calculation in equation (2.15) involves a long position in a binary call, a long OTM call position with strike K^u, and a long position in a collection of OTM calls with strike $K > K^u$.[2]

2.3.2 Elaborating on an $m[S_{t+1}]$ specification and the allied disaster probability

To derive $g[K^d]$, $g'[K^d]$, and $g''[K]$ — the positioning on the spanning securities — in equation (2.13), together with the expressions for conditional expected (logarithmic) return and return volatility used in our empirical work, we feature the following specification of $m[S_{t+1}] \in \mathcal{C}^2$, which guarantees positivity:

$$m[S_{t+1}] = \exp\left(\eta_0 - 1 + \boldsymbol{\eta}'\mathbf{z}_{t+1}[S_{t+1}]\right), \quad \text{where} \quad \underbrace{\mathbf{z}_{t+1}[S_{t+1}]}_{2\times 1} \equiv \begin{pmatrix} \frac{S_{t+1}}{S_t} - R_{\mathbf{f},t+1} \\ \frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1} \end{pmatrix}. \tag{2.16}$$

While the contract with payoff $\{\log(\frac{S_{t+1}}{S_t})\}^2$ is not directly traded, we synthesize its intrinsic value using a portfolio of options. It holds that $\mathrm{v}_t \equiv R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\{\log(\frac{S_{t+1}}{S_t})\}^2) = \int_0^{S_t} \frac{2(1+\log(\frac{S_t}{K}))}{K^2}\, P_t[K]\, dK + \int_{S_t}^{\infty} \frac{2(1-\log(\frac{K}{S_t}))}{K^2}\, C_t[K]\, dK$ (e.g., [78, equation (7)]). One may interpret $\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}$ as the excess return of a volatility contract and $\frac{S_{t+1}}{S_t} - R_{\mathbf{f},t+1}$ as the excess return of equity.

[2] The Heaviside function is a building block of a number of considered objects. Specifically, $\mathbb{E}_t^{\mathbb{P}}(\{\log(\frac{S_{t+1}}{S_t})\}^2\,\theta[K^d - S_{t+1}])$ represents semivariance to the downside, whereas $\mathbb{E}_t^{\mathbb{P}}(\{\log(\frac{S_{t+1}}{S_t})\}^2\,\theta[S_{t+1} - K^u])$ represents semivariance to the upside (e.g., see among others, [75], [76], and [77]).

The following representation of disaster probability follows from equations (2.13) and (2.16) and is employed in our empirical investigation:

$$\mathbb{E}_t^{\mathbb{P}}(\theta[K^d - S_{t+1}]) = g[K^d]\,\mathbb{B}_t^{\mathrm{put}}[K^d] - g[K^d](-\frac{2\eta_2\log(\frac{K^d}{S_t})}{\mathrm{v}_t K^d} - \frac{\eta_1}{S_t})P_t[K^d] + \int_0^{K^d} g[K]\{(\frac{2\eta_2\log(\frac{K}{S_t})}{\mathrm{v}_t K^2} - \frac{2\eta_2}{\mathrm{v}_t K^2}) + (\frac{2\eta_2\log(\frac{K}{S_t})}{\mathrm{v}_t K} + \frac{\eta_1}{S_t})^2\}P_t[K]\,dK \quad (2.17)$$

where $g[K^d] = \frac{1}{m[S_{t+1}]}\big|_{S_{t+1}=K^d} = \exp(-\eta_0 + 1 - \eta_1(\frac{K^d}{S_t} - R_{\mathbf{f},t+1}) - \eta_2(\frac{\{\log(\frac{K^d}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}))$.

The expression for disaster probability in equation (2.17) merits the following comments. Even when the sensitivity of $\log(m[S_{t+1}])$ to $\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}$ is zero, namely $\eta_2 = 0$, the disaster probability in equation (2.17) is distinct from the one implied by the CAPM. A positive and statistically significant η_2 is a robust outcome in our estimations involving S&P 500 equity index as well as futures on the Treasury bond (i.e., see Tables 2.1 and 2.9).

Please insert Table 2.1 here

Please insert Table 2.9 here

The returns $\frac{S_{t+1}}{S_t} - 1$ and $\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - 1$ are distinct. Over the S&P 500 index options expiration cycles of 1/22/1990 to 11/23/2015, the unconditional correlation between $\frac{S_{t+1}}{S_t} - 1$ and $\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - 1$ is -0.45, reflecting that return of equity and return of volatility are negatively associated.

We incorporate the effects of return of volatility in the $m[S_{t+1}]$ specification, and is empirically pertinent. Integral to [70] is a Chebyshev return polynomial approximation for $m[S_{t+1}]$. Our approach to consider the return of volatility in the $m[S_{t+1}]$ specification is in line with [79], and is akin to the variance-dependent SDF in [80, equation (2)]. Our specification differs from [60], who study conditional minimum variance projection of the SDF on a polynomials of gross market returns.

There may be some choices of $m[S_{t+1}]$ for which disaster probabilities and risk premiums may reflect a subset of the relevant pricing information from OTM puts of various strikes. For example, if $m[S_{t+1}] = (\frac{S_{t+1}}{S_t})^{-1}$, in which case, $g[S_{t+1}] = \frac{\theta[K^d - S_{t+1}]}{(\frac{S_{t+1}}{S_t})^{-1}}$, implying that $g[K^d] = \frac{K^d}{S_t} = e^{r^d}$ (a constant), $g'[K^d] = \frac{1}{S_t}$, and $g''[K] = 0$. Hence, $\mathrm{drp}_t[r^d] = \log(\frac{K^d}{S_t} - \frac{P_t[K^d]}{S_t \mathbb{B}_t^{\mathrm{put}}[K^d]}) - \log(R_{\mathbf{f},t+1})$, which is deprived of a contribution from put prices with strikes lower than K^d. This special case of $g''[K] = 0$, that is, *absent term* $\int_0^{K^d} g''[K]\, P_t[K]\, dK$ in equation (2.13), is featured in [81, equation (23)]. In this light, we also consider $m[S_{t+1}]$ specifications that incorporate a nonzero curvature of $g[S_{t+1}]$, that is, $g''[S_{t+1}]$ is *not zero.*

Case 1. *Consider* $m[S_{t+1}] = (1+\upsilon \log(\frac{S_{t+1}}{S_t}))^{-\gamma}$ *for* $\upsilon > 0$ *and* $\gamma > 0$, *where* $-m'/m''$ *is linear in* $r_{t+1} = \log(\frac{S_{t+1}}{S_t})$, *analogous to [82, equation (25)]. It then follows that*

$$\mathbb{E}_t^{\mathbb{P}}(\theta[K^d - S_{t+1}]) = (1+\upsilon \log(\frac{K^d}{S_t}))^{\gamma}\, \mathbb{B}_t^{\mathrm{put}}[K^d] - \frac{\upsilon\, \gamma}{K^d}(1+\upsilon \log(\frac{K^d}{S_t}))^{\gamma-1} P_t[K^d] + \int_0^{K^d} g''[K] P_t[K] dK, \tag{2.18}$$

where $g''[K] = \frac{\upsilon^2(\gamma-1)\gamma(1+\upsilon \log(\frac{K}{S_t}))^{\gamma-2}}{K^2} - \frac{\upsilon\, \gamma(1+\upsilon \log(\frac{K}{S_t}))^{\gamma-1}}{K^2}$.

Case 2. *Let* $m[S_{t+1}] = \psi_a(\log(\frac{S_{t+1}}{S_t}))^2 + \psi_b \log(\frac{S_{t+1}}{S_t}) + \psi_c$. *We assume* $\psi_b < 0$. *To maintain the positivity of* $m[S_{t+1}]$, *we impose* $\psi_a > 0$ *and* $\frac{\psi_b^2}{4\,\psi_a} < \psi_c$. *Let* $g[K^d] =$

$\frac{1}{m[K^d]}$. *Then* $\mathbb{E}_t^{\mathbb{P}}(\theta[K^d - S_{t+1}]) = g[K^d]\,\mathbb{B}_t^{\text{put}}[K^d] - g'[K^d]\,P_t[K^d] + \int_0^{K^d} g''[K]\,P_t[K]\,dK$, *where we determine*

$$
\begin{aligned}
g'[K^d] &= -\frac{1}{K^d}(2\psi_a \log(\frac{K^d}{S_t}) + \psi_b)(m[K^d])^{-2} \quad \text{and} && (2.19)\\
g''[K] &= \frac{2}{K^2}(2\psi_a \log(\frac{K}{S_t}) + \psi_b)^2(m[K])^{-3} - \frac{2}{K^2}(\psi_a - \psi_a \log(\frac{K}{S_t}) - \frac{\psi_b}{2})(m[K])^{-2} && (2.20)
\end{aligned}
$$

Case 3. *Suppose* $m[S_{t+1}] = e^{-\gamma \log(\frac{S_{t+1}}{S_t})}$. *Assuming that the parameter* $\gamma > 1$, *we express*

$$
\mathbb{E}_t^{\mathbb{P}}(\theta[K^d - S_{t+1}]) = (\frac{K^d}{S_t})^{\gamma}\ \mathbb{B}_t^{\text{put}}[K^d] - \frac{\gamma}{S_t}(\frac{K^d}{S_t})^{\gamma-1}P_t[K^d] + \frac{\gamma(\gamma-1)}{S_t^2}\int_0^{K^d}(\frac{K}{S_t})^{\gamma-2}P_t[K]\,dK. \tag{2.21}
$$

These economies indicate the manner in which an $m[S_{t+1}]$ serves to assign weights to the spanning securities: (i) a *long* binary put with strike K^d, (ii) a *short* put with strike K^d, and (iii) a collection of *long* puts with strikes deeper OTM than K^d. While the characterization in equation (2.17) is new, equation (2.21) of Case 3 has been investigated in our earlier version and [83, Proposition 1]. [84] consider $m[S_{t+1}] = e^{-\gamma \log(\frac{S_{t+1}}{S_t})}$, but their goal is to study the properties of the return density under $\mathbb{P}$. The methods developed in [85] and [86] infer the risk aversion function, and they employ option prices and assumptions about the slope of $m[S_{t+1}]$. Finally, our approach to deduce real-world disaster probabilities from put prices differs from [87], [88], [89], and [90], who focus on quantifying *risk-neutral tails* from option prices.

2.3.3 The estimation framework without imposing distributional assumptions

Our identification of disaster probabilities and risk premiums is based on the information contained in options on the S&P 500 index. The horizon over which we measure disaster probabilities is the number of days in the options expiration cycle. This number varies between 23 and 32 calender days, with an average of 27 days. The first expiration cycle started on 1/22/1990 and the last one on 11/23/2015. Thus, we have 311 expiration cycles over 26 years. Crucially, we keep all OTM options up to the minimum tick size.

Equation (2.17) is instrumental to our construction of disaster probabilities and requires an input for $\Theta \equiv (\eta_0, \eta_1, \eta_2) \in \mathbb{R}^3$, the parameter vector that controls the level, slope, and curvature of $m[S_{t+1}]$. Following [56] and [91], we estimate Θ by implementing the following minimum discrepancy problem:

$$\inf_{m} \;\; \mathbb{E}^{\mathbb{P}}(m \, \log(m)) \;\; \text{subject to} \;\; \text{(i)} \;\; \mathbb{E}^{\mathbb{P}}(m) = \mathbb{E}^{\mathbb{P}}(R_{\mathbf{f},t+1}^{-1}), \;\; \text{and} \;\; \text{(ii)} \;\; \mathbb{E}^{\mathbb{P}}(m \, \mathbf{z}) = \mathbf{0}. \tag{2.22}$$

From this problem, it may be deduced that $m[S_{t+1}]$ is of the form $m[S_{t+1}] = \exp\left(\eta_0 - 1 + \boldsymbol{\eta}'\mathbf{z}_{t+1}[S_{t+1}]\right)$, where η_0 and $\boldsymbol{\eta}$ are the Lagrange multipliers associated with the pricing restrictions $\mathbb{E}^{\mathbb{P}}(m) = \mathbb{E}^{\mathbb{P}}(R_{\mathbf{f},t+1}^{-1})$ and $\mathbb{E}^{\mathbb{P}}(m \, \mathbf{z}) = \mathbf{0}$. These restrictions are imposed unconditionally, and $(\eta_0, \boldsymbol{\eta})$ solves the following optimization problem:

$$\inf_{(\eta_0, \boldsymbol{\eta})} \quad - \mathbb{E}^{\mathbb{P}}(R_{\mathbf{f},t+1}^{-1}) \, \eta_0 \; + \; \mathbb{E}^{\mathbb{P}}(\exp(\eta_0 - 1 + \boldsymbol{\eta}'\mathbf{z}_{t+1})). \tag{2.23}$$

Implementing equation (2.23), Table 2.1 presents the estimates of (i) the full specification (2.16) in which η_0, η_1, and η_2 are freely determined parameters, (ii) a restricted specification that imposes $\eta_1 = 0$ (i.e., the projection of the SDF incorporates a stand-alone exposure to the excess returns of volatility), and (iii) a restricted specification that imposes $\eta_2 = 0$ (i.e., the projection of the SDF incorporates a stand-alone exposure to the excess returns of equity).

To select the best and most parsimonious model, and to assess the statistical significance of the estimated $(\eta_0, \boldsymbol{\eta})$, we consider a procedure that randomly selects, with replacement, $(R_{\mathbf{f},t+1}, \mathbf{z}_{t+1})$ jointly via stationary bootstrap. Then we recompute $(\eta_0, \boldsymbol{\eta})$ in each bootstrap draw, exploiting the minimum discrepancy problem in equations (2.22)–(2.23). We perform 1,000 bootstrap draws and record the 95% bootstrap confidence intervals for $(\eta_0, \boldsymbol{\eta})$.

The estimates of η_2 (shown in Table 2.1) are positive and equal 0.192 (when $\eta_1 = 0$) and 0.238 (under the full specification). Whereas η_2 is reliably estimated with bootstrap confidence intervals that do not bracket zero, the η_1 estimates manifest wide confidence intervals, and the hypothesis of $\eta_1 = 0$ is not rejected. Our evidence further suggests that η_0 estimates are close to one.

Complementary evidence that favors the $m[S_{t+1}]$ specification with only excess return of volatility can be garnered from Table 2.2. We evaluate the specification errors of $m[S_{t+1}]$ that arise when pricing a set of returns $\mathbf{R}_{t+1}$ (up to eight returns): gross return of the risk-free bond, gross return of equity, and gross return of 5%, 3%, and 1% OTM puts on the S&P 500 index, as well as 1%, 3%, and 5% OTM calls.

Please insert Table 2.2 here

Exploiting [92], we report $\mathrm{HJ}^{\mathrm{dist}} \equiv \sqrt{\mathbf{e}_{t+1}' \mathbf{I}^{-1} \mathbf{e}_{t+1}}$, where $\mathbf{e}_{t+1} \equiv \mathbb{E}^{\mathbb{P}}(m[S_{t+1}]\, \mathbf{R}_{t+1}) - \mathbf{1}$. For our comparison of nested $m[S_{t+1}]$ specifications, we use the identity matrix, $\mathbf{I}$, as suggested by [93, page 216]. The finding is that the specification $m[S_{t+1}] = \exp(\eta_0 - 1 + \eta_2(\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}))$ yields the lowest $\mathrm{HJ}^{\mathrm{dist}}$ measure, and this finding is robust across the three sets of $\mathbf{R}_{t+1}$. The reported $\frac{\mathrm{HJ}^{\mathrm{dist}}_{\eta_1 \equiv 0}}{\mathrm{HJ}^{\mathrm{dist}}_{\mathrm{Full}}} - 1$, which ranges from -23% to -48%, reflects the improvement in $\mathrm{HJ}^{\mathrm{dist}}$, while moving from the full specification to a frugal one with $\eta_1 \equiv 0$. With the highest bootstrap p-value of 0.095, the hypothesis that $\frac{\mathrm{HJ}^{\mathrm{dist}}_{\eta_1 \equiv 0}}{\mathrm{HJ}^{\mathrm{dist}}_{\mathrm{Full}}} - 1 \geq 0$ is rejected. Our evidence supports the relevance of return of volatility in achieving a better pricing fit to our choice of test assets. Hereon, we restrict $\eta_1 \equiv 0$.

Table 2.3 investigates whether one could improve the workings of our baseline model by adopting a time-varying loading of $\log(m[S_{t+1}])$ on $(\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1})$. Specifically, we study $m[S_{t+1}]$ specifications augmented with a conditioning variable $\mathfrak{u}_t$ as follows:

$$m[S_{t+1}] = \exp(\eta_0 - 1 + \{\eta_2 + \underbrace{\eta_{2,\mathfrak{u}}\, \mathfrak{u}_t}_{\text{time-varying}}\}(\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1})). \tag{2.24}$$

Please insert Table 2.3 here

We consider the following candidates for $\mathfrak{u}_t$: (i) excess return of equity market, (ii) size spread, (iii) value spread, and (iv) momentum spread. Our exercises in

Table 2.3 indicate that the baseline model with $\eta_{2,\mathfrak{u}} = 0$ provides lower or statistically indistinguishable $\mathrm{HJ}^{\mathrm{dist}}$ measures compared to those incorporating interacted lagged variables. In three out of four cases, we are unable to reject the hypothesis that $\eta_{2,\mathfrak{u}} = 0$, as inferred from the bootstrap confidence intervals.

2.3.4 Evidence on real-world probability of disasters

In our implementations and discussions to follow, we define a wealth disaster as the event

$$\{r_{t+1} < r^d\}, \quad \text{for } r^d = -0.05129 \text{ or } r^d = -0.07257. \tag{2.25}$$

This treatment reflects a greater than 5% or 7% decline in the S&P 500 index over an options expiration cycle. Given a 4.53% ($15.7\%/\sqrt{12}$) monthly standard deviation of the S&P 500 index returns over options expiration cycles, a 7% decline can be viewed statistically as a 1.54 sigma event.

The featured breakpoints are compatible with Goldman Sachs EQMOVE. Their model calculates the probabilities of a 5% return movement on the downside and the upside in 30 days, based on free-cash-flow yield, return on equity, Institute for Supply Management data, and capacity utilization (Bloomberg (March 7, 2016)). In contrast, our idea in equation (2.13) is to study probability of disasters that are adapted to information in put option prices. For accuracy, we keep all OTM options up to the minimum tick size.

Please insert Table 2.4 here

Table 2.4 (Panels A and B) provides the estimates of disaster probability (i.e., $\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$ and disaster risk premiums (i.e., $\text{drp}_t[r^d]$),[3] whereas Figure 2.1 depicts the time variation in disaster probabilities.[4] First, the sample average of the probability of a 5% (7%) disaster is 0.082 (0.054). The lower and upper 95% bootstrap confidence intervals for the average of $\mathbb{E}_t^{\mathbb{P}}(\theta[-0.05129 - r_{t+1}])$ are 0.067 and 0.095, respectively. To compute the lower and upper 95% confidence intervals, we randomly draw with replacement $(R_{\mathbf{f},t+1}, \mathbf{z}_{t+1})$. Next, we reestimate $(\eta_0, \boldsymbol{\eta})$. With 1,000 bootstrap trials, we obtain the 1,000 estimates of $(\eta_0, \boldsymbol{\eta})$. The bootstrap inference establishes that our estimates are robust to return perturbations.

Please insert Figure 2.1 here

To examine the relevance of our disaster probabilities, we compute the occurrence probabilities based on actual return realizations over the 311 expiration cycles. We detect 23 (14) declines of magnitude >5% (>7%), resulting in disaster probabilities in the data of 0.074 (0.045). Thus, the unconditional disaster probabilities generated by our method seem aligned with the data attributes.

[3] Options markets imply an average price of a binary put of a greater-than-5% (7%) decline of \$0.124 (\$0.089), with a standard deviation of 0.061 (0.055). Probing further, the bootstrap lower (upper) 95% confidence interval for the average difference between the prices of a 5% and 7% binary put (over the 5/19/1997 to 11/23/2015 period; 223 observations) is 0.045 (0.050).

[4] The accuracy of the calculations in equations (2.13)–(2.14) hinges on the number of strikes $K < K^d$. The sparsity of strikes deeper OTM than 7% precludes us from reliably constructing the price of a 7% binary put prior to 5/19/1997 (often 2 to 4 strikes). In contrast, there is an average of 29 put strikes deeper OTM than 7% after the expiration cycle of 5/19/1997.

Pertinent to our approach is a logistic regression, in which the next-period outcome variable, $\mathbb{Y}_i$, is consistent with a Bernoulli probability function and takes a value 1 with probability $\mathbf{p}_i$ and 0 with probability $1 - \mathbf{p}_i$. We specify $\mathbb{Y}_i \sim \text{Bernoulli}(\mathbb{Y}_i|\mathbf{p}_i)$ and $\mathbf{p}_i = (1 + e^{\psi_0 + \psi_1 x_i})^{-1}$, where $x = \log(\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}]))$. Of interest in Table 2.5 (Panel A) is the rejection of $\psi_1 = 0$, whereby the ψ_1 estimate of 0.98 (0.59) translates into an average marginal effect of 0.062 (0.073) for a 5% (3%) downside return movement.

Please insert Table 2.5 here

Complementing this exercise, we also quantify return upside probabilities as in Result 2. Speaking to the empirical relevance of our approach, the logistic regression results in Table 2.5 (Panel B) reveal that our estimates of upside probabilities have predictive content for 3% and 5% upside return movements. This effect is positive and statistically significant. The takeaway is that the information from the options market is linked to subsequent real-world outcomes.

Our characterization of $\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$ from put options indicates that the probability of a disaster of a fixed size is fluctuating, which is akin to the notion of variable-rate disasters in [94], [95], and [96]. The disaster probabilities are countercyclical. For example, the first principal component of the 17 output and income series in [97, Appendix, Group 1], which includes real income and industrial production, has a correlation of -0.29, with the disaster probability of size 5%. This component inherits the cyclicality of the output series and captures 51.46% of their variance.

The evidence reported in Table 2.4 provides a perspective on the magnitude of the disaster risk premiums. The sample average of $\mathrm{drp}_t[r^d]$ is -63.8% (-73.6%) for a 5% (7%) disaster, and the 95% bootstrap confidence intervals do not bracket zero. We also find that a higher level of the macroeconomic uncertainty index of [98, Figure 1] is associated with higher disaster probabilities (correlation of 0.55). Moreover, $\mathrm{drp}_t[r^d]$ exhibits a correlation of -0.49 with Investor Fears index of [74, equation (32), Figure 5].

Our mechanism for extracting tail risk premiums is different from that of the literature. For example, [99] and [96] assume that jump intensities under the $\mathbb{Q}$ measure are constant multiples of intensities under the $\mathbb{P}$ measure. Moreover, [100] measure jump intensities under $\mathbb{P}$ and $\mathbb{Q}$, and show that their gap is not explained by a constant multiple.

2.3.5 Relevance of disaster uncertainty and economic interpretations

The literature has favored a view that heightened levels of tail uncertainty can percolate through the economy. This has prompted voices that the Federal Reserve act in a timely fashion to remedy tail-risk concerns through forward guidance and macroprudential policies (e.g., *Financial Times* (May 9, 2014)).

Do higher disaster probabilities portend worsening macroeconomy? We explore this question in the context of the following predictive regressions (illustrated by taking $r^d = -0.05129$):

$$y_{\{t \to t+j\}} = \Pi_0 + \Pi_{\text{disaster}} \underbrace{\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])}_{\text{disaster probability via eq. (2.17)}} + \epsilon_{t+j}, \qquad \text{for } j = 1, 3, 6, 12 \text{ months.} \tag{2.26}$$

We employ the database maintained by [97, Appendix] and initially consider the 17, 32, and 14 (properly transformed) variables categorized as output and income; labor markets; and consumption, orders, and inventories, respectively.

For compactness, we feature three dependent variables in our investigation: (i) log change in IP index from "output and income," (ii) log change in all employees: total nonfarm from "labor markets," and (iii) level of the ISM: PMI composite index from "consumption, orders, and inventories." We feature these variables on the grounds that they exhibit high (absolute) correlation with the first principal component of the respective group, with an adjusted R^2 of 94%, 84%, and 67%.

Please insert Table 2.6 here

Table 2.6 indicates that the slope coefficients are negative and statistically significant. For example, the Π_{disaster} estimate is -0.095 (the correlation is -0.41) with IP index regression over the three-month horizon, with a NW[p] of 0.014 and an adjusted R^2 of 16.2%. The H[p] is 0.000, and is corrected for overlapping observations. Additionally, a one standard deviation increase in disaster uncertainty (5.9%) lowers the IP index by 2.24% (400*0.095*0.059) annualized.

Our work implies that efforts by the Federal Reserve to mitigate tail-risk concerns in the equity market can exert a beneficial effect on economic outlook. This positive effect can last up to 12 months on the IP index and total nonfarm. The finding that disaster uncertainty is helpful in predicting macroeconomic outcomes agrees with others, for example, [101], who shows that disaster uncertainty impacts the flow of capital investment and hiring in the labor markets.

One could compare our quantitative results to those from models with time-variation in disaster probabilities, as articulated in [95] and [96], but this analogy is not direct, due to a nontrivial mapping between the parameterizations in equilibrium models with recursive utility versus an estimated reduced-form $m[S_{t+1}]$ specification. While we compute disaster probabilities over the options expiration cycles based on estimated $m[S_{t+1}]$'s, the corresponding one in Wachter's (2013) calibration arises from the probability of a consumption decline larger than 10% annualized. Together with the assumption that dividend equates to levered consumption, the equivalent drop in equity is larger than 20% if a disaster occurs. One could interpret wealth disasters as arising from endogenous time-variation in consumption disaster probabilities, say, in [96], with the observation that we are extracting probabilities from option prices without making distributional assumptions about equity returns.

2.3.6 Gauging the sensibility of our findings

A. Comparing forecast accuracy of conditional return variance (recovered from options) for realized variance. Related to our performance and pricing fit exercises, one may ask how the conditional return variance based on our framework forecasts realized return variance. To study this, we perform the following [102] regression (e.g., [103]), relying on $(\eta_0, \boldsymbol{\eta})$ reported in Table 2.1 (Panel B) to calculate $\mathrm{var}_t^{\mathbb{P}}[r_{t+1}]$:

$$\underbrace{\mathrm{rv}_{\{t \to t+1\}}}_{\text{realized return variance}} = \beta_0 + \beta_v \underbrace{\mathrm{var}_t^{\mathbb{P}}[r_{t+1}]}_{\text{inferred from options at date } t} + \epsilon_{t+1}. \tag{2.27}$$

The realized return variance is computed as the sum of daily squared returns: $\mathrm{rv}_{\{t \to t+1\}} = r^2_{t+\Delta,\Delta} + r^2_{t+2\Delta,\Delta} + \ldots + r^2_{t+\hbar\Delta,\Delta}$, with $r_{t+\hbar\Delta,\Delta} \equiv \log(\frac{S_{t+\hbar\Delta}}{S_{t+(\hbar-1)\Delta}})$, Δ is one day, and $\hbar$ is the number of trading days over an options expiration cycle. Setting $g[S_{t+1}] = \frac{\{\log(\frac{S_{t+1}}{S_t})\}^n}{m[S_{t+1}]}$, for integer n, we compute $\mathrm{var}_t^{\mathbb{P}}[r_{t+1}]$, from options, as

$$\mathrm{var}_t^{\mathbb{P}}[r_{t+1}] \equiv \underbrace{R^{-1}_{\mathbf{f},t+1}\, \mathbb{E}_t^{\mathbb{Q}}(\frac{r^2_{t+1}}{m[S_{t+1}]})}_{\mathbb{E}_t^{\mathbb{P}}\left(r^2_{t+1}\right) \text{ from options}} - \underbrace{\{R^{-1}_{\mathbf{f},t+1}\, \mathbb{E}_t^{\mathbb{Q}}(\frac{r_{t+1}}{m[S_{t+1}]})\}^2}_{\{\mathbb{E}_t^{\mathbb{P}}(r_{t+1})\}^2 \text{ from options}}, \quad \text{where} \tag{2.28}$$

$$\mathbb{E}_t^{\mathbb{P}}\left(r^2_{t+1}\right) = \int_0^{S_t} a[K]\, P_t[K]\, dK + \int_{S_t}^{\infty} a[K]\, C_t[K]\, dK \quad \text{and} \tag{2.29}$$

$$\mathbb{E}_t^{\mathbb{P}}\left(r_{t+1}\right) = 1 - R^{-1}_{\mathbf{f},t+1} + \int_0^{S_t} b[K]\, P_t[K]\, dK + \int_{S_t}^{\infty} b[K]\, C_t[K]\, dK \tag{2.30}$$

With $\eta_1 \equiv 0$, and letting $m[K] \equiv \exp(\eta_0 - 1 + \eta_2(\frac{\{\log(\frac{K}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}))$, we determine

$$\begin{aligned} a[K] &= \{\log(\frac{K}{S_t})\}^2(\frac{4\eta_2^2\{\log(\frac{K}{S_t})\}^2 m[K]^{-1}}{K^2\mathrm{v}_t^2} + \frac{2\eta_2 \log\left(\frac{K}{S_t}\right) m[K]^{-1}}{K^2\mathrm{v}_t} - \frac{2\eta_2 m[K]^{-1}}{K^2\mathrm{v}_t}) \\ &\quad - \frac{8\eta_2\{\log(\frac{K}{S_t})\}^2 m[K]^{-1}}{K^2\mathrm{v}_t} + (\frac{2}{K^2} - \frac{2\log(\frac{K}{S_t})}{K^2})m[K]^{-1} \quad \text{and} \end{aligned} \tag{2.31}$$

$$\begin{aligned} b[K] &= \log(\frac{K}{S_t})(\frac{4\eta_2^2\{\log\left(\frac{K}{S_t}\right)\}^2 m[K]^{-1}}{K^2\mathrm{v}_t^2} + \frac{2\eta_2 \log(\frac{K}{S_t}) m[K]^{-1}}{K^2\mathrm{v}_t} - \frac{2\eta_2 m[K]^{-1}}{K^2\mathrm{v}_t}) \\ &\quad - \frac{4\eta_2 \log(\frac{K}{S_t}) m[K]^{-1}}{K^2\mathrm{v}_t} - \frac{m[K]^{-1}}{K^2}. \end{aligned} \tag{2.32}$$

Our rationale for the Mincer-Zarnowitz regression is that the ex-post value of the realized variance is an unbiased estimator of the conditional variance; that is, $\mathbb{E}_t^{\mathbb{P}}(\mathrm{rv}_{\{t \to t+1\}})$ equates conditional variance ([104, Corollary 1])). While the literature often uses the realized variance to estimate the conditional variance, we investigate whether $\mathrm{var}_t^{\mathbb{P}}[r_{t+1}]$ based on our approach is a good forecast of the subsequent realized return variance.

Please insert Table 2.7 here

Table 2.7 reports that the OLS estimates of β_v and β_0 in equation (2.27) are 0.937 and 0.0004, respectively. Based on [105], with lag automatically selected as in [106], the p-value for the hypothesis of $\beta_v = 0$ ($\beta_0 = 0$) is 0.000 (0.076), and the p-value for the Wald test $\beta_v = 1$ is 0.727 (a low p-value implies rejection). We further evaluate predictability of $\mathrm{var}_t^{\mathbb{P}}[r_{t+1}]$ in an out-of-sample framework, using an expanding window with 120 initial observations. The adjusted mean-squared prediction error (MSPE) statistic of [107] generates a p-value of 0.006, where the hypothesis is that $\mathrm{var}_t^{\mathbb{P}}[r_{t+1}]$ carries no predictability.

To draw further comparisons, we employ $\mathrm{var}_t^{\mathbb{Q}}[r_{t+1}]$ to predict $\mathrm{rv}_{\{t \to t+1\}}$. The β_v estimate is 0.669 and the hypothesis of $\beta_v = 1$ is rejected (p-value is 0.000). Overall, the estimates of conditional return variance based on our method appear compatible with the data counterparts.

Please insert Table 2.8 here

B. Estimates of other objects of interest. Table 2.8 presents the estimates of (i) the conditional expected logarithmic return, (ii) the conditional return variance (its square root), and (iii) the conditional variance risk premium, implied by our approach and option prices. The average conditional expected logarithmic return, which is consistent with equation (2.30), is 11.2% annualized and lies between the bootstrapped 95% confidence intervals of 2.3% and 16.4% of the average *realized* equity returns shown below:

	Mean	Bootstrap 95% CI	SD
Realized equity (logarithmic) return	0.094	[0.023 0.164]	0.157
Realized return volatility	0.161	[0.141 0.182]	0.095

Additionally, the average conditional return volatility is 15.4% annualized and lies between the bootstrapped 95% confidence intervals. The average variance risk premium is -68% and is of the order of disaster risk premium. In sum, our method to recover from option prices provide economically sensible estimates of these considered quantities.

2.4 Studying the spanning hypothesis in the Treasury market

Our aim is to use option-inferred conditional expected excess return of Treasury bond futures under the $\mathbb{P}$ measure to address the empirical controversy associated with the spanning hypothesis. Whereas [65] show that certain variables are relevant to predicting Treasury bond returns, [66] show that only level, slope, and curvature — reflecting the first, second, and third principal components extracted from Treasury yields — matter statistically for subsequent Treasury bond returns.

Other variables are statistically absent of forecasting ability.

Let F_t denote the futures price on the Treasury bond. For a fully collateralized futures position, $\frac{F_{t+1}}{F_t}-1$ represents the excess return of a long position in the futures contract. This implies that

$$\mathbb{E}_t^{\mathbb{Q}}(\frac{F_{t+1}}{F_t}-1)=0. \qquad \text{(the martingale condition holds)} \tag{2.33}$$

The context is that the variable $\mathbb{E}_t^{\mathbb{Q}}(\frac{F_{t+1}}{F_t}-1)$ contains *no* information about forecasting $\frac{F_{t+1}}{F_t}-1$.

In our exercises, we consider an $m[F_{t+1}]$ specification analogous to equation (2.16) of the type

$$m[F_{t+1}]=\exp\left(\eta_0-1+\boldsymbol{\eta}'\mathbf{z}_{t+1}[F_{t+1}]\right), \text{ where } \underbrace{\mathbf{z}_{t+1}[F_{t+1}]}_{2\times 1}\equiv\begin{pmatrix}\overbrace{\frac{F_{t+1}}{F_t}-1}^{\text{excess return of futures}}\\ \frac{\{\log(\frac{F_{t+1}}{F_t})\}^2}{\mathrm{v}_t}-R_{\mathbf{f},t+1}\end{pmatrix}. \tag{2.34}$$

The price of the volatility contract payoff $\{\log(\frac{F_{t+1}}{F_t})\}^2$ can be synthesized as $\mathrm{v}_t\equiv R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\{\log(\frac{F_{t+1}}{F_t})\}^2)=\int_0^{F_t}\frac{2(1+\log(\frac{F_t}{K}))}{K^2}P_t^{\text{futures}}[K]\,dK+\int_{F_t}^{\infty}\frac{2(1-\log(\frac{K}{F_t}))}{K^2}C_t^{\text{futures}}[K]\,dK$. Here, $P_t^{\text{futures}}[K]$ ($C_t^{\text{futures}}[K]$) is the put (call) option price on the futures of the Treasury bond with strike price K.

With our method to achieve recovery, we can derive

$$\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t}-1)=\int_0^{F_t}\mathrm{w}[K]\,P_t^{\text{futures}}[K]\,dK\;+\;\int_{F_t}^{\infty}\mathrm{w}[K]\,C_t^{\text{futures}}[K]\,dK, \tag{2.35}$$

where, letting $m[K] \equiv \exp(\eta_0 - 1 + \eta_1(\frac{K}{F_t} - 1) + \eta_2(\frac{\{\log(\frac{K}{F_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}))$, we obtain

$$\begin{aligned}\mathrm{w}[K] &= \left(\frac{K}{F_t} - 1\right)\left(\left(\frac{2\eta_2 \log(\frac{K}{F_t})}{K^2 \mathrm{v}_t} - \frac{2\eta_2}{K^2 \mathrm{v}_t}\right)(m[K])^{-1}\right. \\ &\quad \left. + \left(-\frac{2\eta_2 \log(\frac{K}{F_t})}{K \mathrm{v}_t} - \frac{\eta_1}{F_t}\right)^2 (m[K])^{-1}\right) + \frac{2\left(-\frac{2\eta_2 \log(\frac{K}{F_t})}{K \mathrm{v}_t} - \frac{\eta_1}{F_t}\right)(m[K])^{-1}}{F_t} \qquad (2.36)\end{aligned}$$

Relying on the minimum discrepancy problem (as before), we estimate the parameters of $m[F_{t+1}]$ in equation (2.34). Table 2.9 shows that the estimate of η_1 is negative, while that of η_2 is positive, and each are statistically significant. The unrestricted specification with $\eta_1 \neq 0$ and $\eta_2 \neq 0$ is empirically relevant in light of the lowest [92] distance measure (see Panel B of Table 2.9).

What do we do differently from others? We consider monthly excess returns based on bond futures and we do not draw inferences based on overlapping 12-month excess Treasury bond returns. Next, we focus on $\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t} - 1)$ extracted from options on futures of Treasury bond using our approach, and explore its ability to predict excess return of futures on a Treasury bond, beyond level, slope, and curvature variables.

Specifically, we work with the following empirical specification to study the spanning hypothesis:

$$\underbrace{\frac{F_{t+1}}{F_t} - 1}_{\text{Excess return of futures on a 10- or 30-year Treasury bond}} = \pi_0 + \underbrace{\sum_{j=1}^{3} \pi_j \, \mathrm{PC}_t^{\lfloor j \rfloor}}_{\text{Level, slope, and curvature variables}} + \pi^{\text{spanning}} \, \mathbb{X}_t + \epsilon_{t+1}, \qquad (2.37)$$

where $\mathrm{PC}_t^{\lfloor j \rfloor}$, for $j = 1, 2, 3$, are the first, second, and third principal components extracted from available Treasury yields (maturity of 1 month, 3 month, 1, 2, 5, 7, 10, 20, and 30 years). We investigate whether

$$\pi^{\text{spanning}} = 0 \text{ for } \textit{any} \text{ predictor } \mathbb{X}_t, \text{ and, specifically for } \mathbb{X}_t = \underbrace{\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t} - 1)}_{\text{recovered from options}} . \tag{2.38}$$

That is, no other variable besides level, slope, and curvature is a significant predictor of excess return of Treasury bond futures.

Table 2.10 highlights the role of level, slope, and curvature variables in jointly predicting excess Treasury bond futures returns. Our inferences about parameter significance rely on NW[p] as well as the p-values based on a parametric bootstrap, denoted as PB[p]. The parametric bootstrap serves as a safeguard and is implemented in line with [108] and [66].

Please insert Table 2.10 here

Our results affirm the predictive ability of level and slope (i.e., $\text{PC}_t^{\lfloor 1 \rfloor}$ and $\text{PC}_t^{\lfloor 2 \rfloor}$) for excess return of Treasury bond futures in the predictive regressions. With a NW[p] of 0.301, the curvature variable is not statistically significant. The sign of the estimated predictive coefficients are the same as those using bond returns (e.g., [66, Table 3]). The χ^2 test that the parameters are jointly zero, depicted by Joint[p], implies that the hypothesis of unpredictability of Treasury bond futures returns is rejected.

On top of yield curve predictors, we find that $\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t} - 1)$, inferred using equation (2.35), is positively associated with next-period excess return of futures on the 10-year (30-year) Treasury bond, manifesting a π^{spanning} of 0.542 (0.788) with a NW[p] of 0.011 (0.076), and a PB[p] of 0.063 (0.222). Furthermore, the reported values of Joint[p] are supportive of joint parameter significance. Thus, our empirical evidence does not favor the spanning hypothesis. The predictor $\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t}-1)$ is also significant, according to NW[p], on a stand-alone basis; that is, in the absence of $\text{PC}_t^{\lfloor 1 \rfloor}$, $\text{PC}_t^{\lfloor 2 \rfloor}$, and $\text{PC}_t^{\lfloor 3 \rfloor}$.

For the futures on the 10-year Treasury bond, our augmented specification increases the adjusted-R^2 from 1.3% to 2.7%. The reported adjusted-R^2's are lower compared to those in [66, Table 2], because we have monthly nonoverlapping observations for Treasury futures returns, as opposed to 12-month overlapping monthly observations for Treasury bond returns.

In sum, our $\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t}-1)$, inferred from options on Treasury bond futures, offers differential information content, and this predictor is mildly associated with $\text{PC}_t^{\lfloor 1 \rfloor}$, $\text{PC}_t^{\lfloor 2 \rfloor}$, and $\text{PC}_t^{\lfloor 3 \rfloor}$ (the largest cross-correlation is -0.27). The relevance of our regressions with $\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t}-1)$ deviates from [109], who use forward rates; from [110], who use variables extracted from a large macroeconomic dataset; from [111], who use Treasury bond supply; from [65], who use measures of inflation and output; and from [112], who use trends in inflation and inflation expectations.

2.5 Conclusion

This paper presents an approach to obtain expressions for the time-t conditional expectation of return functions under the real-world probability measure in terms of the prices of spanning securities known at time-t (the risk-free bond, the asset, and options on the asset). Developed in a discrete-time environment, our method is ingrained in asset pricing theory, is tractable for a broad class of functions of asset returns, and our results obtain without explicitly modeling or estimating the return density function under the real-world probability measure.

What are the steps of our methodology? First, we recognize that the SDF — which connects the real-world probability measure to the risk-neutral probability measure — is not known and explore projecting the SDF onto asset returns. Second, we estimate the projection of the SDF and then invert the Girsanov theorem to link the risk-neutral and real-world probability measures to recover the real-world expectation of return functions. Under this method, the real-world expectation of the return function is the risk-neutral expectation of the same function divided by the projected SDF. Third, we exploit the feature that the risk-neutral expectation of return functions (satisfying certain smoothness properties) can be obtained from the prices of spanning assets. This formulation delivers the conditional expectation of return functions in terms of the (known) prices of the risk-free bond, the underlying asset, and puts and calls on the underlying asset.

We illustrate the workings of our method in the context of formulating wealth disaster probabilities, upside return probabilities, and conditional expected return and volatility. Our investigation employs (i) equity index, and (ii) Treasury bond futures, and the options written on them, and we examine empirical consistency of our formulations. In our implementations, we consider a set of projected SDFs that preserve positivity, correctly price a set of returns, and are parsimonious from the perspective of estimation.

Overall, our methodology improves on approaches to estimating equity tail probabilities and the bond futures risk premium. All in all, our work is consistent with research that assigns a role to the information in traded option prices and their connections to economic phenomena.

2.6 Appendix

2.6.1 Representation of the probability of disasters in equation (2.13)

For the proof to follow,[5] we make our analysis self-contained and clarify some intermediate steps in [113, Appendix 1, page 20].

For the projected SDF, $m[S_{t+1}]$, that is strictly positive and twice-continuously differentiable, we consider the function

$$g[S_{t+1}] \;=\; \frac{\theta[K^d - S_{t+1}]}{m[S_{t+1}]}, \tag{2.39}$$

[5] In his private communications with us, Ngoc-Khanh Tran provided an alternative way to think about the proof and suggested refinements that led to many improvements.

where we recall that $\theta[x]$ is the Heaviside theta function

$$\theta[x] \;\equiv\; \frac{d}{dx}\max(x,0) \;=\; \begin{cases} 0 & \text{for } x < 0, \\ 1 & \text{for } x > 0. \end{cases} \tag{2.40}$$

The derivative of $\theta[x]$, denoted by $\theta'[x]$, is the Dirac delta function

$$\delta[x] = \theta'[x], \qquad \text{implying} \qquad \theta[x] = \int_{-\infty}^{x} \delta[s]\, ds. \tag{2.41}$$

The Dirac delta function, $\delta[x]$, satisfies $1 = \int_{-\infty}^{\infty} \delta[x]\, dx$.

The sifting property asserts that

$$\int_{-\infty}^{\infty} g[x]\ \delta[x - x_0]\ dx \;=\; g[x_0]. \tag{2.42}$$

For concreteness, suppose $x = K - S_{t+1}$. Then

$$\theta[K - S_{t+1}] = \frac{d}{dK}\max(K - S_{t+1}, 0) \quad \text{and} \tag{2.43}$$

$$\delta[K - S_{t+1}] = \frac{d}{dK}\theta[K - S_{t+1}] = \frac{d^2}{dK^2}\{\max(K - S_{t+1}, 0)\}. \tag{2.44}$$

Now consider a test function $g[S_{t+1}]$. By the sifting property in equation (2.42),

$$
\begin{aligned}
g[S_{t+1}] &= \int_0^\infty g[K]\,\delta[K-S_{t+1}]\,dK && (2.45)\\
&= \int_0^\infty g[K]\,\delta[S_{t+1}-K]\,dK \qquad (\text{symmetry } \delta[x]=\delta[-x]) && (2.46)\\
&= \int_0^{K^d} g[K]\,\delta[K-S_{t+1}]\,dK \;+\; \underbrace{\int_{K^d}^\infty g[K]\,\delta[S_{t+1}-K]\,dK}_{=0,\text{ for claims with } S_{t+1}>K^d} && (2.47)\\
&= \int_0^{K^d} g[K]\,d\theta[K-S_{t+1}] \qquad (\text{since } \delta[K-S_{t+1}]\,dK = d\theta[K-S_{t+1}]) && (2.48)
\end{aligned}
$$

Setting $u=g[K]$ and $v=\theta[K-S_{t+1}]$, integration by parts (i.e., $\int u\,dv = uv - \int v du$) implies that

$$
\begin{aligned}
g[S_{t+1}] &= \int_0^{K^d} g[K]\,d\theta[K-S_{t+1}]\\
&= g[K]\,\theta[K-S_{t+1}]\Big|_0^{K^d} - \int_0^{K^d} \theta[K-S_{t+1}]g'[K]\,dK. && (2.49)
\end{aligned}
$$

Substituting $u=g'[K]$ and $v=\max(K-S_{t+1},0)$, a repeated application of integration by parts to $\int_0^{K^d}\theta[K-S_{t+1}]\,g'[K]\,dK$ implies that

$$
\int_0^{K^d} \underbrace{g'[K]}_{u}\,\underbrace{\theta[K-S_{t+1}]\,dK}_{dv} = g'[K]\max(K-S_{t+1},0)\Big|_0^{K^d} - \int_0^{K^d}\max(K-S_{t+1},0)\,g''[K]\,dK. \tag{2.50}
$$

The steps above allow us to write $g[S_{t+1}] = \int_0^{K^d} g[K]\,\delta[K-S_{t+1}]\,dK$ as

$$
g[S_{t+1}] = g[K]\,\theta[K-S_{t+1}]\Big|_0^{K^d} - g'[K]\max(K-S_{t+1},0)\Big|_0^{K^d} + \int_0^{K^d}\max(K-S_{t+1},0)\,g''[K]\,dK. \tag{2.51}
$$

Equation (2.13) then follows.

Finally, the expression for the price of the binary put can be obtained as

$$\begin{aligned} \theta[K^d - S_{t+1}] &= \frac{d}{dK^d} \overbrace{\max(K^d - S_{t+1}, 0)}^{\equiv\, h[K^d]}, \\ &\approx \frac{\max(K^d + \Delta K - S_{t+1}, 0) \; - \; \max(K^d - \Delta K - S_{t+1}, 0)}{2\Delta K} \end{aligned} \tag{2.52}$$

The second-order approximation in (2.52) follows since $h'[K] = \frac{h[K+\Delta K] - h[K-\Delta K]}{2\Delta K} + O[(\Delta K)^2]$, as $\Delta K \to 0$ (e.g., [114, page 208]). Thus,

$$\mathbb{B}_t^{\text{put}}[K^d] \equiv R_{\mathbf{f},t+1}^{-1}\, \mathbb{E}_t^{\mathbb{Q}} \left(\theta[K^d - S_{t+1}]\right) \;\approx\; \frac{P_t[K^d + \Delta K] \; - \; P_t[K^d - \Delta K]}{2\Delta K}. \tag{2.53}$$

The binary put with strike K^d can be synthesized via two adjacent put options. ■

2.6.2 Representation of the probability of upside in equation (2.15)

For functions representing the upside of returns, suppose $x = S_{t+1} - K$ for some K. Then

$$\theta[S_{t+1} - K] \;=\; -\frac{d}{dK} \max(S_{t+1} - K, 0) \quad \text{and} \tag{2.54}$$

$$\delta[S_{t+1} - K] \;=\; -\frac{d}{dK} \theta[S_{t+1} - K] = \frac{d^2}{dK^2} \{\max(S_{t+1} - K, 0)\}. \tag{2.55}$$

Thus, $-\delta[S_{t+1} - K]dK = d\theta[S_{t+1} - K]$.

Consider threshold K^u such that $\frac{S_t}{K^u} - 1 < 0$. By the sifting property in equation (2.42),

$$g[S_{t+1}] = \underbrace{\int_0^{K^u} g[K]\,\delta[K - S_{t+1}]\,dK}_{=0,\text{ for claims with } S_{t+1}<K^u} + \int_{K^u}^{\infty} g[K]\,\delta[S_{t+1} - K]\,dK \tag{2.56}$$

$$= \int_{K^u}^{\infty} g[K]\,\{-d\theta[S_{t+1} - K]\} \qquad (\text{since } \delta[S_{t+1} - K]\,dK = -d\theta[S_{t+1} - K]) \tag{2.57}$$

Setting $u = g[K]$ and $v = -\theta[S_{t+1} - K]$, integration by parts implies that

$$g[S_{t+1}] = \int_{K^u}^{\infty} g[K]\,\{-d\theta[K - S_{t+1}]\} \tag{2.58}$$

$$= -g[K]\,\theta[S_{t+1} - K]\Big|_{K^u}^{\infty} + \int_{K^u}^{\infty} \theta[S_{t+1} - K]\,g'[K]\,dK. \tag{2.59}$$

Substituting $u = g'[K]$ and $v = -\max(S_{t+1}-K, 0)$, which implies $dv = -\frac{d}{dK}\max(S_{t+1}-K, 0)$, a repeated application of integration by parts to $\int_{K^u}^{\infty} \theta[S_{t+1} - K]\,g'[K]\,dK$, implies that

$$\int_{K^u}^{\infty} \underbrace{g'[K]}_{u}\,\underbrace{\theta[S_{t+1} - K]\,dK}_{dv} = -g'[K]\max(S_{t+1}-K, 0)\Big|_{K^u}^{\infty} + \int_{K^u}^{\infty} \max(S_{t+1}-K, 0)\,g''[K]\,dK.$$

Hence, we can write $g[S_{t+1}] = \int_{K^u}^{\infty} g[K]\,\delta[S_{t+1} - K]\,dK$ as

$$g[S_{t+1}] = -g[K]\,\theta[S_{t+1}-K]\Big|_{K^u}^{\infty} - g'[K]\max(S_{t+1}-K, 0)\Big|_{K^u}^{\infty} + \int_{K^u}^{\infty} \max(S_{t+1}-K, 0)\,g''[K]\,dK. \tag{2.60}$$

This is the desired general expression. Additionally,

$$
\begin{aligned}
R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(g[S_{t+1}]) \;&=\; g[K^u]\,R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\theta[S_{t+1}-K^u]) \;+\; g'[K^u]R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\max(S_{t+1}-K^u,0)) \\
&\quad+\; \int_{K^u}^{\infty} g''[K]R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\max(S_{t+1}-K,0))\;dK \qquad (2.61) \\
&=\; g[K^u]\,\mathbb{B}_t^{\text{call}}[K^u] \;+\; g'[K^u]\,C_t[K^u] \;+\; \int_{K^u}^{\infty} g''[K]\,C_t[K]\;dK \qquad (2.62)
\end{aligned}
$$

where $\mathbb{B}_t^{\text{call}}[K^u]$ is the price of the binary call. Finally, we have

$$
\theta[S_{t+1}-K^u] \;=\; -\frac{d}{dK^u}\max(S_{t+1}-K^u,0) \approx \frac{\max(S_{t+1}-\{K^u-\Delta K\},0) \;-\; \max(S_{t+1}-\{K^u+\Delta K\},0)}{2\Delta K} \qquad (2.63)
$$

Thus,

$$
\mathbb{B}_t^{\text{call}}[K^u] \equiv R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}\,(\theta[S_{t+1}-K^u]) \;\approx\; \frac{C_t[K^u-\Delta K] \;-\; C_t[K^u+\Delta K]}{2\Delta K}. \qquad (2.64)
$$

The binary call with strike K^u can be synthesized via two adjacent call options. ■

	Estimates	95% Bootstrap CI Lower	Upper	Bootstrapped values Mean	SD
Panel A: Full specification					
η_0	1.043	1.018	1.117	1.052	0.027
η_1	1.992	-1.290	5.425	1.959	1.621
η_2	0.238	0.140	0.436	0.254	0.077
Panel B: Restricted $\eta_1 = 0$ (only excess return of equity volatility)					
η_0	1.040	1.012	1.114	1.046	0.027
η_2	0.192	0.103	0.389	0.207	0.076
Panel C: Restricted $\eta_2 = 0$ (only excess return of equity market)					
η_0	1.002	0.998	1.023	1.004	0.007
η_1	-1.967	-5.346	0.358	-2.093	1.412

Table 2.1: **Estimates and bootstrap confidence intervals for the parameters of the $\mathbf{m[S_{t+1}]}$ specification**
Reported parameter estimates are based on the following specification of $m[S_{t+1}]$:

$$m[S_{t+1}] = \exp\left(\eta_0 - 1 + \boldsymbol{\eta}'\mathbf{z}_{t+1}[S_{t+1}]\right), \quad \text{where} \quad \underbrace{\mathbf{z}_{t+1}[S_{t+1}]}_{2\times 1} \equiv \begin{pmatrix} \frac{S_{t+1}}{S_t} - R_{\mathbf{f},t+1} \\ \frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1} \end{pmatrix}.$$

The volatility contract, with payoff $\{\log(\frac{S_{t+1}}{S_t})\}^2$, is not directly traded, and its price is synthesized, as $\mathrm{v}_t \equiv R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\{\log(\frac{S_{t+1}}{S_t})\}^2) = \int_0^{S_t} \frac{2(1+\log(\frac{S_t}{K}))}{K^2}\,P_t[K]\,dK + \int_{S_t}^{\infty} \frac{2(1-\log(\frac{K}{S_t}))}{K^2}\,C_t[K]\,dK$. The parameters $(\eta_0, \boldsymbol{\eta}) \in \mathbb{R}^3$ solve the following optimization problem: $\inf_{(\eta_0,\boldsymbol{\eta})} \; -\mathbb{E}^{\mathbb{P}}(R_{\mathbf{f},t+1}^{-1})\,\eta_0 \;+\; \mathbb{E}^{\mathbb{P}}(\exp(\eta_0 - 1 + \boldsymbol{\eta}'\mathbf{z}_{t+1}))$. To obtain the bootstrap distribution of $(\eta_0, \boldsymbol{\eta})$, we consider a procedure that randomly selects, with replacement, $(R_{\mathbf{f},t+1}, \mathbf{z}_{t+1})$ jointly via a stationary bootstrap. Then we recompute η_0, η_1, and η_2 in each bootstrap draw. We perform 1,000 bootstrap draws, yielding the reported mean, standard deviation (SD), and lower and upper 95% bootstrap confidence intervals.

	All eight assets			Excluding 5% OTM options			Excluding calls		
	$\eta_1 \neq 0$ $\eta_2 \neq 0$	$\eta_1 = 0$ $\eta_2 \neq 0$	$\eta_1 \neq 0$ $\eta_2 = 0$	$\eta_1 \neq 0$ $\eta_2 \neq 0$	$\eta_1 = 0$ $\eta_2 \neq 0$	$\eta_1 \neq 0$ $\eta_2 = 0$	$\eta_1 \neq 0$ $\eta_2 \neq 0$	$\eta_1 = 0$ $\eta_2 \neq 0$	$\eta_1 \neq 0$ $\eta_2 = 0$
$\mathrm{HJ}^{\mathrm{dist}}$	0.85	$\boxed{0.47}$	0.77	0.56	$\boxed{0.29}$	0.50	0.48	$\boxed{0.37}$	0.70
Lower CI	0.31	0.21	0.44	0.16	0.09	0.23	0.09	0.08	0.31
Upper CI	2.12	1.43	1.16	1.20	0.85	0.80	1.00	0.93	1.06
$\frac{\mathrm{HJ}^{\mathrm{dist}}_{\eta_1 \equiv 0}}{\mathrm{HJ}^{\mathrm{dist}}_{\mathrm{Full}}} - 1$	-45%			-48%			-23%		
p-value	0.079			0.058			0.095		

Table 2.2: **Specification errors of $m[S_{t+1}]$ using the Hansen-Jagannathan (1997) distance measure** We evaluate the ability of our model in equation (2.16) to price eight assets (or a subset of them): gross return of the risk-free bond ($R_{\mathbf{f},t+1}$), gross return of the equity market index ($R_{t+1} = \frac{S_{t+1}}{S_t}$), gross return of the ℓ% out-of-the-money (OTM) put option, calculated as $R_{t+1,\ell\%,\mathrm{put}} \equiv \max(S_t e^{-\ell} - S_{t+1}, 0)/P_t[S_t e^{-\ell}]$, and gross return of the ℓ% out-of-the-money call option, calculated as $R_{t+1,\ell\%,\mathrm{call}} \equiv \max(S_{t+1} - S_t e^{+\ell}, 0)/C_t[S_t e^{+\ell}]$. The gross return vector employed in our calculations is defined as follows: $\mathbf{R}_{t+1} = \begin{pmatrix} R_{\mathbf{f},t+1} \\ R_{t+1} \\ R_{t+1,5\%,\mathrm{put}} \\ R_{t+1,3\%,\mathrm{put}} \\ R_{t+1,1\%,\mathrm{put}} \\ R_{t+1,1\%,\mathrm{call}} \\ R_{t+1,3\%,\mathrm{call}} \\ R_{t+1,5\%,\mathrm{call}} \end{pmatrix}$. We report the [92] distance as $\mathrm{HJ}^{\mathrm{dist}} \equiv \sqrt{\mathbf{e}'_{t+1}\mathbf{I}^{-1}\mathbf{e}_{t+1}}$, where $\mathbf{e}_{t+1} \equiv \mathbb{E}^{\mathbb{P}}(m[S_{t+1}]\,\mathbf{R}_{t+1}) - \mathbf{1}$. The use of identity matrix, $\mathbf{I}$, follows from [93, page 216], and we report the bootstrap confidence intervals. Specifically, we randomly select, with replacement, $(\mathbf{R}_{t+1}, \mathbf{z}_{t+1})$ jointly via stationary bootstrap. There are 311 options expiration cycles over the sample period of January 1990 to December 2015. The last row reports the bootstrap p-value as the proportion of draws for which $\frac{\mathrm{HJ}^{\mathrm{dist}}_{\eta_1 \equiv 0}}{\mathrm{HJ}^{\mathrm{dist}}_{\mathrm{Full}}} - 1 \geq 0$.

	Equity premium	Size spread	Value spread	Momentum spread
η_0	1.05 [1.02 1.15]	1.04 [1.01 1.12]	1.04 [1.01 1.13]	1.17 [1.10 1.33]
η_2	0.21 [0.11 0.48]	0.19 [0.11 0.45]	0.19 [0.11 0.49]	0.58 [0.40 1.13]
$\eta_{2,\mathfrak{u}}$	1.99 [-2.49 5.67]	0.57 [-2.25 4.25]	-0.08 [-2.80 3.57]	-5.89 [-11.57 -3.91]
$\frac{\mathrm{HJ}^{\mathrm{dist}}_{\eta_{2,\mathfrak{u}}\equiv 0}}{\mathrm{HJ}^{\mathrm{dist}}_{\mathrm{Full}}} - 1$	11%	0.6%	-0.3%	-62%
p-value	0.662	0.593	0.677	0.054

Table 2.3: **Pricing errors and bootstrap confidence intervals for the parameters of $\mathbf{m[S_{t+1}]}$ when the loading on the excess return of volatility is time-varying** Reported parameter estimates are based on the following specification of $m[S_{t+1}]$ that incorporates time-varying loading, as follows: $m[S_{t+1}] = \exp(\eta_0 - 1 + \{\eta_2 + \underbrace{\eta_{2,\mathfrak{u}}\,\mathfrak{u}_t}_{\text{time-varying}}\}\underbrace{(\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1})}_{\text{excess return of volatility}})$, where $\mathfrak{u}_t$ is a conditioning variable that is taken to be (constructed by compounding daily returns over options expiration cycles) — equity premium: excess return of the market index, size spread: return of small-sized stock portfolio minus the return of large-sized stock portfolio, value spread: return of value stock portfolio minus the return of growth stock portfolio, or momentum spread: return of winner stock portfolio minus the return of loser stock portfolio. As in the note to Table 2.1, we compute the price $\mathrm{v}_t \equiv R^{-1}_{\mathbf{f},t+1}\,\mathbb{E}^{\mathbb{Q}}_t(\{\log(\frac{S_{t+1}}{S_t})\}^2) = \int_0^{S_t} \frac{2(1+\log(\frac{S_t}{K}))}{K^2}\,P_t[K]\,dK + \int_{S_t}^{\infty} \frac{2(1-\log(\frac{K}{S_t}))}{K^2}\,C_t[K]\,dK$. The parameters $(\eta_0, \eta_2, \eta_{2,\mathfrak{u}}) \in \mathbb{R}^3$ solve the following optimization problem: $\inf_{(\eta_0,\eta_2,\eta_{2,\mathfrak{u}})} \; -\mathbb{E}^{\mathbb{P}}(R^{-1}_{\mathbf{f},t+1})\,\eta_0 \;+\; \mathbb{E}^{\mathbb{P}}(\exp(\eta_0 - 1 + \{\eta_2 + \eta_{2,\mathfrak{u}}\,\mathfrak{u}_t\}(\frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}))$. To obtain the bootstrap distribution of $(\eta_0, \eta_2, \eta_{2,\mathfrak{u}})$, we consider a procedure that randomly selects, with replacement $(R_{\mathbf{f},t+1}, \frac{\log(\frac{S_{t+1}}{S_t})}{\mathrm{v}_t}, \mathfrak{u}_t)$ jointly via stationary bootstrap. Then we recompute η_0, η_2, and $\eta_{2,\mathfrak{u}}$ in each bootstrap draw. We perform 1,000 bootstrap draws, yielding the lower and upper 95% bootstrap confidence intervals presented in square brackets.

	Average	Bootstrap CI		Across 1,000 bootstraps						
		Lower	Upper	Mean	SD	P5	P25	P50	P75	P95
Panel A: Disaster of size 5% (311 expiration cycles, 1/22/1990–11/23/2015)										
$\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$	0.082	0.067	0.095	0.082	0.007	0.069	0.077	0.082	0.086	0.093
$\text{drp}_t[r^d]$	-0.638	-0.368	-1.195	-0.674	0.213	-0.411	-0.536	-0.648	-0.745	-1.086
Panel B: Disaster of size 7% (233 expiration cycles, 5/19/1997–11/23/2015)										
$\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$	0.054	0.042	0.065	0.054	0.006	0.044	0.051	0.054	0.057	0.063
$\text{drp}_t[r^d]$	-0.736	-1.140	-0.483	-0.759	0.178	-1.059	-0.836	-0.746	-0.644	-0.526

Table 2.4: **Disaster probabilities under the $\mathbb{P}$ measure**The reported results for the disaster probabilities, that is, $\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$, rely on equation (2.13) and further on equation (2.17), when $m[S_{t+1}]$ is based on equation (2.16) and imposing $\eta_1 \equiv 0$. The calculation of the disaster risk premiums uses equation (2.14), that is, $\text{drp}_t[r^d] = \log\left(\frac{\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])}{\mathbb{E}_t^{\mathbb{Q}}(\theta[r^d - r_{t+1}])}\right)$. The wealth disasters of size 5% (7%) correspond to r^d of -0.05129 (-0.07257) over the expiration cycles of the S&P 500 index options, and the results are shown across the 1,000 bootstraps.

Probability	ψ_0	p-val.	ψ_1	p-val.	log-lik.	Pseudo R^2 (%)	$\chi^2(1)$ (LR)	p-val.	Marginal effect
Panel A: Downside return probabilities are the predictor									
5% down	-0.08	0.92	0.98	0.00	-74	6.4	10.2	0.001	0.062
3% down	-0.62	0.36	0.59	0.09	-129	1.3	3.3	0.070	0.073
Panel B: Upside return probabilities are the predictor									
3% up	-0.49	0.50	1.54	0.01	-75	5.7	9.1	0.003	0.098
5% up	0.15	0.76	0.97	0.00	-94	12.2	26.2	0.000	0.087

Table 2.5: **Empirical consistency of downside and upside probabilities implied by our framework** This table presents the results from a logistic regression. The next-period outcome variable, $\mathbb{Y}_t$, is consistent with a Bernoulli probability function and takes a value 1 with probability $\mathbf{p}_t$ and 0 with probability $1 - \mathbf{p}_t$. Assuming $\mathbb{Y}_i \sim \text{Bernoulli}(\mathbb{Y}_t|\mathbf{p}_t)$, our empirical specification is $\mathbf{p}_t = (1 + e^{\psi_0+\psi_1 x_t})^{-1}$, where $x_t = \log(\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}]))$ or $x_t = \log(\mathbb{E}_t^{\mathbb{P}}(\theta[r_{t+1} - r^u]))$. The testable hypothesis is $\psi_1 = 0$. The expression for $\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])$ is displayed in equation (2.17). For the upside probabilities, we choose r^u such that $\frac{S_t}{K^u} - 1 \in \{-0.05, -0.03\}$ where $r^u \equiv \log(\frac{S_t}{K^u})$. Following Result 2, we write $\underbrace{\mathbb{E}_t^{\mathbb{P}}(\theta[r_{t+1} - r^u])}_{\text{Upside probability}} = \mho_0[K^u]\, \underbrace{R_{\mathbf{f},t+1}^{-1}\mathbb{E}_t^{\mathbb{Q}}(\theta[r_{t+1} - r^u])}_{\text{Binary call}} + \mho_1[K^u]\, C_t[K^u] + \underbrace{\int_{K^u}^{\infty} \mho_2[K]\, C_t[K]\, dK}_{\text{calls with strike } K>K^u}$, where $\mho_0[K^u] = \exp(-\eta_0 + 1 - \eta_2(\frac{\{\log(\frac{K^u}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1}))$, $\mho_1[K^u] = (-\frac{2\eta_2 \log(\frac{K^u}{S_t})}{\mathrm{v}_t K^u})\,\mho_0[K^u]$, and $\mho_2[K] = (\frac{2\eta_2 \log(\frac{K}{S_t})}{K^2 \mathrm{v}_t} - \frac{2\eta_2}{K^2 \mathrm{v}_t})\,\mho_0[K] + (-\frac{2\eta_2 \log(\frac{K}{S_t})}{K \mathrm{v}_t})^2\,\mho_0[K]$. "log-lik" represents the log-likelihood from the maximum likelihood estimation. Next, we present the pseudo R^2. Reported alongside is the likelihood ratio (LR), which is distributed $\chi^2(1)$. The p-values correspond to z-statistics, and the final column presents the marginal effect of the predictor.

$y_{\{t\rightarrow t+j\}}$	j	Slope coefficient Π_{disaster}	NW[p]	H[p]	$\overline{R}^2$ (%)	ℓ^*	CORR
Log change of IP index (INDPRO)	1	-0.031	0.031	0.002	8.0	12	-0.29
	3	-0.095	0.014	0.000	16.2	13	-0.41
	6	-0.153	0.014	0.001	13.3	13	-0.37
	12	-0.210	0.013	0.003	8.5	14	-0.30
Log change of total nonfarm (PAYEMS)	1	-0.015	0.004	0.000	27.2	14	-0.52
	3	-0.043	0.004	0.000	29.9	14	-0.55
	6	-0.086	0.001	0.000	31.6	14	-0.56
	12	-0.152	0.000	0.000	27.1	14	-0.52
ISM Index (NAPM)	1	-34.2	0.008	0.000	16.1	13	-0.40
	3	-29.2	0.012	0.000	11.6	13	-0.35
	6	-17.7	0.059	0.001	4.0	13	-0.21
	12	-2.8	0.799	0.565	-0.2	13	-0.03

Table 2.6: **Disaster uncertainty and subsequent output, unemployment, and orders/inventories** We report the results from the predictive regressions with $r^d = -0.05129$: $y_{\{t\rightarrow t+j\}} = \Pi_0 + \Pi_{\text{disaster}} \underbrace{\mathbb{E}_t^{\mathbb{P}}(\theta[r^d - r_{t+1}])}_{\text{disaster probability via eq. (2.17)}} + \epsilon_{t+j}$, for $j = 1, 3, 6, 12$ months. In our specification, we take the dependent variable, $y_{\{t\rightarrow t+j\}}$, to be (i) the log change in the IP index (FRED mnemonic, INDPRO) from "output and income," (ii) the log change in all employees: total nonfarm (FRED mnemonic, PAYEMS) from "labor markets," and (iii) the level of ISM: PMI composite index (FRED mnemonic, NAPM) from "orders and inventories." Each of the three featured variables exhibits high (absolute) correlation with the first principal component of the respective group with an R^2 of 94%, 84%, and 67%. The data is described in [97, Appendix]. To correct for autocorrelation and heteroscedasticity, we use the [105] estimator with automatically selected lag, denoted by ℓ^*, as in [106], and the reported two-sided p-values are denoted by NW[p]. Shown also are [115] two-sided p-values, denoted by H[p]. The [115] p-values are adjusted for overlapping observations when, $j = 3$, $j = 6$, and $j = 12$. The adjusted R^2 is reported as $\overline{R}^2$. The pairwise correlation coefficient between the dependent variable and disaster probability is denoted by CORR. The intercept Π_0 is not reported. The computation of disaster probabilities is based on the expression in equation (2.17).

Predictor	β_0	NW[p] $\beta_0 = 0$	β_v	NW[p] $\beta_v = 0$	Wald NW[p] $\beta_v = 1$	$\overline{R}^2$ (%)	Out-of-sample MSPE[p]	R^2_{OOS} (%)
$\text{var}_t^{\mathbb{P}}[r_{t+1}]$	0.0004	0.076	0.937	0.000	0.727	38.4	0.006	33.9
$\text{var}_t^{\mathbb{Q}}[r_{t+1}]$	0.0000	0.982	0.669	0.000	0.000	40.5	0.008	34.3

Table 2.7: **Forecast accuracy of conditional return variance** Reported are the results from the following Mincer-Zarnowitz regression: $\underbrace{\text{rv}_{\{t\to t+1\}}}_{\text{realized return variance}} = \beta_0 + \beta_v \underbrace{\text{var}_t^{\mathbb{P}}[r_{t+1}]}_{\text{from options via eq. (2.28)}} + \epsilon_{t+1}$. The realized return variance $\text{rv}_{\{t\to t+1\}}$ is calculated as the sum of daily squared returns $\text{rv}_{\{t\to t+1\}} = r^2_{t+\Delta,\Delta} + r^2_{t+2\Delta,\Delta} + \ldots + r^2_{t+\hbar\Delta,\Delta}$, with $r_{t+\hbar\Delta,\Delta} \equiv \log(\frac{S_{t+\hbar\Delta}}{S_{t+(\hbar-1)\Delta}})$, where Δ is one day and $\hbar$ is the total number of trading days over an options expiration cycle. We also report the results from using $\text{var}_t^{\mathbb{Q}}[r_{t+1}]$ as a predictor (inferred from call and put option prices using [78, Section 1.2 and equation (7)]). To correct for autocorrelation and heteroscedasticity, we use the [105] estimator with automatically selected lag l^* (in our case $l^* = 11$), as in [106], and the reported two-sided p-values are denoted by NW[p]. Shown is the p-value corresponding to the Wald test of the hypothesis of $\beta_v = 1$. The adjusted R^2 is reported as $\overline{R}^2$, and the p-value for the MSPE statistic is computed by regressing $\mathfrak{f}_{t+1}$ on a constant, where $\mathfrak{f}_{t+1} = (\text{rv}_{\{t\to t+1\}} - \overline{\text{rv}}_{\{t\to t+1\}})^2 - [(\text{rv}_{\{t\to t+1\}} - \widehat{\text{rv}}_{\{t\to t+1\}})^2 - (\widehat{\text{rv}}_{\{t\to t+1\}} - \overline{\text{rv}}_{\{t\to t+1\}})^2]$. $\widehat{\text{rv}}_{\{t\to t+1\}}$ is an estimate of $\text{rv}_{\{t\to t+1\}}$, using information up to period t, based on equation (2.27), and $\overline{\text{rv}}_{\{t\to t+1\}}$ is obtained from a restricted model imposing $\beta_v = 0$. Low one-sided p-values associated with the adjusted MSPE test statistic indicate that $\text{var}_t^{\mathbb{P}}[r_{t+1}]$ generates additional predictive power. The third term is an adjustment that leads to better small-sample properties of the test. The out-of-sample R^2_{OOS} is calculated as $R^2_{\text{OOS}} \equiv 1 - \frac{\sum_{t=0}^{T-1}(\text{rv}_{\{t\to t+1\}} - \widehat{\text{rv}}_{\{t\to t+1\}})^2}{\sum_{t=0}^{T-1}(\text{rv}_{\{t\to t+1\}} - \overline{\text{rv}}_{\{t\to t+1\}})^2}$.

		Bootstrap CI		Across 1,000 bootstraps						
	Average	Lower	Upper	Mean	SD	P5	P25	P50	P75	P95
$\mathbb{E}_t^{\mathbb{P}}(r_{t+1})$	0.112	0.080	0.153	0.113	0.018	0.087	0.102	0.113	0.122	0.145
$\sqrt{\text{var}_t^{\mathbb{P}}[r_{t+1}]}$	0.154	0.130	0.170	0.153	0.010	0.135	0.148	0.153	0.159	0.167
vrp_t	-0.680	-0.446	-1.114	-0.709	0.180	-0.488	-0.599	-0.690	-0.778	-1.019

Table 2.8: **Empirical evidence on recovered expected (logarithmic) returns and return volatility** The reported properties of conditional expected logarithmic returns is based on equation (2.30), whereas the reported volatility $\sqrt{\text{var}_t^{\mathbb{P}}[r_{t+1}]}$ is based on equation (2.28). We compute variance risk premium $\text{vrp}_t \equiv \log(\text{var}_t^{\mathbb{P}}[r_{t+1}]) - \log(\text{var}_t^{\mathbb{Q}}[r_{t+1}])$, where $\text{var}_t^{\mathbb{Q}}[r_{t+1}]$ is inferred from put and call option prices (e.g., [78, equation (7)]).

Panel A: Model estimates

	η_0	η_1	η_2	η_0	$\eta_1 \equiv 0$	η_2	η_0	η_1	$\eta_2 \equiv 0$
Estimate	1.019	-6.053	0.106	1.004		0.095	1.011	-6.144	
Lower CI	1.002	-11.093	0.010	0.998		-0.002	0.999	-10.726	
Upper CI	1.060	-1.853	0.258	1.029		0.230	1.035	-1.803	

Panel B: Gauging specification errors using Hansen-Jagannathan (1997) distance measure

	All six assets				Excluding put options				Excluding call options			
	$\eta_1 \neq 0$, $\eta_2 \neq 0$	$\eta_1 = 0$, $\eta_2 \neq 0$	$\eta_1 \neq 0$, $\eta_2 = 0$		$\eta_1 \neq 0$, $\eta_2 \neq 0$	$\eta_1 = 0$, $\eta_2 \neq 0$	$\eta_1 \neq 0$, $\eta_2 = 0$		$\eta_1 \neq 0$, $\eta_2 \neq 0$	$\eta_1 = 0$, $\eta_2 \neq 0$	$\eta_1 \neq 0$, $\eta_2 = 0$	
$\mathrm{HJ}^{\mathrm{dist}}$	$\boxed{0.30}$	0.84	0.56		$\boxed{0.26}$	0.61	0.48		$\boxed{0.13}$	0.58	0.28	
Lower CI	0.22	0.42	0.29		0.13	0.23	0.19		0.07	0.17	0.07	
Upper CI	0.92	1.59	0.97		0.76	0.97	0.86		0.75	1.47	0.70	
$\frac{\mathrm{HJ}^{\mathrm{dist}}_{\mathrm{Full}}}{\mathrm{HJ}^{\mathrm{dist}}_{\eta_1 \equiv 0}} - 1$				-65%				-56%				-17%
p-value				0.05				0.07				0.17

Table 2.9: **Estimates and bootstrap confidence intervals for the parameters of the $m[F_{t+1}]$ specification and the associated distance measure** Let F_{t+1} be the price of the futures on the 30-year Treasury bond. Reported parameter estimates are based on the following specification of $m[F_{t+1}]$: $m[F_{t+1}] = \exp(\eta_0 - 1 + \boldsymbol{\eta}'\mathbf{z}_{t+1})$, where $\mathbf{z}_{t+1}[F_{t+1}] = \begin{pmatrix} \frac{F_{t+1}}{F_t} - 1 \\ \frac{\{\log(\frac{F_{t+1}}{F_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1} \end{pmatrix}$. The price of the volatility contract (with payoff $\{\log(\frac{F_{t+1}}{F_t})\}^2$) is $\mathrm{v}_t \equiv R_{\mathbf{f},t+1}^{-1}\,\mathbb{E}_t^{\mathbb{Q}}(\{\log(\frac{F_{t+1}}{F_t})\}^2) = \int_0^{F_t} \frac{2(1+\log(\frac{F_t}{K}))}{K^2}\, P_t^{\mathrm{futures}}[K]\,dK + \int_{F_t}^{\infty} \frac{2(1-\log(\frac{K}{F_t}))}{K^2}\, C_t^{\mathrm{futures}}[K]\,dK$. We obtain $(\eta_0, \boldsymbol{\eta}) \in \mathbb{R}^3$ as the solution to the problem: $\inf_{(\eta_0,\boldsymbol{\eta})} \; - \mathbb{E}^{\mathbb{P}}(R_{\mathbf{f},t+1}^{-1})\,\eta_0 + \mathbb{E}^{\mathbb{P}}[\exp(\eta_0 - 1 + \boldsymbol{\eta}'\mathbf{z}_{t+1})]$. To get the bootstrap distribution of $(\eta_0, \boldsymbol{\eta})$, we consider a procedure that randomly selects, with replacement, $(R_{\mathbf{f},t+1}, \mathbf{z}_{t+1})$ jointly via stationary bootstrap. Then we recompute η_0, η_1, and η_2 in each bootstrap draw. We perform 1,000 bootstrap draws, yielding the reported 95% bootstrap confidence intervals. Let $R_{t+1,\ell\%,\mathrm{put}}$ be the gross return of a ℓ% OTM put option on the Treasury bond futures (and similarly for calls). Panel B evaluates the ability of our model to price six assets (or a subset of them) in the market for futures on the 30-year Treasury bond, as follows: $\mathbf{R}_{t+1} = \left(R_{\mathbf{f},t+1},\; \frac{F_{t+1}}{F_t} + (R_{\mathbf{f},t+1} - 1),\; R_{t+1,3\%,\mathrm{put}},\; R_{t+1,1\%,\mathrm{put}},\; R_{t+1,1\%,\mathrm{call}},\; R_{t+1,3\%,\mathrm{call}}\right)^T$. We report the [92] distance as $\mathrm{HJ}^{\mathrm{dist}} \equiv \sqrt{\mathbf{e}_{t+1}'\mathbf{I}^{-1}\mathbf{e}_{t+1}}$, where $\mathbf{e}_{t+1} \equiv \mathbb{E}^{\mathbb{P}}(m[F_{t+1}]\,\mathbf{R}_{t+1}) - \mathbf{1}$. The use of identity matrix, $\mathbf{I}$, follows from [93, page 216], and we report the bootstrap confidence intervals. Specifically, we randomly select, with replacement, $(\mathbf{R}_{t+1}, \mathbf{z}_{t+1})$ jointly via stationary bootstrap. The sample period is December 1990 to December 2015. The last row reports the bootstrap p-value as the proportion of draws for which $\frac{\mathrm{HJ}^{\mathrm{dist}}_{\mathrm{Full}}}{\mathrm{HJ}^{\mathrm{dist}}_{\eta_1 \equiv 0}} - 1 \geq 0$ (the restricted model has superior pricing ability).

		$\text{PC}_t^{\lfloor 1 \rfloor}$ (level)	$\text{PC}_t^{\lfloor 2 \rfloor}$ (slope)	$\text{PC}_t^{\lfloor 3 \rfloor}$ (curvature)	$\underbrace{\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t}-1)}_{\text{recovered from options}}$	$\overline{R}^2$ (%)	Joint[p]
Panel A: Dependent variable is the excess return of the futures on a 10-year Treasury bond							
Futures	Coef.	0.049	0.342	-0.562		1.3	
(10 year)	NW[p]	0.034	0.004	0.301			0.013
	PB[p]	0.052	0.001	0.421			0.038
	Coef.	0.073	0.293	-0.389	0.542	2.7	
	NW[p]	0.014	0.053	0.566	0.011		0.002
	PB[p]	0.014	0.059	0.683	0.063		0.016
	Coef.				0.519	1.4	
	NW[p]				0.021		
	PB[p]				0.088		
Panel B: Dependent variable is the excess return of the futures on a 30-year Treasury bond							
Futures	Coef.	0.066	0.622	-1.080		1.7	
(30 year)	NW[p]	0.226	0.003	0.312			0.015
	PB[p]	0.238	0.005	0.359			0.030
	Coef.	0.101	0.551	-0.830	0.788	2.7	
	NW[p]	0.081	0.006	0.423	0.076		0.023
	PB[p]	0.086	0.011	0.499	0.222		0.071
	Coef.				0.804	1.2	
	NW[p]				0.034		
	PB[p]				0.118		

Table 2.10: **Assessing the spanning hypothesis using excess return of futures on the Treasury bond** The hypothesis is whether $\pi^{\text{spanning}} = 0$ for an economically motivated predictor $\mathbb{X}_t$ besides the level, slope, and curvature of the Treasury yield curve, as per the predictive regressions below:

$$\underbrace{\frac{F_{t+1}}{F_t} - 1}_{\text{Excess return of futures on a 10- or 30-year Treasury bond}} = \pi_0 + \underbrace{\sum_{j=1}^{3} \pi_j \, \text{PC}_t^{\lfloor j \rfloor}}_{\text{Level, slope, and curvature variables}} + \pi^{\text{spanning}} \, \mathbb{X}_t + \epsilon_{t+1}, \text{ where}$$

$\text{PC}_t^{\lfloor j \rfloor}$, for $j = 1, 2, 3$ are the first, second, and third principal components extracted from Treasury yields (maturities of 1 month, 3 month, 1, 2, 5, 7, 10, 20, and 30 years). In this table, $\mathbb{X}_t$ is option-inferred $\mathbb{E}_t^{\mathbb{P}}(\frac{F_{t+1}}{F_t} - 1)$, as shown in equations (2.35)–(2.36). To correct for autocorrelation and heteroscedasticity, we use the [105] estimator with automatically selected lag (as in [106]), and the reported two-sided HAC p-values are denoted by NW[p]. We further display the p-values based on a parametric bootstrap, denoted by PB[p], following [108] and [66]. In the parametric bootstrap, the predictors follow an ARMA process, selected based on the BIC. The adjusted R^2 is shown as $\overline{R}^2$. The p-values from the χ^2 test, that the parameters are jointly zero, is shown as Joint[p]. Panel A (Panel B) reports the regression results when $\frac{F_{t+1}}{F_t} - 1$ is based on futures prices of the 10- (30)-year Treasury bond. The reported coefficients on $\text{PC}_t^{\lfloor 1 \rfloor}$, $\text{PC}_t^{\lfloor 2 \rfloor}$, and $\text{PC}_t^{\lfloor 3 \rfloor}$ are each multiplied by 10^2. The constant π_0 is included in the regressions but is not reported.

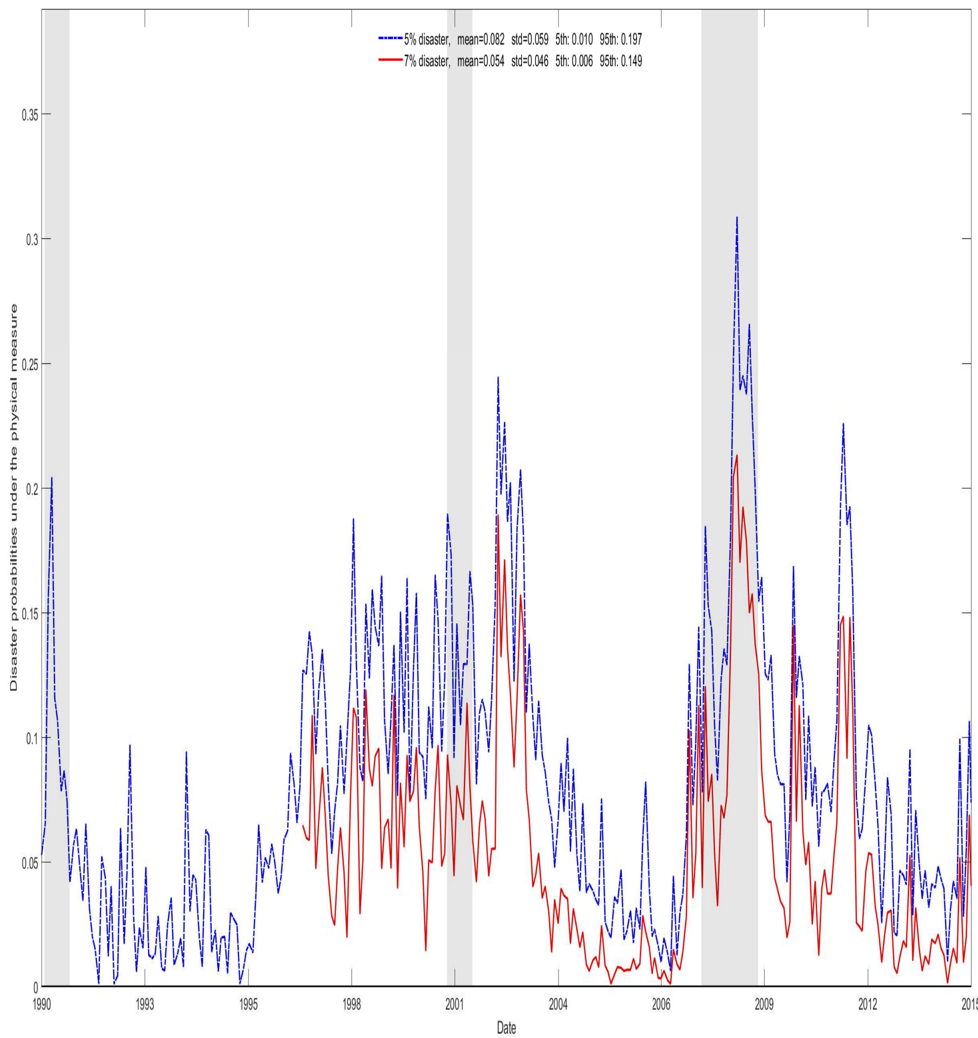

Figure 2.1: **Probability of disasters**
Plotted are the probability of disasters under the $\mathbb{P}$ measure according to equation (2.17), when $m[S_{t+1}] = \exp(\eta_0 - 1 + \boldsymbol{\eta}' \mathbf{z}_{t+1}[S_{t+1}])$, where $\underbrace{\mathbf{z}_{t+1}[S_{t+1}]}_{2\times 1} \equiv \begin{pmatrix} \frac{S_{t+1}}{S_t} - R_{\mathbf{f},t+1} \\ \frac{\{\log(\frac{S_{t+1}}{S_t})\}^2}{\mathrm{v}_t} - R_{\mathbf{f},t+1} \end{pmatrix}$. Our calculations are based on the estimated parameters in Table 2.1. The first point in the sample is 1/22/1990, and the final point is 11/23/2015. The grey shaded regions represent recessions, as dated by the NBER. The probability of a 5% disaster over options expiration cycles is depicted by the dashed-curve (in blue), and the expectation of the 7% disaster is depicted by the solid-curve (in red).

Chapter 3: Oil and Equity Return Predictability: The Importance of Dissecting Oil Price Changes

3.1 Introduction

As the world's major source of energy, oil plays a crucial role in the modern global economy. Not surprisingly, the impact of oil price fluctuations on equity markets and the real economy has been of great interest to academics, policy makers, and market participants alike. Oil price changes could be interpreted in different ways. On the one hand, an oil price increase could be considered bad news for the economy and equity markets as it increases the cost of production in a significant number of sectors and causes consumers to reduce their consumption. Following the same line of thinking, an oil price drop would have the opposite effect and would be perceived as good news. On the other hand, higher oil prices imply higher profits for the oil sector. This would likely cause oil price shares to gain value and, to some extent, boost the aggregate market. Analogously, lower oil prices is bad news for the oil sector and could negatively affect the broader market. The conventional wisdom in the past was that the former effect dominates, implying that an oil price hike is considered to be bad news for equity markets. Accordingly, a possible hypothesis is that positive (negative) oil price changes should predict lower (higher) subsequent stock returns. In line with this hypothesis, [116] document that, based on data until 2003, oil price changes predict Morgan Stanley Capital International (MSCI) equity index returns with a negative and statistically significant predictive slope for a large number of countries. However, the relationship between oil price movements and subsequent stock returns has not been stable over time. Figure 3.1, where we present the two scatter plots of the MSCI World index return versus the one-month lagged log growth rate of West Texas Intermediate (WTI) spot price over the 1982–2003 and 2004–2015 periods, clearly illustrates the shift over time. The correlation has shifted from -0.22 in the former period to 0.25 in the latter period. As a result of this shift, the overall predictive ability of oil price change has dramatically decreased over the sample period covering the last thirty years. This structural change is striking and begs for an explanation.

Please insert Figure 3.1 here

At the most fundamental level, oil prices move when there is a misalignment between supply and demand. Understanding what causes oil price changes, in the first place, can be crucial for determining the potential impact of such a change on equity markets. For instance, lower oil prices due to a slowdown in global economic activity should be viewed as bad news. However, prices could also fall because of excess supply of oil, in which case the message would be different. To provide an explanation for the recent positive correlation between oil price changes and aggregate US market returns, [117] decomposes oil price change into a demand-related component and a residual. He documents that the correlations of the two components with market returns are different and states: "*That's consistent with the idea that when stock traders respond to a change in oil prices, they do so not necessarily because the oil movement is consequential in itself, but because fluctuations in oil prices serve as indicators of underlying global demand and growth.*" Recent academic literature has also discussed the potential differential effects of demand and supply shocks associated with oil price fluctuations. Although oil price shocks were often associated with oil production disruptions in the 1970s and 1980s, it has been argued that the role of global demand for oil, especially from fast-growing emerging economies, should also be emphasized (see [118], [119], and [120]). Furthermore, the last two papers point out that oil supply shocks, global demand shocks, and other types of shocks, all of which can cause oil prices to fluctuate, should have differential effects on the macroeconomy and the stock market.

In this paper, we investigate whether oil price changes contain useful information for predicting future equity returns. Building on the ideas discussed above, we posit that this is indeed the case, once these changes are suitably decomposed into supply and demand shocks. We motivate the research question by first demonstrating that the negative relationship between oil price changes and future G7 country and World MSCI index excess returns, documented using data until 2003, has dramatically changed over the extended sample period ending in 2015. Faced with this empirical result, one might infer oil price changes are useless for forecasting international equity index returns. However, we contend that information contained in oil price changes becomes useful once it is suitably complemented with relevant information about oil supply and global economic activity. The key observation, as made by [119], is that oil price changes are driven by various supply and demand shocks that play fundamentally different roles. Accordingly, we use two distinct structural VAR models that allow us to decompose oil price changes into oil supply, global demand, and oil-specific demand shocks. We, subsequently, illustrate the ability of these three shocks to predict G7 country and World MSCI index excess returns, using metrics of both statistical and economic significance, as well as the structural stability of this predictive relationship over the last thirty years. In addition, we show that the predictability results survive a comprehensive battery of robustness checks.

Our work relates to a growing literature that examines the impact of oil price shocks on the real economy and equity markets. [121], [122], and [120], among others, examine the contemporaneous relationship between the price of oil and stock prices.[1] [120] augment the structural VAR model of [119] by incorporating the US real stock return and study the contemporaneous relationships between shocks embedded in oil price changes and stock returns. They examine cumulative impulse responses of real stock returns to one-time shocks to oil supply, global demand, and oil-specific demand in the crude oil market. Their results, using data over the 1975–2006 period, show that an unexpected decrease in oil production has no significant effect on cumulative US real stock returns and a positive surprise to global

[1] The contemporaneous relationship between the volatility of oil prices and stock returns has been studied by several papers, such as [123] and [124].

demand (oil-specific demand) is associated with a subsequent increase (decrease) in US real stock returns. In addition, several recent papers, including [116], [125], and [126], investigate the ability of oil price shocks to forecast equity index returns and document a negative predictive slope of oil price changes. On a related topic, [127] examine whether crude oil futures returns in the US can forecast equity returns in East Asian markets and point to a structural break in the predictive relationship in the mid 2000s. Unlike [120], and in line with the recent finance literature and [116] in particular, we use a predictive regression framework to examine whether information contained in oil price changes can be used to forecast future stock returns. We approach the question from the perspective of an investor who wishes to use real-time information embedded in oil price changes and captured by the three aforementioned shocks to forecast MSCI equity index excess returns.

This paper is also related to some recent empirical research that focuses on disentangling the intrinsic shocks embedded in oil price changes. [119] identifies oil shocks using a structural VAR model and highlights the importance of disentangling oil supply, global demand, and oil-specific demand shocks. However, the short-run price elasticity of oil supply is assumed to be equal to zero in [119], implying that oil supply shocks only account for a small fraction of oil price variation. [128] and [129] propose structural VAR models that facilitate estimation of both oil supply and oil demand elasticities. They further illustrate that the percentage of variation in oil prices attributed to oil supply shocks critically relies on the oil supply and demand elasticity estimates. [130] and [131] propose to use information from the stock market to identify the underlying types of shocks in oil price changes. [130] identifies shocks specific to the oil market and shocks that affect the overall economy using the sign and magnitude of the correlation between daily oil price changes and aggregate stock market returns, excluding oil companies. [131] uses crude oil futures returns, returns on a global equity index of oil producing firms, and innovations to the VIX index to identify demand and supply shocks. He documents a strong *contemporaneous* relationship between aggregate market returns and the demand/supply shocks from his decomposition based on data from 1986 to 2011, an empirical result that is readily confirmed to remain valid over the sample period extending to 2015. However, the shocks identified by [131] cannot forecast future stock market returns. In contrast, the focus of our paper is the *predictive* relationship between the various shocks embedded in oil price changes and equity index returns. Hence, we find that the structural VAR approach along the lines of [119], which utilizes more direct proxies for oil supply and global demand and does not require stock market information to obtain the decomposition, is more suitable for our purposes.

In this paper, we make a number of contributions to the literature studying the relationship between oil price fluctuations and subsequent international equity returns. First, we document that the ability of oil price changes to forecast G7 country and World MSCI index returns has dramatically changed over the last decade. In particular, using formal structural break tests, we detect a break in the predictive relationship in the third quarter of 2008 for most of the indexes under examination. Second, using two distinct structural VAR models and employing a new proxy for global real economic activity, we obtain real-time decompositions of oil price changes into oil supply, global demand, and oil-specific demand shocks.

The first model is a variant of the [119] model, and the second model is a parsimonious version of the model proposed by [128], which allows for joint estimation of the elasticity of oil supply and demand. The implementation of the structural VAR models is facilitated by the use of suitable proxies for the variables of interest. We use the first principal component of the log growth rates of WTI, Dubai, and Arab Light spot prices as a comprehensive proxy for oil price change. Moreover, we employ two proxies for global real economic activity, namely a shipping cost index and global crude steel production, and use the first principal component of their log growth rates as a comprehensive proxy for global demand growth. Importantly, all the variables that we use in our empirical tests are constructed based on information available in real time, and therefore our results are not subject to look-ahead bias.[2] Third, we illustrate the ability of these three shocks to predict G7 country and World MSCI index returns, denominated in both local currency and US dollars. In particular, using the shocks obtained from both structural VAR models, we find empirical evidence supporting the hypothesis that oil supply shocks and oil-specific demand shocks (global demand shocks) predict equity returns with a negative (positive) slope over the 1986–2015 period. We also demonstrate the advantage of using the oil price decomposition instead of just the oil price change, in economic terms, by the economically substantial and statistically significant improvement in the performance of mean-variance optimal trading strategies. Finally, we examine various other aspects of the predictive relationship. To address real-time data availability

[2] In particular, we make sure to use the vintages of all the required time series that are available at the time. Moreover, as a robustness check, we delay the construction of the MSCI index returns by one and two weeks to completely ensure that the data used are indeed in the investor's information set.

concerns, we construct returns with a delay of one and two weeks and show that the results are essentially identical. We demonstrate that, as the forecasting horizon increases from one to six months, the predictive ability of the three shocks gradually diminishes. For the case of the US, we document that the forecasting ability is present in the cross section of industries and robust in the presence of standard macroeconomic predictors. The estimated conditional expected returns, based on the three shocks, exhibit high volatility and low persistence in comparison to risk premia estimates available in the literature. Finally, these three shocks do not appear to have an effect on conditional return volatility. Collectively, these results do not appear to be consistent with the notion of time-varying risk premia. As pointed out by [132], oil supply and demand shocks are not observable and their estimation is subject to information frictions. A possible explanation for our documented predictability is that, due to information frictions, investors cannot measure oil supply and demand shocks with precision and, therefore, cannot act on such information.

The rest of the paper proceeds as follows. In Section 3.2, we describe the data that we use in our empirical exercises. In Section 3.3, we present evidence on the forecasting ability of oil price changes and how it has changed over the last decade. In Section 3.4, we use structural VAR models to decompose oil price changes into oil supply shocks, global demand shocks, and oil-specific demand shocks. In Section 3.5, we illustrate the ability of these three types of shocks to forecast MSCI equity index returns and provide additional robustness checks. In Section 3.6, we offer some concluding remarks.

3.2 Data

We use five different data sets: returns on international equity indexes, short-term interest rates, oil price proxies, proxies for global economic activity, and global oil production. The full sample period is from January 1982 to December 2015.

We use returns on the G7 country and World MSCI equity indexes, denominated in both local currency and US dollars.[3] We collect monthly short-term interest rates for the G7 countries from the International Monetary Fund (IMF) and the Organisation for Economic Cooperation and Development (OECD). We use IMF Treasury Bill rates whenever available and short-term interest rates obtained from the OECD otherwise.[4]

We further use three proxies for oil price, namely the WTI spot price, the Dubai spot price, and the Arab Light spot price.[5] Note that 75% (83%) of the log growth rates of WTI (Arab Light) prices from October 1973 to September 1981 are zero.[6] Therefore, it is problematic to use WTI and Arab Light prices before September 1981. Therefore, in our empirical analysis, we use oil price data from 1982 onward. Following [116], we use nominal oil prices.[7]

---

[3]Specifically, data on MSCI indexes for the G7 countries, i.e., Canada, France, Germany, Italy, Japan, the United Kingdom, and the United States, as well as the World MSCI index are obtained from Datastream.

[4]For Canada, France, Italy, Japan, and the United Kingdom, we use Treasury Bill rates from the IMF. For Germany, we use Treasury Bill rates from the IMF and, from September 2007, short-term interest rates from the OECD. For the United States, we use the 1-month Treasury Bill rate taken from Kenneth French's website.

[5]Data on the Dubai and Arab Light spot prices are obtained from Bloomberg. Data on WTI spot oil prices for the period of between January 1982 and August 2013 are obtained from the website of St. Louis Fed. Data for the period between August 2013 and December 2015 are obtained from Bloomberg.

[6][133] documents that the US federal control of crude oil prices started in August 1971 and ended in January 1981.

[7]Using real or nominal oil prices is empirically immaterial in our setting. For example, the

We combine the information contained in the three proxies for crude oil spot price into a single proxy using Principal Component Analysis (PCA). The single proxy, denoted by $g^{\mathtt{P}}$ where $\mathtt{P}$ stands for price, is represented by the first principal component of the log growth rates of WTI, Dubai, and Arab Light spot prices. The details of the construction of the single PCA proxy for oil price change are provided in Appendix 3.7.1. To make the proxy $g^{\mathtt{P}}$ comparable to the three individual proxies, we rescale it so that its standard deviation equals 0.09 over the sample period of January 1983 to December 2015. Table 3.1 presents summary statistics, including correlations, for $g^{\mathtt{P}}$ and the log growth rates of the three oil price proxies. Over the subsample period ending in April 2003, which is the last month in the sample used in [116], as well as the full sample period, $g^{\mathtt{P}}$ is highly correlated with the log growth rates of the three individual proxies. The top panel in Figure 3.2 also shows that the four series track each other quite closely.

Please insert Table 3.1 here

Please insert Figure 3.2 here

We use two proxies for global economic activity to capture changes in global demand. The first proxy is a shipping cost index constructed from data on dry cargo single voyage rates and the Baltic Dry Index (BDI). Since the supply of bulk carriers is largely inelastic, fluctuations in dry bulk cargo shipping cost are thought to reflect changes in global demand for transporting raw materials such as metals, grain, and

correlation between log changes of WTI nominal price and log changes of WTI real price is 0.9996.

coal by sea. Therefore, shipping cost is considered to be a useful leading indicator of global economic activity. Data on dry cargo single voyage rates are hand collected from Drewry Shipping Statistics and Economics for the period between January 1982 and January 1985. Rates for seven representative routes are reported each month. We compute the monthly log growth rates of the shipping cost for each route, and then, following [119], obtain the cross-sectional equally-weighted average.[8] Data on the BDI from January 1985 to December 2015 are obtained from Bloomberg.

The second proxy for global economic activity is global crude steel production. [134] argue that world steel production is a good indicator of global real economic activity. Steel is widely used in a number of important industries, such as energy, construction, automotive and transportation, infrastructure, packaging, and machinery. Therefore, fluctuations in world crude steel production reflect changes in global real economic activity. We obtain monthly crude steel production data for the period January 1990 to December 2015 from the World Steel Association's website. The reported monthly figure represents crude steel production in 66 countries and accounts for about 99% of total world crude steel production. In addition, monthly data for the period January 1968 to October 1991 are hand collected from the Steel Statistical Yearbook published by the International Iron and Steel Institute. Crude steel production exhibits strong seasonality and, therefore, we seasonally adjust the data, as we explain in Appendix 3.7.1.

As in the case of oil price proxies, we use PCA to construct a single proxy

[8]It is, however, worth noting that there is an important difference between our proxy for global economic activity and the one constructed in [119]. Specifically, in [119], the average growth rate is cumulated, then deflated using the US CPI, and finally detrended. In that sense, the proxy in [119] is a level variable. In contrast, our proxy is a growth rate.

for global demand growth. The single proxy, denoted by $g^{\mathtt{GD}}$ where GD stands for global demand, is represented by the first principal component of the log growth rates of the shipping cost index and global crude steel production. The details of the construction of the single PCA proxy for global demand growth are provided in Appendix 3.7.1. The correlations between $g^{\mathtt{GD}}$ and the log growth rates of the shipping cost index and global crude steel production are 0.82 and 0.74, respectively. The bottom panel in Figure 3.2 shows that $g^{\mathtt{GD}}$ closely tracks the two individual proxies most of time, except for a few instances in which one of the two proxies takes extreme values.

Finally, we obtain oil production data, covering the period between January 1982 and December 1991, from the website of the US Energy Information Agency.[9] In addition, we hand-collect data on the total supply of crude oil, natural gas liquids, processing gains, and global biofuels, for the period between December 1991 and December 2015, from the monthly Oil Market Report obtained from the website of the International Energy Agency. Combining data from the two sources, we construct a time series of monthly log growth rates of world crude oil production.

3.3 Oil price change as a predictor of MSCI index excess returns

To examine whether oil price changes can predict equity returns, we start by revisiting the evidence documented in [116] who consider a sample period ending in April 2003. In particular, we estimate the following standard predictive regression

[9]Specifically, we use Table 11.1b (World Crude Oil Production: Persian Gulf Nations, Non-OPEC, and World).

model:

$$r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}, \tag{3.1}$$

where $r^{\mathtt{e}}_{t+1}$ is the excess return on an MSCI index and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: WTI, Dubai, and Arab Light. We construct excess returns by subtracting the particular country's short-term rate from each MSCI index return in the case of local currency-denominated indexes and by subtracting the US Treasury Bill rate from each MSCI index return in the case of US dollar-denominated indexes.[10]

We consider the MSCI indexes for the G7 countries as well as the World MSCI index, denominated both in local currencies and US dollars. The oil price change proxy we use is the first principal component obtained from the WTI, Dubai, and Arab Light spot prices, as explained in Section 3.2. [116] document negative and statistically significant estimates of the predictive slope coefficient $\delta^{\mathtt{P}}$ for a large number of countries based on a sample that ends in April 2003. We first run the predictive regression for the sample starting in January 1983 and ending in April 2003 and then consider the extended sample period starting in January 1983 and ending in December 2015. We examine the statistical significance of predictability in terms of p-values and adjusted R-squares. Furthermore, we examine the economic significance of predictability by evaluating the performance of the resulting optimal trading strategies in terms of certainty equivalent returns and Sharpe ratios. Appendix 3.7.2 explains in detail the methods used to evaluate predictive ability.

[10] While [116] use log returns in their empirical analysis, it is more convenient for us to use excess returns for the purpose of assessing the economic significance of the predictive ability of oil price changes. The results for log returns, available upon request, are very similar.

3.3.1 Evidence from the 1983–2003 sample period

In the left panel of Table 3.2, we present statistical significance results for the predictive regression model (3.1) based on MSCI index excess returns, denominated in both local currencies and US dollars over the sample period starting in January 1983 and ending in April 2003. We first focus our analysis on this sample period to facilitate comparison of our results with the evidence presented in [116] who also use a sample ending in April 2003.

Please insert Table 3.2 here

For local currency-denominated returns, the point estimate of the predictive slope δ^{P} in regression (3.1) is negative for all six cases. The null hypothesis H_0 : $\delta^{\mathsf{P}} = 0$ is rejected in six (four) out of six cases at the 10% (5%) level of significance according to [135] standard errors. When [136] standard errors are used, $H_0 : \delta^{\mathsf{P}} = 0$ is rejected in five (four) out of six cases at the 10% (5%) level of significance; Japan yields the highest p-value equal to 0.12. The adjusted R-square is greater than 2% in five out of six cases; Canada yields the lowest adjusted R-square equal to 1%. The results for US dollar-denominated returns are qualitatively similar.[11] Note that our empirical exercise differs from the analysis in [116] in that we use our own oil

[11]Specifically, the point estimate of the predictive slope δ^{P} in regression (3.1) is negative for all eight cases. The null hypothesis $H_0 : \delta^{\mathsf{P}} = 0$ is rejected in seven (six) out of eight cases at the 10% (5%) level of significance according to [135] standard errors. When [136] standard errors are used, $H_0 : \delta^{\mathsf{P}} = 0$ is rejected in six (six) out of eight cases at the 10% (5%) level of significance; Japan again yields the highest p-value equal to 0.27. The adjusted R-square is greater than 2% in six out of eight cases; Canada yields the lowest adjusted R-square equal to 0.7%. Importantly, for the case of the World MSCI index, $H_0 : \delta^{\mathsf{P}} = 0$ is strongly rejected by both methods and the adjusted R-square is equal to 5.6%, which is rather high for monthly returns.

price change proxy, our sample starts at a different point in time, and we use excess returns as opposed to log returns. Despite these differences, our results confirm their evidence on the relationship between oil price changes and subsequent global equity returns for the sample period ending in April 2003.

Please insert Table 3.13 here

In the left panel of Table 3.13 in the Internet Appendix, we present evidence on the ability of oil price changes to predict MSCI index excess returns in terms of economic significance, over the 1983.01–2003.04 period. Specifically, we report results on the certainty equivalent return (CER) and the Sharpe ratio (SR) of the associated optimal trading strategies for a mean-variance investor with a risk aversion coefficient $\gamma = 3$ (see Appendix 3.7.2). These results reinforce the statistical significance results reported in the left panel of Table 3.2. Let $\mathtt{CER}_{\mathtt{IID}}$ and $\mathtt{SR}_{\mathtt{IID}}$ denote the CER and SR achieved by the optimal trading strategy assuming that the MSCI index excess returns are i.i.d., and $\mathtt{CER}_{\mathtt{P}}$ and $\mathtt{SR}_{\mathtt{P}}$ denote the CER and SR achieved by the optimal trading strategy using the predictive regression model (3.1). For local currency-denominated returns, the alternative model using the oil price change proxy $g^{\mathtt{P}}$ as predictor generates significantly higher (point estimates of) CERs and SRs compared to the baseline model that assumes that MSCI index excess returns are i.i.d. across all six countries. More importantly, the null hypothesis $H_0^{\mathtt{CER}} : \mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is rejected in six (five) out of six cases at the 10% (5%) level of significance, respectively, against the one-sided alternative $H_A^{\mathtt{CER}} : \mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$. The null hypothesis

H_0^{SR} : $\mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{P}}$ is rejected against the one-sided alternative H_A^{SR} : $\mathtt{SR}_{\mathtt{IID}} < \mathtt{SR}_{\mathtt{P}}$ in all six cases at the 5% level of significance.[12] Overall, the economic significance results confirm the evidence reported in [116] on the ability of oil price changes to forecast international equity index returns.

3.3.2 Evidence from the 1983–2015 sample period

In this subsection, we extend the sample period to December 2015 and run the same predictive regressions again. As in the previous subsection, we examine both the statistical and economic significance of the predictive ability of oil price changes.

In the right panel of Table 3.2, we present statistical significance results over the 1983.01–2015.12 period. The evidence obtained from the extended sample is quite different: the predictive ability of oil price changes has mostly disappeared.

For local currency-denominated returns, the point estimates of the predictive slope δ^{P} in regression (3.1) are still negative for all six countries. However, they are much smaller in absolute value. For instance, for Japan and the UK, the δ^{P} point estimates obtained over the 1983.01–2003.04 period are -0.11 and -0.11, while they fall to -0.05 and -0.06 over the 1983.01–2015.12 period, respectively. The null

[12] The results for US dollar-denominated returns are qualitatively similar. The alternative model based on the predictive regression model (3.1) still generates significantly higher (point estimates of) CERs and SRs compared to the baseline i.i.d. model for MSCI index excess returns across all eight cases. The null hypothesis H_0^{CER} : $\mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is rejected in five (three) out of eight cases at the 10% (5%) level of significance, respectively, against the one-sided alternative H_A^{CER} : $\mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$, with Canada yielding the highest p-value equal to 0.15. The null hypothesis H_0^{SR} : $\mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{P}}$ is rejected against the one-sided alternative H_A^{SR} : $\mathtt{SR}_{\mathtt{IID}} < \mathtt{SR}_{\mathtt{P}}$ in six (three) out of eight cases at the 10% (5%) level of significance, respectively, with Japan yielding the highest p-value equal to 0.14.

hypothesis $H_0 : \delta^{\mathsf{P}} = 0$ is now rejected in only three (two) out of six cases at the 10% (5%) level of significance according to [135] standard errors. Moreover, we observe a substantial reduction in adjusted R-squares. For instance, for Japan and the UK, the adjusted R-squares obtained over the 1983.01–2003.04 period are 2.9% and 4.1%, while they fall to 0.3% and 1.4% over the 1983.01–2015.12 period, respectively. In the case of Canada, the adjusted R-square even becomes negative.

The results for US dollar-denominated returns are even weaker. While the δ^{P} point estimates are still negative in seven out of eight cases, they are even smaller in absolute value than their counterparts obtained for local currency-denominated returns and, in the case of Canada, the predictive slope estimate becomes positive. The null hypothesis $H_0 : \delta^{\mathsf{P}} = 0$ is rejected only in the case of Italy at the 10% level of significance, regardless of whether we use [135] or [136] standard errors. Importantly, the corresponding p-values for the US and the World MSCI indexes are 0.22 and 0.22, respectively, according to [135] standard errors. In addition, the adjusted R-squares are rather low: they are less than 1% in seven out of eight cases, and even negative in the case of Canada. Hence, our results show substantially weaker statistical evidence on the relationship between oil price changes and subsequent global equity excess returns over the 1983.01–2015.12 period compared to the 1983.01–2003.04 period.

The right panel of Table 3.13 in the Internet Appendix reports results on the economic significance of the predictive ability of oil price changes, over the 1983.01–2015.12 sample period, in terms of the CER and the SR of the associated optimal trading strategies for a mean-variance investor with a risk aversion coefficient $\gamma = 3$ (see Appendix 3.7.2). These results reinforce the message, conveyed by the right panel of Table 3.2, that the forecasting ability of oil price changes has all but disappeared over the extended sample ending in December 2015. For local currency-denominated returns, the null hypothesis $H_0^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is *not* rejected against the one-sided alternative $H_A^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$ in four out of six cases at the 10% level of significance, with the exceptions of Italy and the UK. We obtain the same results when we test $H_0^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{P}}$ against $H_A^{\mathsf{SR}} : \mathtt{SR}_{\mathtt{IID}} < \mathtt{SR}_{\mathtt{P}}$. The results for US dollar-denominated returns are even weaker. The null hypothesis $H_0^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{P}}$ is *not* rejected against the one-sided alternative $H_A^{\mathsf{CER}} : \mathtt{CER}_{\mathtt{IID}} < \mathtt{CER}_{\mathtt{P}}$ in any case, out of eight, at the 10% level of significance. Importantly, the corresponding p-values for the US and the World MSCI indexes are 0.37 and 0.30, respectively.

Collectively, the statistical as well as economic significance results presented in this subsection illustrate that the forecasting ability of oil price changes has been diminished over the extended sample ending in December 2015. In the next subsection, we provide further corroborating evidence by examining the stability, or lack thereof, of the predictive relationship between MSCI index excess returns and past oil price changes.

3.3.3 Instability of the predictive slope coefficients

The empirical evidence gathered in the previous two subsections suggests that the ability of oil price changes to predict MSCI index excess returns is not stable over time. We confirm previous results on the oil price change predictive ability using data until April 2003, consistent with the evidence in [116], but also show that these results do not hold in the extended sample ending in December 2015. While we obtain negative and statistically significant predictive slope estimates in the early sample, these estimates become much closer to zero and lose their statistical significance in the extended sample.

Please insert Figure 3.6 here

As a first attempt to shed some light on these striking findings, we estimate the predictive regression model (3.1) over different samples using an expanding window. The first sample covers the 1983.01–1993.01 period and the last sample covers the 1983.01–2015.12 full period. Figures 3.6 and 3.7, both in the Internet Appendix, present the predictive slope estimates along with 95% confidence intervals, based on [135] standard errors, over the period 1993.01–2015.12 for local currency- and US dollar-denominated MSCI index returns, respectively. The pattern evident in these graphs is rather revealing. For the majority of the cases, the predictive slope estimates are negative and frequently statistically significant until the third quarter of 2008. For most cases after that point in time, however, the estimates start increasing to zero and quickly lose their statistical significance. This effect is more prominent for US dollar-denominated returns.

In addition to the informal analysis based on the predictive slope estimates presented in the aforementioned graphs, we also perform formal structural break tests. Specifically, we employ the methodology developed by [137] to test for multiple structural breaks in the predictive slope coefficients. The Bayesian Information

Criterion (BIC) is used to select the number of breaks.

Please insert Table 3.3 here

Table 3.3 presents [137] structural break tests in the slope coefficient for the predictive regression (3.1). The second column presents the BIC values assuming no break. The third and fourth columns provide the BIC values and the corresponding break dates for the case of the one-break model. The last column shows the number of breaks selected by the BIC. For local currency-denominated index returns, the test identifies the presence of one structural break in five out of six cases, with the only exception of the UK. For US dollar-denominated index returns, the test identifies the presence of one structural break in seven out of eight cases, with the only exception of France. The break dates identified in most cases fall in the third quarter of 2008. However, the break dates for Italy and Japan are October 2003 and September 1990, respectively. Overall, the structural break tests provide additional evidence against the stability of the slope coefficient in the predictive relationship between MSCI index excess returns and past oil price changes.

3.4 Identifying oil price shocks using structural VAR models

In the previous section, we confirm the finding of [116] that oil price changes predict international equity index returns at the monthly frequency with a negative predictive slope based on data up to April 2003. However, we also provide compelling evidence that the predictive power of oil price changes has practically disappeared over the extended sample ending in December 2015. For most of the MSCI indexes, the predictive slope estimates based on expanding windows become closer to zero and turn statistically insignificant after the third quarter of 2008. Moreover, the formal econometric tests of [137] indicate the existence of a structural break in the third quarter of 2008 for the majority of the cases, especially when US dollar-denominated returns are used. The dramatic reduction in the predictive ability of oil price changes, therefore, begs for an explanation.

In this paper, we offer an explanation that emphasizes the differential roles of the various shocks embedded in oil price changes. Specifically, we use a structural Vector Autoregression (VAR) framework that provides a decomposition of oil price changes into oil supply shocks, global demand shocks, and oil-specific demand shocks. As [119] and [120] point out, oil price shocks cannot be treated as strictly exogenous with respect to the global economy. In particular, they argue that oil supply shocks, global demand shocks, and oil-specific demand shocks, the combination of which leads to the observed aggregate oil price changes, should have different effects on the macroeconomy and the stock market.

[120] augment the structural VAR model of [119] by adding the US real stock return in the vector of variables and study the contemporaneous relationships between shocks embedded in oil price changes and stock returns. They further examine cumulative impulse responses of real stock returns to one-time shocks to oil supply, global demand, and oil-specific demand in the crude oil market. Their results, using data over the 1975–2006 period, show that an unexpected decrease in oil production has no significant effect on cumulative US real stock returns and a positive surprise

to global demand (oil-specific demand) leads to a continuous increase (decrease) in US real stock returns. In this paper, following the recent finance literature, [116] in particular, we cast the question in a predictive regression framework using short-horizon, i.e., one-month-ahead, forecasts. Although we utilize the structural VAR framework along the lines of [119], we approach the question from the perspective of an investor who wishes to use the real-time information embedded in oil price changes and captured by the three aforementioned shocks to predict subsequent equity index returns.

To disentangle the supply shocks, demand shocks, and oil-specific demand shocks embedded in the observed oil price changes, we employ two distinct structural VAR models. The first model is a variant of the [119] model which assumes that the short-run elasticity of oil supply is zero. The second model is a parsimonious version of the model proposed by [128] that allows for joint estimation of the elasticities of oil supply and oil demand. Given the ongoing debate in the literature about the relative importance of supply shocks, it is essential to illustrate that our empirical results are robust to different modelling choices for the elasticity of supply and their implications about supply shocks. Both of the structural VAR models describe the joint evolution of three variables capturing changes in the (i) supply of oil; (ii) global economic activity; and (iii) price of oil. The first variable, denoted by $g_t^{\mathtt{S}}$, where $\mathtt{S}$ stands for supply, is the log growth rate of world crude oil production. The second variable, denoted by $g_t^{\mathtt{GD}}$, where $\mathtt{GD}$ stands for global demand, is the first principal component of the log growth rates of the dry bulk cargo shipping cost index and global crude steel production. It has been argued in the literature, e.g.,

[119] and [134], among others, that fluctuations in shipping cost and global crude steel production capture changes in global economic activity growth and demand for oil. The third variable, denoted by $g_t^{\mathtt{P}}$ where $\mathtt{P}$ stands for price, is the first principal component of the log growth rates of West Texas Intermediate, Dubai, and Arab Light spot prices. We provide a detailed explanation of the data sources and construction in Section 3.2.

Our purpose is to employ the two structural VAR models to obtain a decomposition of oil price changes into three types of shocks and use them as predictors of MSCI index returns. We do so by, first, using information available in real time and, second, constructing three variables that are stationary in a consistent way as explained above. As a result, our approach differs from [119] and [128] who use detrended level variables. Importantly, as [138] point out, the variables of global real economic activity and log real oil price used in the structural VAR model in [119] appear to be non-stationary. On the contrary, our variables are, by construction, stationary.[13]

[13] Note that [119] uses the *growth rate* of oil production along with the *detrended level* of the logarithm of the economic activity index and the logarithm of real oil price in his SVAR model. There are several issues with this specification in the context of extracting shocks for forecasting purposes. First, there is no clear justification for mixing growth rates with levels in a VAR system. Second, given that detrending requires knowledge of the full sample, the shocks obtained by [119] cannot be obtained in a real-time fashion. Finally, due to detrending, different vintages of the economic activity index data can vary substantially, as more data become available.

3.4.1 Structural VAR framework

Let $\mathbf{g}_t = [g_t^{\mathtt{S}} \;\; g_t^{\mathtt{GD}} \;\; g_t^{\mathtt{P}}]'$ denote the vector of the three variables described above. The structural VAR model is then stated as:

$$\mathbf{A}_0 \mathbf{g}_t = \mathbf{a} + \sum_{i=1}^{p} \mathbf{A}_i \mathbf{g}_{t-i} + \boldsymbol{\varepsilon}_t, \tag{3.2}$$

where $\mathbf{A}_0$ is the 3×3 matrix specifying the contemporaneous structural relations between the three variables, $\mathbf{a} = [a_{\mathtt{S}} \;\; a_{\mathtt{GD}} \;\; a_{\mathtt{P}}]'$ is a 3×1 vector, $\mathbf{A}_i = [\mathbf{a}_{i,\mathtt{S}} \;\; \mathbf{a}_{i,\mathtt{GD}} \;\; \mathbf{a}_{i,\mathtt{P}}]'$ is a 3×3 matrix, for $i = 1, \ldots, p$, and $\boldsymbol{\varepsilon}_t = [\varepsilon_t^{\mathtt{S}} \;\; \varepsilon_t^{\mathtt{GD}} \;\; \varepsilon_t^{\mathtt{OSD}}]'$ is the vector of orthogonal structural shocks. The interpretation of the fundamental shocks is as follows: $\varepsilon_t^{\mathtt{S}}$ is the oil supply shock, $\varepsilon_t^{\mathtt{GD}}$ is the global demand shock, and $\varepsilon_t^{\mathtt{OSD}}$ is the oil-specific demand shock. The structural innovation vectors $\boldsymbol{\varepsilon}_t$ are, by assumption, serially and cross-sectionally uncorrelated with covariance matrix given by

$$\boldsymbol{\Sigma}_\varepsilon = \mathrm{Var}[\boldsymbol{\varepsilon}_t] = \begin{bmatrix} \sigma_{\mathtt{S}}^2 & 0 & 0 \\ 0 & \sigma_{\mathtt{GD}}^2 & 0 \\ 0 & 0 & \sigma_{\mathtt{OSD}}^2 \end{bmatrix}. \tag{3.3}$$

The reduced-form VAR innovation obtained from the structural VAR model (3.2) is $\mathbf{e}_t = \mathbf{A}_0^{-1}\boldsymbol{\varepsilon}_t$. Letting $\boldsymbol{\Sigma}_e$ denote the covariance matrix of $\mathbf{e}_t$, we obtain $\boldsymbol{\Sigma}_e = \mathbf{A}_0^{-1}\boldsymbol{\Sigma}_\varepsilon(\mathbf{A}_0^{-1})'$ which is the equation that provides identification of the matrices of interest $\mathbf{A}_0$ and $\boldsymbol{\Sigma}_\varepsilon$. Note that the covariance matrix $\boldsymbol{\Sigma}_e$ contains six distinct parameters. Due to the orthogonality of the structural shocks, $\boldsymbol{\Sigma}_\varepsilon$ contains three distinct parameters. Hence, only three distinct parameters in the matrix $\mathbf{A}_0$ can be identified. We consider two different structural VAR models reflected in the form of the matrix $\mathbf{A}_0$. The first model is a variant of the model advanced by [119] and assumes zero short-term oil supply elasticity; it is referred to as the IS model, where IS stands for inelastic supply of oil. The second model is a parsimonious version of the recent model proposed by [128] that facilitates joint identification of oil supply and demand elasticities; it is referred to as the ES model, where ES stands for elastic supply of oil. For now, we assume that the matrix $\mathbf{A}_0$ is fully identified, given knowledge of the covariance matrix $\mathbf{\Sigma}_e$. We elaborate more on how $\mathbf{A}_0$ is determined in the context of both models below.

The estimation of the VAR model and the decomposition of oil price changes proceed as follows. Multiplying both sides of the structural VAR model (3.2) by $\mathbf{A}_0^{-1}$ yields the reduced-form VAR model

$$\mathbf{g}_t = \mathbf{b} + \sum_{i=1}^{p} \mathbf{B}_i \mathbf{g}_{t-i} + \mathbf{e}_t, \tag{3.4}$$

where $\mathbf{b} = \mathbf{A}_0^{-1}\mathbf{a}$ and $\mathbf{B}_i = \mathbf{A}_0^{-1}\mathbf{A}_i$, $i = 1, \ldots, p$. The parameters $\mathbf{b}$ and $\mathbf{B}_i$ are estimated by standard OLS and the VAR order p is selected using the BIC criterion. The matrix $\mathbf{\Sigma}_e$ is estimated by the covariance matrix of the residuals from (3.4). Writing the VAR(p) system in VAR(1) form, we obtain

$$\mathbf{y}_t = \mathbf{C}\mathbf{y}_{t-1} + \mathbf{u}_t, \tag{3.5}$$

where $\mathbf{y}_t$ and $\mathbf{u}_t$ are $3p \times 1$ vectors defined by

$$\mathbf{y}_t = [\ \mathbf{g}_t' - \boldsymbol{\mu}_g' \quad \mathbf{g}_{t-1}' - \boldsymbol{\mu}_g' \quad \cdots \quad \mathbf{g}_{t-p+1}' - \boldsymbol{\mu}_g'\ ]', \tag{3.6}$$

$$\mathbf{u}_t = [\ \mathbf{e}_t' \quad \mathbf{0}_3' \quad \cdots \quad \mathbf{0}_3'\ ]', \tag{3.7}$$

$\boldsymbol{\mu}_g$ is the mean of $\mathbf{g}_t$ and $\mathbf{C}$ is a suitable $3p \times 3p$ matrix (involving the matrices $\mathbf{B}_i$, $i = 1, \ldots, p$). The Wold representation of $\mathbf{y}_t$ reads $\mathbf{y}_t = \sum_{i=0}^{\infty} \mathbf{C}^i \mathbf{u}_{t-i}$. Denoting by $\mathbf{D}_i$ the 3×3 upper-left block of the matrix $\mathbf{C}^i$ and defining the matrix $\mathbf{F}_i = \mathbf{D}_i \mathbf{A}_0^{-1}$, we can express $\mathbf{g}_t$ as $\mathbf{g}_t = \boldsymbol{\mu}_g + \sum_{i=0}^{\infty} \mathbf{F}_i \boldsymbol{\varepsilon}_{t-i}$. The third element of the vector $\mathbf{g}_t$ is the oil price change proxy denoted by $g^{\mathtt{P}}$. Hence, we obtain the following decomposition of $g^{\mathtt{P}}$ into three components

$$g_t^{\mathtt{P}} = \mu_g^{\mathtt{P}} + x_t^{\mathtt{S}} + x_t^{\mathtt{GD}} + x_t^{\mathtt{OSD}}, \tag{3.8}$$

where $\mu_g^{\mathtt{P}}$ is the mean of the oil price change $g_t^{\mathtt{P}}$, $x_t^{\mathtt{S}} = \sum_{i=0}^{\infty} (\mathbf{F}_i)_{31} \varepsilon_{t-i}^{\mathtt{S}}$ is the oil supply shock, $x_t^{\mathtt{GD}} = \sum_{i=0}^{\infty} (\mathbf{F}_i)_{32} \varepsilon_{t-i}^{\mathtt{GD}}$ is the global demand shock, and $x_t^{\mathtt{OSD}} = \sum_{i=0}^{\infty} (\mathbf{F}_i)_{33} \varepsilon_{t-i}^{\mathtt{OSD}}$ is the oil-specific demand shock.

As mentioned above, the reduced-form VAR model (3.4) is estimated using standard OLS. In our implementation, the order p is selected equal to 2 according to BIC.

3.4.2 Inelastic supply (IS) structural VAR model

The identification strategy in [119] rests on the following assumptions: (i) oil production does not respond contemporaneously to either global demand or oil price changes; (ii) global demand responds contemporaneously to oil production but not to oil price changes; and (iii) oil price responds contemporaneously to both oil production and global demand changes. Importantly, according to [119]'s setting, the short-run elasticity of oil supply is assumed to be zero. Note that the above restrictions imply that the matrix $\mathbf{A}_0$ in the IS model is lower triangular:

$$\mathbf{A}_0^{\texttt{IS}} = \begin{bmatrix} 1 & 0 & 0 \\ \vartheta & 1 & 0 \\ \varrho & \varphi & 1 \end{bmatrix}. \tag{3.9}$$

Recall that the matrices $\mathbf{A}_0$ and $\mathbf{\Sigma}_\varepsilon$ are identified through the equation $\mathbf{\Sigma}_e = \mathbf{A}_0^{-1}\mathbf{\Sigma}_\varepsilon(\mathbf{A}_0^{-1})'$. Hence, given the form of $\mathbf{A}_0^{\texttt{IS}}$, one can use the Cholesky factor of the covariance matrix $\mathbf{\Sigma}_e$ to recover both $\mathbf{A}_0^{\texttt{IS}}$ and $\mathbf{\Sigma}_\varepsilon$.

3.4.3 Elastic supply (ES) structural VAR model

While [119] has successfully emphasized the importance of disentangling oil demand and supply shocks, it has also received criticism in the literature regarding the assumption of inelastic supply of oil. One of the implications of the model is that oil supply shocks account for a rather small percentage of variation in oil prices (see [131], [128], and [129], among others). To address these concerns, we also consider a structural VAR model that facilitates estimation of both oil supply and demand elasticities. The model is a parsimonious version of the model proposed by [128]. The three equations describing the joint evolution of $\mathbf{g}_t$ are:

$$g_t^{\mathtt{S}} = \eta_{\mathtt{S}} g_t^{\mathtt{P}} + a_{\mathtt{S}} + \sum_{i=1}^{p} \mathbf{a}'_{i,\mathtt{S}} \mathbf{g}_{t-i} + \varepsilon_t^{\mathtt{S}}, \tag{3.10}$$

$$g_t^{\mathtt{S}} = \eta_{\mathtt{D}} g_t^{\mathtt{P}} + \theta g_t^{\mathtt{GD}} + a_{\mathtt{P}} + \sum_{i=1}^{p} \mathbf{a}'_{i,\mathtt{P}} \mathbf{g}_{t-i} + \varepsilon_t^{\mathtt{OSD}}, \tag{3.11}$$

$$g_t^{\mathtt{GD}} = \lambda g_t^{\mathtt{S}} + a_{\mathtt{GD}} + \sum_{i=1}^{p} \mathbf{a}'_{i,\mathtt{GD}} \mathbf{g}_{t-i} + \varepsilon_t^{\mathtt{GD}}. \tag{3.12}$$

Equations (3.10) and (3.11) characterize the oil market. Equation (3.10) describes the oil supply schedule. We assume that oil production reacts contemporaneously only to oil price changes. The short-run price elasticity of oil supply is captured by the parameter $\eta_{\mathtt{S}}$. Equation (3.11) describes the oil demand schedule. We assume that oil demand responds contemporaneously to both oil price changes and changes in global economic activity. The short-run price elasticity of oil demand is captured by the parameter $\eta_{\mathtt{D}}$. Equation (3.12) describes the evolution of global economic activity. We assume that global economic activity is affected contemporaneously only by changes in oil production. The corresponding $\mathbf{A}_0$ matrix is given by

$$\mathbf{A}_0^{\mathtt{ES}} = \begin{bmatrix} 1 & 0 & -\eta_{\mathtt{S}} \\ -\lambda & 1 & 0 \\ 1 & -\theta & -\eta_{\mathtt{D}} \end{bmatrix}. \tag{3.13}$$

As emphasized by [128], the parameters in $\mathbf{A}_0^{\mathtt{ES}}$ cannot be jointly identified without additional information. However, the ES model is very informative in the sense that it imposes restrictions on the feasible pairs of elasticities $(\eta_{\mathtt{S}}, \eta_{\mathtt{D}})$. Indeed, assuming a value for $\eta_{\mathtt{S}}$ and knowledge of the covariance matrix $\mathbf{\Sigma}_e$, one can recover the rest of the parameters in $\mathbf{A}_0^{\mathtt{ES}}$ according to the equation $\mathbf{A}_0 \mathbf{\Sigma}_e \mathbf{A}_0' = \mathbf{\Sigma}_\varepsilon$. This procedure results in an admissible set of pairs $(\eta_{\mathtt{S}}, \eta_{\mathtt{D}})$ that are consistent with the ES structural VAR model. To determine estimates of supply and demand elasticities consistent with their structural VAR model, [128] propose a two-stage identification strategy. First, they obtain independent estimates $(\eta_{\mathtt{S}}^*, \eta_{\mathtt{D}}^*)$ using an instrumental variable (IV) estimation approach.[14] Then, they select the optimal admissible pair of elasticities $(\widehat{\eta}_{\mathtt{S}}, \widehat{\eta}_{\mathtt{D}})$ by minimizing a suitable distance between the admissible pairs and $(\eta_{\mathtt{S}}^*, \eta_{\mathtt{D}}^*)$, i.e., solving the following problem:

$$\min_{\eta_{\mathtt{S}}} \begin{bmatrix} \eta_{\mathtt{S}} - \eta_{\mathtt{S}}^* \\ \\ \eta_{\mathtt{D}}(\eta_{\mathtt{S}}; \mathbf{\Sigma}_e) - \eta_{\mathtt{D}}^* \end{bmatrix}' \mathbf{V}^{-1} \begin{bmatrix} \eta_{\mathtt{S}} - \eta_{\mathtt{S}}^* \\ \\ \eta_{\mathtt{D}}(\eta_{\mathtt{S}}; \mathbf{\Sigma}_e) - \eta_{\mathtt{D}}^* \end{bmatrix}, \tag{3.14}$$

where $\eta_{\mathtt{D}}(\eta_{\mathtt{S}}; \mathbf{\Sigma}_e)$ is the elasticity of demand consistent with $\eta_{\mathtt{S}}$ and $\mathbf{\Sigma}_e$, and $\mathbf{V}$ is a diagonal matrix of weights that reflect the sampling error in the $(\eta_{\mathtt{S}}^*, \eta_{\mathtt{D}}^*)$ estimates. Adopting the identification procedure advanced by [128] to our setting and using their IV point estimates $(\eta_{\mathtt{S}}^*, \eta_{\mathtt{D}}^*) = (0.077, -0.074)$, we obtain the set of admissible

[14][128] estimate the price elasticity of oil supply and demand by using IV panel regressions. They use a narrative analysis to identify 14 exogenous episodes of large country-specific declines in oil production for 21 countries over the period 1985–2015. In the supply and demand regression equations, for a particular country, the instrumental variables for the price of oil are large exogenous declines in oil production in other countries. In the second stage, country-level IV regressions for supply and demand use changes in crude oil production and petroleum consumption in the particular country, by excluding exogenous episodes involving that country. They obtain point estimates of supply and demand elasticity equal to 0.077 and -0.074, respectively.

pairs as well as the optimal admissible pair of elasticities $(\widehat{\eta}_{\mathtt{S}}, \widehat{\eta}_{\mathtt{D}}) = (0.157, -0.136)$. The results are presented in Figure 3.3 that closely resembles Figure 2 in [128]. Even though we use a more parsimonious model, with three instead of five variables, different proxies and cast our VAR model in terms of growth rates as opposed to levels, our estimates of supply and demand elasticities are quite close to the estimates obtained by [128], i.e., $(\widehat{\eta}_{\mathtt{S}}, \widehat{\eta}_{\mathtt{D}}) = (0.11, -0.13)$. Moreover, our results are comparable to the results in [129] who obtain $(\widehat{\eta}_{\mathtt{S}}, \widehat{\eta}_{\mathtt{D}}) = (0.15, -0.35)$. The rest of the parameters in $\mathbf{A}_0^{\mathtt{ES}}$ and $\mathbf{\Sigma}_e$, namely λ, θ, $\sigma^2_{\mathtt{S}}$, $\sigma^2_{\mathtt{GD}}$, and $\sigma^2_{\mathtt{OSD}}$ are readily recovered from the equation $\mathbf{A}_0 \mathbf{\Sigma}_e \mathbf{A}_0' = \mathbf{\Sigma}_\varepsilon$.

Please insert Figure 3.3 here

3.4.4 Oil price decomposition according to the IS and ES models

We obtain estimates of the matrix $\mathbf{A}_0$, in the context of the IS and ES models, as we explain above. Using these estimates and employing equation (3.8), we decompose oil price changes into oil supply, global demand, and oil-specific demand shocks. The two structural models presented above, IS and ES, have quite different implications in terms of the role of supply shocks in explaining the variation in oil prices. Specifically, our IS model estimates, based on the full sample, imply that 1%, 5%, and 94% of oil price change variation is attributed to oil supply, global demand, and oil-specific demand shocks, respectively. In contrast, our ES model estimates, based on the full sample, imply that 58%, 5%, and 37% of oil price change variation is explained by oil supply, global demand, and oil-specific demand shocks, respectively. In Figure 3.4, we present the percentages of the variance of oil price changes attributed to oil supply and oil-specific demand shocks as a function of the elasticity of supply η_S in the context of the ES model. The pattern in Figure 3.4 illustrates that the percentage of oil price variation explained by oil supply shocks is a monotonically increasing function of the elasticity of supply η_S. The pattern is consistent with the fact that the IS model can be seen a limiting case of the ES model as the elasticity of supply, η_S, approaches zero.[15]

Please insert Figure 3.4 here

For the purposes of predicting future equity returns, we need estimates of the oil supply, global demand, and oil-specific demand shocks based on available data at each point in time. To obtain such estimates, we estimate the reduced-form VAR model (3.4) and obtain the decomposition in equation (3.8) in a real-time fashion for both the IS and ES models. Specifically, for each month in the period between January 1986 and December 2015, we estimate the VAR model using all available data starting in February 1982 and ending in that month. Then, we obtain the

[15] A direct comparison of our results with related findings in the literature might not be so straightforward given that we work with oil price changes and not detrended (log) oil prices. However, in terms of the contribution of global demand shocks to oil price fluctuations, our results are in line with those in [128], as it follows from comparing our Figure 3.4 with the left panel of their Figure 3. Note that the term "oil demand shocks" is used in [128] to describe "oil-specific demand shocks" not "global demand shocks" (see the caption of their Figure 3 for clarification). It follows from their Figure 3 that the contribution of global demand shocks to oil price variation is small, less than 10%, which is consistent with our estimates. In addition, Figure A.5 in [128], replicating [119], shows an even smaller contribution of global demand shocks to oil price variation. Hence, our results are consistent with existing findings in the literature.

time series of three shocks in the decomposition (3.8), but keep the vector of the oil supply, global demand, and oil-specific demand shocks only for the last month. Our *real-time* decomposition is obtained in a fashion that reflects all revisions to past data of crude oil and crude steel production, and, therefore, our results are not subject to look-ahead bias. We plot the real-time estimates of the oil supply, global demand, and oil-specific demand shocks from January 1986 to December 2015 in Figure 3.5. As expected, oil-specific demands shocks are more volatile under the IS model, while oil supply shocks are more volatile under the ES model. The global demand shocks are very similar across the two structural VAR models.

Please insert Figure 3.5 here

3.5 The predictive power of oil supply, global demand, and oil-specific demand shocks

In this section, we examine the ability of the three shocks obtained by the oil price change decomposition (3.8) to forecast next-month MSCI index excess returns over the sample period from January 1986 to December 2015. Specifically, we run the following predictive regression:

$$r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}, \tag{3.15}$$

where $r^{\mathtt{e}}_{t+1}$ denotes excess return on an MSCI index and $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, $x^{\mathtt{OSD}}_t$ denote the oil supply, global demand, oil-specific demand shocks obtained from the decomposition (3.8), respectively. We consider the shocks obtained from both the IS and ES structural VAR models. As in the previous section, we gauge the forecasting ability of the three shocks using measures of both statistical and economic significance. Furthermore, we offer comparisons between model (3.1), which uses oil price change as the sole predictor, and the decomposition-based model (3.15).

The three shocks identified by the decomposition (3.8) are anticipated to have different impacts on future equity returns. Given that demand for oil is less than perfectly elastic, a disruption in oil production would result in an oil price increase. Such a disruption would be potentially bad news for the real economy and the stock market while the corresponding shock $x^{\mathtt{S}}$ would be positive. Hence, one expects $\beta^{\mathtt{S}}$ to have a negative value in the predictive regression (3.15). Second, positive global demand shocks stimulate the global economy as a whole, although the impact might differ across countries. One, therefore, expects that a positive global demand shock would be good news for equity markets. At the same time, a positive global demand shock could drive up the price of oil, which, in turn, could have a slowing-down effect on certain economies. However, the overall effect should be dominated by the first direct impact and, hence, one expects a positive slope $\beta^{\mathtt{GD}}$ in the predictive regression (3.15). Third, following the interpretation in [119], an oil-specific demand shock is thought to capture changes in the demand for oil driven by precautionary motives. Accordingly, a positive oil-specific demand shock is thought to originate from the increased demand for oil due to uncertainty regarding future availability of oil and is, therefore, perceived as bad news for the global economy and the stock market. Hence, one expects $\beta^{\mathtt{OSD}}$ to have a negative value in the predictive regression (3.15).

Summarizing the above discussion, we view an oil price increase due to an oil supply or an oil-specific demand shock as bad news. Recall, however, that the two structural VAR models that we use have different implications regarding the magnitude of these shocks. According to the IS model, the percentages of oil price variation explained by oil supply and oil-specific demand shocks are 1% and 94%, respectively, while, according to the ES model, these percentages are 58% and 37%, respectively. Hence, we anticipate that, when we use the shocks obtained from the IS (ES) model, the oil-specific demand (oil supply) shocks will play a more important role in predicting future equity returns. Indeed, this is confirmed by the evidence presented in the following subsection.

3.5.1 Evidence on the predictive ability of the shocks obtained from the structural VAR models

Please insert Table 3.4 here

Please insert Table 3.5 here

In Tables 3.4 and 3.5, we present statistical significance results for the predictive regression (3.15) over the 1986.01–2015.12 sample period using the shocks obtained from the IS and ES models, respectively. To provide a direct comparison between the predictive regression model (3.1), which uses oil price change as the sole predictor, and the predictive regression model (3.15), we also estimate model (3.1) over the same sample period. We report results for MSCI index excess returns denominated in both local currencies and US dollars and compute standard errors using the [135] method with optimal bandwidth selected as in [139].

As expected, given the evidence presented in Section 3.3, the forecasting power of oil price changes diminishes over the 1986.01–2015.12 sample period. The results are very similar to those presented in Table 3.2 corresponding to the 1983.01–2015.12 sample period. In particular, the adjusted R-square for both the World and the US MSCI index is 0.5%. In stark contrast, we find strong evidence of predictability using the decomposition-based model (3.15) with the shocks obtained from either the IS or the ES model.

The results based on the IS model shocks are reported in Table 3.4 and summarized as follows. We focus on US dollar-denominated returns, as the results for local currency-denominated returns are very similar. The adjusted R-squares for the Canada, France, Germany, Italy, Japan, and UK MSCI indexes are 1.5%, 1.5%, 1.9%, 6.6%, 0.7%, and 2.1%, respectively. The $\beta^{\texttt{GD}}$ estimates are positive in all eight cases and statistically significant in four (six) cases at the 5% (10%) level of significance. The $\beta^{\texttt{OSD}}$ estimates are negative in all eight cases and statistically significant in five (six) cases at the 5% (10%) level of significance. The $\beta^{\texttt{S}}$ estimates are negative in six out of eight cases, although statistically significant only in one case. Importantly, for the US and the World MSCI indexes, the $\beta^{\texttt{GD}}$ estimate is positive, the $\beta^{\texttt{OSD}}$ estimate is negative, and both are statistically significant at the 5% level of significance, while the corresponding R-squares are 4.3% and 3.6%, respectively.

The results based on the ES model shocks, reported in Table 3.5, also show strong predictive ability of the shocks, except that in this case it is the oil supply shock, as opposed to the oil-specific demand shock, that negatively and significantly predicts future equity returns. Next, we briefly summarize the results for US dollar-denominated returns; the results for local currency-denominated returns are very similar. The adjusted R-squares for the Canada, France, Germany, Italy, Japan, and UK MSCI indexes are 1.5%, 1.1%, 1.5%, 6.0%, 0.7%, and 2.3%, respectively. The $\beta^{\mathtt{GD}}$ estimates are positive in all eight cases and statistically significant in four (six) cases at the 5% (10%) level of significance. The $\beta^{\mathtt{S}}$ estimates are negative in all eight cases and statistically significant in four (four) cases at the 5% (10%) level of significance. The $\beta^{\mathtt{OSD}}$ estimates are negative in five out of eight cases, although none of them are statistically significant. Importantly, for the US and the World MSCI indexes, the $\beta^{\mathtt{S}}$ estimate is negative, the $\beta^{\mathtt{GD}}$ estimate is positive, and both are statistically significant at the 5% level of significance, while the corresponding R-squares are 4.2% and 4.1%, respectively.[16]

Please insert Table 3.14 here

Please insert Table 3.15 here

[16] We repeat the above analysis, for both the IS and ES models, computing standard errors according to the [136] method. The results, reported in Tables 3.14 and 3.15 in the Internet Appendix, are similar and convey the same message. Collectively, there is strong statistical evidence supporting the usefulness of the decomposition (3.8) and the ability of the three associated shocks to forecast the World and G7 country MSCI index excess returns.

As mentioned earlier, there are important differences between the IS and ES models in terms of the amount of oil price change variation explained by the three shocks. In the case of the IS model, oil-specific demands shocks are much more important than supply shocks for accounting for oil price fluctuations. The roles are reversed in the case of the ES model. We use both the IS and ES models to illustrate that the predictive power of the three oil price shocks is robust to assumptions about the elasticity of supply.

In addition to the evidence on statistical significance, we also provide evidence on the economic significance of the ability of the oil supply, global demand, and oil-specific demand shocks to predict the G7 country and World MSCI index returns. We refer to the model described by the decomposition-based predictive regression (3.15) as the alternative model and compare it to three baseline models. The first baseline model assumes that the MSCI index excess return $r^{\mathtt{e}}_{t+1}$ is i.i.d. The second baseline model is described by the predictive regression (3.1) that uses the oil price change $g^{\mathtt{P}}$ as predictor. The third baseline model is described by the predictive regression

$$r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{GD}} + \delta^{\mathtt{GD}} g^{\mathtt{GD}}_t + u^{\mathtt{GD}}_{t+1}, \tag{3.16}$$

which uses the global demand growth proxy $g^{\mathtt{GD}}$ as predictor.

Let $\mathtt{CER}_{\mathtt{DEC}}$ and $\mathtt{SR}_{\mathtt{DEC}}$ denote the CER and SR achieved by the optimal trading strategy for a mean-variance investor with a risk aversion coefficient $\gamma = 3$ using the decomposition-based predictive regression model (3.15). Analogously, we denote by $\mathtt{CER}_{\mathtt{IID}}$, $\mathtt{CER}_{\mathtt{P}}$, and $\mathtt{CER}_{\mathtt{GD}}$ ($\mathtt{SR}_{\mathtt{IID}}$, $\mathtt{SR}_{\mathtt{P}}$, and $\mathtt{SR}_{\mathtt{GD}}$) the CERs (SRs) achieved by the optimal trading strategies using the i.i.d. model, the predictive regression (3.1) based on oil price change, and the predictive regression (3.16) based on global demand growth, respectively. To gauge the predictive ability of the shocks $x_t^{\mathtt{S}}$, $x_t^{\mathtt{GD}}$, and $x_t^{\mathtt{OSD}}$, we test the null hypotheses H_0^{CER} : $\mathtt{CER}_{\mathtt{IID}} = \mathtt{CER}_{\mathtt{DEC}}$, H_0^{CER} : $\mathtt{CER}_{\mathtt{P}} = \mathtt{CER}_{\mathtt{DEC}}$, and H_0^{CER} : $\mathtt{CER}_{\mathtt{GD}} = \mathtt{CER}_{\mathtt{DEC}}$ against their one-sided alternatives. Furthermore, in a similar fashion, we test the null hypotheses H_0^{SR} : $\mathtt{SR}_{\mathtt{IID}} = \mathtt{SR}_{\mathtt{DEC}}$, H_0^{SR} : $\mathtt{SR}_{\mathtt{P}} = \mathtt{SR}_{\mathtt{DEC}}$, and H_0^{SR} : $\mathtt{SR}_{\mathtt{GD}} = \mathtt{SR}_{\mathtt{DEC}}$ against their one-sided alternatives. The economic significance test results, based on the shocks obtained from the IS and ES models, are reported in Tables 3.6 and 3.7, respectively. Next, we discuss the results based on the shocks obtained from the IS model. The results based on the shocks obtained from the ES model are similar.

Please insert Table 3.6 here

Please insert Table 3.7 here

For local currency-denominated index returns, the decomposition-based model (3.15) generates CERs that are higher than their counterparts generated by the i.i.d. model in all six cases. The difference is sizable, e.g., more than 3.76%, in annualized terms, for Japan and the UK, and statistically significant in five out of six cases at the 10% level of significance. Moreover, the decomposition-based model (3.15) generates CERs that are higher than their counterparts generated by model (3.1) based on oil price change in five out of six cases, with the exception of France. In the remaining cases, the difference is greater than 1.2%, in annualized terms, and statistically significant in the case of Japan at the 10% level of significance.

The decomposition-based model (3.15) also performs substantially better than the model (3.16) based on global demand growth in terms of CER. It produces CERs that are higher in all six cases and statistically significant in four out of six cases at the 10% level of significance. The SR results are in line with the CER results. The decomposition-based model (3.15) generates SRs that are higher than their counterparts generated by the i.i.d. model in all six cases. The increase in SR is sizable, e.g., from 0.19 to 0.53 for the UK, and the difference is statistically significant in four out of six cases at the 5% level of significance. Moreover, the decomposition-based model (3.15) generates SRs that are at least as high as their counterparts generated by the model (3.1) based on oil price change in all six cases. The differences again are sizable and statistically significant in the case of Japan at the 5% level of significance. The decomposition-based model (3.15) also performs substantially better than the model (3.16) based on global demand growth in terms of SR. It produces SRs that are higher in all six cases and statistically significant in five out of six cases at the 10% level of significance.

The results for US dollar-denominated index returns convey the same message. Importantly, in the case of the US MSCI index, the decomposition-based model (3.15) generates an annualized CER equal to 9.28% compared to 6.52%, 6.08%, and 6.20% generated by the i.i.d. model, the predictive regression model (3.1), and the predictive regression model (3.16), respectively. The corresponding p-values are 0.11, 0.08, and 0.05, respectively. Even stronger results are obtained for the World MSCI index. The decomposition-based model (3.15) generates an annualized CER equal to 7.90% compared to 4.03%, 3.88%, and 3.30% generated by the i.i.d. model, model (3.1), and model (3.16), respectively. The difference is statistically significant in all three comparisons with p-values equal to 0.04, 0.07, and 0.03, respectively. Strong results are obtained in terms of SR as well. In the case of the US MSCI index, the decomposition-based model (3.15) generates an annualized SR equal to 0.65 compared to 0.48, 0.45, and 0.46 generated by the i.i.d. model, model (3.1), and model (3.16), respectively. The corresponding p-values are 0.12, 0.09, and 0.05, respectively. For the World MSCI index, the decomposition-based model (3.15) generates an annualized SR equal to 0.56 compared to 0.31, 0.30, and 0.27 generated by the i.i.d. model, model (3.1), and model (3.16), respectively. The difference is statistically significant in all three comparisons with p-values equal to 0.04, 0.06, and 0.03, respectively.

Thus far, in this subsection, we have provided strong evidence, in terms of statistical as well as economic significance, in support of the ability of the oil supply, global demand, and oil-specific demand shocks to predict the World and G7 country MSCI index excess returns, based on both IS and ES models. Next, we provide further corroborating evidence on the stability of this predictive relation between MSCI index excess returns and these three shocks. As in the case on the predictive regression model (3.1), we use the [137] methodology to test for structural breaks in the decomposition-based model (3.15). The results are presented in Table 3.8. The test identifies breaks in zero (two) out of the 14 cases considered, covering six local currency- and eight US dollar- denominated MSCI indexes, when we use shocks from the IS (ES) model. Collectively, the results illustrate the importance of disentangling oil price changes into oil supply, global demand, and oil-specific demand shocks for the purpose of forecasting international equity returns.

Please insert Table 3.8 here

To further illustrate the importance of using the oil price change decomposition for forecasting purposes, as opposed to the oil price change itself, it is useful to examine the contribution of the three shocks in the overall oil price change variation. Based on the ES model, oil supply shocks, global demand shocks, and oil-specific demand shocks explain 61%, 2%, and 37% of oil price change variation in the 1986-2003 sample period, respectively, while they explain 49%, 17%, and 34% of the variation in the 2004-2015 sample period, respectively. Hence, the bad-news oil supply and oil-specific demand shocks drive more of the variation in the first part of the sample, which is consistent with the top plot in Figure 3.1 and the evidence in [116]. In contrast, the good-news global demand shock explains more oil price change variation in the second part of the sample, which again is consistent with the bottom plot in Figure 3.1. These two facts combined explain the structural break in the predictive relationship between equity returns and oil price change documented in subsection 3.3.3.

In the next subsection, we examine various aspects of the relationship between the three shocks embedded in oil price changes and future stock returns. In particular, we provide evidence suggesting that the documented predictability does not appear to be consistent with time-varying risk premia.

3.5.2 Additional evidence, robustness checks, and discussion

Prior literature has shown that supply and demand shocks are *contemporaneously* correlated with stock market returns. Specifically, [131], using an oil-producing firm portfolio return as a demand proxy, obtains such a result. One natural question is whether our results follow from a continuation of this pattern. The evidence clearly suggests that this is not the case. First, the demand and supply shocks obtained by [131] do not have forecasting ability for stock returns. Second, the contemporaneous relationship between our shocks and stock returns is much weaker than the predictive relationship, as we illustrate in Table 3.16 in the Internet Appendix. For the continuation conjecture to be true, the contemporaneous relationship would have to be stronger than the predictive relationship. However, this conjecture not supported by the empirical evidence.

Please insert Table 3.16 here

To alleviate any concerns regarding the real-time availability of the data required to obtain the oil price change decomposition (3.8), we estimate the predictive regression $r^{\mathtt{d}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{d}}_{t+1}$ is the monthly net return on the World or a G7 country MSCI index constructed with a delay of one or two weeks.[17] The results for the shocks obtained from the IS and ES models, reported in Tables 3.17 and 3.18 in the Internet Appendix, respectively, are in

[17] We do not use excess returns for this exercise due to lack of availability of interest rate data for the relevant time periods.

line with the baseline evidence reported in Tables 3.4 and 3.5. Overall, the results illustrate the robustness of the forecasting ability of oil supply, global demand, and oil-specific demand shocks with respect to one- or two-week delays.

Please insert Table 3.17 here

Please insert Table 3.18 here

Please insert Table 3.9 here

Please insert Table 3.10 here

When studying return predictability, one natural question that emerges is whether the predictors under examination can forecast asset returns over horizons longer than one month. We present statistical evidence on the predictive ability, at the three-month and six-month horizons, of the shocks obtained using the IS and ES models in Tables 3.9 and 3.10, respectively. The results show that the predictive ability of the three shocks gradually diminishes as the horizon gets longer. Moreover, the statistical significance of the slope corresponding to the global demand shock decreases as we move from the three-month to the six-month horizon. This evidence is reinforced by the adjusted R-squares over horizons up to six months that we report in Table 3.11. The adjusted R-squares at the six-month horizon are typically much lower than their three-month counterparts. As argued by [140] and [116], among others, predictability typically associated with time-varying risk premia is long-lived and persists over long horizons. We document that this is not the case for the oil supply, global demand, and oil-specific demand shocks and, hence, we conclude that the documented predictability is not consistent with time-varying risk premia.

Please insert Table 3.11 here

As we detail in subsection 3.4.3, to estimate the ES model and obtain full identification, we rely on the approach of [128] and their auxiliary IV estimates of supply and demand elasticities. To illustrate the robustness of our results, we also estimate the ES model by setting the supply elasticity η_S equal to three fixed values, namely 0.10, 0.15, and 0.20. Based on each value of η_S, the elasticity of demand η_D, along with the rest of the parameters in $\mathbf{A}_0^{\mathbf{ES}}$ and $\mathbf{\Sigma}_e$, is recovered from the equation $\mathbf{A}_0 \mathbf{\Sigma}_e \mathbf{A}_0' = \mathbf{\Sigma}_\varepsilon$. The three sets of results, reported in Table 3.19 in the Internet Appendix, show that our predictability evidence is robust with respect to the level of supply elasticity.

Please insert Table 3.19 here

[132], using a theoretical framework, argue that participants in commodity markets face severe information frictions and, as a result, demand shocks are likely to be overestimated. To understand whether our predictability results are subject to such a concern, we focus on the ES model and perform the following simulation exercise. For each time period, we reduce the estimated demand shock by a factor of up to 50%, drawn from a uniform distribution, and then adjust the supply and

oil-specific demands proportionately so that the sum of the three shocks is still equal to the observed oil price change. We perform the predictive regression (3.15) using the perturbed time series of the three shocks and repeat the exercise 10,000 times. The quantiles of the resulting distributions of the p-values associated with the three predictive slopes and the adjusted R-square from the predictive regression are reported in Table 3.20 in the Internet Appendix, for the World and US MSCI indexes. The results show that, even under this sizable distortion, the supply and global demand shocks maintain their statistically significant predictive power and the adjusted R-square ranges between 2.8% and 4.8%, for at least 90% of the cases. We conclude that even if investors used crude methods to disentangle oil supply and global demand shocks, they could still benefit from the oil price decomposition and improve their return forecasts.

Please insert Table 3.20 here

Next, we examine whether the results on the predictability of aggregate equity index returns are robust in the cross section of US industries. Specifically, we use the 17 Fama-French value-weighted industry portfolios.[18]

First, we conduct [137] structural break tests for (i) the predictive regres-

[18]We use monthly returns on the 17 Fama-French value-weighted industry portfolios from Kenneth French's website. The abbreviations (descriptions) of the 17 industries are Food (Food), Mines (Mining and Minerals), Oil (Oil and Petroleum Products), Clths (Texiles, Apparel and Footware), Durbl (Consumer Durables), Chems (Chemicals), Cnsum (Drugs, Soap, Perfumes, Tobacco), Cnstr (Construction and Construction Materials), Steel (Steel Works Etc), FabPr (Fabricated Products), Machn (Machinery and Business Equipment), Cars (Automobiles), Trans (Transportation), Utils (Utilities), Rtail (Retail Stores), Finan (Banks, Insurance Companies, and Other Financials), and Other (Other).

sion using oil price change as the predictor, for the 1983.01–2015.12 sample period, and (ii) the predictive regression using the oil supply, global demand, and oil-specific demand shocks, obtained from both the IS and ES models, as predictors, for the 1986.01–2015.12 sample period. Table 3.21 in the Internet Appendix shows that the tests identify the presence of one structural break in 14 of 17 industry portfolios when oil price change is used as the sole predictor, with the exception of Mining and Minerals, Oil and Petroleum Products, and Utilities. In contrast, the tests do not identify a break for any industry when the three shocks embedded in oil price changes are used as predictors. Collectively, these results are consistent with the evidence from the G7 country and World MSCI indexes reported in subsection 3.5.1.

Please insert Table 3.21 here

Please insert Table 3.23 here

Second, we examine the ability of the three shocks, obtained from the IS and ES models, to forecast industry portfolio excess returns. The results are presented in Tables 3.22 and 3.23 in the Internet Appendix, where we also provide the results of the predictive regression using oil price change as the sole predictor, for the purposes of comparison. Given the results of the structural break tests discussed above, the results of the oil price change regression are meaningful only for the three industries that do not exhibit a break. As expected, the estimated predictive slope on oil price change is positive for the Oil and Petroleum Products industry, although not statistically significant. Overall, there is no evidence of predictability based on oil price change alone, with 15 out of 17 adjusted R-squares being less than 1%. In contrast, there exists strong evidence of predictability across the various industries based on the three shocks embedded in oil price changes, according to [135] standard errors. The results for the IS model shocks (Table 3.22) are summarized as follows. The $\beta^{\mathtt{GD}}$ estimates are positive for all 17 industries and statistically significant for 12 (13) industries at the 5% (10%) level. The $\beta^{\mathtt{OSD}}$ estimates are negative for 15 industries and statistically significant for seven (eight) industries at the 5% (10%) level. The $\beta^{\mathtt{S}}$ estimates are negative for all 17 industries, although statistically significant only for four industries at the 10% level. Moreover, the adjusted R-square is greater than 1.5% for 12 out of 17 industries. The results based on the ES model shocks (Table 3.23) convey the same message, except that, as in the case of MSCI index returns, it is the oil supply shock, and not the oil-specific demand shock, that negatively predicts industry returns. The $\beta^{\mathtt{GD}}$ estimates are positive for all 17 industries and statistically significant for 12 (14) industries at the 5% (10%) level. The $\beta^{\mathtt{S}}$ estimates are negative for all 17 industries and statistically significant for seven (eight) industries at the 5% (10%) level. The $\beta^{\mathtt{OSD}}$ estimates are negative for 10 industries but not statistically significant for any of them. Finally, the adjusted R-square is greater than 1.5% for 14 out of 17 industries. Overall, these results demonstrate that the predictive ability of the oil supply, global demand, and oil-specific demand shocks is strong not only for the aggregate equity index, but also across different US industry portfolios.

Please insert Table 3.22 here

Please insert Table 3.24 here

Another natural question in the context of equity return predictability is: how do the proposed predictors relate to macroeconomic variables that have been extensively used in the extant literature to model time-varying expected equity returns? Due to data limitations, we examine this issue only for the case of the US. Table 3.24 in the Internet Appendix presents the contemporaneous correlations between the oil supply, global demand, and oil-specific demand shocks, obtained from the IS and ES models in a real-time fashion, and four macroeconomic variables: the log dividend yield, the term spread, the default yield spread, and the one-month Treasury Bill rate. The correlations are rather low in magnitude with the largest (in absolute value) being the correlation between the global demand shock and the default yield equal to -0.17. In addition, we examine whether the forecasting ability of the oil supply, global demand, and oil-specific demand shocks is robust to the presence of the macroeconomic predictors in the case of the US. We examine the following linear predictive regressions $r^{\mathtt{e}}_{t+1} = \gamma^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + \boldsymbol{\theta}'\mathbf{z}_t + v^{\mathtt{P}}_{t+1}$ and $r^{\mathtt{e}}_{t+1} = \gamma^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + \boldsymbol{\lambda}'\mathbf{z}_t + v^{\mathtt{DEC}}_{t+1}$ over various sample periods, where $r^{\mathtt{e}}_{t+1}$ is the US MSCI index excess return and $\mathbf{z}_t$ is the vector of the four macroeconomic variables mentioned above. The results, based on shocks from both the IS and ES models, are reported in Table 3.12. According to the evidence, the inference results we have reported thus far in the paper are robust to the presence of the macroeconomic variables. In particular, the forecasting ability of oil price change over the early 1982.01–2003.04 sample period is unaffected. Moreover, over the 1986–2015 sample period, the slopes of the global demand shock and the oil supply (oil-specific demand) shock are positive and negative, respectively, and significant at the 5% level of significance for the ES (IS) model, consistent with our baseline results.

Please insert Table 3.12 here

Please insert Table 3.25 here

In our next empirical exercise, we examine the descriptive statistics of the conditional expected excess returns based on the oil supply, global demand, and oil-specific demand shocks obtained from the IS and ES models. In particular, we focus on the mean, the standard deviation, and the first three autocorrelations. This evidence can shed light on the issue of whether the documented predictive ability of the three shocks is consistent with time-varying risk premia. For the purposes of comparison, we use two benchmarks. The first benchmark is the predicted MSCI US index excess return based on the aforementioned macroeconomic variables in terms of descriptive statistics. We report the results for the time period 1986.01–2015.12 in Table 3.25 in the Internet Appendix. Overall, the conditional expected excess returns predicted by the three shocks are more volatile and much less persistent compared to their analogues obtained from the macroeconomic variables. One might argue that the predicted excess returns based on the macroeconomic variables are just too persistent, given the nature of these macroeconomic predictors. To address this concern, in our second comparison, we use as a benchmark the equity risk premium estimates obtained by [141] based on option prices over different maturities, ranging from one month to one year.[19] Table 3.26 in the Internet Appendix reports the results for the time period 1996.01–2012.01, which is the sample period used by [141]. The main message from the second comparison remains the same. In particular, the second and third order autocorrelations of the conditional expected excess returns predicted by the three shocks are much lower than their counterparts obtained from either the predicted excess returns based on the macroeconomic predictors or the risk premium estimates of [141]. Collectively, this evidence suggests that the forecasting ability of the three shocks is not consistent with time-varying risk premia, in line with the evidence of predictability diminishing over longer horizons as reported above.

Please insert Table 3.26 here

Next, we investigate whether there is a more direct link between time variation in expected returns and changes in risk, as captured by return volatility. Such an exercise can shed more light on the question of whether the predictive ability of the

[19] We thank Ian Martin for making the data available on his website.

oil supply, global demand, and oil-specific demand shocks is associated with changes in risk premia. To this end, we employ an augmented EGARCH(1,1) model that includes these three shocks in the volatility equation as exogenous regressors. If the variation of expected returns is to be attributed to time-varying risk premia, we expect that any of these three shocks would have the same effect on both the drift and the volatility. In the context of the EGARCH model, we expect the coefficient on any of these shocks to have the same sign as in the drift equation and be statistically significant. The econometric specification is:

$$r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}, \tag{3.17}$$

$$u^{\mathtt{DEC}}_{t+1} = \sigma_t z_{t+1}, \quad z_{t+1} \sim \text{i.i.d.}(0,1), \tag{3.18}$$

$$\log(\sigma^2_t) = \tau_0 + \tau_1|z_t| + \tau_2 z_t + \tau_3 \log(\sigma^2_{t-1}) + \zeta^{\mathtt{S}} x^{\mathtt{S}}_t + \zeta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \zeta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t. \tag{3.19}$$

As argued above, time-varying risk premia would be consistent with $\zeta^{\mathtt{S}} < 0$, $\zeta^{\mathtt{GD}} > 0$, and $\zeta^{\mathtt{OSD}} < 0$. We estimate the model using monthly excess returns on the MSCI indexes for the G7 countries as well as the World MSCI index, denominated both in local currencies and US dollars, over the 1986.01–2015.12 sample period. We consider three distributions for the disturbances z_{t+1}: Normal, Student-t, and GED. The Student-t distribution was selected according to the Bayesian Information Criterion.[20] The results for the IS and the ES models are presented in Tables 3.27 and 3.28 in the Internet Appendix, respectively. In the majority of the cases, the estimates $\zeta^{\mathtt{S}} < 0$, $\zeta^{\mathtt{GD}} > 0$, and $\zeta^{\mathtt{OSD}} < 0$ are statistically insignificant at conventional

[20]The results are very similar across all three distributional assumptions.

levels. Moreover, for the IS model, whenever there is significance, the sign is the opposite of what would be consistent with time-variation of risk premia; e.g., in five out of 14 instances, the estimates of $\zeta^{\mathtt{GD}}$ are statistically significant but negative. This evidence is inconsistent with the notion of time-varying risk premia, reinforcing the message conveyed by the evidence documented earlier.[21]

Please insert Table 3.27 here

Several pieces of the empirical evidence presented above collectively indicate that the documented predictability cannot be explained by time variation of risk premia. Hence, one important remaining question is what might explain such predictability. [116] suggest that the gradual diffusion of information hypothesis might explain the predictive ability of oil price change, as investors appear to react to readily available information with a delay. However, there is an important distinction between their context and ours. On the one hand, while oil price changes are directly observable, assessing their impact on the future trajectory of the real economy and the value of stocks is a formidable task for investors. On the other hand, while supply and demand shocks, embedded in oil price changes, have clear implications for the global economy and, by extension, international stock markets, these shocks are not directly observable and their estimation poses a serious challenge to investors.

[21]It is important to note that our results do not imply the existence of an arbitrage opportunity. What we document is that there is time variation of conditional expected equity returns explained by the three shocks from the oil price change decomposition and that this variation does not appear to be related to changes in risk. The trading strategies generated by using the three shocks as predictors are obviously risky; they do, however, offer a better risk-return trade-off compared to the strategies based on other predictors.

As pointed out by [142], economists and investors alike ask whether oil price fluctuations are caused by supply or demand shocks. This distinction is quite important as the bad-news supply shocks have very different implications from the good-news demand shocks. However, these shocks are not observable and their estimation is subject to a number of frictions. First, as emphasized by [132], the globalization of the crude oil market, with new supply from shale oil and demand stemming from many fast-growing emerging economies, exposes market participants to heightened informational frictions regarding the global supply and demand of crude oil. As a result, taking a stand on observable proxies for oil supply and demand becomes a non-trivial task. The second friction is the uncertainty regarding the elasticities of oil supply and demand, which are crucial for the estimation of oil supply and demand shocks, as illustrated by [128]. Finally, extracting the underlying fundamental shocks embedded in the observed oil prices requires sophisticated econometric modelling and is subject to much recent academic debate (e.g., [119], [128], and [129]). Given that a decomposition of oil price changes into fundamental supply and demand shocks relies on a number of choices regarding data, economic models, and econometric methods, it is conceivable that different agents end up with different estimates of these underlying shocks. Hence, the documented predictability can be potentially explained by these informational frictions that prevent investors from measuring supply and demand shocks with precision. At the same time, as better models and techniques become available and more investors become aware of the importance of oil supply and demand shocks, it is also conceivable that their predictive power diminishes over time. This issue should be the subject of future research.

3.6 Conclusion

As the modern global economy heavily depends on oil, the price of oil is widely thought to affect global real economic activity and, consequently, the global equity market. An oil price drop, in the past, has been considered as good news, as it lowers the cost of production in a significant number of sectors and allows consumers to boost their consumption. Accordingly, one could hypothesize that negative (positive) oil price changes should predict higher (lower) subsequent equity returns. [116] document that this is indeed the case for a large number of MSCI equity indexes based on data until 2003. However, this predictive relationship has dramatically changed over the last ten years. Specifically, the correlation between the World MSCI index return and the lagged one-month log growth rate of West Texas Intermediate spot price has increased from -0.22 over the 1982–2003 period to 0.25 over the 2004–2015 period. As a result, the ability of oil price change to forecast future equity returns has diminished over the sample period that extends to 2015. Furthermore, using the formal econometric test of [137], we detect a structural break in the predictive relationship in the third quarter of 2008 for most of the G7 country MSCI index returns.

In this paper, we suggest that oil price changes do, in fact, contain useful information for forecasting subsequent equity indexes, provided that these changes are suitably disentangled into supply and demand shocks. Using two distinct structural

VAR models, we obtain an oil price change decomposition into an oil supply shock, a global demand shock, and an oil-specific demand shock and argue that these three different types of shocks should have different effects on equity markets. The first model is a variant of the model proposed by [119], and assumes zero elasticity of oil supply. The second model is a parsimonious version of the model advanced by [128] that facilitates joint estimation of oil supply and demand elasticities. The hypothesis that oil supply shocks and oil-specific demand shocks (global demand shocks) predict equity returns with a negative (positive) slope is supported by the empirical evidence over the 1986–2015 sample period, using shocks obtained from both structural VAR models. Using the oil price decomposition obtained from the first structural VAR model, instead of just oil price change, leads to an increase of the annualized certainty equivalent return and Sharpe ratio of a mean-variance optimal trading strategy for the World MSCI index from 3.88% to 7.90% and from 0.30 to 0.56, respectively, with the differences being statistically significant. When we use the shocks obtained from the second structural VAR model, the corresponding increases are from 3.88% to 7.67% for the annualized certainty equivalent return and from 0.30 to 0.55 for the Sharpe ratio. These results survive a comprehensive battery of robustness checks and, in general, do not appear to be consistent with time-varying risk premia.

3.7 Appendices

3.7.1 Data construction

The single oil price change proxy $g^{\mathtt{P}}$ is constructed in a real-time fashion using PCA. Specifically, for each month t between January 1983 and December 2015, we use data on three proxies for oil price change starting in February 1982 and ending in month t. We first rescale the three log growth rates, obtained from the West Texas Intermediate, the Dubai, and the Arab Light spot prices, so they all have variances equal to one over the given sample period and then perform PCA. The first PCA corresponding to month t is kept each time and the process is repeated using expanding windows until December 2015 is reached.

To address the strong seasonality of the global crude steel production data, we use X-13ARIMA-SEATS to compute seasonally-adjusted level data from which we compute log growth rates in a real-time fashion.[22] Specifically, for each month in the period between February 1982 and December 2015, we perform seasonal adjustment on the level data starting in January 1968 and ending in that month, compute the log growth rates of the seasonally adjusted level data, and, finally, keep the log growth rate over the last month.

The single global demand growth proxy $g^{\mathtt{GD}}$ is also constructed in a real-time fashion using PCA. Specifically, for each month t between January 1983 and December 2015, we use data on two proxies for global economic activity starting in

[22] We use the X-13 Toolbox for Matlab, written by Yvan Lengwiler, to perform seasonal filtering. The source codes are retrieved from `http://www.mathworks.com/matlabcentral/fileexchange/49120-x-13-toolbox-for-seasonal-filtering/content/x13tbx/x13.m`.

February 1982 and ending in month t. We first rescale the two log growth rates, obtained from the shipping cost index and the global crude steel production data, so they all have variances equal to one over the given sample period and then perform PCA. The first PCA corresponding to month t is kept each time and the process is repeated using expanding windows until December 2015 is reached.

3.7.2 Evaluation of predictive ability

In this paper, we examine the ability of (i) oil price changes and (ii) the oil supply, global demand, and oil-specific demand shocks embedded in these changes to forecast MSCI index excess returns. Following the literature, we employ linear predictive regressions of the type:

$$r^{\mathsf{e}}_{t+1} = \alpha + \boldsymbol{\beta}'\mathbf{x}_t + u_{t+1}, \tag{3.20}$$

where r^{e}_{t+1} is the MSCI index excess return, $\mathbf{x}_t = [x_{1,t} \ \cdots \ x_{n,t}]'$ is a vector of predictors, $\boldsymbol{\beta} = [\beta_1 \ \cdots \ \beta_n]'$ is the vector of predictive slope coefficients, and u_{t+1} is a zero-mean random disturbance. The predictor $\mathbf{x}_t$ could be a scalar $(n = 1)$, e.g., when we use a sole predictor such as oil price change, or a multidimensional vector $(n > 1)$, e.g., when we use multiple predictors such as the three different shocks and/or additional controls. In some instances, we also consider the i.i.d. model for r^{e}_{t+1} in which case the vectors $\boldsymbol{\beta}$ and $\mathbf{x}_t$ are null and equation (3.20) reduces to $r^{\mathsf{e}}_{t+1} = \alpha + u_{t+1}$. We evaluate predictive ability in terms of both statistical and economic significance.

The question we wish to address is whether $\mathbf{x}_t$ can forecast the MSCI index excess return r^{e}_{t+1}. Hence, we are interested in testing the null hypotheses H_0 : $\beta_i = 0$, for $i = 1, \ldots, n$. We evaluate the statistical significance of predictive ability of $\mathbf{x}_t$ using standard metrics. Specifically, we obtain two-sided p-values for the null hypotheses $H_0 : \beta_i = 0$, $i = 1, \ldots, n$ based on standard errors computed according to two well-established approaches: (i) the [135] method, where the optimal bandwidth is selected following the approach in [139], and (ii) the [136] method that imposes the no-predictability condition. Finally, we also report adjusted R-squares.

To gauge the economic significance of the predictive ability of $\mathbf{x}_t$, we consider a mean-variance investor who can invest in an MSCI index and the corresponding short-term Treasury Bill. The investor uses the regression model (3.20) to forecast MSCI index excess returns. An optimal trading strategy is then developed based on the resulting estimates of the conditional mean and variance of excess returns. We evaluate economic significance in terms of two commonly used metrics: (i) the certainty equivalent return (CER) and (ii) the Sharpe ratio (SR) of the associated optimal portfolio returns.

Following [143], we assume that the risk aversion coefficient of the mean-variance investor is $\gamma = 3$. At the end of each period t, the investor uses all available data to estimate the parameters of the linear predictive regression (3.20). Using these parameter estimates, the investor then obtains estimates of the mean and the variance of the MSCI index excess return r^{e}_{t+1} at time t, denoted by $\widehat{\mu}_{t+1}$ and $\widehat{v}_{t+1}$, respectively. These estimates give rise to the following optimal portfolio weight on the MSCI index:

$$\omega_t = \frac{1}{\gamma}\frac{\widehat{\mu}_{t+1}}{\widehat{v}_{t+1}}. \tag{3.21}$$

The rest of the investor's wealth is invested in the short-term Treasury Bill. We assume that the portfolio weight on the MSCI index is constrained between a minimum and maximum feasible weight, denoted by $\underline{\omega}$ and $\overline{\omega}$, respectively. The minimum weight $\underline{\omega}$ is set equal to zero so that short-selling is precluded. Following [143], we set the maximum weight, $\overline{\omega}$, equal to 150% so that the investor is allowed to borrow up to 50% and invest the proceeds in the MSCI index. Optimal weights are determined according to equation (3.21) and then the realized portfolio returns are computed. Below, we describe the two metrics, CER and SR, used in our evaluation of economic significance of predictability.

The CER of the resulting optimal portfolio from period 1 to period T based on the predictive regression (3.20) is given by

$$\widehat{\mathtt{CER}} \; = \; \widehat{\mu}_p \; - \frac{\gamma}{2}\widehat{v}_p, \tag{3.22}$$

where the mean $\widehat{\mu}_p$ and the variance $\widehat{v}_p$ of the realized optimal portfolio net returns are defined by

$$\widehat{\mu}_p = \frac{1}{T}\sum_{t=0}^{T-1}(r^{\mathtt{f}}_{t+1} + \omega_t r^{\mathtt{e}}_{t+1}) \quad \text{and} \quad \widehat{v}_p = \frac{1}{T}\sum_{t=0}^{T-1}\left((r^{\mathtt{f}}_{t+1} + \omega_t r^{\mathtt{e}}_{t+1}) - \widehat{\mu}_p\right)^2, \tag{3.23}$$

and $r^{\mathtt{e}}_{t+1}$ and $r^{\mathtt{f}}_{t+1}$ denote the excess return on the MSCI index and the corresponding Treasury Bill rate at time $t+1$, respectively.

The SR of the resulting optimal portfolio from period 1 to period T based on the predictive regression (3.20) is given by

$$\widehat{\mathtt{SR}} \;=\; \frac{\widehat{\mu}_p^{\mathsf{e}}}{\sqrt{\widehat{v}_p^{\mathsf{e}}}}, \tag{3.24}$$

with the mean and the variance of the realized optimal portfolio excess returns defined by

$$\widehat{\mu}_p^{\mathsf{e}} = \frac{1}{T}\sum_{t=0}^{T-1} \omega_t r_{t+1}^{\mathsf{e}} \quad \text{and} \quad \widehat{v}_p^{\mathsf{e}} = \frac{1}{T}\sum_{t=0}^{T-1} (\omega_t r_{t+1}^{\mathsf{e}} - \widehat{\mu}_p^{\mathsf{e}})^2. \tag{3.25}$$

We express the CERs obtained from equation (3.22) in annualized percentages by multiplying by 1,200 and annualize the monthly SRs from equation (3.24) by multiplying by $\sqrt{12}$. CER represents the equivalent risk-free rate of return that a mean-variance investor would require in exchange of a risky portfolio return series, while SR measures the average portfolio excess return per unit of risk as measured by the portfolio excess return standard deviation.

If the variables in a vector $\mathbf{x}_t$ have nontrivial predictive ability, then using the predictive regression model (3.20) is expected to generate a higher CER (SR) than using the i.i.d. model for the MSCI index excess returns r_{t+1}^{e}. In this context, we refer to the i.i.d. model as the baseline model (Model 1) and the predictive regression model using the vector of predictors $\mathbf{x}_t$ as the alternative model (Model 2). We are interested in providing a formal comparison of the two models in terms of CER and SR. Denote by $\mathtt{CER}_j$ and $\mathtt{SR}_j$ the CER and SR of Model j, for $j = 1, 2$. Even if the point estimate of the CER and/or SR generated by the alternative model is higher than its counterpart generated by the baseline model, i.e., $\widehat{\mathtt{CER}}_1 < \widehat{\mathtt{CER}}_2$, one might be concerned whether this is due to genuine predictive ability of $\mathbf{x}_t$ or simply due to sample variability. Therefore, it is important to test the statistical significance of any differences in CER and SR. To this end, we develop asymptotic tests for the null hypothesis $H_0^{\mathsf{CER}} : \mathtt{CER}_1 = \mathtt{CER}_2$ against the one-sided alternative $H_A^{\mathsf{CER}} : \mathtt{CER}_1 < \mathtt{CER}_2$ and similarly the null hypothesis $H_0^{\mathsf{SR}} : \mathtt{SR}_1 = \mathtt{SR}_2$ against the one-sided alternative $H_A^{\mathsf{SR}} : \mathtt{SR}_1 < \mathtt{SR}_2$. Our purpose is to evaluate the incremental value of the alternative Model 2 compared to the baseline Model 1 and, hence, we focus on one-sided alternative hypotheses. The same framework can be used to facilitate more general comparisons. For instance, an important comparison in terms of predictive ability is between the oil price change and the vector of oil supply, global demand, and oil-specific demand shocks. For this comparison, Model 1 (baseline) corresponds to the predictive regression (3.20) with $\mathbf{x}_t$ consisting of the oil price change, while Model 2 (alternative) corresponds to the predictive regression (3.20) with $\mathbf{x}_t$ consisting of the three shocks.

Next, we provide the details about the aforementioned asymptotic tests and the computation of p-values. Let $\mathbf{r}_t = (r_{1,t}, r_{2,t})'$ denote the pair of returns on the portfolios generated by Models 1 and 2 at time t. These returns would be either net or excess depending on whether we focus on the CER or the SR. Note that, in the context of our mean-variance framework, the CER of a portfolio is expressed as a function of the first two moments of net portfolio returns, while the SR of a portfolio is expressed as a function of the first two moments of the portfolio excess returns. Denote the mean, variance, and non-central second moment of $r_{j,t}$ by μ_j, σ_j^2, and

ν_j, respectively for the portfolios $j = 1, 2$. Note that $\sigma_j^2 = \nu_j - \mu_j^2$. It follows that, in the case of net returns, the CERs for an investor with mean-variance preferences and a risk aversion coefficient equal to γ are given by $\mathtt{CER}_j = \mu_j - \frac{\gamma}{2}\left(\nu_j - \mu_j^2\right)$, $j = 1, 2$. Similarly, in the case of excess returns, the SRs are given by $\mathtt{SR}_j = \frac{\mu_j}{\sqrt{\nu_j - \mu_j^2}}$, $j = 1, 2$. Therefore, the relevant hypotheses can be stated using a suitable function of the parameter vector $\boldsymbol{\theta} = (\mu_1, \mu_2, \nu_1, \nu_2)'$. We estimate $\boldsymbol{\theta}$ by the sample analogue $\widehat{\boldsymbol{\theta}} = (\widehat{\mu}_1, \widehat{\mu}_2, \widehat{\nu}_1, \widehat{\nu}_2)'$, where $\widehat{\mu}_j = \frac{1}{T}\sum_{t=1}^{T} r_{j,t}$ and $\widehat{\nu}_j = \frac{1}{T}\sum_{t=1}^{T} r_{j,t}^2$, for $j = 1, 2$. Under regularity conditions, such as stationarity and ergodicity, $\widehat{\boldsymbol{\theta}}$ asymptotically follows a normal distribution described by

$$\sqrt{T}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}) = \frac{1}{\sqrt{T}}\sum\nolimits_{t=1}^{T} \mathbf{y}_t \lim N(\mathbf{0}, \boldsymbol{\Psi}), \tag{3.26}$$

where $\boldsymbol{\Psi}$ is the long-run variance-covariance matrix of

$$\mathbf{y}_t = \left(r_{1,t} - \mu_1, r_{2,t} - \mu_2, r_{1,t}^2 - \nu_1, r_{2,t}^2 - \nu_2\right)'. \tag{3.27}$$

The matrix $\boldsymbol{\Psi}$ is given by $\boldsymbol{\Psi} = \boldsymbol{\Gamma}_0 + \sum_{\ell=1}^{\infty}(\boldsymbol{\Gamma}_\ell + \boldsymbol{\Gamma}_\ell')$, where $\boldsymbol{\Gamma}_\ell = \mathrm{E}\left[\mathbf{y}_t \mathbf{y}_{t-\ell}'\right]$, for $\ell = 0, 1, \ldots$ and is estimated by a heteroscedasticity and autocorrelation consistent (HAC) estimator of the form

$$\widehat{\boldsymbol{\Psi}} = \widehat{\boldsymbol{\Gamma}}_0 + \sum_{\ell=1}^{T} \kappa\left(\frac{\ell}{b_T}\right)(\widehat{\boldsymbol{\Gamma}}_\ell + \widehat{\boldsymbol{\Gamma}}_\ell'), \tag{3.28}$$

where

$$\widehat{\mathbf{\Gamma}}_\ell = \frac{1}{T-\ell}\sum_{t=\ell+1}^{T}\widehat{\mathbf{y}}_t\widehat{\mathbf{y}}'_{t-\ell}, \quad \widehat{\mathbf{y}}_t = \left(r_{1,t}-\widehat{\mu}_1, r_{2,t}-\widehat{\mu}_2, r_{1,t}^2-\widehat{\nu}_1, r_{2,t}^2-\widehat{\nu}_2\right)', \tag{3.29}$$

$\kappa(\cdot)$ is a kernel function, and b_T is the bandwidth. HAC estimators have been developed by several authors including [135], [144], [145], and [139]. We report p-values based on the [135] approach with the Bartlett kernel and the optimal bandwidth computed as suggested in [139].

Consider testing the null hypothesis $H_0 : f(\boldsymbol{\theta}) = 0$ against the alternative hypothesis $H_A : f(\boldsymbol{\theta}) < 0$, where $f(\boldsymbol{\theta})$ is a smooth real-valued function of $\boldsymbol{\theta}$. Applying the delta method, we obtain

$$\sqrt{T}\left(f(\widehat{\boldsymbol{\theta}}) - f(\boldsymbol{\theta})\right) \lim N\left(0, \nabla' f(\boldsymbol{\theta})\boldsymbol{\Psi}\nabla f(\boldsymbol{\theta})\right), \tag{3.30}$$

where $\nabla f(\cdot)$ is the gradient of f. For large T, the standard error of $f(\widehat{\boldsymbol{\theta}})$ is given by

$$se(f(\widehat{\boldsymbol{\theta}})) = \sqrt{\frac{1}{T}\nabla' f(\widehat{\boldsymbol{\theta}})\widehat{\boldsymbol{\Psi}}\nabla f(\widehat{\boldsymbol{\theta}})}, \tag{3.31}$$

and, therefore, the corresponding t-statistic is $t(f,\widehat{\boldsymbol{\theta}}) = \frac{f(\widehat{\boldsymbol{\theta}})}{se(f(\widehat{\boldsymbol{\theta}}))}$, yielding the one-sided p-value $p(f,\widehat{\boldsymbol{\theta}}) = \Phi(t(f,\widehat{\boldsymbol{\theta}}))$, where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution.

To test for equality of CERs, we use net returns and the function f takes the form

$$f_{\mathsf{CER}}\left(\boldsymbol{\theta}\right) = \left(\mu_1 - \frac{\gamma}{2}\left(\nu_1 - \mu_1^2\right)\right) - \left(\mu_2 - \frac{\gamma}{2}\left(\nu_2 - \mu_2^2\right)\right), \tag{3.32}$$

with gradient equal to

$$\nabla f_{\mathsf{CER}}(\boldsymbol{\theta}) = \left(1 + \gamma\mu_1, -1 - \gamma\mu_2, -\frac{\gamma}{2}, \frac{\gamma}{2}\right)'. \tag{3.33}$$

To test for equality of SRs, we use excess returns and the function f takes the form

$$f_{\mathsf{SR}}\left(\boldsymbol{\theta}\right) = \frac{\mu_1}{\sqrt{\nu_1 - \mu_1^2}} - \frac{\mu_2}{\sqrt{\nu_2 - \mu_2^2}}, \tag{3.34}$$

with gradient equal to

$$\nabla f_{\mathsf{SR}}(\boldsymbol{\theta}) = \left(\frac{\nu_1}{(\nu_1 - \mu_1^2)^{\frac{3}{2}}}, -\frac{\nu_2}{(\nu_2 - \mu_2^2)^{\frac{3}{2}}}, -\frac{1}{2}\frac{\mu_1}{(\nu_1 - \mu_1^2)^{\frac{3}{2}}}, \frac{1}{2}\frac{\mu_2}{(\nu_2 - \mu_2^2)^{\frac{3}{2}}}\right)'. \tag{3.35}$$

3.7.3 Additional results and robustness checks

In this Internet Appendix, we provide additional results and robustness checks. In Table 3.13, we report results on the economic significance of the predictive ability of oil price change in terms of the CER and the SR of the associated optimal trading strategies for a mean-variance investor with a risk aversion coefficient $\gamma = 3$. In Tables 3.14 and 3.15, we present predictive regression estimation results for MSCI indexes based on (i) the oil price change and (ii) the oil supply, global demand, and oil-specific demand shocks obtained from the IS and the ES models, respectively, using [136] standard errors. In Table 3.16, we report the results of the contemporaneous regression of MSCI index excess returns on oil supply, global demand, and oil-specific demand shocks, based on the IS and ES models. In Tables 3.17 and 3.18, to address any concerns regarding the real-time data availability, we illustrate the robustness of the forecasting ability of oil supply, global demand, and oil-specific demand shocks with respect to one- and two-week delays. In Table 3.19, to illustrate the robustness of our inference with respect to the estimation of the elasticity of oil supply η_S, we present three sets of results where η_S takes the values 0.10, 0.15, and 0.20, respectively. In Table 3.20, to address the concern that global demand shocks might be overestimated, we report the predictive regression results based the three shocks obtained from the oil price change decomposition, after randomly perturbing the demand shocks downwards in 10,000 simulation repetitions. In Table 3.21, we report [137] structural break test results for the regressions using (i) oil price change and (ii) the oil supply, global demand, and oil-specific demand shocks to forecast the Fama-French 17 industry portfolio excess returns. In Tables 3.22 and 3.23, we present predictive regression estimation results for the Fama-French 17 industry portfolios based on (i) the oil price change and (ii) the oil supply, global demand, and oil-specific demand shocks obtained from the IS and the ES models, respectively. In Table 3.24, we report correlations between oil supply, global demand, and oil-specific demand shocks and several US macroeconomic variables. In Table 3.25, we present descriptive statistics for predicted excess returns on MSCI indexes, based on US macroeconomic predictors as well as the three shocks obtained from the oil price change decomposition, according to the IS and ES models, over the 1986.01–2015.12 sample period. For purposes of comparison, in Table 3.26, we also present the descriptive statistics for predicted excess returns on MSCI indexes over the 1996.01–2012.01 sample period along with the descriptive statistics for the US risk premium estimates of [141]. In Tables 3.27 and 3.28, we report estimation results for an augmented EGARCH(1,1) model that includes the oil supply, global demand, and oil-specific demand shocks in the volatility equation as exogenous regressors for the IS and the ES models, respectively. Finally, in Figures 3.6 and 3.7, we present the slope estimates for the predictive regression model (3.1) over different samples using an expanding window for local currency- and US dollar-denominated MSCI index returns, respectively.

1983.01–2003.04 Sample Period

	WTI	Dubai	Arab Light	g^{P}
Min	−39.60	−37.76	−48.51	−49.85
Max	37.71	53.68	48.73	38.16
Mean	−0.05	−0.10	−0.08	−0.26
Std. dev.	8.17	10.51	10.94	9.18
# of obs.	244	244	244	244

Correlation Matrix

	WTI	Dubai	Arab Light	g^{P}
WTI	1.00	0.76	0.72	0.90
Dubai	0.76	1.00	0.90	0.94
Arab Light	0.72	0.90	1.00	0.92
g^{P}	0.90	0.94	0.92	1.00

1983.01–2015.12 Sample Period

	WTI	Dubai	Arab Light	g^{P}
Min	−39.60	−49.71	−48.51	−49.85
Max	37.71	53.68	48.73	38.16
Mean	0.04	0.02	0.01	−0.09
Std. dev.	8.45	10.16	10.34	9.00
# of obs.	396	396	396	396

Correlation Matrix

	WTI	Dubai	Arab Light	g^{P}
WTI	1.00	0.76	0.72	0.90
Dubai	0.76	1.00	0.91	0.95
Arab Light	0.72	0.91	1.00	0.93
g^{P}	0.90	0.95	0.93	1.00

Table 3.1: **Oil price change summary statistics**
This table presents summary statistics for three oil spot price log growth rates, i.e., West Texas Intermediate (WTI), Dubai, and Arab Light, and their first principal component g^{P}. Results are presented for the 1983.01–2003.04 and the 1983.01–2015.12 sample periods. All reported numbers are in percentages.

	1983.01–2003.04				1983.01–2015.12			
	δ^{P}	NW[p]	H[p]	$\bar{R}^2$	δ^{P}	NW[p]	H[p]	$\bar{R}^2$
Local currency								
Canada	−0.06	0.06*	0.08*	1.0	−0.01	0.86	0.82	−0.2
France	−0.15	0.00***	0.01**	4.9	−0.08	0.08*	0.06*	1.4
Germany	−0.17	0.00***	0.01***	4.8	−0.08	0.11	0.10*	1.0
Italy	−0.32	0.00***	0.00***	16.4	−0.19	0.01***	0.00***	6.2
Japan	−0.11	0.07*	0.12	2.9	−0.05	0.29	0.34	0.3
UK	−0.11	0.00***	0.01***	4.1	−0.06	0.03**	0.03**	1.4
US dollar								
Canada	−0.06	0.09*	0.13	0.7	0.01	0.78	0.72	−0.2
France	−0.13	0.00***	0.02**	3.3	−0.06	0.31	0.22	0.4
Germany	−0.15	0.00***	0.02**	3.7	−0.06	0.24	0.24	0.4
Italy	−0.31	0.00***	0.00***	14.3	−0.17	0.01**	0.00***	4.1
Japan	−0.09	0.14	0.27	0.9	−0.04	0.34	0.42	0.1
UK	−0.10	0.00***	0.04**	2.3	−0.03	0.49	0.40	0.0
US	−0.12	0.00***	0.00***	5.2	−0.05	0.22	0.14	0.6
World	−0.12	0.00***	0.01***	5.6	−0.04	0.22	0.20	0.6

Table 3.2: **Oil price change as a predictor of MSCI index excess returns: statistical significance over the 1983.01–2003.04 and 1983.01–2015.12 sample periods**

This table presents in-sample statistical significance results for the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where r^{e}_{t+1} is the excess return on an MSCI index and the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. We report the two-sided p-values for the null hypotheses $H_0 : \delta^{\mathsf{P}} = 0$, denoted by NW[$p$] and H[$p$], based on [135] standard errors, with optimal bandwidth selected as in [139], and [136] standard errors, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	No break model	One break model		# of breaks
	BIC	BIC	Break date	selected by BIC
Local currency				
Canada	−6.281	**−6.283**	2008.08	1
France	−5.789	**−5.790**	2008.08	1
Germany	−5.588	**−5.592**	2008.07	1
Italy	−5.514	**−5.554**	2003.10	1
Japan	−5.760	**−5.763**	1990.09	1
UK	**−6.232**	−6.228	2008.08	0
US dollar				
Canada	−5.810	**−5.815**	2008.07	1
France	**−5.601**	−5.598	2008.07	0
Germany	−5.419	**−5.421**	2008.07	1
Italy	−5.301	**−5.335**	2003.10	1
Japan	−5.570	**−5.575**	1990.09	1
UK	−5.929	**−5.933**	2008.07	1
US	−6.290	**−6.318**	2008.08	1
World	−6.283	**−6.308**	2008.07	1

Table 3.3: **Bai-Perron structural break tests: oil price changes for the 1983.01–2015.12 sample period**
This table presents [137] structural break tests for the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where r^{e}_{t+1} is excess return on an MSCI index and the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. Bold numbers indicate the lowest value assumed by the Bayesian Information Criterion (BIC).

	Oil Price Change			IS Model OPC Decomposition						
	$\delta^{\mathtt{P}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local currency										
Canada	−0.00	0.91	−0.3	−0.18	0.60	0.27	0.01***	−0.03	0.39	1.5
France	−0.08	0.09*	1.5	−0.43	0.39	0.23	0.06*	−0.10	0.01**	2.9
Germany	−0.08	0.12	1.0	−0.50	0.39	0.31	0.02**	−0.11	0.01**	3.1
Italy	−0.18	0.01**	6.2	−0.75	0.07*	0.32	0.03**	−0.22	0.00***	9.5
Japan	−0.05	0.32	0.3	0.21	0.56	0.31	0.00***	−0.08	0.06*	1.9
UK	−0.06	0.03**	1.4	0.17	0.70	0.16	0.13	−0.09	0.00***	2.3
US dollar										
Canada	0.02	0.74	−0.2	−0.26	0.56	0.35	0.01***	−0.01	0.82	1.5
France	−0.06	0.30	0.4	−0.57	0.21	0.23	0.24	−0.07	0.08*	1.5
Germany	−0.06	0.26	0.3	−0.66	0.20	0.30	0.10*	−0.08	0.04**	1.9
Italy	−0.16	0.02**	4.1	−0.91	0.07*	0.32	0.12	−0.20	0.00***	6.6
Japan	−0.04	0.37	0.1	0.61	0.18	0.19	0.06*	−0.07	0.11	0.7
UK	−0.03	0.49	0.0	0.01	0.98	0.32	0.05**	−0.06	0.04**	2.1
US	−0.04	0.27	0.5	−0.36	0.36	0.32	0.00***	−0.07	0.01**	4.3
World	−0.04	0.26	0.5	−0.19	0.58	0.31	0.01**	−0.07	0.02**	3.6

Table 3.4: **Oil supply, global demand, and oil-specific demand shocks, based on the IS model, as predictors of MSCI index excess returns: statistical significance results, based on [135] standard errors, over the 1986.01–2015.12 sample period**
The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	Oil Price Change			ES Model OPC Decomposition						
	$\delta^{\mathtt{P}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local currency										
Canada	−0.00	0.91	−0.3	−0.05	0.30	0.30	0.02**	−0.00	0.98	1.5
France	−0.08	0.09*	1.5	−0.18	0.00***	0.31	0.08*	−0.00	1.00	3.3
Germany	−0.08	0.12	1.0	−0.18	0.01**	0.40	0.01**	−0.01	0.88	3.2
Italy	−0.18	0.01**	6.2	−0.25	0.00***	0.39	0.06*	−0.18	0.14	9.0
Japan	−0.05	0.32	0.3	−0.16	0.04**	0.43	0.00***	0.04	0.64	3.1
UK	−0.06	0.03**	1.4	−0.15	0.00***	0.23	0.07*	0.02	0.72	3.2
US dollar										
Canada	0.02	0.74	−0.2	−0.04	0.44	0.40	0.02**	0.03	0.72	1.5
France	−0.06	0.30	0.4	−0.11	0.11	0.28	0.24	−0.03	0.77	1.1
Germany	−0.06	0.26	0.3	−0.12	0.12	0.35	0.09*	−0.04	0.69	1.5
Italy	−0.16	0.02**	4.1	−0.19	0.01***	0.36	0.12	−0.21	0.11	6.0
Japan	−0.04	0.37	0.1	−0.12	0.16	0.26	0.01***	0.02	0.81	0.7
UK	−0.03	0.49	0.0	−0.11	0.03**	0.37	0.07*	0.01	0.85	2.3
US	−0.04	0.27	0.5	−0.10	0.01***	0.38	0.01***	−0.03	0.55	4.2
World	−0.04	0.26	0.5	−0.12	0.03**	0.38	0.01**	−0.01	0.88	4.1

Table 3.5: **Oil supply, global demand, and oil-specific demand shocks, based on the ES model, as predictors of MSCI index excess returns: statistical significance results, based on [135] standard errors, over the 1986.01–2015.12 sample period**
The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	CER							SR						
	DEC-IS	IID	p-value	P	p-value	GD	p-value	DEC-IS	IID	p-value	P	p-value	GD	p-value
Local currency														
Canada	6.68	4.40	0.09*	4.90	0.17	5.62	0.26	0.45	0.25	0.10	0.31	0.22	0.37	0.27
France	6.03	4.24	0.24	6.10	0.51	2.93	0.11	0.41	0.23	0.14	0.41	0.49	0.18	0.08*
Germany	7.44	3.58	0.08*	5.77	0.23	3.78	0.09*	0.51	0.17	0.03**	0.40	0.18	0.24	0.06*
Italy	10.70	4.32	0.03**	8.73	0.12	3.04	0.01***	0.61	−0.06	0.00***	0.50	0.13	0.00	0.00***
Japan	3.56	−0.20	0.05**	−1.67	0.07*	−1.37	0.01**	0.42	−0.11	0.00***	0.02	0.02**	−0.14	0.00***
UK	8.52	4.74	0.04**	7.32	0.25	3.94	0.03**	0.53	0.19	0.02**	0.45	0.28	0.13	0.02**
US dollar														
Canada	4.06	3.31	0.37	1.96	0.25	4.12	0.51	0.34	0.26	0.30	0.19	0.23	0.34	0.48
France	0.90	3.46	0.81	2.48	0.76	1.82	0.63	0.22	0.26	0.64	0.25	0.63	0.20	0.47
Germany	3.50	3.27	0.47	2.37	0.37	3.29	0.47	0.33	0.21	0.20	0.25	0.27	0.26	0.33
Italy	5.30	2.27	0.19	3.20	0.27	0.48	0.10*	0.44	0.05	0.01***	0.34	0.23	0.05	0.02**
Japan	4.11	1.46	0.04**	0.60	0.02**	−0.69	0.01***	0.30	−0.11	0.00***	−0.05	0.00***	−0.26	0.00***
UK	5.33	3.38	0.19	2.74	0.21	2.67	0.12	0.41	0.25	0.12	0.23	0.18	0.22	0.08*
US	9.28	6.52	0.11	6.08	0.08*	6.20	0.05**	0.65	0.48	0.12	0.45	0.09*	0.46	0.05**
World	7.90	4.03	0.04**	3.88	0.07*	3.30	0.03**	0.56	0.31	0.04**	0.30	0.06*	0.27	0.03**

Table 3.6: **Oil supply, global demand, and oil-specific demand shocks, based on the IS model, as predictors of MSCI index excess returns: economic significance over the 1986.01–2015.12 sample period**

This table presents evidence on the performance of the optimal trading strategies using oil supply, global demand, and oil-specific demand shocks as predictors in terms of two metrics: the certainty equivalent return (CER) and the Sharpe ratio (SR). The investor forms optimal mean-variance portfolios between an MSCI index and the corresponding Treasury Bill and has a risk aversion coefficient $\gamma = 3$. The weight on the MSCI index is constrained between $\underline{\omega} = 0$ and $\overline{\omega} = 150\%$. We compare the alternative model (DEC-IS) described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8) to three baseline models. The first baseline model (IID) assumes that the MSCI index excess return $r^{\mathtt{e}}_{t+1}$ is i.i.d. The second baseline model (P) is described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The third baseline model (GD) is described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{GD}} + \delta^{\mathtt{GD}} g^{\mathtt{GD}}_t + u^{\mathtt{GD}}_{t+1}$, where the global demand growth proxy $g^{\mathtt{GD}}_t$ is the first principal component obtained from the log growth rates of shipping cost index and global crude steel production. The portfolio strategies use the mean and variance estimates resulting from each model respectively. Reported are the CERs in annualized percentage points and the annualized SRs for the alternative and the three baseline models. We also report the one-sided p-values for the null hypothesis that the alternative model does not improve the CER and SR obtained from the optimal trading strategies based on each of the three baseline models. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. The results are based on an expanding window with 60 initial observations.

	CER							SR						
	DEC-ES	IID	p-value	P	p-value	GD	p-value		DEC-ES	IID	p-value	P	p-value	GD	p-value
Local currency														
Canada	7.05	4.40	0.06*	4.90	0.15	5.62	0.21	0.48	0.25	0.08*	0.31	0.18	0.37	0.22
France	6.00	4.24	0.25	6.10	0.52	2.93	0.11	0.41	0.23	0.13	0.41	0.49	0.18	0.08*
Germany	7.56	3.58	0.10*	5.77	0.22	3.78	0.08*	0.52	0.17	0.02**	0.40	0.21	0.24	0.05*
Italy	10.47	4.32	0.04**	8.73	0.12	3.04	0.01***	0.59	−0.06	0.00***	0.50	0.14	0.00	0.00***
Japan	1.92	−0.20	0.14	−1.67	0.16	−1.37	0.04**	0.30	−0.11	0.00***	0.02	0.09*	−0.14	0.00***
UK	9.14	4.74	0.00***	7.32	0.15	3.94	0.01***	0.59	0.19	0.00***	0.45	0.16	0.13	0.00***
US dollar														
Canada	4.30	3.31	0.34	1.96	0.24	4.12	0.47	0.36	0.26	0.28	0.19	0.20	0.34	0.44
France	1.61	3.46	0.74	2.48	0.63	1.82	0.53	0.25	0.26	0.55	0.25	0.52	0.20	0.38
Germany	4.55	3.27	0.32	2.37	0.27	3.29	0.32	0.38	0.21	0.13	0.25	0.20	0.26	0.23
Italy	5.18	2.27	0.20	3.20	0.24	0.48	0.11	0.44	0.05	0.01***	0.34	0.19	0.05	0.02**
Japan	1.85	1.46	0.39	0.60	0.26	−0.69	0.07*	0.09	−0.11	0.09*	−0.05	0.18	−0.26	0.02**
UK	5.84	3.38	0.14	2.74	0.19	2.67	0.07*	0.44	0.25	0.08*	0.23	0.14	0.22	0.06*
US	9.78	6.52	0.07*	6.08	0.05*	6.20	0.02**	0.68	0.48	0.07*	0.45	0.06*	0.46	0.02**
World	7.67	4.03	0.05*	3.88	0.09*	3.30	0.03**	0.55	0.31	0.05**	0.30	0.06*	0.27	0.02**

Table 3.7: **Oil supply, global demand, and oil-specific demand shocks, based on the ES model, as predictors of MSCI index excess returns: economic significance over the 1986.01–2015.12 sample period**

This table presents evidence on the performance of the optimal trading strategies using oil supply, global demand, and oil-specific demand shocks as predictors in terms of two metrics: the certainty equivalent return (CER) and the Sharpe ratio (SR). The investor forms optimal mean-variance portfolios between an MSCI index and the corresponding Treasury Bill and has a risk aversion coefficient $\gamma = 3$. The weight on the MSCI index is constrained between $\underline{\omega} = 0$ and $\overline{\omega} = 150\%$. We compare the alternative model (DEC-ES) described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8) to three baseline models. The first baseline model (IID) assumes that the MSCI index excess return $r^{\mathtt{e}}_{t+1}$ is i.i.d. The second baseline model (P) is described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The third baseline model (GD) is described by the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{GD}} + \delta^{\mathtt{GD}} g^{\mathtt{GD}}_t + u^{\mathtt{GD}}_{t+1}$, where the global demand growth proxy $g^{\mathtt{GD}}_t$ is the first principal component obtained from the log growth rates of shipping cost index and global crude steel production. The portfolio strategies use the mean and variance estimates resulting from each model respectively. Reported are the CERs in annualized percentage points and the annualized SRs for the alternative and the three baseline models. We also report the one-sided p-values for the null hypothesis that the alternative model does not improve the CER and SR obtained from the optimal trading strategies based on each of the three baseline models. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. The results are based on an expanding window with 60 initial observations.

	IS Model				ES Model			
	No break model	One break model		# of breaks	No break model	One break model		# of breaks
	BIC	BIC	Break date	selected by BIC	BIC	BIC	Break date	selected by BIC
Local currency								
Canada	**−6.298**	−6.257	1994.02	0	**−6.297**	−6.259	2008.08	0
France	**−5.794**	−5.774	1995.03	0	**−5.799**	−5.774	1998.09	0
Germany	**−5.598**	−5.573	1994.06	0	**−5.600**	−5.574	1999.03	0
Italy	**−5.537**	−5.534	1994.06	0	−5.532	**−5.544**	1995.10	1
Japan	**−5.737**	−5.703	1990.09	0	**−5.749**	−5.712	1990.09	0
UK	**−6.240**	−6.197	2008.08	0	**−6.250**	−6.209	2010.03	0
US dollar								
Canada	**−5.803**	−5.765	2008.07	0	**−5.803**	−5.766	2008.07	0
France	**−5.597**	−5.561	1999.02	0	**−5.593**	−5.566	2010.03	0
Germany	**−5.436**	−5.407	1999.03	0	**−5.432**	−5.402	1999.03	0
Italy	**−5.313**	−5.296	1990.09	0	−5.306	**−5.306**	1990.09	1
Japan	**−5.551**	−5.531	1990.09	0	**−5.551**	−5.522	1990.09	0
UK	**−5.976**	−5.944	2010.01	0	**−5.978**	−5.951	2010.03	0
US	**−6.299**	−6.270	2008.08	0	**−6.298**	−6.268	2008.08	0
World	**−6.276**	−6.249	2008.07	0	**−6.281**	−6.258	2010.03	0

Table 3.8: **Bai-Perron structural break tests: oil supply, global demand, and oil-specific demand shocks, based on the IS and ES models, as predictors of MSCI index returns for the 1986.01–2015.12 sample period**
This table presents [137] structural break tests for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, oil-specific demand shocks obtained in the oil price change decomposition (3.8), respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. Bold numbers indicate the lowest value assumed by the Bayesian Information Criterion (BIC).

	Three-month return							Six-month return						
	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local currency														
Canada	−0.79	0.21	0.67	0.00***	0.01	0.80	3.4	−1.17	0.24	0.45	0.11	0.03	0.70	0.5
France	−1.18	0.12	0.76	0.05**	−0.11	0.11	4.0	−1.58	0.25	0.72	0.07*	−0.19	0.06*	2.3
Germany	−1.53	0.09*	0.92	0.01***	−0.09	0.19	4.6	−1.29	0.40	0.72	0.08*	−0.17	0.08*	1.4
Italy	−2.22	0.01***	1.04	0.00***	−0.29	0.00***	10.9	−1.97	0.13	0.92	0.02**	−0.36	0.00***	5.2
Japan	−0.97	0.26	0.63	0.01***	−0.11	0.15	2.3	−1.98	0.14	0.34	0.26	−0.25	0.06*	2.0
UK	−0.19	0.77	0.48	0.06*	−0.04	0.30	1.4	−0.09	0.93	0.59	0.06*	−0.05	0.53	0.7
US dollar														
Canada	−0.80	0.31	0.89	0.01***	0.05	0.52	3.7	−1.29	0.28	0.56	0.22	0.06	0.52	0.4
France	−1.46	0.11	0.92	0.02**	−0.01	0.88	4.6	−2.10	0.14	0.72	0.10	−0.13	0.24	2.1
Germany	−1.84	0.05**	1.10	0.00***	0.00	0.99	5.7	−1.75	0.19	0.72	0.11	−0.13	0.15	1.3
Italy	−2.45	0.01**	1.21	0.00***	−0.20	0.04**	9.0	−2.77	0.03**	0.91	0.06*	−0.33	0.03**	4.9
Japan	−0.32	0.79	0.71	0.00***	−0.10	0.27	1.7	−1.71	0.31	0.66	0.03**	−0.26	0.19	2.1
UK	−0.25	0.75	0.95	0.00***	−0.01	0.89	5.1	−0.67	0.53	0.87	0.02**	−0.03	0.73	1.7
US	−0.38	0.57	0.79	0.01***	−0.02	0.62	4.9	−1.02	0.27	0.72	0.02**	−0.03	0.71	1.9
World	−0.66	0.33	0.89	0.00***	−0.05	0.33	6.5	−1.25	0.22	0.78	0.02**	−0.12	0.17	3.0

Table 3.9: **Oil supply, global demand, and oil-specific demand shocks, based on the IS model, as predictors of long-horizon MSCI index returns: statistical significance over the 1986.01–2015.12 sample period**
This table presents results for the predictive regression $r_{t,t+h} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{GD}} x_t^{\mathtt{GD}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + u_{t+h}^{\mathtt{DEC}}$, where $r_{t,t+h}$ is the h-month return on an MSCI index and $x_t^{\mathtt{S}}$, $x_t^{\mathtt{GD}}$, and $x_t^{\mathtt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. Results based on three-month and six-month returns are presented in the left and right panel, respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Three-month return							Six-month return						
	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local currency														
Canada	−0.08	0.39	0.72	0.05*	0.12	0.27	2.6	0.05	0.76	0.37	0.26	−0.04	0.79	−0.3
France	−0.31	0.01**	0.93	0.07*	0.13	0.41	4.2	−0.33	0.09*	0.82	0.14	−0.04	0.87	1.7
Germany	−0.31	0.04**	1.14	0.02**	0.17	0.25	4.6	−0.36	0.10*	0.89	0.12	0.08	0.72	1.4
Italy	−0.46	0.00***	1.26	0.01**	−0.12	0.60	9.7	−0.53	0.00***	1.14	0.03**	−0.16	0.54	4.9
Japan	−0.20	0.13	0.74	0.03**	−0.01	0.97	2.0	−0.31	0.12	0.46	0.26	−0.21	0.31	1.4
UK	−0.13	0.21	0.49	0.17	0.09	0.46	1.2	−0.01	0.94	0.46	0.25	−0.06	0.70	0.1
US dollar														
Canada	−0.08	0.49	0.97	0.03**	0.19	0.21	3.1	0.11	0.56	0.44	0.34	−0.06	0.76	−0.2
France	−0.19	0.12	1.09	0.03**	0.18	0.27	3.9	−0.23	0.23	0.80	0.13	−0.06	0.78	1.0
Germany	−0.20	0.13	1.31	0.01***	0.21	0.16	5.0	−0.28	0.15	0.88	0.11	0.04	0.84	0.9
Italy	−0.35	0.01***	1.43	0.00***	−0.08	0.73	7.4	−0.48	0.01***	1.14	0.04**	−0.20	0.51	4.0
Japan	−0.09	0.45	0.71	0.01***	−0.09	0.62	1.4	−0.15	0.59	0.60	0.13	−0.44	0.10*	1.8
UK	−0.06	0.57	0.94	0.02**	0.08	0.59	3.9	0.04	0.87	0.70	0.10*	−0.11	0.55	0.8
US	−0.15	0.09*	0.89	0.04**	0.15	0.16	5.1	−0.03	0.85	0.69	0.10*	−0.04	0.78	1.0
World	−0.17	0.08*	1.00	0.01***	0.11	0.38	6.2	−0.12	0.45	0.77	0.05**	−0.14	0.42	1.9

Table 3.10: **Oil supply, global demand, and oil-specific demand shocks, based on the ES model, as predictors of long-horizon MSCI index returns: statistical significance over the 1986.01–2015.12 sample period**
This table presents results for the predictive regression $r_{t,t+h} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{GD}} x_t^{\mathtt{GD}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + u_{t+h}^{\mathtt{DEC}}$, where $r_{t,t+h}$ is the h-month return on an MSCI index and $x_t^{\mathtt{S}}$, $x_t^{\mathtt{GD}}$, and $x_t^{\mathtt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. Results based on three-month and six-month returns are presented in the left and right panel, respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

IS Model OPC Decomposition

	One-month	Two-month	Three-month	Four-month	Five-month	Six-month
Local currency						
Canada	1.6	3.4	3.4	3.5	1.7	0.5
France	2.9	3.9	4.0	4.8	3.4	2.3
Germany	3.0	5.1	4.6	4.2	2.6	1.4
Italy	9.6	11.1	10.9	9.9	6.7	5.2
Japan	1.9	2.7	2.3	2.5	2.1	2.0
UK	2.2	2.6	1.4	3.0	2.0	0.7
US dollar						
Canada	1.5	3.8	3.7	3.5	1.6	0.4
France	1.4	3.1	4.6	4.6	2.9	2.1
Germany	1.9	4.7	5.7	4.4	2.4	1.3
Italy	6.6	8.6	9.0	8.0	5.7	4.9
Japan	0.6	1.8	1.7	2.4	2.3	2.1
UK	2.1	5.3	5.1	5.5	3.6	1.7
US	4.3	5.5	4.9	5.6	3.6	1.9
World	3.7	6.4	6.5	6.9	4.7	3.0

ES Model OPC Decomposition

	One-month	Two-month	Three-month	Four-month	Five-month	Six-month
Local currency						
Canada	1.6	3.1	2.6	2.6	0.9	−0.3
France	3.3	3.9	4.2	5.2	3.9	1.7
Germany	3.1	4.8	4.6	5.6	4.2	1.4
Italy	9.1	10.2	9.7	9.3	7.4	4.9
Japan	3.1	2.7	2.0	2.4	2.2	1.4
UK	3.2	2.7	1.2	2.1	1.5	0.1
US dollar						
Canada	1.5	3.5	3.1	2.7	0.9	−0.2
France	1.1	2.9	3.9	4.6	3.1	1.0
Germany	1.5	4.4	5.0	5.4	3.8	0.9
Italy	5.9	7.8	7.4	7.2	6.0	4.0
Japan	0.7	1.5	1.4	2.2	2.2	1.8
UK	2.3	4.8	3.9	4.2	2.7	0.8
US	4.2	6.0	5.1	4.6	2.9	1.0
World	4.1	6.4	6.2	6.3	4.4	1.9

Table 3.11: **Oil supply, global demand, and oil-specific demand shocks as predictors of long-horizon MSCI index returns: adjusted R-squares over the 1986.01–2015.12 sample period**
This table presents adjusted R^2 for the predictive regression $r_{t,t+h} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{GD}} x_t^{\mathtt{GD}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + u_{t+h}^{\mathtt{DEC}}$, where $r_{t,t+h}$ is the h-month return on an MSCI index and $x_t^{\mathtt{S}}$, $x_t^{\mathtt{GD}}$, and $x_t^{\mathtt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). The adjusted R^2, stated as a percentage, based on one-month to six-month returns are presented. The top and bottom panels contain results, for local currency- and US dollar-denominated index returns, based on the oil price change decompositions resulting from the IS and ES models, respectively.

				IS Model	ES Model
Sample Period	1983.01–2003.04	1983.01–2015.12	1986.01–2015.12	1986.01–2015.12	1986.01–2015.12
$g_t^{\mathtt{P}}$	-0.11	-0.04	-0.04		
	[0.00***]	[0.28]	[0.25]		
$x_t^{\mathtt{S}}$				-0.34	-0.09
				[0.38]	[0.03**]
$x_t^{\mathtt{GD}}$				0.34	0.38
				[0.00***]	[0.00***]
$x_t^{\mathtt{OSD}}$				-0.06	-0.03
				[0.03**]	[0.55]
dy_t	0.03	0.03	0.03	0.03	0.03
	[0.04**]	[0.00***]	[0.00***]	[0.00***]	[0.00***]
tms_t	-0.78	-0.43	-0.49	-0.57	-0.56
	[0.14]	[0.08*]	[0.08*]	[0.04**]	[0.05**]
dfy_t	-0.32	-1.16	-1.36	-0.99	-0.97
	[0.74]	[0.25]	[0.15]	[0.17]	[0.16]
tbl_t	-0.45	-0.29	-0.33	-0.35	-0.34
	[0.14]	[0.02**]	[0.05**]	[0.03**]	[0.03**]
$\bar{R}^2$	6.3	2.0	2.0	5.7	5.6

Table 3.12: **Robustness checks of the US MSCI index excess return predictive regressions: the role of macroeconomic variables**

This table presents results for the predictive regressions $r_{t+1}^{\mathtt{e}} = \gamma^{\mathtt{P}} + \delta^{\mathtt{P}} g_t^{\mathtt{P}} + \boldsymbol{\theta}'\mathbf{z}_t + v_{t+1}^{\mathtt{P}}$ and $r_{t+1}^{\mathtt{e}} = \gamma^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{GD}} x_t^{\mathtt{GD}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + \boldsymbol{\lambda}'\mathbf{z}_t + v_{t+1}^{\mathtt{DEC}}$ over various sample periods. The variables in these regressions are as follows: (i) $r_{t+1}^{\mathtt{e}}$ is excess return on the US MSCI index, (ii) the oil price change proxy $g_t^{\mathtt{P}}$ is the first principal component obtained from three oil spot price log growth rates, West Texas Intermediate, Dubai, and Arab Light, (iii) $x_t^{\mathtt{S}}$, $x_t^{\mathtt{GD}}$, and $x_t^{\mathtt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8) based on the IS and ES models, and (iv) $\mathbf{z}_t$ is the vector of macroeconomic variables including the log dividend yield (dy), the term spread (tms), the default yield spread (dfy), and the one-month Treasury Bill rate (tbl). The numbers in square brackets represent two-sided p-values for the null hypotheses that the slope coefficients are zero, based on [135] standard errors with optimal bandwidth selected as in [139]. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	1983.01–2003.04						1983.01–2015.12					
	CER			SR			CER			SR		
	P	IID	p-value	P	IID	p-value	P	IID	p-value	P	IID	p-value
Local currency												
Canada	9.65	5.56	0.04**	0.49	0.07	0.02**	5.59	5.10	0.39	0.29	0.23	0.37
France	12.26	6.19	0.05**	0.61	0.28	0.04**	7.08	5.24	0.23	0.44	0.31	0.22
Germany	13.11	4.10	0.01**	0.72	0.22	0.01***	7.21	4.59	0.18	0.48	0.30	0.15
Italy	19.11	6.64	0.00***	0.82	−0.06	0.00***	9.30	4.50	0.06*	0.49	−0.01	0.00***
Japan	−1.22	−6.66	0.06*	0.13	−0.23	0.00***	−2.24	−2.83	0.43	0.06	−0.08	0.19
UK	11.91	5.81	0.00***	0.54	0.08	0.00***	7.93	5.21	0.06*	0.42	0.18	0.04**
US dollar												
Canada	7.28	4.20	0.15	0.38	0.12	0.11	3.25	3.85	0.60	0.23	0.25	0.55
France	10.67	5.47	0.06*	0.59	0.36	0.06*	4.78	4.32	0.43	0.38	0.35	0.41
Germany	11.70	3.80	0.02**	0.64	0.23	0.01**	4.73	4.09	0.42	0.38	0.30	0.30
Italy	13.51	2.53	0.00***	0.72	0.07	0.00***	3.97	1.86	0.29	0.37	0.09	0.04**
Japan	−3.83	−6.30	0.11	−0.10	−0.25	0.14	−2.19	−2.27	0.48	−0.07	−0.13	0.25
UK	8.17	4.12	0.02**	0.45	0.21	0.02**	3.91	3.69	0.46	0.29	0.27	0.40
US	11.61	7.81	0.11	0.66	0.45	0.08*	7.74	7.06	0.37	0.52	0.48	0.36
World	7.79	2.26	0.06*	0.43	0.19	0.06*	4.70	3.54	0.30	0.34	0.29	0.36

Table 3.13: **Oil price change as a predictor of MSCI index excess returns: economic significance over the 1983.01–2003.04 and 1983.01–2015.12 sample periods.**
This table presents evidence on the performance of the optimal trading strategies using oil price change as a predictor in terms of two metrics: the certainty equivalent return (CER) and the Sharpe ratio (SR). The investor forms optimal mean-variance portfolios between an MSCI index and the corresponding Treasury Bill and has a risk aversion coefficient $\gamma = 3$. The weight on the MSCI index is constrained between $\underline{\omega} = 0$ and $\overline{\omega} = 150\%$. We compare the baseline model (IID), according to which the MSCI index excess return r^{e}_{t+1} is i.i.d., to the alternative model (P) described by the predictive regression $r^{\mathsf{e}}_{t+1} = \alpha^{\mathsf{P}} + \delta^{\mathsf{P}} g^{\mathsf{P}}_t + u^{\mathsf{P}}_{t+1}$, where the oil price change proxy g^{P}_t is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The portfolio strategies use the mean and variance estimates resulting from each model, respectively. Reported are the CERs in annualized percentage points and the annualized SRs for both baseline and alternative models. We also report the one-sided p-values for the null hypothesis that the alternative model does not improve the CER and SR obtained from the optimal trading strategy based on the baseline model. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. The results are based on an expanding window with 60 initial observations.

	Oil Price Change			IS Model OPC Decomposition						
	$\delta^{\mathtt{P}}$	H[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	H[p]	$\beta^{\mathtt{GD}}$	H[p]	$\beta^{\mathtt{OSD}}$	H[p]	$\bar{R}^2$
Local currency										
Canada	−0.00	0.87	−0.3	−0.18	0.64	0.27	0.02**	−0.03	0.38	1.5
France	−0.08	0.07*	1.5	−0.43	0.40	0.23	0.10*	−0.10	0.03**	2.9
Germany	−0.08	0.11	1.0	−0.50	0.41	0.31	0.07*	−0.11	0.02**	3.1
Italy	−0.18	0.00***	6.2	−0.75	0.15	0.32	0.03**	−0.22	0.00***	9.5
Japan	−0.05	0.36	0.3	0.21	0.67	0.31	0.03**	−0.08	0.13	1.9
UK	−0.06	0.04**	1.4	0.17	0.70	0.16	0.14	−0.09	0.01***	2.3
US dollar										
Canada	0.02	0.67	−0.2	−0.26	0.56	0.35	0.03**	−0.01	0.82	1.5
France	−0.06	0.22	0.4	−0.57	0.30	0.23	0.26	−0.07	0.11	1.5
Germany	−0.06	0.26	0.3	−0.66	0.30	0.30	0.23	−0.08	0.09*	1.9
Italy	−0.16	0.01***	4.1	−0.91	0.11	0.32	0.11	−0.20	0.00***	6.6
Japan	−0.04	0.44	0.1	0.61	0.28	0.19	0.17	−0.07	0.21	0.7
UK	−0.03	0.40	0.0	0.01	0.98	0.32	0.03**	−0.06	0.10*	2.1
US	−0.04	0.17	0.5	−0.36	0.39	0.32	0.03**	−0.07	0.03**	4.3
World	−0.04	0.23	0.5	−0.19	0.63	0.31	0.03**	−0.07	0.04**	3.6

Table 3.14: **Oil supply, global demand, and oil-specific demand shocks, based on the IS model, as predictors of MSCI index excess returns: statistical significance results, based on [136] standard errors, over the 1986.01–2015.12 sample period**
The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by H[p], based on [136] standard errors. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	Oil Price Change			ES Model OPC Decomposition						
	$\delta^{\mathtt{P}}$	H[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	H[p]	$\beta^{\mathtt{GD}}$	H[p]	$\beta^{\mathtt{OSD}}$	H[p]	$\bar{R}^2$
Local currency										
Canada	−0.00	0.87	−0.3	−0.05	0.30	0.30	0.02**	−0.00	0.98	1.5
France	−0.08	0.07*	1.5	−0.18	0.01**	0.31	0.04**	−0.00	1.00	3.3
Germany	−0.08	0.11	1.0	−0.18	0.04**	0.40	0.03**	−0.01	0.88	3.2
Italy	−0.18	0.00***	6.2	−0.25	0.00***	0.39	0.02**	−0.18	0.08*	9.0
Japan	−0.05	0.36	0.3	−0.16	0.06*	0.43	0.01***	0.04	0.62	3.1
UK	−0.06	0.04**	1.4	−0.15	0.00***	0.23	0.06*	0.02	0.71	3.2
US dollar										
Canada	0.02	0.67	−0.2	−0.04	0.45	0.40	0.03**	0.03	0.71	1.5
France	−0.06	0.22	0.4	−0.11	0.13	0.28	0.21	−0.03	0.76	1.1
Germany	−0.06	0.26	0.3	−0.12	0.19	0.35	0.18	−0.04	0.66	1.5
Italy	−0.16	0.01***	4.1	−0.19	0.03**	0.36	0.11	−0.21	0.08*	6.0
Japan	−0.04	0.44	0.1	−0.12	0.18	0.26	0.10*	0.02	0.82	0.7
UK	−0.03	0.40	0.0	−0.11	0.07*	0.37	0.03**	0.01	0.85	2.3
US	−0.04	0.17	0.5	−0.10	0.03**	0.38	0.02**	−0.03	0.56	4.2
World	−0.04	0.23	0.5	−0.12	0.04**	0.38	0.02**	−0.01	0.89	4.1

Table 3.15: **Oil supply, global demand, and oil-specific demand shocks, based on the ES model, as predictors of MSCI index excess returns: statistical significance results, based on [136] standard errors, over the 1986.01–2015.12 sample period.** The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an MSCI index and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by H[p], based on [136] standard errors. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	IS Model OPC Decomposition							ES Model OPC Decomposition						
	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local currency														
Canada	−0.20	0.60	0.29	0.04**	0.08	0.01***	4.5	0.06	0.25	0.32	0.04**	0.09	0.07*	4.3
France	−0.48	0.31	0.18	0.36	−0.04	0.38	0.6	−0.11	0.13	0.29	0.14	0.03	0.69	1.0
Germany	−0.05	0.93	0.23	0.22	−0.03	0.59	−0.0	−0.06	0.40	0.29	0.11	0.01	0.86	0.2
Italy	−0.55	0.21	0.14	0.51	−0.02	0.74	−0.1	−0.14	0.16	0.28	0.23	0.11	0.11	0.8
Japan	−0.26	0.59	0.38	0.02**	0.03	0.46	2.0	0.03	0.71	0.44	0.01***	0.02	0.81	2.3
UK	−0.60	0.14	0.11	0.42	−0.00	0.93	0.5	−0.01	0.81	0.13	0.32	−0.02	0.77	−0.4
US dollar														
Canada	−0.36	0.39	0.45	0.00***	0.14	0.00***	8.0	0.12	0.10*	0.48	0.01**	0.13	0.04**	7.6
France	−0.79	0.16	0.24	0.30	−0.00	0.94	0.9	−0.06	0.51	0.35	0.18	0.02	0.80	0.6
Germany	−0.40	0.56	0.29	0.16	0.01	0.86	0.4	−0.01	0.85	0.35	0.10*	0.02	0.86	0.4
Italy	−0.85	0.08*	0.21	0.37	0.00	0.98	0.4	−0.11	0.37	0.37	0.16	0.11	0.22	0.6
Japan	−0.14	0.83	0.20	0.28	0.02	0.71	−0.2	0.01	0.88	0.26	0.17	0.02	0.82	0.1
UK	−0.65	0.18	0.34	0.06*	0.04	0.44	3.0	0.04	0.54	0.37	0.06*	0.00	0.97	2.1
US	−0.40	0.33	0.27	0.11	0.01	0.72	1.8	−0.03	0.63	0.33	0.08*	0.05	0.41	1.7
World	−0.45	0.22	0.30	0.08*	0.02	0.59	2.5	−0.01	0.83	0.37	0.02**	0.04	0.54	2.3

Table 3.16: **Contemporaneous regression of MSCI index excess returns on oil supply, global demand, and oil-specific demand shocks, based on the IS and ES models: statistical significance results, based on [135] standard errors, over the 1986.01–2015.12 sample period**

The table presents in-sample results for the contemporaneous regression $r_t^{\mathtt{e}} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x_t^{\mathtt{S}} + \beta^{\mathtt{GD}} x_t^{\mathtt{GD}} + \beta^{\mathtt{OSD}} x_t^{\mathtt{OSD}} + u_t^{\mathtt{DEC}}$, where $x_t^{\mathtt{S}}$, $x_t^{\mathtt{GD}}$, and $x_t^{\mathtt{OSD}}$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	One-week delay							Two-week delay						
	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local currency														
Canada	−0.21	0.55	0.29	0.00***	−0.04	0.22	1.7	−0.14	0.68	0.30	0.00***	−0.03	0.35	1.7
France	0.01	0.98	0.27	0.06*	−0.12	0.00***	2.7	−0.16	0.69	0.26	0.01**	−0.09	0.02**	2.2
Germany	−0.18	0.76	0.35	0.05*	−0.13	0.00***	3.2	−0.50	0.32	0.37	0.00***	−0.11	0.00***	3.7
Italy	−0.36	0.44	0.47	0.00***	−0.21	0.00***	7.7	−0.30	0.56	0.50	0.00***	−0.19	0.00***	7.0
Japan	−0.09	0.81	0.37	0.00***	−0.08	0.02**	2.5	−0.13	0.76	0.40	0.00***	−0.08	0.02**	2.6
UK	0.39	0.50	0.24	0.02**	−0.09	0.00***	2.4	0.27	0.56	0.24	0.02**	−0.07	0.02**	1.8
US dollar														
Canada	−0.29	0.52	0.40	0.01***	−0.02	0.68	1.8	−0.28	0.51	0.41	0.00***	−0.01	0.83	2.3
France	0.03	0.95	0.32	0.02**	−0.08	0.01***	1.3	0.05	0.91	0.35	0.01***	−0.06	0.09*	1.6
Germany	−0.16	0.75	0.40	0.01***	−0.10	0.00***	2.2	−0.28	0.55	0.46	0.00***	−0.09	0.01***	3.3
Italy	−0.35	0.47	0.51	0.00***	−0.19	0.00***	5.5	−0.02	0.97	0.59	0.00***	−0.18	0.00***	5.6
Japan	0.19	0.71	0.27	0.02**	−0.08	0.02**	1.1	0.20	0.70	0.37	0.00***	−0.09	0.01***	2.0
UK	0.33	0.54	0.39	0.00***	−0.07	0.04**	2.4	0.40	0.36	0.41	0.00***	−0.06	0.04**	3.2
US	−0.04	0.92	0.34	0.00***	−0.07	0.01***	3.2	0.09	0.79	0.34	0.00***	−0.03	0.23	2.6
World	−0.06	0.86	0.37	0.00***	−0.07	0.00***	3.8	0.03	0.91	0.38	0.00***	−0.06	0.02**	4.0

Table 3.17: **Oil supply, global demand, and oil-specific demand shocks, based on the IS model, as predictors of delayed monthly MSCI index returns: statistical significance over the 1986.01–2015.12 sample period**
This table presents results for the predictive regression $r^{\mathtt{d}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}}x^{\mathtt{S}}_t + \beta^{\mathtt{GD}}x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}}x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{d}}_{t+1}$ is the delayed monthly return on an MSCI index and $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. Results based on a one-week and two-week delay are presented in the left and right panel, respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	One-week delay							Two-week delay						
	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Local currency														
Canada	−0.06	0.14	0.33	0.00***	−0.01	0.85	1.6	−0.09	0.04**	0.37	0.00***	0.06	0.37	2.4
France	−0.21	0.00***	0.36	0.04**	0.01	0.85	3.5	−0.20	0.00***	0.37	0.01***	0.07	0.35	3.7
Germany	−0.21	0.00***	0.44	0.00***	−0.01	0.92	3.7	−0.20	0.00***	0.48	0.00***	0.01	0.92	4.4
Italy	−0.25	0.00***	0.55	0.02**	−0.15	0.11	7.8	−0.27	0.00***	0.65	0.00***	−0.07	0.44	8.0
Japan	−0.11	0.05**	0.42	0.00***	−0.04	0.54	2.6	−0.13	0.02**	0.48	0.00***	−0.00	0.96	2.9
UK	−0.13	0.00***	0.30	0.00***	−0.01	0.86	2.5	−0.10	0.06*	0.28	0.00***	0.00	1.00	1.8
US dollar														
Canada	−0.03	0.59	0.44	0.00***	−0.01	0.90	1.7	−0.08	0.13	0.50	0.00***	0.09	0.27	2.9
France	−0.14	0.05**	0.38	0.02**	0.00	0.97	1.6	−0.18	0.01***	0.48	0.00***	0.10	0.21	3.3
Germany	−0.14	0.07*	0.46	0.01***	−0.03	0.75	2.3	−0.18	0.01**	0.60	0.00***	0.04	0.67	4.4
Italy	−0.19	0.00***	0.56	0.01***	−0.18	0.11	5.5	−0.27	0.00***	0.77	0.00***	−0.04	0.67	6.9
Japan	−0.06	0.31	0.27	0.01**	−0.08	0.19	1.0	−0.09	0.15	0.40	0.00***	−0.07	0.34	2.0
UK	−0.08	0.20	0.43	0.00***	−0.03	0.69	2.4	−0.10	0.17	0.46	0.00***	0.01	0.92	3.1
US	−0.14	0.01***	0.42	0.00***	0.02	0.70	3.9	−0.13	0.02**	0.44	0.00***	0.11	0.09*	4.8
World	−0.12	0.01**	0.43	0.00***	−0.01	0.86	4.1	−0.14	0.01***	0.49	0.00***	0.07	0.32	5.8

Table 3.18: **Oil supply, global demand, and oil-specific demand shocks, based on the ES model, as predictors of delayed monthly MSCI index returns: statistical significance over the 1986.01–2015.12 sample period**

This table presents results for the predictive regression $r^{\mathtt{d}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{d}}_{t+1}$ is the delayed monthly return on an MSCI index and $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8). Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. Results based on a one-week and two-week delay are presented in the left and right panel, respectively. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
$\eta_{\mathtt{S}} = 0.10$							
US	−0.11	0.09*	0.35	0.01***	−0.06	0.13	4.1
World	−0.11	0.16	0.34	0.01**	−0.05	0.21	3.7
$\eta_{\mathtt{S}} = 0.15$							
US	−0.09	0.04**	0.36	0.01**	−0.06	0.20	4.0
World	−0.10	0.11	0.36	0.02**	−0.05	0.33	3.7
$\eta_{\mathtt{S}} = 0.20$							
US	−0.08	0.02**	0.38	0.01**	−0.06	0.28	4.0
World	−0.09	0.08*	0.37	0.02**	−0.05	0.47	3.7

Table 3.19: **Oil supply, global demand, and oil-specific demand shocks, based on the ES model with a set of fixed oil supply elasticities, as predictors of US and World MSCI index excess returns: statistical significance results, based on [135] standard errors, over the 1986.01–2015.12 sample period**

This table presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8) based on the ES model with fixed oil supply elasticities equal to 0.10. 0.15, and 0.20. Shown are the estimates of the predictive slope coefficients β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$, as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	q_5	q_{25}	q_{50}	q_{75}	q_{95}
US					
p-value of $\beta^{\texttt{S}}$	0.00	0.01	0.02	0.04	0.08
p-value of $\beta^{\texttt{GD}}$	0.00	0.00	0.01	0.01	0.02
p-value of $\beta^{\texttt{OSD}}$	0.17	0.24	0.31	0.40	0.55
Adjusted R^2	3.1	3.6	3.9	4.3	4.9
World					
p-value of $\beta^{\texttt{S}}$	0.01	0.02	0.04	0.05	0.08
p-value of $\beta^{\texttt{GD}}$	0.00	0.01	0.01	0.02	0.04
p-value of $\beta^{\texttt{OSD}}$	0.34	0.45	0.54	0.64	0.79
Adjusted R^2	2.8	3.4	3.8	4.2	4.8

Table 3.20: **Perturbed oil supply, global demand, and oil-specific demand shocks, based on the ES model as predictors of US and World MSCI index excess returns: 5, 25, 50, 75, and 95 quantiles of the p-values of beta coefficients and the adjusted R^2, over the 1986.01–2015.12 sample period**

To address the concern that global demand shocks might be overestimated, as pointed out by [132], we perform the following simulation exercise. For each time period, the estimated global demand shocks are reduced by a factor of up to 50%, drawn from a uniform distribution, and the oil supply and oil-specific demand shocks are adjusted proportionately so that the sum of the three shocks is still equal to the observed oil price change. We run the predictive regression $r^{\texttt{e}}_{t+1} = \alpha^{\texttt{DEC}} + \beta^{\texttt{S}} x^{*\texttt{S}}_t + \beta^{\texttt{GD}} x^{*\texttt{GD}}_t + \beta^{\texttt{OSD}} x^{*\texttt{OSD}}_t + u^{\texttt{DEC}}_{t+1}$, where $x^{*\texttt{S}}_t$, $x^{*\texttt{GD}}_t$, and $x^{*\texttt{OSD}}_t$ are the perturbed oil supply, global demand, and oil-specific demand shocks, described above, and repeat the exercise 10,000 times. This table presents the quantiles of the resulting distributions of the p-values associated with the three predictive slopes and the adjusted R^2. Reported are two-sided p-values for the null hypotheses that the slope coefficients are zero, based on [135] standard errors with optimal bandwidth selected as in [139]. The adjusted R^2 is presented as a percentage.

	Oil Price Change		IS Model OPC Decomposition	ES Model OPC Decomposition
	# of breaks selected by BIC	Break date	# of breaks selected by BIC	# of breaks selected by BIC
Food	1	1991.12	0	0
Mines	0	–	0	0
Oil	0	–	0	0
Clths	1	2008.08	0	0
Durbl	1	2008.06	0	0
Chems	1	2008.07	0	0
Cnsum	1	2000.05	0	0
Cnstr	1	2008.08	0	0
Steel	1	2008.07	0	0
FabPr	1	2008.08	0	0
Machn	1	2008.08	0	0
Cars	1	2008.08	0	0
Trans	1	2008.06	0	0
Utils	0	–	0	0
Rtail	1	2008.08	0	0
Finan	1	2008.06	0	0
Other	1	2008.08	0	0

Table 3.21: **Bai-Perron structural break tests for excess returns on Fama-French 17 value-weighted industry portfolios: oil price change as the predictor over the 1983.01–2015.12 sample period, and oil supply, global demand,and oil-specific demand shocks, based on the IS and the ES models, as predictors over the 1986.01–2015.12 sample period**
The first column presents [137] structural break test results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The second and third columns present [137] structural break test results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, oil-specific demand shocks obtained in the oil price change decomposition (3.8), based on the IS and ES models, respectively. The number of breaks is selected by the Bayesian Information Criterion (BIC).

	Oil Price Change			IS Model OPC Decomposition						
	$\delta^{\mathtt{P}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Food	−0.028	0.45	0.1	−0.36	0.39	0.15	0.11	−0.04	0.23	0.9
Mines	0.018	0.74	−0.2	−0.69	0.28	0.34	0.06*	0.01	0.92	0.6
Oil	0.004	0.90	−0.3	−0.04	0.93	0.12	0.30	−0.00	0.91	−0.6
Clths	−0.077	0.22	1.0	−1.13	0.05*	0.52	0.00***	−0.12	0.01**	7.5
Durbl	−0.025	0.66	−0.1	−0.39	0.44	0.53	0.00***	−0.07	0.11	4.9
Chems	0.018	0.76	−0.2	−0.59	0.22	0.56	0.00***	−0.02	0.72	4.7
Cnsum	−0.042	0.29	0.5	−0.37	0.37	0.13	0.15	−0.05	0.16	1.0
Cnstr	−0.085	0.09*	1.5	−0.78	0.15	0.34	0.02**	−0.11	0.00***	4.7
Steel	−0.001	0.98	−0.3	−0.56	0.38	0.59	0.01**	−0.04	0.51	2.5
FabPr	−0.020	0.65	−0.2	−0.39	0.41	0.30	0.01***	−0.04	0.34	1.5
Machn	−0.064	0.32	0.4	−0.24	0.70	0.53	0.00***	−0.12	0.02**	3.9
Cars	−0.058	0.32	0.4	−0.85	0.09*	0.43	0.01**	−0.09	0.04**	3.6
Trans	−0.028	0.50	−0.0	−0.31	0.52	0.43	0.00***	−0.07	0.08*	4.1
Utils	0.005	0.87	−0.3	−0.55	0.08*	0.07	0.33	0.01	0.76	0.4
Rtail	−0.094	0.01***	2.6	−0.40	0.40	0.26	0.03**	−0.12	0.00***	5.1
Finan	−0.041	0.51	0.2	−0.80	0.09*	0.54	0.00***	−0.08	0.04**	7.0
Other	−0.055	0.22	0.7	−0.24	0.59	0.37	0.01***	−0.09	0.01***	4.2

Table 3.22: **Oil supply, global demand, and oil-specific demand shocks, based on the IS model, as predictors of excess returns on Fama-French 17 value-weighted industry portfolios: statistical significance over the 1986.01–2015.12 sample period**

The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8), based on the IS model. Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	Oil Price Change			ES Model OPC Decomposition						
	$\delta^{\mathtt{P}}$	NW[p]	$\bar{R}^2$	$\beta^{\mathtt{S}}$	NW[p]	$\beta^{\mathtt{GD}}$	NW[p]	$\beta^{\mathtt{OSD}}$	NW[p]	$\bar{R}^2$
Food	−0.028	0.45	0.1	−0.11	0.00***	0.22	0.08*	0.04	0.41	1.5
Mines	0.018	0.74	−0.2	−0.03	0.75	0.39	0.13	0.02	0.90	0.2
Oil	0.004	0.90	−0.3	−0.06	0.25	0.18	0.24	0.07	0.30	−0.2
Clths	−0.077	0.22	1.0	−0.14	0.03**	0.58	0.00***	−0.11	0.13	5.9
Durbl	−0.025	0.66	−0.1	−0.07	0.21	0.54	0.01***	−0.06	0.45	4.1
Chems	0.018	0.76	−0.2	−0.07	0.31	0.63	0.00***	0.03	0.76	4.2
Cnsum	−0.042	0.29	0.5	−0.13	0.00***	0.22	0.08*	0.04	0.52	1.8
Cnstr	−0.085	0.09*	1.5	−0.15	0.01***	0.41	0.02**	−0.08	0.31	4.1
Steel	−0.001	0.98	−0.3	−0.05	0.54	0.64	0.01***	−0.05	0.68	2.3
FabPr	−0.020	0.65	−0.2	−0.09	0.17	0.37	0.01***	0.01	0.94	1.5
Machn	−0.064	0.32	0.4	−0.10	0.14	0.56	0.00***	−0.14	0.16	4.0
Cars	−0.058	0.32	0.4	−0.07	0.29	0.43	0.03**	−0.13	0.18	2.6
Trans	−0.028	0.50	−0.0	−0.08	0.13	0.47	0.00***	−0.05	0.41	4.0
Utils	0.005	0.87	−0.3	−0.07	0.10*	0.16	0.18	0.09	0.12	0.8
Rtail	−0.094	0.01***	2.6	−0.15	0.00***	0.31	0.03**	−0.09	0.16	4.9
Finan	−0.041	0.51	0.2	−0.14	0.01**	0.64	0.00***	−0.02	0.74	6.4
Other	−0.055	0.22	0.7	−0.11	0.01***	0.43	0.00***	−0.06	0.41	4.3

Table 3.23: **Oil supply, global demand, and oil-specific demand shocks, based on the ES model, as predictors of excess returns on Fama-French 17 value-weighted industry portfolios: statistical significance over the 1986.01–2015.12 sample period**

The left panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{P}} + \delta^{\mathtt{P}} g^{\mathtt{P}}_t + u^{\mathtt{P}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and the oil price change proxy $g^{\mathtt{P}}_t$ is the first principal component obtained from three oil spot price log growth rates: West Texas Intermediate, Dubai, and Arab Light. The right panel presents in-sample results for the predictive regression $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, where $r^{\mathtt{e}}_{t+1}$ is excess return on an industry portfolio and $x^{\mathtt{S}}_t$, $x^{\mathtt{GD}}_t$, and $x^{\mathtt{OSD}}_t$ are the oil supply, global demand, and oil-specific demand shocks obtained in the oil price change decomposition (3.8), based on the ES model. Shown are the estimates of the predictive slope coefficients $\delta^{\mathtt{P}}$ (left panel) and β^i, $i = \mathtt{S}, \mathtt{GD}, \mathtt{OSD}$ (right panel), as well as two-sided p-values for the null hypotheses that the slope coefficients are zero, denoted by NW[p], based on [135] standard errors with optimal bandwidth selected as in [139]. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The adjusted R^2, presented as a percentage, is denoted by $\bar{R}^2$.

	IS Model OPC Decomposition			
	dy	tms	dfy	tbl
$x^{\mathtt{S}}$	−0.04	0.01	−0.01	−0.02
$x^{\mathtt{GD}}$	−0.03	0.05	−0.17	0.04
$x^{\mathtt{OSD}}$	−0.08	−0.00	−0.08	0.04
	ES Model OPC Decomposition			
	dy	tms	dfy	tbl
$x^{\mathtt{S}}$	−0.08	0.03	−0.07	0.03
$x^{\mathtt{GD}}$	−0.03	0.04	−0.17	0.04
$x^{\mathtt{OSD}}$	−0.06	−0.03	−0.07	0.03

Table 3.24: **Correlations between oil supply, global demand, and oil-specific demand shocks and several US macroeconomic variables**

This table presents correlations between the oil supply ($x_t^{\mathtt{S}}$), global demand ($x_t^{\mathtt{GD}}$), and oil-specific demand ($x_t^{\mathtt{OSD}}$) shocks, obtained from the IS and ES models in a real-time fashion, and the log dividend yield (*dy*), the term spread (*tms*), the default yield spread (*dfy*), and the one-month Treasury Bill rate (*tbl*). Results are presented for the 1986.01–2015.12 sample period. The top and bottom panels show results of the three shocks obtained by oil price change decomposition (3.8) based on the IS and ES models, respectively.

	Mean	STD	AC(1)	AC(2)	AC(3)
Macroeconomic Variables					
US	7.76	2.56	0.97	0.93	0.90
IS Model OPC Decomposition					
Local currency					
Canada	4.79	2.32	0.56	0.14	−0.09
France	6.04	3.76	0.32	−0.03	−0.11
Germany	5.74	4.25	0.36	−0.01	−0.12
Italy	2.50	7.36	0.28	−0.03	−0.09
Japan	3.62	3.28	0.39	0.05	−0.09
UK	4.21	2.74	0.27	0.00	−0.09
US dollar					
Canada	6.79	2.94	0.61	0.19	−0.07
France	7.96	3.22	0.38	−0.03	−0.13
Germany	7.15	3.81	0.41	−0.02	−0.13
Italy	5.24	6.90	0.30	−0.04	−0.10
Japan	3.29	2.65	0.31	−0.02	−0.07
UK	6.68	3.02	0.46	0.09	−0.10
US	7.77	3.44	0.44	0.04	−0.12
World	6.35	3.24	0.43	0.06	−0.11
ES Model OPC Decomposition					
Local currency					
Canada	4.79	2.30	0.52	0.16	−0.01
France	6.01	3.97	0.34	0.12	−0.02
Germany	5.72	4.33	0.36	0.11	−0.03
Italy	2.49	7.19	0.29	0.03	−0.06
Japan	3.60	3.94	0.41	0.14	−0.01
UK	4.18	3.13	0.32	0.13	−0.01
US dollar					
Canada	6.78	2.92	0.56	0.20	0.02
France	7.95	2.98	0.37	0.09	−0.04
Germany	7.13	3.50	0.39	0.09	−0.05
Italy	5.24	6.59	0.27	−0.01	−0.08
Japan	3.28	2.65	0.37	0.13	−0.01
UK	6.66	3.14	0.45	0.14	−0.02
US	7.76	3.40	0.43	0.10	−0.05
World	6.33	3.40	0.43	0.12	−0.03

Table 3.25: **Descriptive statistics for predicted excess returns on MSCI indexes: 1986.01–2015.12**

This table presents the annualized mean and standard deviation (STD), and the first three autocorrelations (AC(k), $k = 1, 2, 3$) of MSCI index predicted excess returns using two sets of predictors: (i) four macroeconomic variables, i.e., the log dividend yield, the term spread, the default yield spread, and the one-month Treasury Bill rate and (ii) the oil supply, global demand, and oil-specific demand shocks obtained from the oil price decomposition. The top panel contains the results for the MSCI US index based on the macroeconomic variables. The second and third panels contain the results for MSCI local currency- and US dollar-denominated indexes based on the oil price change decomposition resulting from the IS and ES models, respectively.

	Mean	STD	AC(1)	AC(2)	AC(3)
[141]'s US Equity Premium					
M1	4.93	3.99	0.78	0.55	0.45
M2	4.97	3.65	0.82	0.61	0.51
M3	4.93	3.35	0.85	0.67	0.57
M6	4.87	2.85	0.88	0.73	0.64
M12	4.64	2.38	0.90	0.80	0.72
Macroeconomic Variables					
US	4.89	3.42	0.94	0.87	0.81
IS Model OPC Decomposition					
Local currency					
Canada	6.98	4.25	0.65	0.19	−0.09
France	5.18	5.23	0.51	−0.00	−0.17
Germany	5.75	6.78	0.47	−0.06	−0.19
Italy	2.55	6.21	0.42	−0.03	−0.16
Japan	−2.19	3.66	0.65	0.22	−0.04
UK	2.52	3.21	0.38	−0.01	−0.12
US dollar					
Canada	10.19	5.72	0.67	0.23	−0.06
France	5.57	5.20	0.54	0.05	−0.15
Germany	6.09	6.62	0.49	−0.02	−0.18
Italy	4.19	6.21	0.46	0.01	−0.14
Japan	−2.60	3.04	0.54	0.04	−0.03
UK	4.34	4.76	0.62	0.17	−0.08
US	4.73	5.36	0.54	0.06	−0.14
World	3.61	4.84	0.57	0.11	−0.11
ES Model OPC Decomposition					
Local currency					
Canada	6.96	4.38	0.65	0.23	−0.02
France	5.12	5.22	0.51	0.14	−0.04
Germany	5.65	6.27	0.46	0.08	−0.07
Italy	2.52	6.66	0.42	0.10	−0.05
Japan	−2.17	3.73	0.64	0.21	−0.03
UK	2.53	4.40	0.37	0.13	−0.01
US dollar					
Canada	10.15	5.72	0.69	0.28	0.00
France	5.52	5.08	0.55	0.14	−0.06
Germany	6.01	6.18	0.47	0.06	−0.08
Italy	4.17	6.41	0.46	0.10	−0.06
Japan	−2.52	2.15	0.67	0.27	−0.01
UK	4.34	5.38	0.58	0.21	−0.01
US	4.68	5.49	0.55	0.15	−0.05
World	3.59	4.98	0.57	0.17	−0.04

Table 3.26: **Descriptive statistics for predicted excess returns on MSCI indexes: 1996.01–2012.01**

This table presents the mean, the standard deviation (STD), and the first three autocorrelations (AC(k), $k = 1, 2, 3$) of MSCI index predicted excess returns using two sets of predictors: (i) four macroeconomic variables, i.e., the log dividend yield, the term spread, the default yield spread, and the one-month Treasury Bill rate and (ii) the oil supply, global demand, and oil-specific demand shocks obtained from the oil price decomposition. For purposes of comparison, the top panel reports the descriptive statistics for the US equity premium estimates of [141], obtained from option prices, for one-month (M1), two-month (M2), three-month (M3), six-month (M6), and one-year (M12) maturities. The second panel contains the results for the MSCI US index based on the macroeconomic variables. The third (fourth) panel contains the results for MSCI local currency-denominated (US dollar-denominated)

	τ_1	p-value	τ_2	p-value	τ_3	p-value	$\zeta^{\mathtt{S}}$	p-value	$\zeta^{\mathtt{GD}}$	p-value	$\zeta^{\mathtt{OSD}}$	p-value
Local currency												
Canada	0.16	0.03**	0.07	0.11	0.97	0.00***	−2.98	0.59	−3.16	0.04**	−0.03	0.96
France	0.33	0.00***	−0.20	0.00***	0.81	0.00***	3.05	0.75	0.07	0.98	−0.04	0.95
Germany	0.28	0.01**	−0.09	0.15	0.87	0.00***	6.73	0.45	−2.21	0.36	0.18	0.81
Italy	0.13	0.11	−0.07	0.04**	0.96	0.00***	7.24	0.24	−2.46	0.15	−0.02	0.97
Japan	−0.00	0.97	−0.04	0.43	−0.79	0.00***	6.13	0.70	−8.56	0.05*	−0.02	0.98
UK	0.28	0.03**	−0.17	0.05**	0.78	0.00***	0.19	0.98	1.02	0.69	0.14	0.86
US dollar												
Canada	0.13	0.04**	0.08	0.09*	0.96	0.00***	−4.59	0.38	−3.49	0.02**	−0.40	0.49
France	0.14	0.11	−0.21	0.00***	0.89	0.00***	5.37	0.45	−0.60	0.71	−0.27	0.58
Germany	0.21	0.04**	−0.08	0.17	0.92	0.00***	5.43	0.46	−1.85	0.32	−0.20	0.73
Italy	−0.07	0.03**	−0.02	0.54	0.98	0.00***	11.79	0.00***	−4.52	0.00***	0.26	0.52
Japan	0.00	0.95	−0.07	0.00***	0.99	0.00***	−4.64	0.30	−2.43	0.04**	0.29	0.33
UK	0.14	0.12	−0.11	0.06*	0.94	0.00***	−1.03	0.88	−0.20	0.91	−0.23	0.70
US	0.25	0.03**	−0.18	0.03**	0.86	0.00***	−3.60	0.66	−0.37	0.86	−0.29	0.66
World	0.16	0.12	−0.18	0.02**	0.85	0.00***	4.40	0.61	−0.63	0.75	−0.22	0.73

Table 3.27: **Effect of oil supply, global demand, and oil-specific demand shocks, based on the IS model, on conditional return volatility: evidence from an EGARCH(1,1) model**
This table presents results of an augmented EGARCH(1,1) model that includes the three shocks in the volatility equation as exogenous regressors. The econometric specification is: $r^{\mathtt{e}}_{t+1} = \alpha^{\mathtt{DEC}} + \beta^{\mathtt{S}} x^{\mathtt{S}}_t + \beta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \beta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t + u^{\mathtt{DEC}}_{t+1}$, with $u^{\mathtt{DEC}}_{t+1} = \sigma_t z_{t+1}$, $z_{t+1} \sim$ i.i.d.$(0,1)$, and $\log(\sigma^2_t) = \tau_0 + \tau_1|z_t| + \tau_2 z_t + \tau_3 \log(\sigma^2_{t-1}) + \zeta^{\mathtt{S}} x^{\mathtt{S}}_t + \zeta^{\mathtt{GD}} x^{\mathtt{GD}}_t + \zeta^{\mathtt{OSD}} x^{\mathtt{OSD}}_t$. The reported results are based on the Student-t distribution for the disturbances z_{t+1}. The model is estimated using monthly excess returns on the MSCI indexes for the G7 countries as well as the World MSCI index over the 1986.01–2015.12 sample period. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	τ_1	p-value	τ_2	p-value	τ_3	p-value	$\zeta^{\texttt{S}}$	p-value	$\zeta^{\texttt{GD}}$	p-value	$\zeta^{\texttt{OSD}}$	p-value
Local currency												
Canada	0.18	0.02**	0.06	0.27	0.95	0.00***	−0.04	0.96	−2.80	0.11	0.05	0.97
France	0.33	0.00***	−0.20	0.00***	0.80	0.00***	0.20	0.79	0.04	0.99	−0.23	0.86
Germany	0.31	0.01***	−0.10	0.14	0.86	0.00***	0.43	0.61	−2.48	0.32	0.09	0.95
Italy	0.15	0.07*	−0.08	0.01***	0.96	0.00***	0.19	0.71	−2.79	0.13	0.20	0.84
Japan	0.02	0.76	−0.06	0.03**	0.93	0.00***	0.81	0.19	−2.68	0.04**	−0.71	0.45
UK	0.27	0.03**	−0.20	0.01**	0.73	0.00***	1.77	0.13	−0.26	0.93	−3.12	0.09*
US dollar												
Canada	0.14	0.03**	0.07	0.18	0.96	0.00***	−0.37	0.65	−3.22	0.06*	−0.44	0.75
France	0.14	0.11	−0.20	0.00***	0.89	0.00***	0.37	0.46	−1.15	0.45	−1.02	0.36
Germany	0.23	0.03**	−0.08	0.15	0.91	0.00***	−0.00	1.00	−1.86	0.34	−0.47	0.72
Italy	0.08	0.16	−0.10	0.01**	0.95	0.00***	0.24	0.64	−3.57	0.04**	−0.18	0.85
Japan	0.01	0.83	−0.06	0.01***	1.00	0.00***	−0.13	0.81	−2.02	0.09*	0.50	0.44
UK	0.14	0.11	−0.11	0.06*	0.95	0.00***	−0.16	0.83	−0.01	1.00	−0.28	0.82
US	0.30	0.02**	−0.26	0.00***	0.77	0.00***	1.06	0.29	−1.39	0.54	−2.96	0.08*
World	0.17	0.12	−0.21	0.01***	0.78	0.00***	1.30	0.23	−1.69	0.40	−2.87	0.10*

Table 3.28: **Effect of oil supply, global demand, and oil-specific demand shocks, based on the ES model, on conditional return volatility: evidence from an EGARCH(1,1) model**
This table presents results of an augmented EGARCH(1,1) model that includes the three shocks in the volatility equation as exogenous regressors. The econometric specification is: $r^{\texttt{e}}_{t+1} = \alpha^{\texttt{DEC}} + \beta^{\texttt{S}} x^{\texttt{S}}_t + \beta^{\texttt{GD}} x^{\texttt{GD}}_t + \beta^{\texttt{OSD}} x^{\texttt{OSD}}_t + u^{\texttt{DEC}}_{t+1}$, with $u^{\texttt{DEC}}_{t+1} = \sigma_t z_{t+1}$, $z_{t+1} \sim \text{i.i.d.}(0,1)$, and $\log(\sigma^2_t) = \tau_0 + \tau_1|z_t| + \tau_2 z_t + \tau_3 \log(\sigma^2_{t-1}) + \zeta^{\texttt{S}} x^{\texttt{S}}_t + \zeta^{\texttt{GD}} x^{\texttt{GD}}_t + \zeta^{\texttt{OSD}} x^{\texttt{OSD}}_t$. The reported results are based on the Student-t distribution for the disturbances z_{t+1}. The model is estimated using monthly excess returns on the MSCI indexes for the G7 countries as well as the World MSCI index over the 1986.01–2015.12 sample period. The top and bottom panels contain results for local currency- and US dollar-denominated index returns, respectively. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

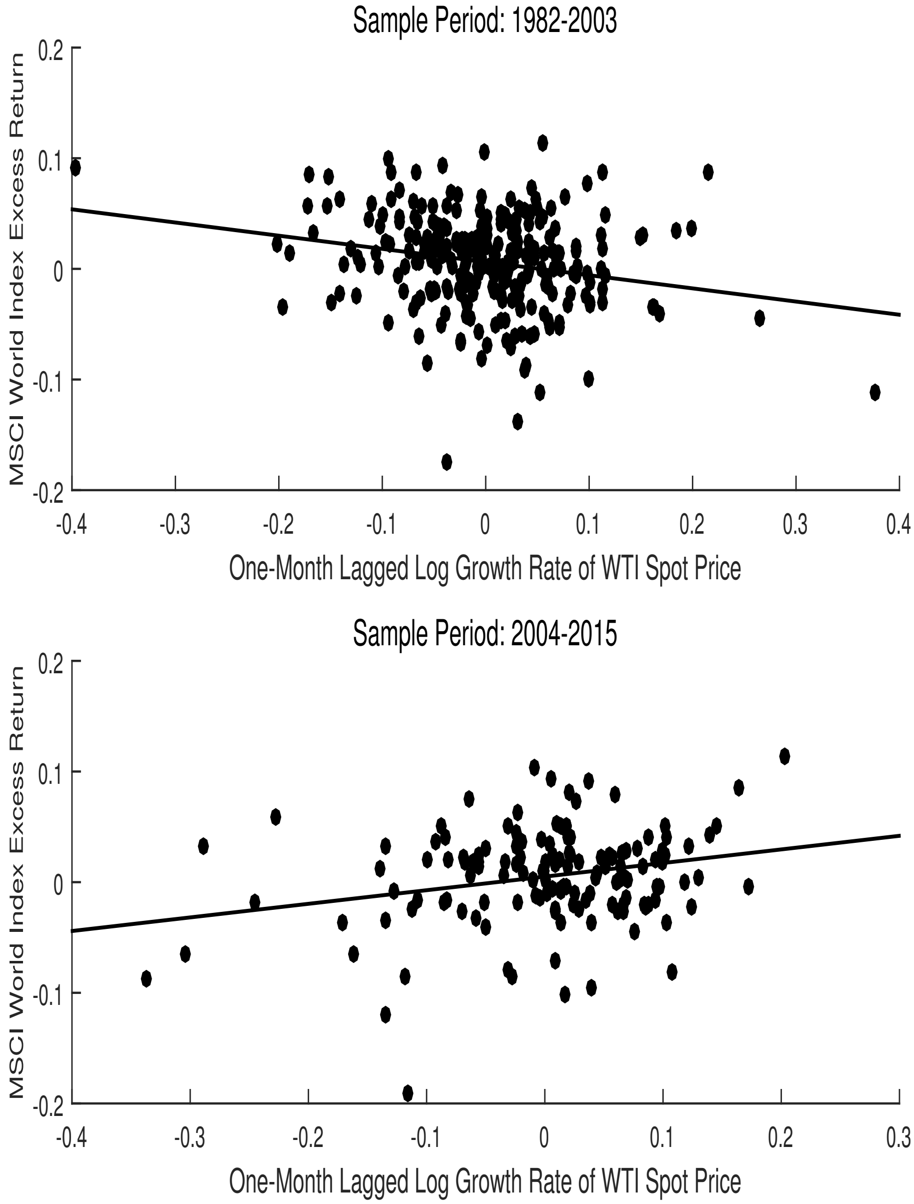

Figure 3.1: **Scatter plots of US dollar-denominated MSCI World index excess return versus the one-month lagged log growth rate of West Texas Intermediate (WTI) spot price**

In this figure, we present the scatter plots of the US dollar-denominated MSCI World index return versus the one-month lagged log growth rate of WTI spot price over the 1982.01–2003.12 and 2004.01–2015.12 sample periods. The solid lines represent the fitted least-squares regression lines. The correlation between the MSCI World index excess return and the one-month lagged log growth rate of WTI spot price is -0.22, 0.25, and -0.04 over the 1982.01–2003.12, 2004.01–2015.12, and 1982.01–2015.12 sample periods, respectively.

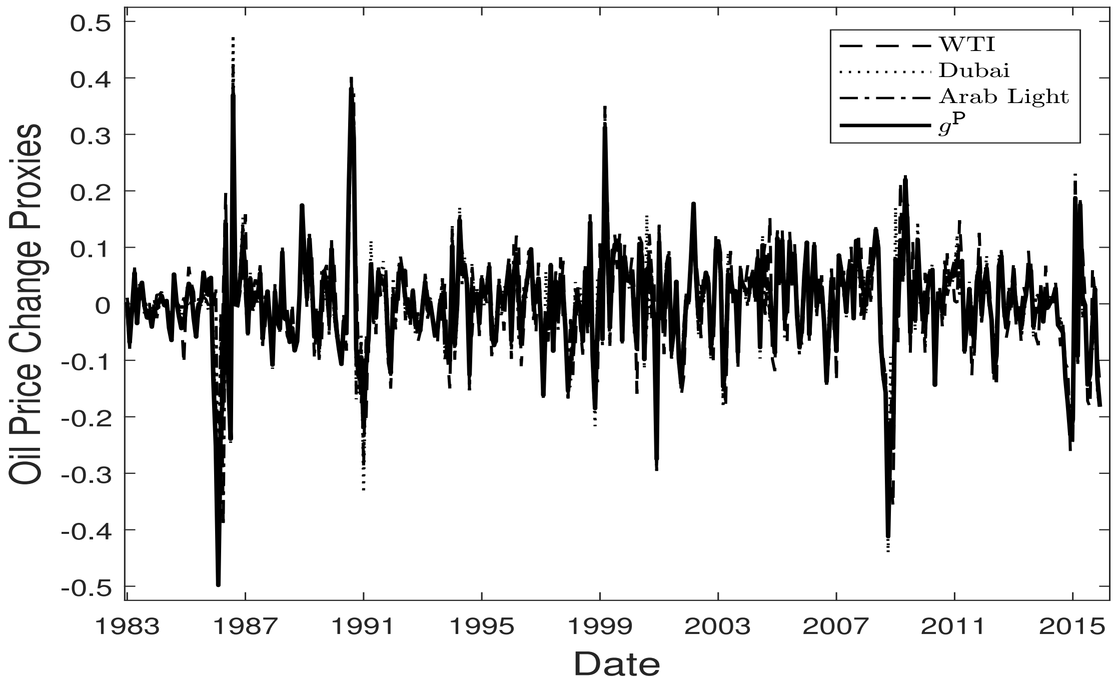

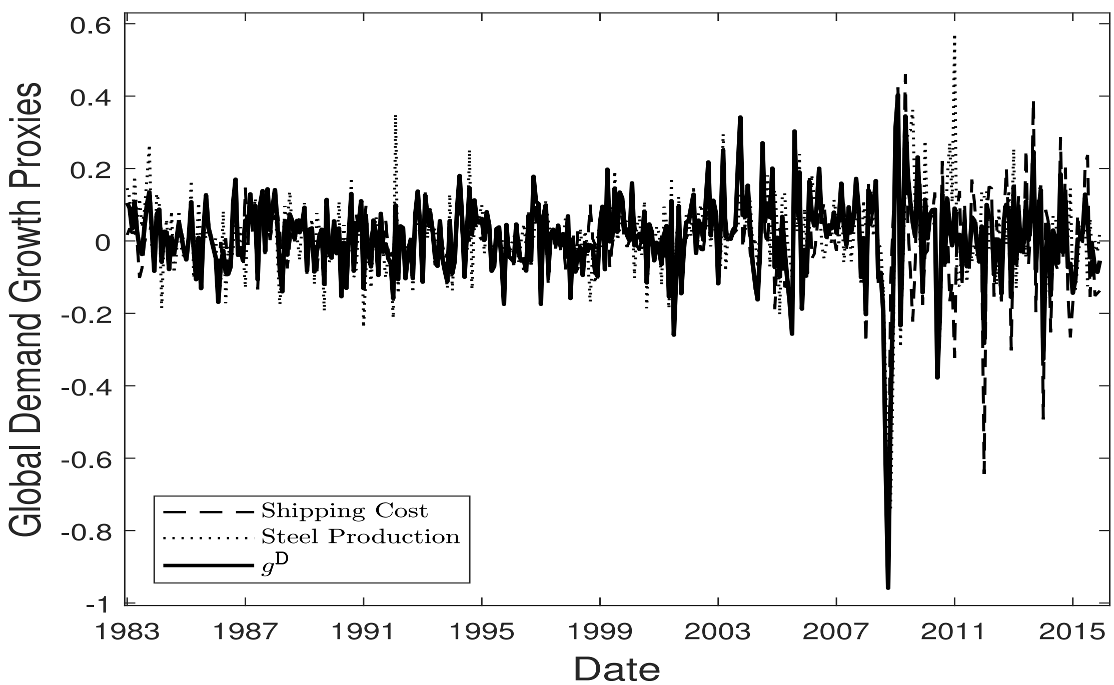

Figure 3.2: **Time series of oil price change and global demand growth proxies**
In the first figure, we plot the log growth rates of three oil spot price proxies, i.e., WTI, Dubai, and Arab Light, along with their first principal component, g^{P}, over the 1983.01–2015.12 sample period. The series are rescaled so that they have a standard deviation equal to 0.09. In the second figure, we plot the log growth rates of the shipping cost index and the seasonally-adjusted crude steel production, along with their first principal component, g^{GD}, over the 1983.01–2015.12 sample period. The series are rescaled so that they have a standard deviation equal to 0.12.

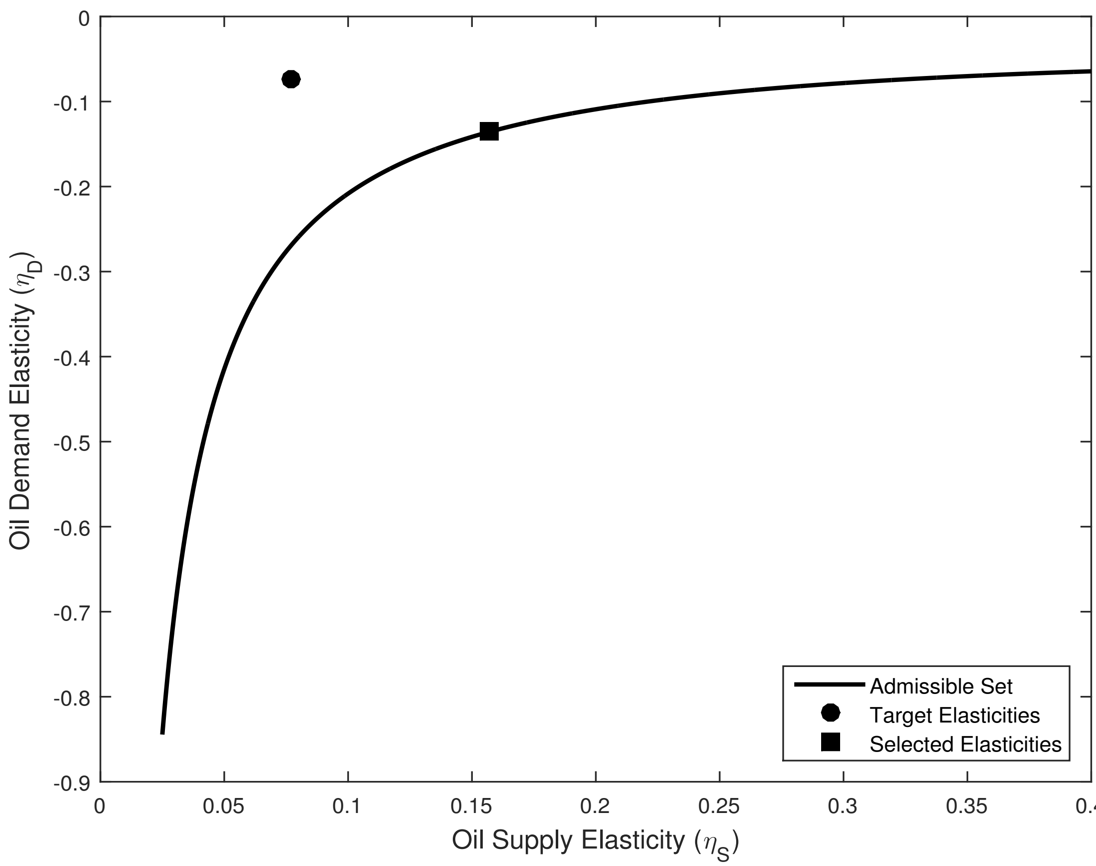

Figure 3.3: **Oil supply and demand elasticities consistent with the estimates of the reduced-form VAR model**
In this figure, we plot the pairs of supply and demand elasticities $(\eta_{\mathtt{S}}, \eta_{\mathtt{D}})$ that are consistent with the estimates of the reduced-form VAR model (3.4). The circle corresponds to the independent IV estimates $(\eta^*_{\mathtt{S}}, \eta^*_{\mathtt{D}}) = (0.077, -0.074)$ of [128]. The square corresponds to the pair of optimal admissible elasticities $(\widehat{\eta}_{\mathtt{S}}, \widehat{\eta}_{\mathtt{D}}) = (0.157, -0.136)$.

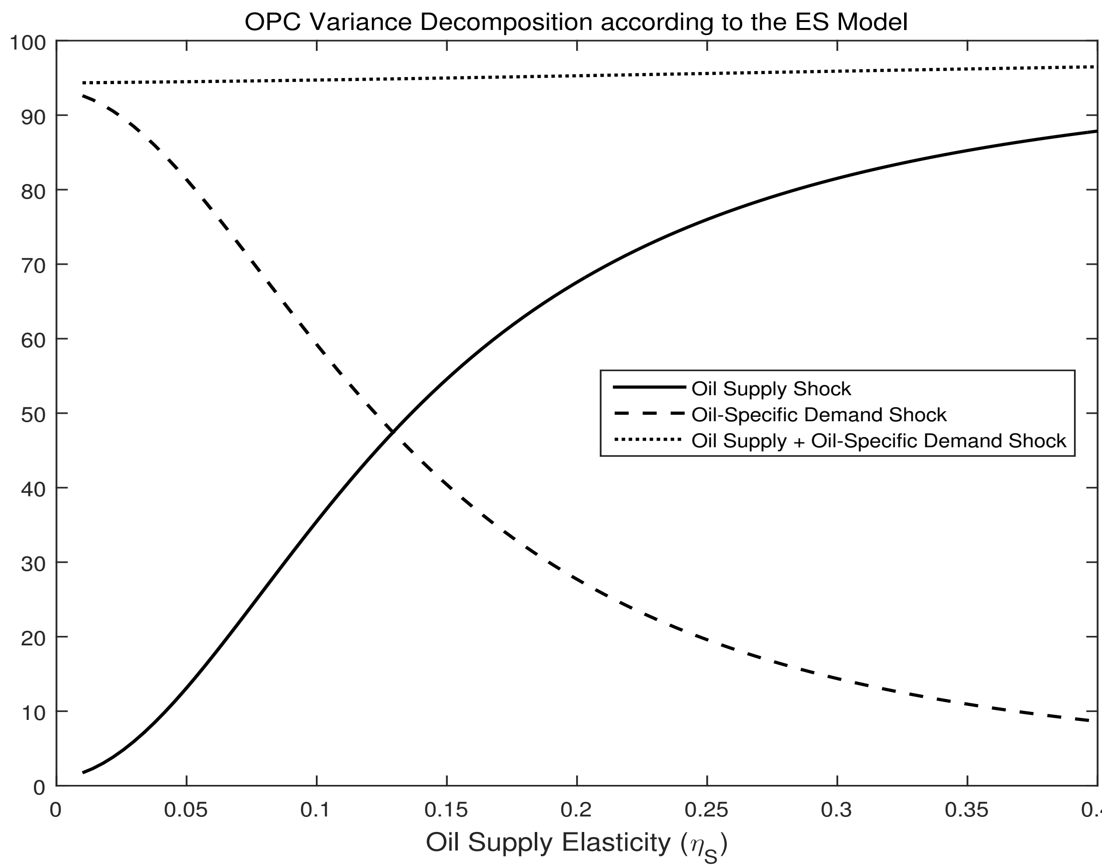

Figure 3.4: **Variance decomposition of oil price changes according to the ES model**
In this figure, we plot the percentages of the oil price variance attributed to the oil supply and oil-specific demand shocks based on the estimates of the ES model.

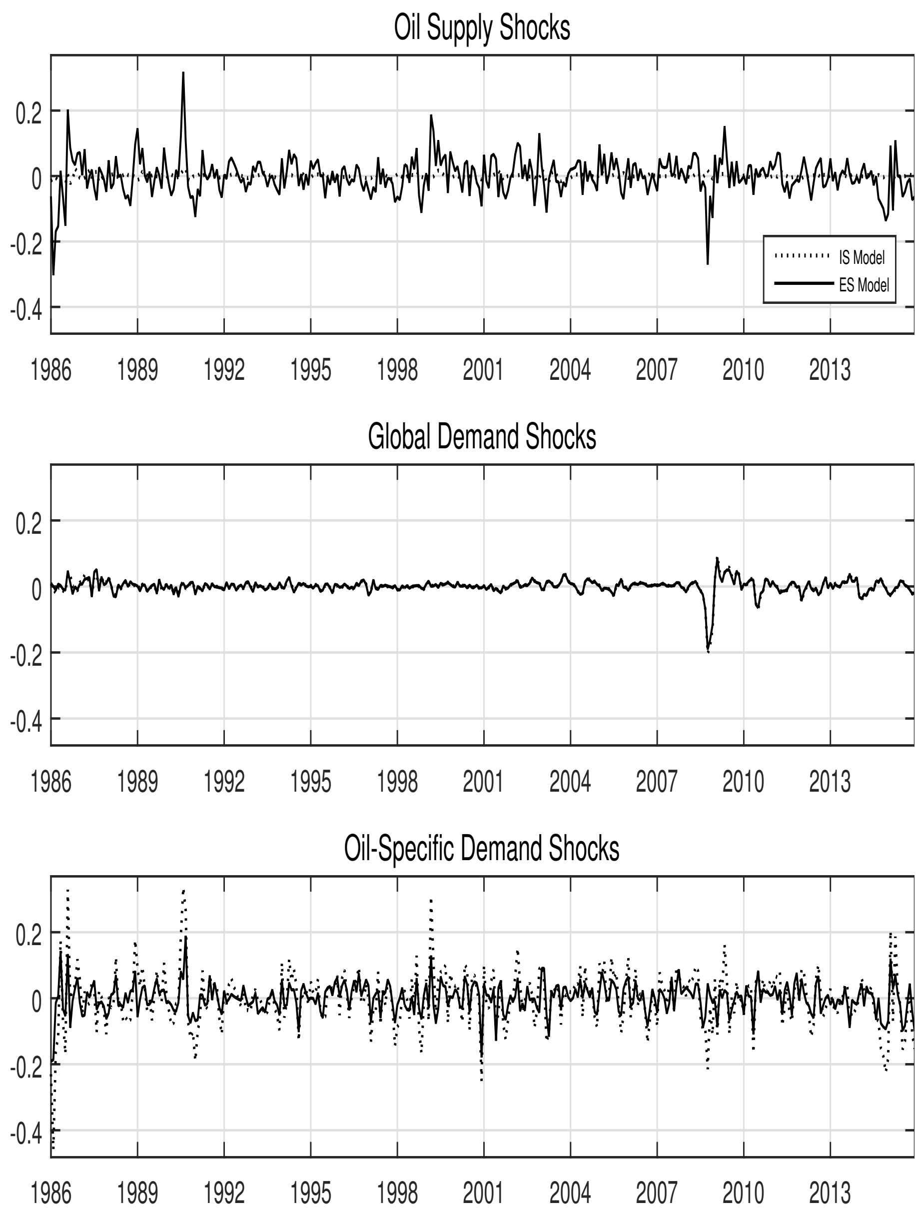

Figure 3.5: **Time series of oil supply, global demand, and oil-specific demand shocks**

In this figure, we plot the time series of the oil supply, global demand, and oil-specific demand shocks obtained using the decomposition in equation (3.8), based on the IS and ES models, over the 1986.01–2015.12 sample period. The shocks are obtained in a real-time fashion as described in section 3.4.

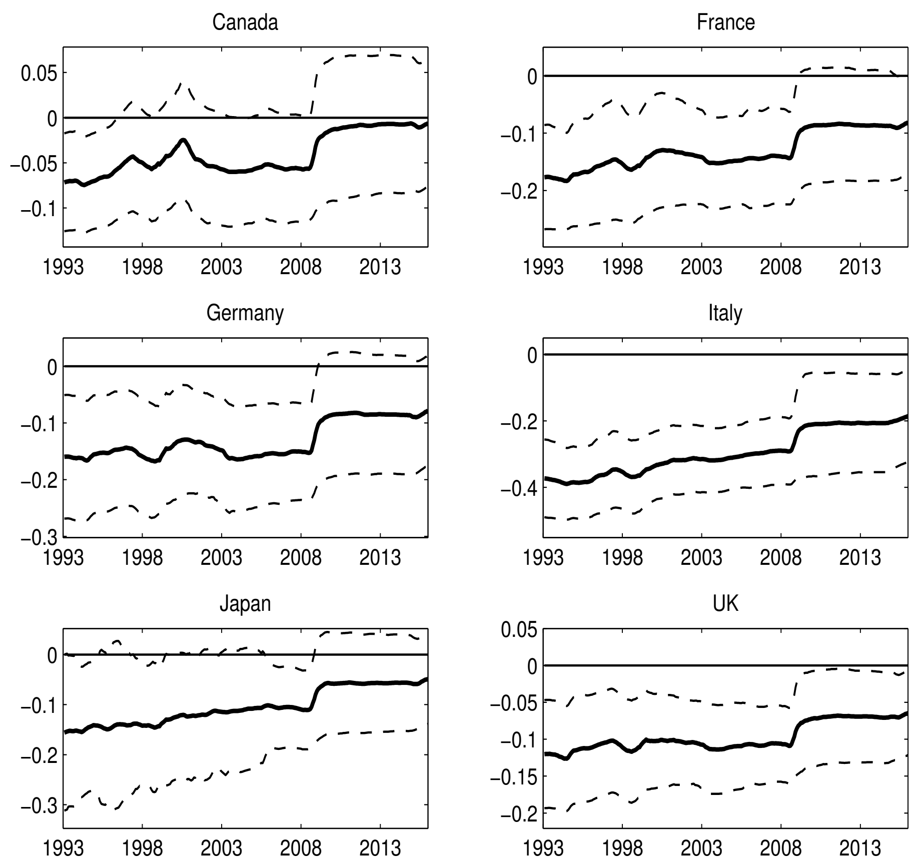

Figure 3.6: **Oil price change as a predictor of local currency-denominated MSCI index excess returns: slope estimates over expanding sample periods**

In this figure, we plot the time series of the slope estimates, along with the corresponding 95% confidence intervals based on [135] standard errors, from the predictive regression model (3.1) over different samples using an expanding window with the first sample being 1983.01–1993.01 and the last sample being 1983.01–2015.12.

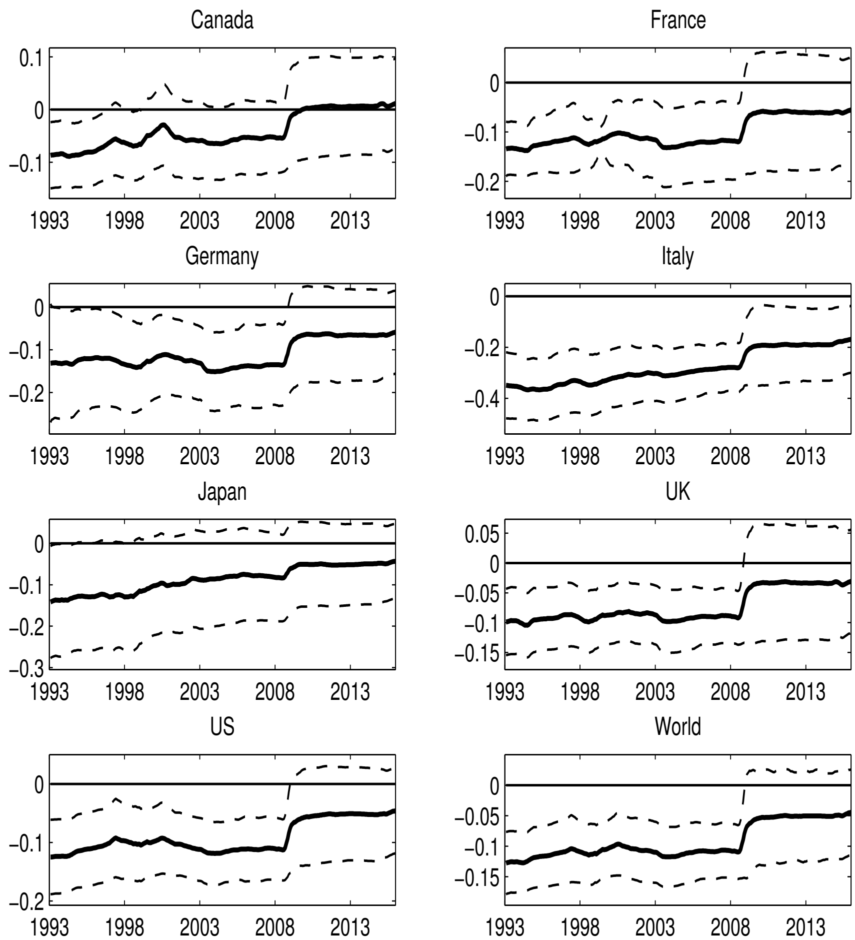

Figure 3.7: **Oil price change as a predictor of US dollar-denominated MSCI index excess returns: slope estimates over expanding sample periods**
In this figure, we plot the time series of the slope estimates, along with the corresponding 95% confidence intervals based on [135] standard errors, from the predictive regression model (3.1) over different samples using an expanding window with the first sample being 1983.01–1993.01 and the last sample being 1983.01–2015.12.

Bibliography

[1] Hendrik Bessembinder, Chester Spatt, and Kumar Venkataraman. A Survey of the Microstructure of Fixed-Income Markets. *Journal of Financial and Quantitative Analysis*, 2019. ISSN 17566916. doi: 10.1017/S0022109019000231. Forthcoming.

[2] Jaewon Choi and Yesol Huh. Customer Liquidity Provision: Implications for Corporate Bond Transaction Costs. available at SSRN 2848344, August 2019.

[3] Richard Roll. A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *The Journal of Finance*, 39(4):1127–1139, sep 1984.

[4] Yakov Amihud. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of financial markets*, 5(1):31–56, 2002.

[5] Rainer Jankowitsch, Amrut Nashikkar, and Marti G Subrahmanyam. Price dispersion in otc markets: A new measure of liquidity. *Journal of Banking & Finance*, 35(2):343–357, 2011.

[6] Jack Bao, Jun Pan, and Jiang Wang. The Illiquidity of Corporate Bonds. *The Journal of Finance*, 66(3):911–946, jun 2011.

[7] Nils Friewald, Rainer Jankowitsch, and Marti Subrahmanyam. Illiquidity or Credit Deterioration: A Study of Liquidity in the US Corporate Bond Market during Financial Crises. *Journal of Financial Economics*, 105(1):18–36, 2012. ISSN 0304405X. doi: 10.1016/j.jfineco.2012.02.001. URL `http://dx.doi.org/10.1016/j.jfineco.2012.02.001`.

[8] Peter Feldhütter. The same bond at different prices: Identifying search frictions and selling pressures. *The Review of Financial Studies*, 25(4):1155–1206, April 2012.

[9] Jack Bao, Maureen O'Hara, and Xing Zhou. The Volcker Rule and Corporate Bond Market Making in Times of Stress. *Journal of Financial Economics*, 130 (1):95–113, oct 2018.

[10] Hendrik Bessembinder, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman. Capital Commitment and Illiquidity in Corporate Bonds. *The Journal of Fiannce*, 73(4):1615–1661, 2018.

[11] Gary Gorton and Andrew Metrick. Securitized Banking and the Run on Repo. *Journal of Financial Economics*, 104(3):425–451, 2012.

[12] Maria Chaderina, Alexander Mürmann, and Christoph Scheuch. The Dark Side of Liquid Bonds in Fire Sales. available at SSRN 2995544, 2019.

[13] Maureen O'Hara and Xing Zhou. The electronic evolution of corporate bond dealers. *Working paper*, 2019.

[14] Jens Dick-Nielsen, Peter Feldhütter, and David Lando. Corporate Bond Liquidity before and after the Onset of the Subprime Crisis. *Journal of Financial Economics*, 103(3):471–492, 2012. ISSN 0304405X. doi: 10.1016/j.jfineco.2011.10.009.

[15] Gideon Saar, Jian Sun, Ron Yang, and Haoxiang Zhu. From Market Making to Matchmaking: Does Bank Regulation Harm Market Liquidity? available at SSRN 3399063, 2019.

[16] Amber Anand, Chotibhak Jotikasthira, and Kumar Venkataraman. Do Buy-Side Institutions Supply Liquidity in Bond Markets? Evidence from Mutual Funds. 2018.

[17] Viral Acharya, Yakov Amihud, and Sreedhar Bharath. Liquidity Risk of Corporate Bond Returns: Conditional Approach. *Journal of Financial Economics*, 110(2):358–386, 2013. ISSN 0304405X. doi: 10.1016/j.jfineco.2013.08.002. URL `http://dx.doi.org/10.1016/j.jfineco.2013.08.002`.

[18] Harald Uhlig. What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. *Journal of Monetary Economics*, 52(2):381–419, 2005. ISSN 03043932. doi: 10.1016/j.jmoneco.2004.05.007.

[19] Darrell Duffie. Market Making under the Proposed Volcker Rule. Rock Center for Corporate Governance at Stanford University Working Paper 106, jan 2012.

[20] Mike Anderson and Rene Stulz. Is Post-Crisis Bond Liquidity Lower? No. w23317. National Bureau of Economic Research, April 2017.

[21] Francesco Trebbi and Kairong Xiao. Regulation and Market Liquidity. *Management Science*, 65(5):1949–1968, 2019. ISSN 15265501. doi: 10.1287/mnsc.2017.2876.

[22] Jonathan Goldberg. Liquidity supply by broker-dealers and real activity. *Journal of Financial Economics*, 2019. forthcoming.

[23] Jonathan Goldberg and Yoshio Nozawa. Liquidity supply and demand in the corporate bond market. 2019. available at SSRN 3209658.

[24] Jaewon Choi and Or Shachar. Did Liquidity Providers Become Liquidity Seekers? *Federal Reserve Bank of New York Staff Reports*, 2013.

[25] Hendrik Bessembinder, William Maxwell, and Kumar Venkataraman. Market Transparency, Liquidity Externalities, and Institutional Trading Costs in Corporate Bonds. *Journal of Financial Economics*, 82(2):251–288, nov 2006.

[26] Nils Friewald and Florian Nagler. Over-the-Counter Market Frictions and Yield Spread Changes. *The Journal of Finance*, 2019. doi: 10.1111/jofi.12827. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12827`.

[27] Erin Lu, Gene Lai, and Qingzhong Ma. Organizational Structure, Risk-Based Capital Requirements, and the Sales of Downgraded Bonds. *Journal of Banking and Finance*, 74:51–68, jan 2017.

[28] Helmut Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer Science & Business Media, 2005.

[29] Jens Dick-Nielsen. How to Clean Enhanced TRACE Data. available at SSRN 2337908, dec 2014.

[30] Michael Goldstein and Edith Hotchkiss. Providing Liquidity in an Illiquid Market: Dealer Behavior in US Corporate Bonds. *Journal of Financial Economics*, 2019.

[31] Marco Di Maggio, Amir Kermani, and Zhaogang Song. The Value of Trading Relations in Turbulent Times. *Journal of Financial Economics*, 124(2):266–284, may 2017.

[32] CGFS. Fixed Income Market Liquidity. available at `https://www.bis.org/publ/cgfs55.pdf`, January 2016.

[33] Tobias Adrian, Michael Fleming, Or Shachar, and Erik Vogt. Market Liquidity after the Financial Crisis. *Annual Review of Financial Economics*, 9(1):43–83, 2017. ISSN 1941-1367. doi: 10.1146/annurev-financial-110716-032325.

[34] Bank for International Supervision. Basel III: The Liquidity Coverage Ratio and Liquidity Risk Monitoring Tools. *Basel Committee on Banking Supervision*, pages 1–75, 2013. URL `www.bis.org`.

[35] SIFMA. SIFMA Electronic Bond Trading Report: US Corporate & Municipal Securities, feb 2016. available at `https://www.sifma.org/wp-content/uploads/2017/05/sifma-electronic-bond-trading-report-us-corporate-and-municipal-securities.pdf`.

[36] Terrence Hendershott and Ananth Madhavan. Click or Call? Auction versus Search in the Over-the-Counter Market. *Journal of Finance*, 70(1):419–447, 2015. ISSN 15406261. doi: 10.1111/jofi.12164.

[37] Pierre Collin-Dufresne, Robert Goldstein, and Spencer Martin. The Determinants of Credit Spread Changes. *Journal of Finance*, 56(6):2177–2207, 2001. ISSN 00221082. doi: 10.1111/0022-1082.00402.

[38] Gurdip Bakshi and Dilip Madan. Spanning and Derivative-Security Valuation. *Journal of Financial Economics*, 55(2):205–238, 2000. ISSN 0304405X. doi: 10.1016/S0304-405X(99)00050-1.

[39] Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock Return Characteristics, Skew Laws, and the Differential Pricing. *The Review of Financial Studies*, 16(1):101–143, 2003.

[40] Yakov Amihud and Haim Mendelson. Asset Pricing and the Bid-Ask Spread. *Journal of Financial Economics*, 17(2):223–249, 1986. ISSN 0304405X. doi: 10.1016/0304-405X(86)90065-6.

[41] Long Chen, David Lesmond, and Jason Wei. Corporate Yield Spreads and Bond Liquidity. *Journal of Finance*, 62(1):119–149, 2007. ISSN 00221082. doi: 10.1111/j.1540-6261.2007.01203.x.

[42] Francis Longstaff and Eduardo Schwartz. A Simple Approach to Valuing Risky Fixed and Floating Rate Debt. *The Journal of Finance*, 50(3):789–819, 1995. URL `https://pdfs.semanticscholar.org/ce18/29a7e6c80a35153a9541044b43b5d5f54302.pdf{%}0Ahttps://dl-web.dropbox.com/get/LiteraturIlmiah/Keuangan/RiskPremiums/Sovereign/1995LonstaffSchwartzJoF.pdf?w=c4013668`.

[43] Christian Danne. The VARsignR Package, 2015. URL `https://cran.r-project.org/web/packages/VARsignR/vignettes/VARsignR-vignette.html`. available at `https://cran.r-project.org/web/packages/VARsignR/vignettes/VARsignR-vignette.html`.

[44] Julien Hugonnier, Benjamin Lesterz, and Pierre-Olivier Weill. Heterogeneity in Decentralized Asset Markets. No. w20746. National Bureau of Economic Research, feb 2016.

[45] Darrell Duffie, Nicolae Gârleanu, and Lasse Pedersen. Over-the-Counter Markets. *Econometrica*, 73(6):1815–1847, nov 2005.

[46] Dan Li and Norman Schürhoff. Dealer Networks. *The Journal of Finance*, 74 (1):91–144, 2019.

[47] Burton Hollifield, Artem Neklyudov, and Chester Spatt. Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations. *The Review of Financial Studies*, 30(9):3048–3085, sep 2017.

[48] Terrence Hendershott, Dan Li, Dmitry Livdan, and Norman Schürhoff. Relationship Trading in OTC Markets. Swiss Finance Institute Research Paper No. 17-30, 2017.

[49] Hong Liu and Yajun Wang. Market Making with Asymmetric Information and Inventory Risk. *Journal of Economic Theory*, 163:73–109, may 2016.

[50] Semyon Malamud and Marzena Rostek. Decentralized Exchange. *American Economic Review*, 107(11):3320–3362, nov 2017.

[51] Sriketan Mahanti, Amrut Nashikkar, Marti Subrahmanyam, George Chacko, and Gaurav Mallik. Latent Liquidity: A New Measure of Liquidity, with an Application to Corporate Bonds. *Journal of Financial Economics*, 88(2): 272–298, may 2008.

[52] Sanford Grossman and Merton Miller. Liquidity and Market Structure. *The Journal of Finance*, 43(3):617–633, jul 1988.

[53] Wen Chen and Yajun Wang. Dynamic Market Making and Asset Pricing. available at SSRN 3224495, aug 2018.

[54] Whitney Newey and Kenneth West. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–708, 1987.

[55] Stephen Ross. The recovery theorem. *Journal of Finance*, 70:615–648, 2015.

[56] Jaroslav Borovička, Lars Hansen, and Jose Scheinkman. Misspecified recovery. *Journal of Finance*, 71(6):2493–2544, 2016.

[57] Gurdip Bakshi, Fousseni Chabi-Yo, and Xiaohui Gao. A recovery that we can trust? Deducing and testing the restrictions of the recovery theorem. *Review of Financial Studies*, 31(2):532–555, 2018.

[58] Likuan Qin, Vadim Linetsky, and Yutian Nie. Long forward probabilities, recovery and the term structure of bond risk premiums. *Review of Financial Studies*, 31(12):4863–4883, 2018.

[59] Christian Jensen, David Lando, and Lasse Pedersen. Generalized recovery. *Journal of Financial Economics*, 133(1):154–174, 2019.

[60] Paul Schneider and Fabio Trojani. (Almost) model-free recovery. *Journal of Finance*, 74(1):323–370, 2019.

[61] Ngoc-Khanh Tran and Shixiang Xia. Specified recovery. Working paper, Virginia Tech, 2019.

[62] Francesco Audrino, Robert Huitema, and Markus Ludwig. An empirical implementation of the Ross recovery theorem as a prediction device. *Journal of Financial Econometrics*, (forthcoming), 2019.

[63] Yannick Dillschneider and Raimond Maurer. Functional Ross recovery: Theoretical results and empirical tests. *Journal of Economic Dynamics and Control*, (forthcoming), 2019.

[64] Jens Jackwerth and Marco Menner. Does the Ross recovery theorem work empirically? *Journal of Financial Economics*, (forthcoming), 2020.

[65] Scottt Joslin, Marcel Priebsch, and Kenneth Singleton. Risk premiums in dynamic term structure models with unspanned macro risks. *Journal of Finance*, 69:1197–1233, 2014.

[66] Michael Bauer and James Hamilton. Robust bond risk premia. *Review of Financial Studies*, 31(2):399–448, 2018.

[67] Bernt Økesendal and Anges Sulem. *Applied Stochastic Control of Jump Diffusions*. Springer, New York, 2007.

[68] Homer Reid. Smoothness of functions. Lecture Notes, Department of Mathematics, MIT (http://homerreid.com/teaching/18.305/Notes/SmoothnessOfFunctions.pdf), 2015.

[69] Gurdip Bakshi and Dilip Madan. Spanning and derivative-security valuation. *Journal of Financial Economics*, 55(2):205–238, February 2000.

[70] Joshua Rosenberg and Robert Engle. Empirical pricing kernels. *Journal of Financial Economics*, 64(3):341–372, 2002.

[71] Ranadeb Chaudhuri and Mark Schroder. Monotonicity of the stochastic discount factor and expected option returns. *Review of Financial Studies*, 28(5): 1462–1505, 2015.

[72] Maurice Kendall and Alan Stuart. *The Advanced Theory of Statistics.*, volume 1. Charles Griffin, London, 4 edition, 1977.

[73] R Kanwal. *Generalized Functions: Theory and Technique*. Birkhauser, Boston, MA, 1998.

[74] Tim Bollerslev and Viktor Todorov. Tails, fears, and risk premia. *Journal of Finance*, 66:2165–2211, 2011.

[75] Bruno Feunou, Mohammad Jahan-Parvar, and Romeo Tedongap. Which parametric model for conditional skewness? *European Journal of Finance*, 22(3): 1237–1271, 2016.

[76] Bruno Feunou and Cedric Okou. Good volatility, bad volatility and option pricing. *Journal of Financial and Quantitative Analysis*, 54(2):695–727, 2019.

[77] M. Kilic and I. Shaliastovich. Good and bad variance premia and expected returns. *Management Science*, 65(6):2445–2945, 2019.

[78] Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143, 2003.

[79] Andrew Ang, Robert Hodrick, Yuhang Xing, and Xiaoyan Zhang. The cross-section of volatility and expected returns. *Journal of Finance*, 61:259–299, 2006.

[80] Peter Christoffersen, Kris Jacobs, and Steven Heston. Capturing option anomalies with a variance-dependent pricing kernel. *Review of Financial Studies*, 26:1963–2006, 2013.

[81] Ian Martin. What is the expected return on the market? *Quarterly Journal of Economics*, 150:367–433, 2017.

[82] William Sharpe. Expected utility asset allocation. Working paper, Stanford University, 2007.

[83] Niels Gormsen and Christian Jensen. Higher-moment risk. Working paper, Copenhagen Business School, 2019.

[84] Robert Bliss and Nikolaos Panigirtzoglou. Option-implied risk aversion estimates. *Journal of Finance*, 59(1):407–446, 2004.

[85] Y. Aït-Sahalia and A. Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94:9–51, 2000.

[86] Jens Jackwerth. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2):433–451, 2000.

[87] Yacine Aït-Sahalia, Yubo Wang, and Francis Yared. Do options markets correctly price the probabilities of movement of the underlying asset? *Journal of Econometrics*, 102:67–110, 2001.

[88] Jian Du and Nikunj Kapadia. The tail in the volatility index. Working paper, University of Massachusetts, 2013.

[89] G. Gao and Z. Song. Rare disaster concerns everywhere. *Management Science*, 65(7):2947–3448, 2019.

[90] Emil Siriwardane. The probability of rare disasters: Estimation and implications. Working paper, Harvard University, 2015.

[91] Anisha Ghosh, Christian Julliard, and Alex Taylor. What is the consumption-CAPM missing? An information-theoretic framework for the analysis of asset pricing models. *Review of Financial Studies*, 30(2):442–504, 2017.

[92] Lars Hansen and Ravi Jagannathan. Assessing specification errors in stochastic discount factor models. *Journal of Finance*, 52(2):557–590, 1997.

[93] John Cochrane. *Asset Pricing*. Princeton University Press, Princeton, NJ, 2005.

[94] Hui Chen, Scott Joslin, and Ngoc-Khanh Tran. Rare disasters and risk sharing with heterogeneous beliefs. *Review of Financial Studies*, 25:2189–2224, 2012.

[95] Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, 127:645–700, 2012.

[96] Jessica Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, 68:987–1035, 2013.

[97] Michael McCraken and Serena Ng. FRED-MD: A monthly database for macroeconomic research. Working paper # 012B, Federal Reserve Bank of St. Louis, 2015.

[98] K. Jurado, S. Ludvigson, and S. Ng. Measuring uncertainty. *American Economic Review*, 117:1177–1216, 2015.

[99] Itamar Drechsler and Amir Yaron. What's vol got to do with it. *Review of Financial Studies*, 24:1–45, 2011.

[100] Torben Andersen, Nicola Fusari, and Viktor Todorov. The risk premia embedded in index options. *Journal of Financial Economics*, 117:558–584, 2015.

[101] Franois Gourio. Disaster risk and business cycles. *American Economic Review*, 102(6):2734–2766, 2012.

[102] J. Mincer and V. Zarnowitz. The evaluation of economic forecasts. In *In Economic forecasts and expectations: Analysis of forecasting behavior and performance*, pages 3–46. NBER, 1969.

[103] Andrew Patton and Allan Timmermann. Testing forecast optimality under unknown loss. *Journal of the American Statistical Association*, 102(480):1172–1184, 2007.

[104] Torben Andersen, Tim Bollerslev, Francis Diebold, and Paul Labys. Modeling and forecasting realized volatility. *Econometrica*, 71(2):529–626, 2003.

[105] Whitney Newey and Kenneth West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.

[106] Whitney Newey and Kenneth West. Automatic lag selection in covariance matrix estimation. *Review of Economic Studies*, 61(4):631–653, 1994.

[107] Todd Clark and Kenneth West. Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138:291–311, 2007.

[108] Yakov Amihud, Clifford Hurvich, and Yi Wang. Multiple-predictor regressions: Hypothesis testing. *Review of Financial Studies*, 22(1):413–434, 2009.

[109] John Cochrane and Monika Piazzesi. Bond risk premia. *American Economic Review*, 95(1):138–160, 2005.

[110] Sydney Ludvigson and Serena Ng. Macro factors in bond risk premia. *Review of Financial Studies*, 22(12):5027–5067, 2009.

[111] Robin Greenwood and Dimitri Vayanos. Bond supply and excess bond returns. *Review of Financial Studies*, 27:663–713, 2014.

[112] A. Cieslak and P. Povala. Expected returns in Treasury bonds. *Review of Financial Studies*, 28:2859–2901, 2015.

[113] Peter Carr and Dilip Madan. Towards a theory of volatility trading. Working paper, NYU and University of Maryland (http://www.math.nyu.edu/research/carrp/papers/pdf/twrdsfig.pdf), 2002.

[114] Dan Stefanica. *A Primer for the Mathematics of Financial Engineering*. FE Press, New York, 2011.

[115] Robert Hodrick. Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies*, 5: 141–161, 1992.

[116] Gerben Driesprong, Ben Jacobsen, and Benjamin Maat. Striking oil: Another puzzle? *Journal of Financial Economics*, 89:307–327, 2008. URL `https://ideas.repec.org/a/eee/jfinec/v89y2008i2p307-327.html`.

[117] Ben Bernanke. The relationship between stocks and oil prices. Article, The Brookings Institution, 2016.

[118] James D. Hamilton. What is an oil shock? *Journal of Econometrics*, 113(2): 363–398, April 2003.

[119] Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99:1053–1069, 2009. URL `https://ideas.repec.org/a/aea/aecrev/v99y2009i3p1053-69.html`.

[120] Lutz Kilian and Cheolbeom Park. The impact of oil price shocks on the U.S. stock market. *International Economic Review*, 50:1267–1287, 2009. URL `https://ideas.repec.org/a/ier/iecrev/v50y2009i4p1267-1287.html`.

[121] Nai-Fu Chen, Richard Roll, and Stephen Ross. Economic forces and the stock market. *Journal of Business*, 59(3):383–403, 1986.

[122] Charles Jones and Gautam Kaul. Oil and the stock markets. *Journal of Finance*, 51:463–491, 1996.

[123] I-hsuan Ethan Chiang, W. Keener Hughen, and Jacob S. Sagi. Estimating oil risk factors using information from equity and derivatives markets. *Journal of Finance*, 70(2):769–804, 2015. ISSN 1540-6261. doi: 10.1111/jofi.12222. URL `http://dx.doi.org/10.1111/jofi.12222`.

[124] Peter Christoffersen and Xuhui Pan. Oil volatility risk and expected stock returns. *Journal of Banking and Finance*, forthcoming, 2017.

[125] Jaime Casassus and Freddy Higuera. Short-horizon return predictability and oil prices. *Quantitative Finance*, 12(12):1909–1934, December 2012.

[126] Paresh Kumar Narayan and Rangan Gupta. Has oil price predicted stock returns for over a century? *Energy Economics*, 48(C):18–23, 2015. URL `https://ideas.repec.org/a/eee/eneeco/v48y2015icp18-23.html`.

[127] Conghui Hu and Wei Xiong. Are commodity futures prices barometers of the global economy? In E. L. Glaeser, T. Santos, and E. G. Weyl, editors, *After the Flood: How the Great Recession Changed Economic Thought*, pages 209–242. University of Chicago Press, 2017.

[128] Dario Caldara, Michele Cavallo, and Matteo Iacoviello. Oil price elasticities and oil price fluctuations. Working Paper, Federal Reserve Board, 2017.

[129] Christiane J.S. Baumeister and James D. Hamilton. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. NBER Working Paper #24167, 2017.

[130] Avihai Rapaport. Supply and demand shocks in the oil market and their predictive power. Working Paper, University of Chicago, 2014.

[131] Robert C. Ready. Oil prices and the stock market. *Review of Finance*, 2: 155–176, 2018.

[132] Michael Sockin and Wei Xiong. Informational frictions and commodity markets. *Journal of Finance*, 70:2063–2098, 2015.

[133] James D. Hamilton. Historical oil shocks. In R. E. Parker and R. Whaples, editors, *Handbook of Major Events in Economic History*, pages 239–265. Routledge, 2013.

[134] Francesco Ravazzolo and Joaquin L. Vespignani. A new monthly indicator of global real economic activity. Working Paper, BI Norwegian Business School and University of Tasmania, 2015.

[135] Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708, 1987.

[136] Robert Hodrick. Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies*, 5: 357–386, 1992.

[137] Jushan Bai and Pierre Perron. Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1):1–22, 2003.

[138] Nicholas Apergis and Stephen M. Miller. Do structural oil-market shocks affect stock prices? *Energy Economics*, 31(4):569–575, July 2009. URL `https://ideas.repec.org/a/eee/eneeco/v31y2009i4p569-575.html`.

[139] Whitney Newey and Kenneth West. Automatic lag selection in covariance matrix estimation. *Review of Economic Studies*, 61:631–653, 1994.

[140] Eugene F. Fama and Kenneth R. French. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25:23–49, 1989.

[141] Ian Martin. What is the expected return on the market? *Quarterly Journal of Economics*, 132:367–433, 2017.

[142] Olivier Blanchard. World economic and market developments. Speech, Shanghai Development Research Foundation, 2015.

[143] J.Y. Campbell and S.B. Thompson. Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21:1509–1531, 2008.

[144] Donald Andrews. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59:817–858, 1991.

[145] Donald W. K. Andrews and J. Christopher Monahan. An improved heteroscedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica*, 60:953–966, 1992.