

ABSTRACT

Title of dissertation: **ESSAYS ON ISSUES IN PUBLIC AND ENVIRONMENTAL ECONOMICS**
Daniel Kraynak, Doctor of Philosophy, 2023

Dissertation directed by: **Professor Roberton Williams III**
Department of Agricultural and Resource Economics

This dissertation is composed of three applied economics essays about important topics in public and environmental economics. The first is an analysis of the distributional effects of demand shocks or demand-shifting policies in the context of energy markets and climate policy. The second focuses on the use of remote monitoring technology and its effects on the provision of the local public good of public safety. The third analyzes the effect of imperfect real-world carbon pricing policies on carbon emissions.

Chapter 1 studies the impact of declining coal demand on local labor markets in coal mining regions of the US. I separate the effect of a recent contraction in the coal industry from other factors driving economic trends in coal country by constructing an instrument for coal demand from producing counties. The instrument combines a regional model of coal plant dispatch with variation in the exposure of producing counties to demand shocks from the electricity sector. My estimates demonstrate that demand-driven declines in the value of coal produced eliminate jobs primarily in coal mining and adjacent industries, with the largest effects occurring in Appalachia and the West. I also estimate decreases in in-migration, home values, and expenditures on public education, and increases in poverty.

Applied in a stylized spatial equilibrium model of location choice, my estimates imply an aggregate decline of \$0.5-1 billion in the economic welfare of coal country residents resulting from a net decline of \$3.7 billion in thermal coal production value from 2007-2017.

In Chapter 2, using a novel data set on CCTV cameras in Chandigarh, India, we test whether police officers' effort changes in response to the presence of traffic cameras. Although the cameras are useful in sanctioning drivers, they can also capture the passive (shirking) or active (rent-seeking) corruption of officers. Accounting for the spatial and temporal variations in the operation of the cameras, we find that the presence of a functioning camera results in an increase in *on-the-ground* tickets. Although we do not rule out possible decreases in rent-seeking behavior, a decline in passive corruption appears to be driving the increase in officer ticketing behavior, particularly for the most common vehicles and violations that can be observed from the CCTV cameras. Our findings indicate that remote monitoring technology can serve, if not a substitute for, then as a complement to *on-the-ground* enforcement.

In Chapter 3, we contribute to a growing body of empirical evidence on the efficacy of carbon pricing policies, much of which finds that carbon pricing has produced only modest reductions in emissions. We hypothesize that a complex policy environment and political uncertainty are two possible mechanisms behind the limited effects measured in the literature. We focus on the experience of Australia which substantially expanded subsidies for renewable energy in 2009 and also implemented a controversial carbon "tax" from 2012-2014 before it was repealed. Using synthetic control and recent extensions, we estimate the joint effect of the subsidy expansion and the tax to be substantial. We explore the dynamic nature of the treatment effect as it relates to the changing political environment and explore the mechanisms for the observed reduction in emissions.

ESSAYS ON ISSUES IN PUBLIC AND ENVIRONMENTAL
ECONOMICS

by

Daniel Kraynak

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
2023

Advisory Committee:
Professor Roberton Williams III, Chair/Advisor
Professor Joshua Linn
Assistant Professor Louis Preonas
Assistant Professor James Archsmith
Professor Melissa Kearney

© Copyright by
Daniel Kraynak
2023

Foreword

The first chapter is sole-authored work. The second chapter is jointly authored with Emily Conover and Prakarsh Singh, a pre-publication draft of our article in the *Journal of Development Economics*. The third chapter is jointly authored with Govinda Timilsina and Anna Alberini. The Dissertation Committee acknowledges that Daniel Kraynak made substantial contributions to relevant aspects of both chapters as affirmed by the coauthors on each project.

Dedication

To my family, but especially to my Pops, the other social scientist in the family.

Acknowledgments

There are a lot of people who I should thank for being in my life and helping me arrive at a finished dissertation. This section can highlight at least a few of them.

Starting with the academics, thanks to Rob Williams for being my chair, advisor, mentor, and advocate. Also thanks to Josh Linn for being a great example and for a tremendous amount of invaluable feedback on my work. Thanks to Louis Preonas and Jim Archsmith, who were important young faculty mentors in various respects. Special thanks also to all my coauthors (attempted, successful, and TBD), who are great people. I owe thanks to several professors at Hamilton College too, especially to Emily Conover. She was my first mentor in economics who brought me into the field and has helped me in so many ways.

Keeping it professional, thanks to the folks at Resources for the Future and the Federal Energy Regulatory Commission for my part-time stints while in graduate school. I can't overstate how useful it was to learn from my colleagues (broadly defined) at both places.

More personally, thanks to my Maryland family and friends, including the students at UMD AREC (especially my cohort and surrounding cohorts), my family, and my "Hamilton" people. Thanks for believing in me, for not believing me sometimes, for setting a high standard whether we do or don't see each other a lot, and for always broadening my horizons. Finally, thanks so much to Cata and Moon for all you do and are.

And of course, thanks to the University and the AREC department as a whole for instruction, financial support, and an opportunity.

Table of Contents

Foreword	ii
Dedication	iii
Acknowledgements	iv
Table of Contents	v
List of Tables	viii
List of Figures	ix
1 The Local Economic and Welfare Consequences of Demand Shocks for Coal Country	1
1.1 Introduction	1
1.2 Background and Data	5
1.3 Empirical Strategy	8
1.4 Details of the Coal Demand Instrument	12
1.4.1 Dispatch Models	13
1.4.2 Mapping from Load Regions to Coal Counties	16
1.4.3 Instrument Relevance	17
1.4.4 Instrument Validity	18
1.5 The Impact of Demand Declines on Employment and Wages	21
1.5.1 Benchmarking the Employment Estimates	23
1.6 Impacts Beyond Employment and Wages	24
1.6.1 Population and the Labor Force	25
1.6.2 Public Finance and Public Goods	27
1.6.3 Regional Heterogeneity	28
1.7 Changes in Economic Welfare	29
1.7.1 Capitalization	29
1.7.2 Calculations based on Spatial Equilibrium	31

1.8	Conclusion	39
1.9	Figures for Chapter 1	42
1.10	Tables for Chapter 1	52
2	The Effect of Traffic Cameras on Police Effort: Evidence from India	60
2.1	Introduction	60
2.2	Background	65
2.3	Data	67
2.4	Conceptual Framework	70
2.5	Empirical Strategy	72
2.5.1	Estimation	74
2.5.2	Placebo Tests	75
2.6	Results	77
2.6.1	Mechanisms	79
2.7	Robustness and Additional Results	83
2.8	Conclusion	85
2.9	Figures for Chapter 2	88
2.10	Tables for Chapter 2	93
3	The Effect of Constrained Climate Policies: Evidence from Australia	99
3.1	Introduction	99
3.2	Policy Details	102
3.3	Methods: Synthetic Control and Extensions	104
3.4	Data	106
3.5	Estimation Details	107
3.6	Results	109
3.7	Mechanisms	113
3.8	Exports	114
3.9	Conclusion	115
3.10	Figures for Chapter 3	118
3.11	Tables for Chapter 3	123
A	Appendices for Chapter 1	124
A.1	Supplementary Tables	124
A.2	Data Appendix	125
A.2.1	Energy Market Data	125
A.2.2	Outcomes	125
A.2.3	Other (including controls)	126
B	Appendices for Chapter 2	127
C	Appendices for Chapter 3	128
C.1	Donor Pool	128

C.2	Alternate Placebo-in-time	129
C.3	Alternate Permutation Gaps Plots	130
C.4	Alternate Permutation Distributions	131
C.5	More Detail on Mechanisms	132
Bibliography		133

List of Tables

1.1	Summary statistics	52
1.2	Reduced-Form Effects of Demand Declines by Period	53
1.3	Results on Employment and Wages	54
1.4	Results on Population and Migration	55
1.5	Results on Labor Force Components	55
1.6	Results on Public Finances	56
1.7	Results on Local Public Goods	56
1.8	Heterogeneity Across Regions	57
1.9	Summary of Spatial Equilibrium Effects	58
1.10	Welfare Estimates	59
2.1	Summary Statistics	93
2.2	Effect of CCTV Cameras <i>Prior</i> to the Installation of the Cameras (Pre-Trends)	94
2.3	Effect of CCTV Cameras on Officers <i>On-the-Ground</i> Ticketing Behavior	95
2.4	Violations by Type (Levels)	96
2.5	Violations by Vehicle Type (Levels)	97
2.6	Main Results - Robustness	98
2.7	Main Results by Distance to CCTV Camera	98
3.1	Emissions per Capita Predictors Before Policy Changes	123
3.2	Nonzero Country Weights in Synthetic Australia	123
A.1	Supplemental Summary Statistics	124
B.1	Effect of CCTV Cameras <i>Prior</i> to the Installation of the Cameras (Pre-Trends)	127
C.1	Donor Pool for Australia	128

List of Figures

1.1	Aggregate Fossil Fuel Generation	42
1.2	Coal Production and Powerplant Deliveries	43
1.3	US Coal Miners	43
1.4	Coal Supply Regions	44
1.5	Subregions and Coal County Centroids	45
1.6	Example Supply Stacks	46
1.7	Dispatch Predictions	47
1.8	Market Shares	48
1.9	Predicted Coal Demand vs. Actual Coal Production	49
1.10	Reduced-Form placebo regressions, 2002-2007	50
1.11	Wage Effects	51
2.1	Map of Chandigarh and Location of CCTV Cameras	88
2.2	Location of CCTV Cameras and 300m Buffer Circles	89
2.3	Daily Number of <i>On-the-Ground</i> Violations	90
2.4	Event Study Type Figure Using Months Pre/Post Camera Installation	91
2.5	Event Study Type Figure Using Months Pre/Post Camera Installation on Driver Behavior	91
2.6	Permutation Inference Results	92
3.1	Emissions per Capita in Australia and Donor Pool	118
3.2	Path plot	119
3.3	Leave-one-out Plot	119
3.4	Placebo-in-time Test	120
3.5	Permutation Test Plot	120
3.6	Permutation Test Distribution	121
3.7	Australia's Emissions by Sector	121
3.8	Australia's Electric Generation Capacity	122
3.9	Australia's Coal Production and Exports	122
C.1	Placebo-in-time (2000)	129
C.2	Placebo-in-time (1996)	129
C.3	Placebo-in-space Plot (MSPE ratio threshold: 20)	130
C.4	Placebo-in-space Plot (MSPE ratio threshold: 2)	130

C.5 Placebo-in-space Distribution (MSPE ratio threshold: 20) 131

C.6 Placebo-in-space Distribution (MSPE ratio threshold: 2) 131

C.7 Australia’s Energy Consumption by Source 132

C.8 Australia’s Electricity Generation 132

Chapter 1: The Local Economic and Welfare Consequences of Demand Shocks for Coal Country

1.1 Introduction

In the past 15 years, coal-fired electricity generation in the US has declined dramatically, and approximately 25,000 (34%) of domestic coal mining jobs have been eliminated.¹ This trend is expected to continue, spurring significant interest in policies to limit the welfare impact of structural changes in electricity markets on the coal producing communities that have supplied them for over 100 years (Morris, 2016). For instance, in February 2022 the Biden administration announced more than \$700 million to “catalyze economic revitalization in coal communities”.² This announcement expands on existing policies that have already distributed some \$410 million dollars through various initiatives to applicant counties (Shelton et al, 2021).

The impacts of declining coal demand on coal country are of interest to economists for several reasons. First, they are a prominent distributional consequence of the transition to a

¹Data: <https://fred.stlouisfed.org/series/CEU1021210001>. Difference calculated Jan. 2005 to Jan. 2020.

²<https://www.doi.gov/pressreleases/biden-administration-announces-nearly-725-million-create-good-paying-union-jobs>

lower carbon economy in response to the threats posed by climate change. Second, they are a significant demand-side shock to labor markets in producing regions, which relates to a vast literature that studies trade and natural resource shocks (Autor, Dorn, and Hanson, 2013; Bartik, Currie, Greenstone and Knittel, 2019). Third, in the US and elsewhere, some of the policy response is likely to take the form of place-based policies, such as geographically targeted grants (Furnaro et al. 2021). The success of these policies ultimately depends on whether they address the real impacts of a declining coal industry.

Identifying the effects of demand shocks on coal producers, however, is challenging. Coal producing regions in the US differ from the rest of the country in important ways, making it difficult to construct counterfactuals for economic indicators in these regions. This difficulty stems from the fact that many mining regions in the US were settled from 1870 to 1930 when coal was the dominant source of fuel in the US (Matheis, 2016). These regions developed economies closely tied to the coal industry and have endured many cyclical booms and busts in population, employment, wages, and economic growth, despite coal mining's long-term decline in employment at the national level.

This paper aims to present a more complete and policy-relevant picture of the effects of declining coal production on residents of coal country by exploiting downstream variation from the power sector. I construct an instrument for coal demand from each coal-producing county in the US that maps changes in electricity markets up the supply chain to coal producers, using comprehensive data on contracts between large coal-fired power plants and mines. To isolate aggregate shocks to coal demand from electricity markets that are

exogenous to labor markets in mining counties, I predict demand by constructing regional electricity models using hourly generation data from thermal powerplants.³ The motivation for the empirical strategy is a comparison of producing counties with greater or lesser exposure to recent declines in coal demand from the electricity sector, obviating the need to compare historical producers with the rest of the country.

The instrument is applied to estimate the effect of production declines at mines on a broad set of county-level outcomes in a long-difference, within-coal-country specification using data from 2007-2017. I find that a typical decline in coal production value per capita at the county level reduces mining employment and causes modest employment and wage spillovers across industries, on the order of 30% of the initial impact on the mining sector. The largest effects occur in Appalachia and the West. On the 5-10 year time horizon studied, I find evidence consistent with the idea that many out-of-work coal miners find new jobs and unemployment effects are small. I also provide causal estimates of the effect of demand shocks on several knock-on effects previously undocumented in the literature. I find that local school district expenditures and state aid to general purpose governments decline, and that poverty increases. Motivated by estimated declines in population and immigration, I summarize the causal estimates in the context of a theoretical model of location choice. My estimates of household-level declines in economic welfare average around 5% of initial wages per household (3.2% of initial adjusted gross income per household), and imply a total welfare impact on the order of \$500 million resulting from a \$3.7 billion

³I employ the EPA's Continuous Emissions Monitoring Systems (CEMS) data and model static, least-cost dispatch of generators following a large economic literature dating at least to Borenstein et al. (2002).

decline in thermal coal production (15.4% of baseline production value).

The primary contribution of the paper is the development of credible new estimates of the effect of a recent contraction in coal industry on coal country's residents, perhaps the most broadly publicized distributional consequence of the transition away from coal in the US electricity sector.⁴ My results confirm some of the analysis in Weber (2020), which finds that coal-producing regions have experienced declines in aggregate earnings and employment, while contrasting its finding of nonexistent local spillovers across industries.⁵ My results on public goods also complement those of Metcalf and Wang (2019), who find little evidence for the hypothesis that coal mining layoffs lead to opioid usage. I also provide causal estimates of the phenomena suggested in Morris, Kaufman, and Doshi (2020), who provide descriptive evidence suggesting that there have been declines in local public revenue related to declines in coal production.

Second, by exploring a broader set of outcomes than previous studies, estimating capitalization effects, and explicitly estimating economic welfare changes using a model of spatial equilibrium, this paper contributes a more complete picture of the effect of demand shocks on coal country's residents. In doing so it connects to other literature that has explored broad effects of economic shocks on labor markets in other contexts (Feler and

⁴In using policy-relevant variation from the electricity sector, I also connect labor markets in coal country to a sizeable environmental and industrial organization literature studying the effects of environmental regulations and the natural gas boom on electricity markets (e.g. Johnsen, LaRiviere, Wolff, 2019; Linn and McCormack, 2019; Davis, Holladay and Sims, 2021; Watson, Lange, Linn, 2023).

⁵This finding aligns more closely with literature measuring the effects of demand and price shocks in the 1970s and '80s in Appalachia, which demonstrates that the economies of coal producing areas follow trends in coal demand, with employment and wage spillovers mostly to the locally traded goods sector. The same literature also documents the effects that booms and busts have on migration and educational decision-making, suggesting that non-wage channels are also important in the transmission of demand shocks to the local economy (Black et al. 2005 A,B).

Senses, 2017; Allcott and Keninston, 2018) as well as literature examining place-based policy (Busso, Gregory and Kline; 2013, Kline and Moretti, 2014). Ideally, the analysis presented here is useful in understanding the consequences of shocks to local labor markets more generally through a new application, as well as useful to policymakers as evidence of the real (present and likely future) costs of climate policy and the energy transition to coal communities.

The next section, Section 1.2, motivates the rest of the paper with background on coal markets and the data available to perform the analysis. Section 1.3 outlines the econometric strategy. Section 1.4 describes construction of the instrumental variable in detail. Sections 1.5 and 1.6 present 2SLS estimates of the effects of declining coal production on a variety of outcomes. Section 1.7 ties together the causal estimates of the paper and calculates approximate changes in economic welfare. Section 1.8 concludes.

1.2 Background and Data

From 2002 to 2007, the share of coal-fired electricity generation in the US was relatively constant.⁶ Over the next 10 years it fell to nearly half of its previous market share, reducing demand for coal from mines

Both during and prior to the period of declining coal demand from the power sector, coal deliveries to power plants tracked closely with total production reported from mines

⁶Cicala (2015) argues that a flurry of state-level electricity market restructuring came to a halt in the wake of the California electricity crisis of 2000-2001, suggesting that state-level regulation during this period was also relatively stable.

as demonstrated by Figure 1.1. The figure shows mine-reported production from the Mine Safety and Health Administration (MSHA) and Energy Information Administration (EIA) form 7A alongside plant-reported coal deliveries from a combination of EIA 423 and EIA 923. Both data series are reported in short tons and benchmarked against EIA published aggregates. All of the data sources show a significant decline from the early and mid-2000s to the later 2010s.⁷ Coal mining employment, also reported in the MSHA data and EIA form 7A, follows a similar pattern over time as mine production and deliveries to power plants (Figure 1.3). These simple figures are highly suggestive that labor demand and coal production at mines are responsive to coal consumption by power plants, which links aggregate power plant demand to local labor markets.

The core empirical task of the paper is to credibly estimate the effects the trends outlined above on labor markets in coal country. My measures of local labor market outcomes come from several publicly available sources of data. I draw data on wages, employment, and unemployment from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages and Local Area Unemployment Statistics data series. I also examine effects on local public revenues using data from the Census of Governments. A full Census of Governments only occurs in years ending in 2 and 7, so I construct a harmonized panel of 5-year differences at the county level to maintain a consistent estimation strategy throughout.

Finally, I present evidence of effects on population, migration, home values, and local

⁷Over the longer run, 90-95% of thermal coal produced in the US ends up at a US power plant. However, some care is required in interpreting Figure 1.2, because coal is stored relatively cheaply at both the minemouth and onsite at power plants, so not all production and delivery is contemporaneous even at the annual level.

public goods. These measures come from a variety of sources which are outlined in detail in the data appendix.⁸

The estimation sample includes all counties with a coal mine recorded in the MSHA data as of 2000, 255 counties. One dominant producer and extreme outlier in terms of productivity (Campbell County, Wyoming of the Powder River Basin) is omitted from all estimates, as it accounts for around 35% of coal production by short tons for most years in the sample - a figure which remains stable from 2007-2017. Summary statistics for the estimation sample in 2007 are shown in Table 1.1. On average, coal counties are relatively rural, with an average population of around 62,000. Though the share of coal miners in the workforce averaged around 3%, the maximum in the sample is more than 40%, and coal counties produced an average of \$94.3 million of thermal coal in 2007. Notably, the average mining sector job in coal country had about 20% higher wages than the average across all jobs.⁹

One of the key contributions of the paper is a new estimation strategy that more credibly identify the effects of demand shocks to the coal industry on coal communities, described in further detail in the Empirical Strategy section. Here I briefly outline the data sources I use to instrument for coal demand from the electricity sector. Data on electricity generation by fossil fuel plants comes from the Environmental Protection Agency's (EPA) Continuous Emissions Monitoring Systems (CEMS) data, which contains hourly gross generation and

⁸Population data comes from the US Census and migration from the IRS county-to-county migration files. Several indicators (such as educational attainment and median home values) are drawn from the ACS. The data appendix also contains further detail on data processing and lists other data sources used to construct control variables described with the estimation strategy.

⁹Appendix A.1 shows statistics on variables relevant to the study but not covered in Table 1.1.

emissions for most coal-fired boilers in the US. I use the EPA’s eGRID data as my source of metadata on boilers and power plants. The EIA 423 and EIA 923 data for power plants contains information on coal characteristics and origin reported by power plants. These data are used to infer a long run average heat content of coal from each coal producing county, and to link between coal mines and the coal-fired power plants they supply. I value the coal using state-level fuel prices from the EIA’s State Energy Data System (SEDS).¹⁰ Production at the mine-mouth and employment of coal miners is reported by the MSHA and EIA Form 7.

1.3 Empirical Strategy

The empirical strategy is designed to estimate the effect of recent declines in coal production by mines on coal producing communities. Using nearly all¹¹ of the 255 counties with a coal mine present in MSHA data since 2000, I start by estimating the coefficients of models with the form of Equation (1.1). Counties are indexed by c and times are indexed by $t \in \{2002, 2007, 2012, 2017\}$, while the difference operator Δ indicates a 5-year difference

¹⁰While the natural physical unit of measurement in the EPA CEMS and EIA 423 and 923 data is Btus of coal, the MSHA/EIA supply-side data is measured in short tons, meaning that conversion to dollar amounts for continuity with the coal demand measure is not trivial. To do so, I use the EIA 423 and 923 data to infer a long-run average Btu content of coal produced in each county. Where sufficient data is not available, I impute the average of the five nearest neighbor counties (based on centroid distance), allowing me to use state-level prices from the SEDS database to convert to dollar amounts.

¹¹As previously mentioned, Campbell County, WY is excluded from all estimation due to its status as a very large producer and an extreme outlier. Estimation results are sensitive to its inclusion.

between time t and $t - 5$.¹²

$$\frac{\Delta Y_{ct}}{Pop_{c,t-1}} = \beta_1 \left[\frac{-\Delta D_{ct}}{Pop_{c,t-1}} \right] + X'_{ct} \beta_2 + \gamma_t + \varepsilon_{ct} \quad (1.1)$$

The variable ΔY_{ct} is a county-level change in local outcomes (such as employment in different industries), and is normalized by start-of-period t population estimates from the US Census. The key independent variable is ΔD_{ct} , the change in total value of coal produced at mines in county c , similarly normalized. This variable captures both declining production (for which miners might be laid off) and mine closures (where the number of miners goes to 0). Production values are smoothed with two lags at the start of each period (e.g. $t = 2007$ is averaged with 2006 and 2005) to reduce measurement error due to coal storage both at mines and plants, as well as frictions that might slow labor market adjustment. The coefficient of interest is β_1 , which measures the relationship between changes in coal production value per capita and county outcomes.

I use Equation (1.1) to estimate labor market outcomes for which the object of interest is an aggregate change within the county, such as the total number of jobs lost or the total decrease in wage earnings by sector. The model is estimated using two stacked five-year differences, 2007-2012 and 2012-2017, which breaks the ten-year period of rapidly declining coal dispatch into an earlier period with small impacts (in which coal production and

¹²I use 5-year differences ending in years 2 and 7 due to the data restrictions imposed by the Census of Governments. They also align relatively conveniently with trends in coal demand - relative stability from 2002-2007, declining dispatch but with some modestly increasing production from 2007-2012, and rapid declines from 2012-2017.

employment actually increased in many counties) and a later period in which production and employment in mining rapidly declined. The period fixed effect (γ_t) is allowed to vary at the regional level to control for transportation cost differences across regions and other cultural and economic differences that might independently drive trends in local outcomes, with coal-producing regions adapted from Stoker et al. (2005) and depicted in Figure 1.4.

The unit of analysis is the county, but estimation is structured around the commuting zone to acknowledge the important economic interactions between counties within a commuting zone (Tolbert and Sizer, 1996). Standard errors are clustered at the commuting zone level, and the X_{ct} vector contains several relevant controls measured at the start-of-period t computed at the commuting zone level (including non-producing counties) that might independently impact mining employment, wages, and local labor market conditions more generally. Specifically, I include the baseline share of mining employment in the commuting zone, the share of manufacturing employment, share of employment in the transportation and utility sectors, and share of the population with a bachelor's degree. Additionally, I include the baseline value of oil and gas production calculated using Enverus and EIA data (aggregated to the commuting zone level) to mitigate concerns that co-location of oil and gas production with coal reserves influences the estimates.

I also use Equation (1.2), which shares almost the exact same structure as Equation (1.1) with one key difference - the dependent variable is typically a log change, and any normalization by population or employment is *contemporaneous*. While this specification has the disadvantage of omitting zero-valued outcomes, it has the advantage of directly estimating

an approximate proportional change in the outcome variable. This is useful to estimate changes in quantities that have most meaning with contemporaneous normalizations, such as average wages (per job) or average public expenditures per capita.

$$\Delta \log(Y_{ct}) = \beta_1 \left[\frac{-\Delta D_{ct}}{Pop_{c,t-1}} \right] + X'_{ct} \beta_2 + \gamma_t + \varepsilon_{ct} \quad (1.2)$$

Coefficients β_1 from Equations (1.1) and (1.2) might have a causal interpretation if one is willing to assume that changes in coal production value are exogenous or as-good-as random with respect to the characteristics of the county of production. It is not difficult to imagine channels through which this assumption is potentially violated. Some mines may face a secular decline in productivity that independently drives lower production and lower wages for miners. Some commuting zones may have better amenities or tighter labor markets that affect the ability of mines to compete on output price and decrease their production. Put another way, if local economic conditions impact mine output, then ordinary least squares estimates based on Equations (1.1) and (1.2) suffer from reverse causality. To overcome these issues, I instrument for changes in coal production value using predicted demand from a model of the electricity sector mapped back to coal counties, described in greater detail in the next section.

1.4 Details of the Coal Demand Instrument

The instrumental variables strategy of the paper is to identify the proportion of declining coal production at mines that is driven by factors reducing generation by coal-fired power plants in electricity markets (most importantly, declining natural gas prices). To parsimoniously summarize the influence of these factors, I estimate yearly variation in aggregate coal demand from power plants using simplified models of electricity generation at the regional level covering the continental US.¹³ Long run changes in demand are mapped to coal producing counties using market shares calculated during the period of relative stability from 2002-2007. The instrument thus follows a shift-share logic, where market shares are held fixed and are multiplied by exogenous variations in demand, yielding an instrument for coal demand from each coal mining county in the US.¹⁴

There are two key components to the instrument - the dispatch models and the mapping between coal-fired power plants and mines. The dispatch models summarize aggregate coal demand from different regions. The inputs to the dispatch models are region-level electricity load, technology mixes, and state-level fuel prices, so identification relies on the assumption that those inputs are sufficiently “large” relative to mining counties to be treated as exogenous. The supply relationships between power plants and coal mines provide a mapping between labor demand in coal producing counties and coal demand in the power

¹³Specifically, I use dispatch models based on the EPA’s eGRID Subregions (discussed further in the next section).

¹⁴In this case however, there are many aggregate shocks and the shares that define exposure are not local employment ratios, but county-to-market shares of coal.

sector. Supply relationships are held constant prior to the estimation period from 2007-2017 to eliminate the potential influence of endogenous contract changes between plants and mines during the estimation period from 2007-2017. By construction then, mining counties are exposed to demand shocks only through an immutable characteristic - past relationships to prior customers. This section describes construction of the instrument in more detail - the dispatch models, the shares, and the formal definition of the instrument; before moving on to direct and indirect tests of its validity.

1.4.1 Dispatch Models

The dispatch models are least-cost economic dispatch models. Generating units are dispatched according to their costs, a very simple approximation of actual electricity market operation in the US.¹⁵ Low-cost units operate when total load (electricity usage) is low and higher-cost units operate only at high levels of load. The model closely follows the approach of Johnsen, LaRiviere and Wolff (2019) or JLW, where supply curves are constructed for each electricity load region and year and dispatch is run hourly using observed state-level fuel prices and region-level demand. Inputs are the EPA's CEMS data and the EIA's state-level fuel prices (from SEDS).

The main departure of this paper's dispatch model from that of JLW is the use of the EPA's eGRID Subregions (polygons depicted in Figure 1.5) rather than North American Electric Reliability Corporation (NERC) regions. EPA eGRID subregions divide the con-

¹⁵This is sometimes achieved through a coordinated market mechanism as in the ISO/RTO regions, and sometimes by regulated utility companies managing wholesale generation.

tinental US into 20 electricity load regions within which electricity flows relatively freely, rather than 8 as with NERC. This choice was motivated to balance the goal of choosing the smallest electricity market definition possible (increasing predictive power of the instrument) with ensuring that each coal-producing county only enjoyed a small market share in terms of fuel delivered to market participants (strengthening the argument for exogeneity).¹⁶ Market shares of mining counties and the instrument's predictive power are discussed in the following subsections.

Within a market and year, a model run consists of 4 main steps. First, using only CEMS data, an average heat rate for each thermal generating unit is calculated by averaging heat input in millions of British thermal units (mmBtu) per megawatt-hour (MWh) of electricity generation. Second, unit capacity is calculated using the observed maximums of generation on days in which the unit produced for more than 1.5 hours. Third, a supply curve is constructed by multiplying unit heat rates by state-level fuel input prices, yielding an estimated marginal cost for each unit. Ranking units by marginal cost and calculating a running sum of observed capacities yields the supply curve. Fourth, for each hour of the year, the supply curve is intersected with total observed generation by fossil fuel plants in the load region, and the lowest cost plants required to meet this demand are “dispatched” for the full hour.¹⁷

¹⁶In reality, neither NERC regions nor EPA eGRID Subregions are the actual boundaries at which electricity dispatch occurs, which are known as balancing authority areas (BAAs). In most cases BAAs are contained within eGRID subregions. However, because electricity is traded across BAAs, EPA subregions, and even NERC regions, a model that simulates least-cost dispatch using NERCs or subregions as an approximate market definition is reasonably capable of predicting true generation by thermal plants (albeit with significant prediction errors if transmission and dynamic constraints are not modelled).

¹⁷Though no data similar to CEMS is publicly available for non-CEMS generators, this method approximates electricity supply assuming that non-CEMS generators are price insensitive. This not always, but often

To illustrate, Figure 1.6 shows an example of the supply curve modelled for a market in the southern US (SRMV, a subregion of SERC) in years 2008 and 2016. The gray shape behind each estimated supply curve shows the (arbitrarily rescaled) density of hourly fossil load in the region. In other words, the height of the gray curve is proportional to the number of hours that total regional load was equal to the corresponding value on the x axis. In 2008, the lowest cost units were all coal-fired power plants. By 2016 there is almost complete overlap in the marginal cost of gas units and coal units, and an oil-fired “peaker” plant also appears. Note also the overall downward shift in the cost curve for the same quantity of generation, which is almost entirely due to declining natural gas prices which shifted them lower in the merit order.

The dispatch model yields a value of predicted generation for nearly the universe of coal-fired generators in the US, derived from the interaction of generation technologies, market-level load, and state-level fuel prices. Figure 1.7 plots predicted versus actual generation for the regions whose supply curves are depicted in Figure 1.6. The model has significant predictive power, which is sufficient for use in the instrument, though it fails to capture dynamic and congestion constraints that shift the true merit order. Annual market-level coal demand is calculated by converting each coal plant’s predicted generation to the mmBtus required to meet its dispatch requirements (via the estimated heat rate), and summing across all coal plants in the market.

the case (i.e. most renewables are non-dispatchable, nuclear plants face very costly constraints in altering output).

1.4.2 Mapping from Load Regions to Coal Counties

To map changes in aggregate demand for coal derived from the electricity sector model to mining counties, I use producing-county-to-load-region shares that do not vary during the estimation period following Watson, Lange, and Linn (2023). For each load region b , I compute the share of coal Btus delivered by each supply county c over the entire period from 2002-2007.¹⁸ Figure 1.8 shows a strip plot that displays the entire distribution of market shares separated by eGRID Subregion. No coal producing county in the estimation sample enjoys a market share in the period from 2002 to 2007 of more than 40% in terms of Btus of coal delivered from county to load region and only 3 counties enjoy a market share of more than 30% in any load region.¹⁹ This fact supports the assumption that EPA's eGRID Subregions are sufficiently large as a power market definition such that no one coal producing county can exert a dominant influence on load region dispatch, which is required for the regional shocks to serve as valid instruments.

Predicted demand for coal deliveries from county c is derived by mapping yearly variation in region-level coal demand during the years of falling dispatch back to previous suppliers. Mathematically, predicted demand for coal from county c in year t , $\widetilde{D}_{ct}$ is given by

¹⁸This calculation relies on EIA 423 and 923 data, which has a plant reporting threshold of 50 or more MW of nameplate capacity, but still contains about 95% of all coal deliveries Figure 1.2.

¹⁹Campbell County, Wyoming is the only exception to this condition, and is excluded from all calculations and estimation other than previous descriptive analysis of trends in national production.

$$\widetilde{D}_{ct} = \sum_b \left[\frac{Deliv_{cb0}}{Deliv_{b0}} \right] \widetilde{Demand}_{bt} \quad (1.3)$$

for $t > 2007$. Reading Equation (1.3) from right to left, for each electricity load region b , annual coal demand $\widetilde{Demand}_{bt}$ is aggregated from the dispatch model in years $t > 2007$. These load region demands are mapped back to supply county c via its share of total deliveries to load region b , $\frac{Deliv_{cb0}}{Deliv_{b0}}$, where $t = 0$ denotes the entire period from 2002-2007. Summing across shocks from the 20 load regions yields the expected demand for coal from county c . All market share calculations are performed in Btus and later converted to dollars using the EIA SEDS data.

1.4.3 Instrument Relevance

The predicted demand instrument is highly correlated with county level production across different specifications and is robust to controls. Figure 1.9 plots the fully transformed IV (5-year changes in predicted coal demand value per capita) against observed changes in coal production value in USD per capita. Each county has two observations corresponding to the periods 2007-2012 and 2012-2017. Figure 1.9a differentiates points by region suggesting that it is highly (but imperfectly) predictive across all coal regions, while Figure 1.9b differentiates points by period showing that the majority of declines in demand and production occurred in the latter period from 2012-2017.

1.4.4 Instrument Validity

Valid IV estimation of the coefficient β_1 in Equations (1.1) and ((1.2)) requires that my measure of coal demand only operates on local labor market outcomes through its direct, contemporaneous impact on the coal industry. This section presents reduced form estimates that explore whether the instrument follows patterns of correlation consistent with this condition.

Table 1.2 shows reduced-form estimates of β_1 from Equation (1.4) below against employment in different sectors in different time periods. Equation (1.4) follows the same structure as Equation (1.1), except commuting zone controls are omitted and outcomes are regressed directly against the instrument. The key regressor ($\widetilde{\Delta D_{ct}}$ per capita) is standardized to have mean 0 and standard deviation 1, so that coefficients measure the effect of a one standard deviation decline in predicted coal demand per capita on local employment across sectors.

$$\frac{\Delta Y_{ct}}{Pop_{c,t-1}} = \beta_1 \left[\frac{-\widetilde{\Delta D_{ct}}}{Pop_{c,t-1}} \right] + \gamma_t + \varepsilon_{ct} \quad (1.4)$$

Panels A-C of Table 1.2 show reduced-form estimates of β_1 from Equation (1.4) by time period over the period of declining coal demand (2007-2017). Panel A shows the results of estimating with both of the later periods in the stacked difference model. The coefficient of -0.53 in column (1) of Panel A means that a one standard deviation exogenous decline

in coal demand is associated a 0.53 percentage point decrease in coal miners per capita in the county. All other coefficients in the table are interpreted similarly. Panels B and C of Figure 1.2 estimate Equation (1.4) separately for the periods 2012 – 2017 (panel B) and 2007 – 2012 (panel C), where $\widetilde{D}_{ct}$ is the *contemporaneous* value of demand declines from power plants. The results in panel B are similar to those in panel A, suggesting that the dominant effect of demand declines occurred in this later half-decade. In panel C the results remain negative but become statistically insignificant.

Panel D of Table 1.2 serves as a placebo exercise of the reduced form specification. In this panel, *past* changes in employment are regressed against the average of *future* changes in predicted demand for coal from the electricity sector. This is done to alleviate concerns that the estimates capture a long-run secular decline in either productivity or general labor market conditions, rather than the contemporaneous effect of declining demand from coal power plants. The key right-hand-side regressor is the average value of long-run demand declines in the half-decades from 2007-2012 and 2012-2017. The results show no evidence of negative correlation between future demand declines and past (2002-2007) trends in employment across sectors. In fact, the correlations are relatively small and positive, which is evidence against such reverse causality and suggests that counties that fared worse from 2007-2017 fared (if anything) slightly better from 2002-2007. In this context, Panel B also alleviates statistical power concerns that might arise in the falsification exercise - there is indeed sufficient variation to identify effects even in a 1-period estimation, though the main results of the paper use stacked differences to incorporate more data and increase precision.

The results in Panel D of Table 1.2 show that the measure of predicted coal demand used as an instrumental variable is uncorrelated with pre-existing trends in local labor markets, which indirectly supports a valid exclusion restriction. That is, no effect is found prior to the demand shocks that the IV is meant to measure. While this is comforting for direct estimates on employment and wages, it is possible that supply counties that would later suffer demand declines have other pre-existing trends that separate them from other coal counties.

To address the concern that the instrument may capture spurious trends in locations that would later suffer demand shocks, Figure 1.10 show coefficients from regressions similar to panel A of Table 1.2 that test this assumption for several outcomes relevant later in the paper.²⁰

That is, Figure 1.10 shows regressions of past trends (2002-2007) against future declines in coal demand. While technically Figure 1.10 does not show strong statistical evidence of correlation between future demand declines and trends in total employment and wages, labor force participation, home values, migration, or local government revenue from 2002-2007, some of the point estimates are large and 95% confidence intervals are wide (a coefficient of 0.1 indicates a 10% difference). Comfortingly, Figure 1.10b suggests that adjusting for start-of-period commuting zone controls does a reasonable job of adjusting for pre-period differences in these variables, moving the largest coefficients close

²⁰The specification for these regressions is a reduced-form version of Equation (1.2) given below:

$$\Delta \log(Y_{ct}) = \beta_1 \left[\frac{-\Delta \widetilde{D}_{ct}}{Pop_{c,t-1}} \right] + \gamma_t + \varepsilon_{ct}$$

to zero. Eight of 10 coefficients are indistinguishable from zero (and more generally close to zero relative to the tails of their 95% confidence interval) bolstering the argument for a conditional-on-observables interpretation of the instrumental variables results in the rest of the paper.

1.5 The Impact of Demand Declines on Employment and Wages

Table 1.3 presents instrumental variables estimates of β_1 from Equation (1.1) where the dependent variable is employment or total earnings per capita in coal mining or other sectors within the county according to BLS definitions. The key independent variable ΔD_{ct} per capita is the 5-year change in the total value of coal produced by mines per capita according to MSHA and EIA data, a supply-side measure of production. It is instrumented with $\widetilde{\Delta D_{ct}}$, changes in predicted demand for coal by power plants per capita as defined in Equation (1.3). For ease of interpretation, total production value per capita is standardized to have mean 0 and standard deviation 1. The coefficients measure the effect of a one standard deviation decline in production value per capita, or about \$4,800 per capita (Table 1.1b) on local outcomes.

The first two columns in Panel A of Table 1.3 show that a one standard deviation decline in coal production per capita is expected to reduce the number of miners per capita by 0.94 percentage points as measured by the MSHA, or 1.16 percentage points by the BLS measure of mining employment. The MSHA data is filtered to include only employees classified as miners at coal mines and so is potentially a more conservative measure. Col-

umn (3) of Panel A suggests that there are indeed employment spillover effects on the order of 39% according to BLS employment categories. Columns (4) and (5) show negative but statistically insignificant point estimates on manufacturing employment and employment in non-tradeable (or locally consumed) sectors. Column (6) clarifies that most of the effects on the non-tradeable goods sectors appear to be dominated by sectors directly in business with coal mining such as utilities and transportation. These findings are most consistent with the idea that coal mining mostly affects the local economy through a direct business-to-business demand channel.

Panel B tells a similar story, where in column (2) the results imply that a decrease of \$4,800 per capita of coal production decreased total mining wages by \$1,048 per capita and total wages by \$1,349, an estimated spillover of roughly 29%. Column (4) shows insignificant results across Panels A and B on manufacturing employment and wages, while non-tradeables show a decline of \$250, about 70-80% of which is driven by adjacent industries, in particular those that directly transport or use coal. Results in column (5) are somewhat nuanced. While the evidence for an employment effect in non-tradeables is statistically weaker than the wage effect, Panel B suggests that coal mining also affects the economy through the demand channel of mining wages, though it is not large relative the impact on the utilities and transport sector.²¹

²¹Figure 1.11a displays these coefficients visually. Figure 1.11b shows complementary evidence, estimates of changes in log wages per job estimated with Equation (1.2). The patterns are similar across figures, but suggest that the decline in total wage earnings by miners is a result of fewer workers rather than lower wages per worker, while in other sectors it appears to be both fewer workers and lower wages per job.

1.5.1 Benchmarking the Employment Estimates

One way to understand the economic importance of the employment effects is to benchmark them using the overall decline in coal production and employment during the sample period. Put another way, one can ask: How much of the observed decline in mining employment do the estimates explain? An important assumption greatly simplifies this calculation - that the estimates measure absolute changes in employment, not solely relative changes, so that a decrease in the number of miners in a county means a similar decrease in the total number of miners nationally. This assumption is supported by the fact that the number of miners nationally was indeed rapidly declining as demand from the power sector declined (Figure 1.3).

The simplest way to perform this exercise is to read the coefficient in Panel A of Table 1.3 as the effect of a \$4,808 exogenous decline in coal production value on mining employment. This is also an appropriate interpretation given that the normalization by population is lagged on both sides of the equation. Normalizing the total reduction in thermal coal production value from coal country between 2007 and 2017 (3.7 billion dollars) by \$4,808 and multiplying by the estimated effect size (0.0094) yields an estimate of about 7,234 miners, or approximately 41% of the observed decline (17,678 miners).²² Replacing the

²²Equivalently, one can use the original interpretation of the coefficients as follows, though the calculation is messier. The estimated effect on MSHA-measured coal miners in column (1), Panel A of Table 3 suggests that a \$4,808 decline in coal production value per capita at mines reduces coal mining employment per capita by 0.94 percentage points. Normalizing the 3.7 billion dollar decline in production value by the 2007 population of coal country yields \$232 per capita. Normalizing this figure by 4808, multiplying by the estimated coefficient (0.0094) yields a similar 41% when compared to 17,678 miners (also normalized by the 2007 population of coal country).

base year (2007) with the year of maximal production value and employment in the sample (2011), the change in production value increases to over \$10 billion, the change in mining employment to over 28,000, and the portion of explained variation to just over 80%.

Aside from the fact that meaningful proportions of US-produced coal are exported and used for metallurgical purposes, it is intuitively reasonable that demand declines from the power sector may not explain a significant portion or even a majority of recent employment declines at coal mines. Secular productivity declines at mines, compliance or other cost pressures, and substitution of labor for capital in the industry more broadly may also be behind these trends. Nonetheless, the back-of-the-envelope calculation presented here suggests that declining power plant demand is an important, if not dominant factor in recent declines in coal mining employment. This fact in-and-of itself is important for understanding the broader impacts of environmental policies (or other policies) that advantage some electricity generation technologies over others.²³

1.6 Impacts Beyond Employment and Wages

This section looks beyond direct labor market impacts in the coal mining sector and even spillovers to local industries via demand channels. The main reason to extend the analysis in this way is that estimated effects on wages and employment may not adequately characterize even the *economic* effects of demand shocks on coal communities. For instance, out-of-work coal miners have several margins on which they can respond to a job

²³For a detailed exploration of the causes of recent changes in coal dispatch and retirement, see Linn and McCormack (2019), among others.

loss, such as remaining unemployed, leaving the labor force through retirement, or migrating. Moreover, indirect evidence on the price of locally consumed goods such as housing can provide evidence on the magnitude of the demand spillovers caused by lower employment. Finally, the effect of large demand shocks can have ambiguous effects on local amenities and publicly-provided services. Reduced tax revenue may cause declines in public service levels, but outmigration can simultaneously reduce congestion pressures on local public goods.

1.6.1 Population and the Labor Force

Table 1.4 shows the results of IV estimation using the full 2-period differenced model with population (from the Census) and migration (inferred by the IRS from tax receipts) as the outcome variables. Panel A presents estimates from Equation (1.1) with the dependent variable normalized by lagged population, and panel B from Equation (1.2), with the dependent variable log-differenced. Coefficients report the estimated effect of an exogenous one standard deviation decline in coal production value per capita in the county. Estimates in columns (1) and (2) suggest that reduced coal production may accelerate population decline, though perhaps in an unexpected way. In column (1) of Panel B, total population declined by approximately 0.9% following a decline in production value relative to a counterfactual in which the decline did not occur. In columns (2) and (3) of panel B, IRS county-to-county migration files show a 5.5% decrease in *in-migration* for each standard deviation decline in coal production value per capita, with a very small (statistically zero)

estimated increase in outmigration. The statistical strength of these results are robust to changing the specification as shown in Panel A, where a one standard deviation decline in coal production value decreases in-migration by 0.2 persons per capita and population growth by 0.89 percent.

Table 1.5 presents further evidence on the activities of the working age population in a similar format to Table 1.4, with estimates from Equation (1.1) in panel A and from Equation (1.2) in panel B. Columns (1) and (2) reproduce results on employment with slightly different aggregation, with panel B demonstrating that mining employment decreased almost 14% while nonmining employment decreased only 1.8%. Columns (3) and (4) show a somewhat surprising result - very small and insignificant effects of declines in coal production value on both the number of individuals not in the labor force (NILF) and the number of unemployed individuals.²⁴ Both estimates show small but noisily estimated increases. Combined with the results of Table 4, these estimates suggest that laid-off coal miners may re-integrate into the labor force quickly. Column (5) provides some evidence on the economic meaning of shocks to residents of coal country, suggesting that the number of poor individuals increased by an average of 3.6% in response to a one standard deviation decline in coal production value per capita in the county.

²⁴I estimate the number not in the labor force (NILF) by subtracting the BLS labor force estimates from census estimates of population from ages 15-64 in each county.

1.6.2 Public Finance and Public Goods

Table 1.6 examines local public finances, split into general purpose (GP) expenditures and revenues in Panel A and school district (SD) expenditures and revenues in panel B. General purpose governments include counties as well as municipalities and township governments with corresponding boundaries according to the Census of Governments. Panel A shows that expenditures and revenues collected by local general purpose governments show statistically insignificant changes, while state support shows a striking 20.2% decrease. This is likely due to the distribution of severance tax revenues collected at the state level, which are generally allocated by formula. Thus as production declines, state support declines. This result certainly begs the question of whether this is sound policy, because state aid is reduced when its marginal value may be relatively high. On the other hand, general purpose governments appear positioned to cushion this blow to total revenue, perhaps by raising taxes, consistent with the positive point estimate in column (3). Total school district revenues and expenditures on the other hand declined by a substantial 4.4% and 4.1%, respectively. This result is supported by the even larger proportional decrease estimated in local revenues from property taxes per standard deviation decline in coal production value per capita, because school district revenue is generally closely tied to property taxes.

Table 1.7 explores effects on two public goods commonly studied in the economics literature for which data is available, crime rates and education. Columns (1) and (2) show no measureable effect on either violent or property crime rates per 1000 residents. Both point estimates are negative, but have large standard errors. Column (3) suggests that an

exogenous decrease in coal production value per capita may substantially increase teacher to student ratios, but the estimate is too noisy to make a definitive statement about the result, while column (4) suggests that the number of teachers declines by 2.2%, consistent with lower school funding. The overall message from Table 1.7 is that traditional measures of effects on public goods or “local amenities” show relatively weak evidence of positive or negative effects over the time horizon studied.

1.6.3 Regional Heterogeneity

Table 1.8 examines effect heterogeneity at the regional level. Each row contains 4 coefficients with estimates of the effect of a one standard deviation decline in coal production value per capita on changes in the outcome listed on the left hand side of the table. The table is split into three main categories - the labor market, home values and migration, and public finance. In general Appalachia appears have the largest effect sizes across most categories (and the most precise estimates due to sample size). Several patterns emerge from the labor market results in Panel A. First, the largest and most precisely estimated employment effects in the coal industry are in the Appalachian/East and Illinois basin, while the largest point estimates on wages are in Appalachia and the West. Second, significant impacts on the total labor force show up in the Western region, but the largest effects occur in Appalachia. Interestingly, home value declines appear to be proportionally largest in the West and reductions in in-migration are comparable between Appalachia and the West. A few unexpected results emerge in the migration categories, where the Illinois basin expe-

rienced increased in-migration correlated with declining production values and the Interior region counties apparently experiencing a very large decrease in both in-migration and out-migration. In panel C, public revenue impacts are again driven by school district revenue declines both in Appalachia and the West, with the largest (in percentage terms) occurring in the West.

1.7 Changes in Economic Welfare

This paper estimates the effect of demand-driven declines in the produced value of coal on a variety of outcomes in coal communities. The primary effects are declines in employment and wages, a decline in home values and local public revenues and expenditures, and declining in-migration. While these effects can be summarized and aggregated across different regions or characteristics of individuals, it remains unclear what the relative valuation of such changes mean to individuals in the economy, and therefore difficult to make statements about the welfare of the residents of coal country. This section offers several calculations to tie together the key estimates in the paper.

1.7.1 Capitalization

One way to understand the effects in the paper is to consider their incidence across in- and out-of-county wage earners, home owners and renters, and households with and without children. An in-county wage earner that owns a home in a household with children bears the incidence of the three largest estimated effects: a decline in wages, a decline in

home equity, and a decline in education spending per capita. Households with no children, renters, and non-resident landlords or workers who commute out of the county avoid some of the incidence of these shocks.

Table 1.9 summarizes these effects with estimates following the standard format. Column (1) shows that a standard deviation decrease in coal production per capita is predicted to reduce average wages by 4.2%, while columns (2) and (3) show that a one standard deviation production decline per capita is capitalized into a 3% decline in median home values, while rents decline to a lesser extent, a statistically insignificant 1.3%. Column (4) replicates the result from Table 1.6 showing a 4.1% decline in school district expenditures per standard deviation decline in coal production value per capita.

In a strict interpretation of a Roback (1982) style spatial equilibrium model with inelastic housing supply and elastic migration, the object of most interest is the change in property values, which incorporates all changes in the willingness-to-pay of individuals to live in a particular location. A ballpark estimate of the aggregate effect on home values might use the median as a conservative estimate of mean combined with the \$3.7 billion dollar decline in coal production across coal country, similar to the benchmarking exercise in Section 5. Normalizing by population, scaling to the effect of one standard deviation, and multiplying by the total number of “households” (identified by tax returns in Appendix A.1) yields an estimate of about \$1.03 billion in home value declines. It is likely correct to think of the observed change in school district expenditures on public education as a knock on effect of this change, and it is important to note that some of this decline is likely a

transfer from property owners to renters.

1.7.2 Calculations based on Spatial Equilibrium

To complement the capitalization effects described above, in this section I examine a model that focuses on the choice of mobile households in the rental market. Combining the insights of the model with my causal estimates and parameters from the literature, the analysis provides a summary evaluation of the effects of declining coal demand on economic welfare in coal country.

The model is a representation of location choice by homogeneous agents that value real wages and public goods or amenities. Using the concept of spatial equilibrium, it suggests a mapping from a proposed utility function to observable migration decisions, and simplifying assumptions imply that both the utility model and the model of migration decisions contain the same parameters (for example, higher wages in county c both increase utility and encourage migration to the county). The intuition is that migration responses reveal the value of marginal changes in the bundle of wage and amenity characteristics of different locations. Drawing on key parameters estimated in the literature that link the characteristics of locations (wages, rents, and amenities) to long-run migration decisions allows me to combine the observed changes in wages with money-metric valuations of changes in non-wage housing costs and school expenditures in the context of demand shocks to coal country.

The model characterizing location choice draws mostly on Bartik, Greenstone, Currie,

Knittel (2019), which uses a similar approach to estimate the willingness-to-pay of individuals and communities to permit hydraulic fracturing. My derivation is quite similar to theirs, and also draws on Kline and Moretti's (2014) analysis of place-based policy. Parameter estimates of migration elasticities with respect to wages, housing costs, and a vector of amenities (including education expenditures) are drawn from the structural modelling of Diamond (2016).

The model economy consists of many homogeneous individuals $1, \dots, i, \dots, N$ who value wages, rental rates, and local public goods and receive an individual-county specific preference shock. Individuals choose a region c to maximize their indirect utility U_{ic} given by Equation (1.5), where w_c , r_c , and G_c are the wage rate, cost of housing, and local amenities in c . While their choice of region is endogenous, they take the characteristics of each region as exogenous.

$$\ln U_{ic} = \alpha \ln w_c + \beta \ln r_c + \gamma \ln G_c + \varepsilon_{ic} \quad (1.5)$$

Conditional on location choice, individuals inelastically supply one unit of labor. Consider the comparison that individual i makes between region c and another location k . The individual chooses region c if Inequality (1.6) holds.

$$\alpha [\ln w_c - \ln w_k] + \beta [\ln r_c - \ln r_k] + \gamma [\ln G_c - \ln G_k] + [\varepsilon_{ic} - \varepsilon_{ik}] > 0 \quad (1.6)$$

Intuitively, Inequality (1.6) holds when an individual's idiosyncratic preference for region c

exceeds the utility premium the individual would receive from living in k . If the differences $\varepsilon_{ic} - \varepsilon_{ik}$ are draws from an exponential distribution with shape parameter one, then the log share of individuals residing in region c is approximated by Equation (1.7).

$$\ln\left(\frac{N_c}{N}\right) = \alpha[\ln w_c - \ln w_k] + \beta[\ln r_c - \ln r_k] + \gamma[\ln G_c - \ln G_k] \quad (1.7)$$

An elastic housing supply function given by Equation (1.8) closes the model and shrinks the range of parameter values that yield corner solutions.²⁵ Equilibrium is achieved when all agents choose their preferred location, and the housing market adjusts.

$$\ln r_j = \eta_{1j} + \eta_{2j} N_j, \quad j \in \{c, k\} \quad (1.8)$$

To reach a usable expression for changes in willingness to pay for the bundle of characteristics of region c , assume that region c is small relative to location k ²⁶ and take the derivative of $\ln(\frac{N_c}{N})$ from Equation (1.7) with respect to $\ln w_c$. This leads to Equation (1.9), in which α in the utility model is equal to the elasticity of (N_c/N) with respect to the wage.

Analogously, β and γ are also equal to the elasticities of population share in region c to

²⁵While the distributional assumption on $\varepsilon_{ic} - \varepsilon_{ik}$ is analytically convenient, it indicates that all agents have a “relative” unobserved preference for region c . Endogenizing rents creates a feedback between population shares and home prices captured by the semi-elasticity η_{2j} that pushes the desirability of locations towards equality.

²⁶The assumption of the “smallness” of location c follows Bartik et al. and implies no general equilibrium effects - e.g. marginal changes in prices in region c do not affect prices in k . It can be thought of as letting k represent “everywhere else” and greatly simplifies the relationship between the migration function and the utility function.

rents and local amenities in c .

$$\frac{d\ln(N_c/N)}{d\ln w_c} = \alpha \quad (1.9)$$

Next, let wages, rents, and public good expenditure each be a function of a latent index θ that may represent global demand for a good produced only in region c . Taking the derivative of U_{ic} for a resident of location c with respect to θ and dividing by α yields Equation (1.10), a money-metric expression for the change in indirect welfare from residing in location c (equivalently, the change in willingness to pay for changes in the bundle of characteristics associated with location c).

$$\frac{1}{\alpha} \frac{d\ln U_{ic}}{d\theta} = \frac{d\ln w_c}{d\theta} + \frac{\beta}{\alpha} \frac{d\ln r_c}{d\theta} + \frac{\gamma}{\alpha} \frac{d\ln G_c}{d\theta} \quad (1.10)$$

The value of expressions (1.9) and (1.10) is that they relate elasticities of migration with respect to the characteristics of region c to a utility function in a simple economic model of location decisions. For a marginal change in the latent index θ , migration decisions reveal effects on individuals' willingness to pay for the bundle of wages and amenities they enjoy in location c , and hence reveal changes in the economic welfare of its residents. Appealing to the envelope theorem argument in Busso, Gregory, and Kline (2013) suggests that to assess the aggregate effect of a change in θ on the model economy, we need only sum expression (1.10) over individuals in region c . Intuitively, this is because individuals just on the margin between locations are indifferent between alternatives, implying that the

relevant welfare effects of the shock to θ are only the price changes it induces, almost as though the population were immobile.

I combine my estimates of changes in wages, housing costs, public good expenditures with parameters from Diamond (2016), who estimated 10-year migration elasticities in terms of wages, housing costs, and local amenities (including public education expenditures), to approximate economic welfare changes from demand shocks to coal country. I estimate the total wage-equivalent welfare impact of observed declines in coal production value using the expression given by equation (1.11),

$$\frac{1}{\hat{\alpha}} \frac{dU}{d\theta} = N_c * \left(\frac{d \ln w_c}{d\theta} + \frac{\hat{\beta}}{\hat{\alpha}} \frac{d \ln r_c}{d\theta} + \frac{\hat{\gamma}}{\hat{\alpha}} \frac{d \ln G_c}{d\theta} \right) * \overline{w_{c0}} \quad (1.11)$$

where $\overline{w_{c0}}$ is the average baseline wage, which monetizes the proportional changes implied by Equation (1.10) following Bartik et al. (2019). The set of estimated of parameter values are drawn from Diamond (2016), one set of migration elasticities estimated for college-educated workers and one for noncollege workers. An unweighted average of the parameter vector is $(\alpha, \beta, \gamma) = (4.1185, -2.5515, 0.7045)$.²⁷ Normalizing β and γ by α implies that housing cost changes are valued at approximately 60% of wage changes and amenity changes are valued at roughly 20%. In practice this approach allows me to create a coherent, combined valuation of the three largest components of welfare changes revealed by the causal estimates - the effect on wages, housing prices, and public expenditures on

²⁷These parameters are taken from Panel A of Table 5, and reflect migration elasticities estimated under the least restrictive possible model. Noncollege-worker parameter estimates on wages, housing costs, and Diamond's amenities index are $(3.261, -2.944, 0.771)$ and college-worker parameters are $(4.976, -2.159, 0.638)$.

education.²⁸

Table 1.10 presents the results of applying the parameters described above to Equation (1.11) to estimate welfare effects on the average household in coal country and to estimate aggregate welfare effects on all of coal country. Column (1) monetizes the welfare effect of the shock estimated throughout the paper - a one standard deviation or approximately \$4,800 decline in coal production value per capita. Estimated effects on average wages (4.2%) and public education expenditures (4.1%) are drawn from Table 1.9 (and re-scaled to the household level where appropriate). Each panel of Table 1.10 has 3 rows that explore the sensitivity of the estimates to different assumed effect sizes on housing costs, varying between no change and 3% change based on differing interpretations of the estimates in Table 1.9 (no effect on rental prices, 1.3% effect imprecisely estimated, or 3% coinciding with home values.) The two panels use different sets of migration elasticities, one for college-educated workers and one for non-college workers as presented in the Diamond paper.

Column (2) aggregates the per capita effects using the total decline in coal production per capita and total population of coal country, similar to the employment effect benchmarks in section 5. Column (3) follows Bartik et al. and scales the estimate in column (2)

²⁸Since G_c is typically a *vector* of local amenities (of which education is just one component), if other components of the amenity vector change this affects the overall welfare expression. Equation (1.11) in this case captures the welfare change net of other changes in local public goods. In my context, I find strong evidence only of changing education expenditures, and therefore use Equation (1.11) as a method of valuing changes measured in the empirical section. To the extent that other amenities change, this may increase or decrease total welfare changes, but does not change the economic welfare impact of the measured changes. In applying parameters estimated in the literature, I assume that the elasticity of migration to public expenditures is the same as the elasticity estimated on a bundle of amenities. This likely overstates the effect of school expenditures, which in any event have a relatively minor impact on the welfare estimates. I am not aware of published estimates of migration elasticities with respect to education expenditures.

assuming the observed changes in wages, housing costs, and public good expenditures are constant and remain the same in perpetuity using a 5% discount rate. Using the parameters for non-college workers in Panel A, the estimates of changes in economic welfare in coal country in column (1) range from a decline of \$885 per capita to \$1,859 per capita. Assuming permanent declines in wages, rental rates, and school expenditures yields estimates of \$6.24 to \$13.11 billion. Assuming a larger decline in housing costs significantly offsets some of the estimated welfare losses, while for smaller decreases in housing costs the combined effect of wage declines and reductions in public expenditures make residents worse off. Estimates in the second row of each panel using the estimated decline in rental rates roughly coincide for both sets of parameters, suggesting a total aggregate decline in economic welfare of just above \$500 million.

How can we benchmark the magnitudes in Table 1.10? By the valuation techniques described in the data section, I estimate that US coal country produced just under \$24 billion dollars of thermal coal in 2007 before reaching a maximum value in the sample of more than \$31 billion dollars of thermal coal (in 2011), and \$20.3 billion in 2017. Compared to the decline from 2007 to 2017, the effects in column (2) range from about 8.4% of the value of produced coal at the smallest to 17.8% at the largest. These estimates are on a similar order of magnitude to the Bartik et al. welfare estimates for fracking counties as a proportion of new oil and gas production value.

The model is highly stylized and requires several important caveats. The most important of which is that the model derives changes in willingness-to-pay to live in coal country

from the perspective of marginal decision makers in a mobile base of residents. Changes in wages drive these effects, and they are either mitigated or reinforced by reductions in housing costs or declines in public expenditure, while the model is relatively silent on inframarginal residents (e.g. owners of property). Another caveat is that several assumptions were required to bring the parameters estimated in Diamond (2016) to the present context - it could be the case that some qualitatively different migration parameters would be better suited to the context of coal country. Finally, unobserved changes in amenities may push in either direction, so it remains unclear whether the estimates presented are liberal or conservative measures of *total* welfare changes. It is possible that congestion and environmental amenities improve in the case of coal demand shocks, but it is extremely difficult to observe and measure the vector of all possible amenities. Consequently, the approach taken in this paper is to estimate several important economic effects and attach a model-based welfare interpretation to those effects using parameters from the literature. In defense of this strategy, evidence of major negative spillovers to public goods has so far not appeared, with Metcalf and Wang's (2018) results not supporting an association between mining layoffs and opioid usage, and the results of Table 1.7 showing relatively weak evidence on crime. Environmental effects are a natural area for future research.

The value of the calculations in Table 1.10 is that they offer a coherent summary of changes in economic welfare derived from exogenous changes to the coal sector. The migration elasticities from the literature are important adjustments, and suggest that wages are the main factor driving welfare changes. This seems sensible in the context of coal

country - housing prices and rents are already low relative to the rest of the US (Appendix A.1), and it is likely that congestion effects in local education (and other local public goods) are low, while non-college employment opportunities are relatively highly valued.

1.8 Conclusion

Many economic changes and policies that are welfare-enhancing in the long run carry acute and important costs to certain groups or regions. It was certainly true in the 1990s and early 2000s when US manufacturing faced fierce competition from China as studied by Autor et al. (2013) and Feler and Senses (2017). This paper contributes to the study of local economic shocks by examining US coal communities in a down decade of coal demand driven by competition from alternative generation fuels and environmental policy. I find evidence of substantial labor market reallocation, migration responses, and knock-on effects to home values and public revenues. These findings share some commonalities with other contexts but differ in important ways, suggesting that the lessons from such shocks may not be so easily carried from one context to the next.

To identify the effect of declining coal production on coal country, I track demand shocks from the power sector up the supply chain to producers in coal country. I construct a dispatch model of the electricity sector at the regional level and use predicted changes in coal dispatch as exogenous variation to predict shocks to the coal mines that previously supplied them. Applying this demand measurement as an instrument for production value by mines, I find that mining employment and wages make up the bulk of the direct la-

bor market effects, but that total effects are on the order of 30% larger, with spillovers mostly occurring in directly connected industries such as transportation. I estimate that state support for general purpose governments decline, and that public school revenues and expenditures fall (4.4 and 4.1% respectively), driven by a significant decline in home values (3% at the median). Correspondingly, I show that in-migration and medium-run trends in population are significantly impacted by declining coal production. Interpreted as a capitalization effects, the home value estimates imply a total decline of around \$1 billion in individuals' valuation of life in coal country. Alternatively, by bringing together the estimates of wage, rents, and public school expenditures in a valuation exercise motivated by a theoretical model of spatial equilibrium, I calculate that mobile households showed declining valuations for living in coal country of \$1,400-1,500 per household and about \$500 million in the aggregate in my preferred specifications, driven mostly by the prospect of lower wages.

In summary, my results suggest that the effect of declines in coal demand from the power sector on producing areas has thus far been meaningfully negative but not yet overwhelming. The (negative) multipliers on mining jobs and wages appear to be moderate and consistent with literature exploring previous cycles in coal demand. I do not find strong evidence of high or persistent unemployment, increasing crime, rapid outmigration, or declining expenditures by general purpose governments, despite reduced support from the state. On the other hand, it seems beyond doubt that coal demand shocks reduce the overall welfare of coal country's residents as evidenced by reduced in-migration and declining

home values. Of potential concern for policymakers, I find that substantial numbers of potential in-migrants sort away from living in coal country, school district expenditures on public education decline, and poverty rates increase.

1.9 Figures for Chapter 1




Figure 1.1: **Aggregate Fossil Fuel Generation.** US fossil fuel generation by source since 2002, (author reproduction) from EIA data at <https://www.eia.gov/todayinenergy/detail.php?id=43675>.




Figure 1.2: **Coal Production and Powerplant Deliveries.** Blue bars are derived from the MSHA and EIA from 7A. Red bars are derived from EIA 423 and 923. The solid and dotted lines are EIA published aggregates from annual coal reports.




Figure 1.3: **US Coal Miners.** Coal mining employment (total miners) reported in MSHA/EIA-7A represented by vertical bars and FRED data series (CEU1021210001) by the dotted line.




Figure 1.4: **Coal Supply Regions from Stoker et al. (2005).** Boundaries of coal supply regions drawn at the county level.






Figure 1.5: Subregions and Coal County Centroids. This figure shows approximate boundaries of EPA eGRID Subregions as colored polygons and the centroids of coal-supplying counties as black triangles. eGRID Subregions are an intermediate division of electricity zones between NERC (North American Electric Reliability Corporation) regions and BAAs (balancing area authorities).



(a) SRMV region, 2008



(b) SRMV region, 2016

Figure 1.6: Example Supply Stacks. These figures show two examples of estimated electricity load region supply curves for the SRMV eGRID subregion. The grey shape in each figure shows the distribution of total fossil fuel load from CEMS plants at the hourly level. Marginal costs for CEMS plants are estimated as the unit-level average heat rate times fuel prices drawn from the EIA SEDS database.




Figure 1.7: Dispatch Predictions. This figure shows predicted generation from a least-cost dispatch model of the SRMV load region against actual observed gross generation in 2008 and 2016. Predicted values are derived by constructing the supply curves depicted in the previous figures and “intersecting” them with observed load from CEMS plants in each hour (modelling hourly demand as perfectly inelastic). All units with costs below the intersection of supply and demand are dispatched for the full hour. Summing over all hours in the year yields annual predicted generation for each unit as a function of state-level fuel prices, the region’s fuel mix, and total fossil load.






Figure 1.8: Market Shares. This figure is a strip plot of the share of Btus of coal delivered by each coal producing county to each eGRID Subregion in the continental US over the entire period from 2002-2007. Each row on the y-axis represents an eGRID Subregion and the x-axis measures the share of fuel deliveries. Campbell County, WY is excluded from all calculations.




(a) By region




(b) By period

Figure 1.9: **Predicted Coal Demand vs. Actual Coal Production.** These figures show the first stage relationship between the predicted coal demand instrument and observed coal production value in dollars. See section 1.4 for details.




(a) No adjustment for commuting zone characteristics




(b) Adjusted for commuting zone characteristics

Figure 1.10: **Reduced-Form placebo regressions, 2002-2007.** See section 1.4 for details. GP Rev. = General Purpose Revenue (county, municipality, township). SD = School District Revenue.



(a) Total wages per capita



(b) Log of wages per job

Figure 1.11: **Wage Effects.** See section 1.5

1.10 Tables for Chapter 1

Table 1.1: Summary statistics

	(1) N	(2) mean	(3) sd	(4) min	(5) max
Production (MM short tons)	254	2.625	4.957	0	34.47
Production (MM Btus)	254	56.94	107.3	0	804.3
Production value (MM USD)	254	94.30	176.7	0	1,456
Coal mining emp. share	254	0.0305	0.0614	0	0.430
Mining emp. share	254	0.0553	0.0708	0	0.452
Avg. Mining Wage (000s)	240	46.28	15.42	17.53	92.92
Wage per job (000s)	254	38.03	6.393	26.83	66.44
AGI per HH (000s)	254	45.34	11.39	24.56	92.23
Population (000s)	254	62.73	150.6	1.865	1,848
Labor Force (000s)	254	30.77	82.82	1.131	1,067
Total Employment (000s)	254	27.18	90.15	0.660	1,173
Unemployed (000s)	254	1,416	3,011	48	34,024
GP Exp. per capita (000s)	254	1.718	1.117	0.145	7.989
SD Exp. per capita (000s)	254	1.629	0.730	0	4.297

(a) Summary statistics for coal counties in 2007. Campbell County, Wyoming excluded.

	p5	p25	p50	p75	p95	iqr	sd
Appalachia	-9,206	-413	-5	43	3,037	456	4,239
Illinois	-4,465	-162	0	1,628	10,857	1,790	5,893
Interior	-2,405	-381	0	197	2,363	579	1,791
West	-3,557	-94	0	2,022	13,705	2,116	6,639
Total	-5,804	-348	0	211	5,676	559	4,808

(b) Distribution of $\Delta D_{ct}/Pop_{c,t-1}$, the change in production value per capita by mines in county c.

Table 1.2: Reduced-Form Effects of Demand Declines by Period

Dependent variable: % point change in emp/capita over 5-year difference

	(1)	(2)	(3)	(4)	(5)
	Coal Miners	Mining emp. (BLS)	Tot. Emp.	Manufacturing	Non-Tradeables
<u>Panel A: 2007-2017</u>					
Demand Decline (Δ s.d.)	-0.53*** (0.12)	-0.63*** (0.12)	-0.93*** (0.16)	-0.01 (0.03)	-0.19* (0.09)
Obs	508	508	508	508	508
<u>Panel B: 2012-2017</u>					
Demand Decline (Δ s.d.)	-0.60*** (0.13)	-0.67*** (0.11)	-0.97*** (0.18)	0.01 (0.03)	-0.22* (0.10)
Obs	254	254	254	254	254
<u>Panel C: 2007-2012</u>					
Demand Decline (Δ s.d.)	-0.14 (0.13)	-0.46 (0.38)	-0.69* (0.29)	-0.11 (0.06)	-0.07 (0.29)
Obs	254	254	254	254	254
<u>Panel D: 2002-2007 (placebo)</u>					
Demand Decline (Δ s.d.)	0.03 (0.12)	0.08 (0.15)	0.25 (0.26)	0.03 (0.10)	0.13 (0.21)
Obs	254	254	254	254	254
Region-time FE	yes	yes	yes	yes	yes
CZ controls	no	no	no	no	no

Notes: This table reports reduced form regressions of 5-year changes in employment per capita on the instrumental variable $\Delta \widehat{D}_{ct}$ per capita. Outcome variables come from the MSHA (column 1) and the BLS (all other columns). Each coefficient comes from a separate regression. Panel A shows coefficients estimated using Equation (1.4) with two periods of data in the stacked difference model. Panels B and C estimate one-period regressions of contemporaneous changes in employment against changes in the instrument. Panel (D) regresses changes from 2002-2007 against the average of future values of the instrument. Non-tradeables are defined as all employment except mining and agriculture, manufacturing, and government. The instrument is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.3: Results on Employment and Wages

Dependent variable: % point change in emp/capita over 5-year difference

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Employment						
	Coal miners (MSHA)	Mining	Total	Manufacturing	Non-Tradeable	Util./Transp.
Coal Value (Δ s.d.)	-0.94*** (0.14)	-1.16*** (0.23)	-1.61*** (0.32)	-0.07 (0.07)	-0.22 (0.22)	-0.17*** (0.06)
F excl.	20.47	20.47	20.47	20.47	20.47	20.47
Obs	508	508	508	508	508	508
Panel B. Wages						
		Mining	Total	Manufacturing	Non-Tradeable	Util./Transp.
Coal Value (Δ s.d.)		-1,048*** (214)	-1,349*** (272)	-24 (51)	-250** (119)	-174** (79)
F excl.		20.47	20.47	20.47	20.47	20.47
Obs		508	508	508	508	508
Region-time FE	yes	yes	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in employment and wages per capita on the treatment variable ΔD_{ct} , coal production value changes per capita. Outcome variables come from the MSHA (column 1) and the BLS (all other columns). Each coefficient comes from a separate regression. Panels A and B both show coefficients estimated using Equation (1.1) with two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the “Empirical Strategy” section. Non-tradeables is defined as all employment except mining and agriculture, manufacturing, and government. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.4: Results on Population and Migration

	(1)	(2)	(3)
Panel A: $(\Delta Y_{ct}/Pop_{c,t-1})$	Population	Inmigr.	Outmigr.
Coal Value (Δ s.d.)	-0.89*** (0.23)	-0.20** (0.09)	0.01 (0.07)
F excl.	20.47	20.47	20.47
Obs	508	508	508
Panel B: $(\Delta \ln Y_{ct})$	Population	Inmigr.	Outmigr.
Coal Value (Δ s.d.)	-0.009*** (0.002)	-0.055*** (0.020)	0.003 (0.013)
F excl.	20.47	20.47	20.47
Obs	508	508	508
Region-time FE	yes	yes	yes
CZ controls	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in population and migration on the treatment variable ΔD_{ct} , coal production value changes per capita. Outcome variables come from the US Census (column 1) and the IRS (columns 2 and 3). Each coefficient comes from a separate regression. Panel A shows coefficients estimated using Equation (1.1) and panel B both show coefficients estimated using Equation (1.2) with log differenced outcomes. Both panels use two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the “Empirical Strategy” section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.5: Results on Labor Force Components

	(1)	(2)	(3)	(4)	(5)
Panel A: $(\Delta Y_{ct}/Pop_{c,t-1})$	Mining	NonMining	NILF	Unemployed	Poor
Coal Value (Δ s.d.)	-1.16*** (0.23)	-0.45* (0.25)	0.33 (0.38)	0.04 (0.07)	0.67** (0.31)
F excl.	20.47	20.47	20.47	20.47	20.47
Obs	508	508	508	508	508
Panel B: $(\Delta \ln Y_{ct})$	Mining	NonMining	NILF	Unemployed	Poor
Coal Value (Δ s.d.)	-0.139*** (0.042)	-0.018** (0.009)	0.004 (0.023)	0.018 (0.025)	0.036** (0.016)
F excl.	16.29	20.47	20.20	20.47	20.47
Obs	455	508	506	508	508
Region-time FE	yes	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in employment, labor force participation, and poverty on the treatment variable ΔD_{ct} , coal production value changes per capita. Outcome variables come from the BLS (columns 1-4) and the Census’ SAIPE (column 5). Each coefficient comes from a separate regression. Panel A shows coefficients estimated using Equation (1.1) and panel B both show coefficients estimated using Equation (1.2) with log differenced outcomes. Both panels use two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the “Empirical Strategy” section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.6: Results on Public Finances

	(1)	(2)	(3)	(4)
Panel A: General Purpose	GP Expend.	GP Rev.	GP Rev. (local)	GP Rev. (state)
Coal Value (Δ s.d.)	0.015 (0.036)	-0.008 (0.035)	0.026 (0.044)	-0.202*** (0.043)
F excl.	20.50	20.50	20.50	20.50
Obs	507	507	507	507
Panel B: School District	SD Expend.	SD Rev.	SD Rev. (local)	SD Rev. (state)
Coal Value (Δ s.d.)	-0.041** (0.017)	-0.044*** (0.013)	-0.076*** (0.026)	-0.034 (0.021)
F excl.	19.59	19.59	19.59	19.59
Obs	469	469	469	469
Region-time FE	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in per-capita local public finance on the treatment variable ΔD_{ct} , coal production value changes per capita. Outcome variables come from the Census of Governments. Each coefficient comes from a separate regression. Panels A and B both show coefficients estimated using Equation (1.2) with log differenced outcomes. Both panels use two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the “Empirical Strategy” section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.7: Results on Local Public Goods

	(1)	(2)	(3)	(4)
Panel A:	Crime Rate (viol.)	Crime Rate (prop.)	Students/Teachers	Teachers
Coal Value (Δ s.d.)	-0.011 (0.031)	-0.035 (0.043)	0.039 (0.026)	-0.022*** (0.005)
F excl.	20.47	20.47	17.16	15.82
Obs	508	508	468	462
Region-time FE	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in measures of public goods on the treatment variable ΔD_{ct} , coal production value changes per capita. Outcome variables come from the FBI’s Uniform Crime Reporting Data (columns 1 and 2) and the National Center for Education Statistics (columns 3 and 4). Each coefficient comes from a separate regression. Coefficients are estimated using Equation (1.2) with log differenced outcomes and two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the “Empirical Strategy” section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.8: Heterogeneity Across Regions

	(1) East/Appalachia	(2) Illinois	(3) Interior	(4) West
Panel A: Employment and Wages				
miners per cap	-1.18*** (0.07)	-0.96*** (0.10)	-0.25 (0.25)	-0.27** (0.11)
mining wages per cap	-1,187.81*** (222.70)	-507.08** (247.59)	496.72 (397.00)	-812.85 (502.74)
ln(employment)	-0.07*** (0.02)	-0.01 (0.01)	0.06 (0.07)	-0.04*** (0.01)
ln(avg wage)	-0.06*** (0.02)	-0.01 (0.01)	0.04 (0.05)	-0.03** (0.01)
Panel B: Real Estate and Migration				
ln (med. home value)	-0.026*** (0.009)	-0.018* (0.011)	0.073 (0.046)	-0.070** (0.031)
ln (med. rent)	-0.006 (0.009)	-0.039 (0.025)	-0.083 (0.077)	-0.015 (0.027)
ln (inmigr.)	-0.065** (0.028)	0.048*** (0.015)	-0.253*** (0.067)	-0.063** (0.032)
ln (outmigr.)	0.009 (0.014)	0.008 (0.022)	-0.175*** (0.032)	-0.006 (0.031)
Panel C: Public Revenue				
ln (avg GP exp.)	0.015 (0.044)	0.022 (0.078)	-0.082 (0.261)	0.004 (0.060)
ln (avg SD exp.)	-0.041*** (0.014)	0.050 (0.045)	-0.036 (0.068)	-0.091*** (0.026)
No. Counties	149	41	31	33
Time FE	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes

Notes: This table reports region-level 2SLS regressions of 5-year changes in labor market outcomes, home values, migration, and local public finance on the treatment variable ΔD_{ct} , coal production value changes per capita. Coefficients in the first two rows are estimated using Equation (1.1) and all other rows using Equation (1.2) with log differenced outcomes. All rows use two periods of data in the stacked difference model. Time FE is an indicator for the later period. Control variables are described in the “Empirical Strategy” section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.9: Summary of Spatial Equilibrium Effects

	(1)	(2)	(3)	(4)
Panel A:	Wages	Med. Home Vals	Med. Rents	SD Exp.
Coal Value (Δ s.d.)	-0.042*** (0.008)	-0.030*** (0.010)	-0.013 (0.011)	-0.041** (0.017)
F excl.	20.47	20.47	20.47	19.59
Obs	508	508	508	469
Region-time FE	yes	yes	yes	yes
CZ controls	yes	yes	yes	yes

Notes: This table reports 2SLS regressions of 5-year changes in measures of wages, housing costs, and public expenditures on education on the treatment variable ΔD_{ct} , coal production value changes per capita. Each coefficient comes from a separate regression. Coefficients are estimated using Equation (1.2) with log differenced outcomes and two periods of data in the stacked difference model. Region-time fixed effects follow Stoker et al. (2005), and control variables are described in the “Empirical Strategy” section. Treatment is standardized to have mean 0 and standard deviation 1. SE clustered at commuting zone.

Table 1.10: Welfare Estimates

	(1)	(2)	(3)
	Avg. (\$/cap)	Agg (\$ B)	Agg Inf. Disc. (\$ B)
Panel A: Non-College parameters			
Δ housing costs 0%	-1,859	-0.66	-13.11
Δ housing costs 1.3%	-1,437	-0.51	-10.13
Δ housing costs 3%	-885	-0.31	-6.24
Panel B: College parameters			
Δ housing costs 0%	-1,700	-0.6	-11.98
Δ housing costs 1.3%	-1,497	-0.53	-10.55
Δ housing costs 3%	-1,231	-0.43	-8.68

Notes: The table reports estimates of the effect of demand declines on the economic welfare of residents of coal country. Column (1) reports estimates per household, while column (2) aggregates per capita estimates by scaling the effect to the \$3.7 billion decline in coal production value observed over sample and the 7,300,375 households in coal producing counties inferred from IRS tax data. Column (3) inflates column (2) under the assumption that the observed changes in wages, housing costs, and public spending on education are one-time and permanent. The rows of the panels explore the sensitivity of the estimates to different effect sizes on home values. The calculations are converted to dollars using a mean household wage and salary of \$35,963 inferred from BLS and IRS data.

Chapter 2: The Effect of Traffic Cameras on Police Effort: Evidence from India

2.1 Introduction

Video and electronic surveillance are becoming increasingly common in the workplace (Chyi 2021, Lohr 2014, West 2021). Employee behavior has been monitored and studied as a way to improve worker productivity in firms, restaurants, and call centers (Pierce et al. 2015, Scott et al. 2021). Monitoring has been also used in health care facilities and schools, and studies indicate that performance is improved when strict surveillance is available (Borcan et al. 2017, Duflo et al. 2012, Cilliers et al. 2018). Among police, in light of increased use of body-cameras,¹ several studies have looked at how video surveillance impacts the use of force (Ariel et al. 2015, Çubukçu et al. 2021) and racial disparities (Makofske 2020), but relatively little is known about how monitoring public enforcement officials via cameras can impact their productivity.² This is the case despite the fact that

¹In the U.S., The Urban Institute reports that 17 states have laws indicating the use of body cameras by police, and 6 states have pending legislation (<https://apps.urban.org/features/body-camera-update/index.html>).

²Related literature has looked at monitoring of officials for corruption through audits or randomized controlled trials (Bertrand et al. 2007, Bobonis et al. 2016, Olken 2007, Zamboni and Litschig 2018); or phone base monitoring to improve service delivery (Muralidharan et al. 2021).

there has been widespread adoption of CCTV cameras for traffic and crime monitoring in many cities around the world.³ In this paper, using an administrative dataset of traffic tickets (“challans”) issued by officers before and after the installation of closed-circuit TV (CCTV) cameras in the Indian city of Chandigarh, we investigate whether the use of video cameras can impact the productivity of on-the-ground police officers.⁴

Previous literature has identified three likely effects of monitoring employees. First, there is an incentive for employees to work harder because they could be subject to disciplinary actions if caught shirking (Hubbard 2000, Duflo et al. 2012).⁵ In our case, on-the-ground traffic officers were aware when the cameras were working through communication with a satellite office. This suggests that increased monitoring would translate into officers issuing more tickets. Second, there might be a crowding-out effect resulting from the loss of trust between principle and agent, which could diminish intrinsic motivation and thus productivity (Falk and Kosfeld 2004). The extent of this effect may vary depending on whether monitoring is interpersonal, where output is higher when monitored in a distant relationship rather than an interpersonal one (Frey 1993, Dickinson and Villeval 2008). In our case, the monitoring is not interpersonal. The monitoring of principals

³According to industry reports, by mid 2021, there were 770 million cameras in use globally (<https://www.comparitech.com/studies/surveillance-studies/the-worlds-most-surveilled-cities/>).

⁴In Chandigarh alone, there has been a striking increase in the use of CCTV cameras since our period of study. News articles reported in February 2022, that 1,800 of 2,000 proposed cameras had been installed (<https://www.tribuneindia.com/news/chandigarh/chandigarh-no-challans-to-be-issued-via-new-cctv-cameras-for-now-368250>). We thank an anonymous referee for bringing this to our attention.

⁵Shirking is a significant problem among officials in many countries, for instance, Bandiera et al. 2009 find that shirking accounts for 83% of total corruption in Italian public bodies. Chaudhury et al. 2006 report that in India absence rates of government teachers and health workers was 25 and 40% respectively.

(traffic officers), happens along with the monitoring of agents (drivers), and is carried out remotely in a satellite office. Nevertheless, it is possible that fewer tickets are recorded if, as Chang and Lai (1999) argue, the crowding-out effect dominates the disciplining effect. Third, monitoring could result in a decline of active corruption (extortion or acceptance of bribes) where public office is used for private gains (Schulze and Frank 2003). In our context, because the cameras are useful in providing tickets to traffic offenders, they may capture passive (shirking) or active (rent-seeking) corruption of officers. While it is possible that the intrinsic motivation of officers was eroded by the installation of the cameras, our baseline hypothesis is that increased monitoring of traffic officers translates into more tickets issued through a reduction in active and/or passive corruption.

To motivate our empirical estimates, we provide a conceptual model of the data generating process for observed tickets, taking into account officers' incentives to issue tickets for violations, to solicit or accept bribes, and to shirk. To identify the effect of cameras on officer behavior, we use exogenous variation in the functionality of the CCTV cameras. Cameras were installed at twenty key intersections in Chandigarh, but as is often the case with new technologies, there was variation in when and how often the cameras worked.⁶ We exploit this variation with a difference-in-differences methodology and together with date and location fixed effects, we measure the number of tickets recorded by *on-the-ground* police officers located within 100 meters (or approximately 1 block) of the cameras.⁷ We find that as the intensity of monitoring increases, the number of *on-the-*

⁶Our understanding is that all the CCTV camera models and quality were similar across locations.

⁷Recent literature has highlighted concerns on biased estimates in two-way fixed effects models arising from negative weights when there is staggered timing and heterogeneous treatment effects (Goodman-Bacon

ground tickets issued also increases substantially. Different model specifications, controls, and alternative definitions of key variables also indicate an increase in ticketing behavior. We run placebo tests prior to the date when the cameras were installed and find no monitoring impacts of the cameras. We also use placebo locations and compare them to locations with non-functioning cameras to ensure that our results are not being driven by changes in driver behavior. To alleviate concerns of few clusters, we estimate our results using block bootstrapping and do a robustness check using permutation inference.

We find consistent support for the proposition that changes in officer effort result from a decline in shirking, particularly for violations that are observable by the CCTV cameras and those issued to cars and motorcycles registered locally. We test for changes in violations issued to the types of vehicles that are anecdotally more likely to be affected by bribe-taking (out-of-state vehicles and big cars) and find a statistically insignificant increase in violations in categories observable by the CCTV cameras. Interestingly, we do find a statistically significant increase in the number of tickets issued to big cars for violations not observable by the CCTV cameras, which could be consistent with officers pocketing fewer bribes for these higher risk vehicles. In theory, this could be driven by the cameras reducing the potential payoff of engaging in corrupt behavior, though our data prevents us from making strong claims with regards to corruption. Overall, our results indicate that remote

2021, Chaisemartin and d’Haultfoeuille 2020, Callaway and Sant’Anna 2020). Our model includes two-way fixed effects but the timing is not staggered since the installation of the cameras occurred on May 7, 2010. The treatment variable in our preferred model, “mean intensity” (described in section 2.3), remains constant for each camera throughout the entire time period. Nevertheless, to address possible concerns, as described in section 2.3, we present our main results using three definitions of treatment each of which incorporates variation in treatment over time in a different way.

monitoring technology can serve as, if not a substitute for, then a complement to *on-the-ground* enforcement.

Within India, Chandigarh is a positive outlier in terms of its reputation for the maintenance of law and order. Our results may therefore be generalizable for aspiring “smart cities” and metropolitan areas in developing countries where the local population is well educated, earning relatively high wages, and where the police are equipped with skilled manpower and technology. More generally, our results are relevant to policy makers who are considering adopting similar methods of traffic enforcement, and indicate that technology can serve as a complement to on-the-ground monitoring, even when such interventions may have a theoretically ambiguous or empirically negligible effect on other outcomes, such as traffic accidents (Gallagher and Fisher, 2020). This paper expands on economic literature exploring the consequences of monitoring of public officials, in particular as it relates to the use of cameras and traffic police. Further, we show evidence of complementarities between *on-the-ground* and remote traffic enforcement, documenting an example of a new technology typically adopted for one purpose (monitoring drivers) which has positive spillover effects on other agents (traffic officers) engaging with it.

The rest of the paper is structured as follows: in the next section we provide some background information and a conceptual framework for our analysis. In section 2.3 we describe the datasets used in our empirical models and in section 2.4 we model police incentives and relate them to observed ticketing behavior. Sections 2.5 and 2.6 present the empirical model and the results including a discussion of potential mechanisms. Section

2.7 shows results addressing possible data challenges. We conclude in section 2.8.

2.2 Background

The CCTV cameras, installed on May 7 of 2010, are located at twenty different intersections in Chandigarh and are operated remotely by officers in the Traffic Control Room, a central office. The cameras capture routine violations, but they also serve as monitoring devices for interactions between traffic officers and motorists. Ensuring fair, transparent dealings between officers and drivers and curbing corruption at traffic checkpoints were explicitly listed as reasons for the adoption of the cameras.⁸ Tellez et al. 2020 note that Indian police have much discretion in the solicitation of bribes “including helping people to avoid legitimate traffic violations.” A visual of where the cameras are located in Chandigarh can be seen in Figures 2.1 and 2.2. The gray dots indicate approximate locations where tickets were issued (derived from text descriptions), and the black triangles the location of the CCTV cameras. Gray filled triangles denote “placebo” locations used to examine driver behavior - we return to this in the empirical section.

Officers are selected based on their educational qualifications, age, physical qualifications, and test scores. Typically, *on-the-ground* officers have the rank of constable, while camera operators are senior administrative officials or officials who report to senior officers, allowing them to monitor those *on-the-ground* even though they do not have explicit

⁸See page 6 in <http://chandigarhpolice.gov.in/pdf/coffee-table-book/section-iiii/section-iiii.pdf> last accessed on October 31, 2021. See also The Times of India article from Oct. 13, 2006 <https://timesofindia.indiatimes.com/city/chandigarh/rla-fits-cctvs-to-curb-corruption/articleshowprint/2162415.cms> last accessed on October 31, 2021.

supervisory authority over them. Camera operators can give instructions or feedback to *on-the-ground* police personnel. Officers salaries depend on their seniority level and officers can be promoted based on overall performance, performance on certain assigned cases, or because of tenure. The police operate in a reward and punishment system, where certificates are given for meritorious service or excellent performance. A supervising officer makes a written annual evaluation of each constable.

Officers in Chandigarh face severe penalties for shirking or bribing, although not all cases come to light because drivers do not always complain. Tellez et al. 2020 indicate that a common mechanism used to fight corruption is by citizens exposing an incident in a newspaper or a television station. They also resort to approaching local elected officials. News articles indicate that officers accused of taking bribes are taken to court and if found guilty are suspended and/or arrested.⁹ Inquiries are conducted within the department when complaints are made, and action is taken.

Information provided by the head of the traffic control room at the time of our study indicates the Chandigarh Police headquarters made a monthly plan by using software that randomly assigned constables to locations. A senior officer then made a duty roster, and constables had no choice in location assignment. Within a day, constables generally stayed in one location, but could be moved to different locations depending on the orders given to

⁹See for example an article from the *Times of India* March 31, 2012 “Cop Caught on Facebook for Graft” indicates how a video of a constable accepting a Rs 500 (approximately USD \$10) bribe was uploaded to Facebook. The article also refers to 13 incidents issued against police personnel for violating traffic rules. <https://timesofindia.indiatimes.com/city/chandigarh/Cop-caught-on-Facebook-for-graft/articleshowprint/12475768.cms> accessed on May 5, 2021. An extensive search of news sources and the website ipaidabribe.com indicate that the average bribe paid in Chandigarh during the time period that we study was Rs 400 (approx. USD \$8).

them. Across different days, constables were frequently posted to different locations. These changes were announced with one day's notice to administrative staff, and to constables via wireless message. The head of the traffic control room was clear that the allocation of constables was not related to the functionality of the cameras on any given day, but to traffic density. He also indicated that typically, 2-3 constables were assigned on small junctions (which include CCTV locations), 3-4 on bigger junctions, and sometimes more if there was some special event that increased the traffic flow.

2.3 Data

The Police Department of Chandigarh provided the two data sources described below.

(1) *On-the-ground traffic violations or tickets*: This dataset contains information on all traffic violations that took place in the north Indian city of Chandigarh between January 1, 2009 and March 28, 2012. The raw data file we received includes just under 600,000 observations. Each observation contains information on: date, time of day, approximate location, type of violation, and vehicle committing the violation. We use only *on-the-ground* tickets as our outcome variable, and restrict the sample to tickets issued within a given radius of the CCTV camera installation sites (typically 100 meters, or approximately one city block).

(2) *CCTV cameras*: The CCTV cameras were introduced in 20 locations throughout the city on May 7, 2010. We received a dataset containing information on CCTV recorded violations between the installation date and May of 2012. The raw data includes approx-

imately 85,000 observations, and each observation contains information on: date, time of day, the relevant CCTV camera, and the type of violation.

Examination of the two data sources indicates that CCTV cameras issue tickets contemporaneously with *on-the-ground* officers. We were told that the cameras and officers do not replicate tickets, and see no obvious duplication in the two data sources. There appear to be complementarities between the tickets issued by CCTV cameras and those issued by *on-the-ground* officers, as the types of violations enforced by the cameras is not as extensive as those enforced by the *on-the-ground* officers. Importantly, significant overlap in violation types between the two data sources allows us to group violation types into several intuitive categories and clearly delineate which violation types are observed and enforced by CCTV cameras and which are not.

We also use the CCTV violations data to identify when and where the cameras worked, creating three different “treatment” measures: (1) *Functional=1* when a CCTV camera recorded any violations and zero otherwise; (2) “*Months*” *Running* is the cumulative number of days a CCTV camera was working after May 7, 2010, the day when the cameras were installed (divided by 30); and (3) *Mean Intensity* is the average percentage of time a CCTV camera was functioning since the installation. This last measure is constant for a given camera, while “*Months*” *Running* could change.

Because we analyze *on-the-ground* tickets, our period of study runs from January 1, 2009 to March 28, 2012. We were informed that the cameras were only functional during the daytime, so we restrict the data to include only *on-the-ground* violations recorded be-

tween 8 am and 7 pm.¹⁰ The unit of observation is the CCTV location (20) by date (1,183), for a total of 23,660 observations in the estimation sample. Summary statistics are provided in Table 2.1, and the raw variation in total tickets and cumulative days of camera function is plotted in Figure 2.3.

As Panel A of Table 2.1 indicates, at least one *on-the-ground* violation was recorded within 100 meters of a CCTV location on about 65% of days in the sample, and on average there were approximately 5.3 violations recorded in a particular location and day. The ‘Post’ variable indicates that our data are fairly evenly split before and after camera installation, with about 58% of observations from the post period, while the *Mean Intensity* variable suggests that on average the cameras worked 41% of the days after their installation.

Panel B of Table 2.1 reports summary statistics according to vehicles’ observable characteristics. The average number of out-of-state license plates tickets issued is 2.5, while 2.2 tickets were issued on average to cars, and around 1.6 tickets were issued on average for motorcycles. Panel C reports summary statistics by type of violation. We further split this panel to identify violations that are observable by the CCTV cameras, and those that are not. The most common violations are those broadly categorized as breaking “Major road-rules” which include disobeying signs, stopping on pedestrian crossings, and running yellow and/or red lights. Another common violation is motorcycle safety violations (“No or unauthorized helmet/unsafe motorcycle”), most often a failure to wear a helmet, not

¹⁰About 95% of CCTV violations were recorded between these hours in the raw data we received. Results including nighttime data and conditioning on a day/night indicator are consistent with those reported here. Later in the paper we explore nighttime observations in a robustness check.

wearing a certified helmet, or carrying excess passengers unsafely. Violations that involve documentation, such as those reported as “License/insurance/registration (papers)” are not observable remotely with the cameras, but are fairly common. The second most common non-CCTV observable violation is dangerous driving¹¹, which can involve officer judgement or discretion.

2.4 Conceptual Framework

In the data we observe the ticketing behavior of officers in approximate locations, which is an imperfect measure of effort and/or corruption. We propose a simple model of officer ticketing behavior to motivate the empirical results under two simplifying assumptions. First, drivers are confronted only given some type of violation. In other words, an officer stops a vehicle only if he has a pretext for doing so. Second, “a bribe changing hands” implies that money changes hands from driver to officer and a ticket is not recorded.¹²

Constable i observes violation v at a given intersection, and takes one of three qualitatively different choices: to work, to engage in active corruption (extort or accept a bribe), or to avoid work (shirk), denoted $\{W, X, S\}$. In all cases, the constable receives a fixed wage w_i with certainty. If he works, he receives his wage. If he engages in active corruption and solicits or accepts a bribe of size b , his expected penalty is f_x . If he shirks he faces an expected penalty f_s . Allowing for a preference shock η_{ivc} , denote his indirect utility and

¹¹Dangerous driving includes: rash/dangerous driving, failure to lower lights, dangerous overtaking, failure to keep a safe distance, abrupt breaking, failure to yield at right-of-way and reversing more than necessary.

¹²This assumption rules out other types of bribery, such as paying a bribe to reduce a large challan to a small challan which local knowledge and studies suggest are unlikely (Jauregui 2016).

probability of taking action c as:

$$\begin{aligned}\pi_{ivc} &= V(w_i, b, f_x, f_s, \eta_{ivc}) \\ P_{ivc}(b, f_x, f_s) &= \Pr(\text{i choose } c \mid b, f_x, f_s).\end{aligned}$$

Applying the definitions above and summing across violation-specific shocks yields an expected number of tickets N_i honestly cited by officer i . Introducing cameras that affect f_x, f_s changes the expected number of tickets predictably:

$$\begin{aligned}N_i &= \sum_v (1 - P_{ivX} - P_{ivS}) \\ \frac{\partial N_i}{\partial f_j} &= \sum_v \left(1 - \frac{\partial P_{ivX}}{\partial f_j} - \frac{\partial P_{ivS}}{\partial f_j}\right) \text{ for } j \in \{x, s\}.\end{aligned}$$

Importantly, the model suggests potential leakage between shirking and active corruption and indicates its effect on the average number of tickets issued. Standard assumptions on preferences imply that increasing the expected penalty for extortion both decreases the likelihood of extortion ($\partial P_{ivX} / \partial f_x < 0$) while potentially increasing the likelihood of shirking ($\partial P_{ivS} / \partial f_x \geq 0$). Intuitively, a corrupt officer threatened with greater sanctions may choose simply to shirk rather than work.

As a non-interpersonal continuous monitoring technology, traffic cameras may affect the expected penalties for both extortion and shirking. A sufficient condition for the total number of tickets to increase is that the sum of direct incentive effects dominates the leakage effects (cross-partial derivatives). This is the natural primary question for the incentive

effects of traffic cameras on police ticketing, and is the focus of the empirical analysis. Secondly, we subset the data based on vehicle characteristics such as out-of-state license plates and big cars (anecdotally more likely targets for corrupt officers) to explore what drives the results.

2.5 Empirical Strategy

The twenty CCTV cameras, installed at major intersections of the city, feed data to a central control room in the police headquarters. The controller in this room is in charge of recording remote violations and in contact with the police stationed at the intersections with CCTV cameras. After visiting the field site, it was clear that *on-the-ground* traffic officers were aware, through their communication with officers in the control room, which cameras were functional and which were not. As is often the case with new technologies, there was variation in when and how often the cameras worked during our sample, with some cameras operating as little as 0.1% of the time, and others as much as 94% of the time. We exploit this variation across time and space to isolate the impact of the cameras on officers' ticketing behavior and control for driver-side reactions to camera installations.

We estimate the impact of the CCTV cameras in a difference-in-differences framework tailored to deal with two primary challenges to identification. First, cameras were strategically placed throughout the city, indicating that a comparison between locations with and without cameras may not yield a causal estimate. Second, the mere presence of a camera may impact driver behavior, confounding an estimate of the effect of the cameras on po-

lice behavior. To address these challenges, our estimation strategy compares locations with low- and high-functioning cameras, netting out the effect of cameras on driver behavior.

The key added assumption in our context is that driver behavior does not change in response to camera *functionality*, instead responding only to CCTV camera *presence* at an intersection. We do not have data to test this assumption directly, but as an indirect test we can check whether drivers behave differently in the presence of a (unknown to them) non-functioning camera relative to a location without a camera.¹³ There are 8 locations where cameras function less than 10% of the time ('Non-Functioning' Cameras), and we chose 5 nearby placebo locations with comparable ticketing patterns that did not have a camera installed. We expect that drivers are if anything, more careful in the presence of a camera, and thus the number of *on-the-ground* violations with a CCTV camera should weakly decrease.

If drivers are unaware of when the cameras function - but police are - the comparison of working and non-working cameras isolates the effect of a functioning monitoring device on police behavior. Any increase in tickets at locations with highly-functioning cameras is likely only to operate through changes in behavior on the supply-side (here, the traffic officers), as opposed to the demand-side (the drivers). Below, we introduce our general estimation strategy before turning to a discussion of placebo tests to explore the validity of the strategy.

¹³Note that we cannot compare functioning from non-functioning camera locations since police officers *are* aware whether cameras are working, and thus can change their behavior.

2.5.1 Estimation

We aim to estimate the effect of monitoring by cameras on the *on-the-ground* ticketing behavior of police officers at CCTV camera locations. A violation occurring “at” a CCTV location is defined as occurring at most 100m (approximately 1 block) from a CCTV camera. Using all of the administrative ticketing data both pre- and post-camera installation, our specification for the reduced form impact of CCTV cameras on *on-the-ground* police behavior is as follows:

$$y_{lt} = \beta_0 + \beta_1 [Treat_{lt} \times Post_t] + \delta_l + \gamma_t + \rho_{lt} + \varepsilon_{lt} \quad (2.1)$$

where: l corresponds to one of twenty location codes; t corresponds to a particular date. The outcome variable, y_{lt} , is either an indicator variable that takes a value of 1 if there were any tickets recorded or the total number of tickets issued by officers at location l and date t (using only daytime observations). The coefficient of interest is β_1 , which captures the interaction of camera installment ($Post_t$) and the camera’s *functionality*, as defined in section 2.3 by the three different measures of treatment: (1) the indicator variable *Functional=1*; (2) the cumulative measure for “*Months*” *Running*; and (3) the average measure using *Mean Intensity*.

We estimate effects after controlling for location fixed-effects δ_l , date fixed-effects γ_t , and location-specific “seasonal” controls ρ_{lt} (location-by-day-of-week and location-by-month). All standard errors are estimated by block bootstrap to effectively cluster standard errors at the CCTV-location level and alleviate concerns due to a low number of clusters.

More precisely, we randomly re-sample CCTV locations with replacement 500 times and use the distribution of estimated coefficients to construct standard errors and confidence intervals. To ensure our approach is robust to different model specifications, we report results from OLS regressions using an indicator for any police-recorded tickets (Linear Probability Model or LPM), estimates by Poisson regression to account for the dependent variable being a count variable (estimated by maximum likelihood), and results from OLS regressions where the dependent variable is the number of violations.

2.5.2 Placebo Tests

Since our identification strategy relies on estimated differences between intersections with low- and high-function cameras, we check for parallel trends across these groups in the pre-period (prior to the installation of the cameras). We conduct a placebo test that uses only pre-treatment data and checks whether *on-the-ground* ticketing behavior was affected by a fake $Post(k)$ variable defined k months before the actual installation of the cameras. We test whether cameras with higher average camera functionality in the post period (the *Mean Intensity* variable) have differential trends in ticketing prior to the installation of the cameras. A violation occurring “at” a CCTV location is defined as occurring at most 100m (approximately 1 block) from a CCTV camera as above. Formally, the placebo specifications for the reduced form impact of CCTV cameras prior to the installation of the cameras is as follows:

$$y_{lt} = \beta_0 + \beta_{1k}[Treat_l \times FakePost(k)_t] + \delta_l + \gamma_t + \rho_{lt} + \varepsilon_{lt} \quad (2.2)$$

The outcome variable, y_{lt} is the total number of tickets in a particular location, date. The coefficient of interest is β_{1k} , which measures the treatment effect of camera intensity k periods *prior* to the installation of the cameras. All other details of the specification and estimation are identical to those described above in relation to 2.1. For brevity, we report estimates using only the count of tickets as the dependent variable in Panel A of Table 2.2. In Appendix B, we add to this specification estimates from a model using an indicator variable that takes a value of 1 if there were any tickets recorded and a Poisson regression coefficients estimated by maximum likelihood.

To ensure that our indirect test of driver behavior will lead to reliable inferences about the effect of cameras on driver behavior, we also estimate a version of equation 2.2 where $Treat_l$ is replaced by a dummy variable indicating the presence of a ‘Non-functioning’ camera on the sample of cameras functioning less than 10% of the time and the 5 nearby intersections we chose as a “peer-group”. As before, we report estimates using only the count of tickets as the dependent variable in Panel B of Table 2.2 for brevity, but include alternate model specifications in Appendix B.

The top panel of Table 2.2 indicates that the estimated coefficients β_{1k} on the interaction term are always insignificant. This is also the case for the alternative models reported in Appendix B. This is consistent with the idea that pre-trends in locations where cameras

worked more often are not different from those in areas where cameras worked less often. Panel B of Table 2.2 similarly indicates that the locations with non-functioning cameras do not have differential trends to our hand-picked “peer-group” of locations without cameras prior to their installation. These results support a causal interpretation of estimates of β_1 from equation 2.1 presented in the following sections.¹⁴

2.6 Results

Table 2.3 shows the paper’s main results - estimates of β_1 from equation 2.1, indicating the effect of the CCTV cameras on the officers’ *on-the-ground* ticketing behavior. The dependent variable in column (1) is an indicator variable that takes a value of 1 when any violation is recorded at that particular location, in any given day, and 0 otherwise. In panel A, the key explanatory variable is the interaction between *Post* and the *Functional* variable, an indicator when a CCTV camera recorded any violations and zero otherwise. Results in panel A column (1) indicate that on days when the cameras work, relative to when they do not, there is an increase of 11.8 percentage points in the probability of any violation being recorded. This corresponds to a 19% increase from the pre-period mean. The dependent variable in columns (2) and (3) use the number of violations recorded. Exponentiating

¹⁴As an additional check for pre-trends, Figure 2.4 uses an event study type specification that estimates the following regression:

$$y_{lt} = \beta_0 + \sum_{m=-7}^{m=7} \beta_m \text{treat}_{lm} * D_m + \delta_l + \gamma_t + \rho_{lt} + \epsilon_{lt}$$

where m indexes months pre/post camera installation. Coefficients are normalized in relation to April 2010, the month prior to the onset of the CCTV cameras. The figure shows no pre-trends prior to the onset of the cameras and positive coefficients after. Figure 2.5 shows the analogous figure for the driver behavior test, with negative coefficients after.

the coefficient in column (2) estimated using the Poisson model indicates that the ratio of tickets issued for when the camera is functional relative to when it is not is 1.56, indicating a 56% increase in daily tickets issued. Column (3) indicates that there are on average 2.469 additional tickets issued when the camera is functional. This corresponds to a 42% increase from the mean number of pre-period tickets. Overall, all results in panel A, regardless of how the outcome variable is defined, indicate substantial increases in the number of tickets issued by *on-the-ground* officers.

Panel B shows an interaction between *Post* and the '*Months*' *running* variable, the cumulative number of days a CCTV camera was working after the day when the cameras were installed divided by 30, a variable that could change daily in the post period in any given location. Like panel A, these results also indicate an increase in the number of *on-the-ground* tickets issued. The magnitude ranges between 4% and 14% being somewhat lower than those presented in panel A using the *Functional* variable.

Results presented in panel C use the interaction between *Post* and the *Mean Intensity* variable, a constant within each location equal to the percentage of days each CCTV camera was functioning since installation. Reported coefficients estimate the effect of changing mean intensity from 0 (a camera that never works) to 1 (a camera that always works). Scaling these coefficients by the standard deviation reported in Table 2.1 imply that a one standard deviation increase in the *Mean Intensity* variable results in an 11% increase in the probability of any violation being recorded (column 1). For the Poisson model in column (2), a one standard deviation increase in the *Mean Intensity* variable results in the number of

violations being 1.69 times higher. Results for column (3) are consistent, and indicate that a one standard deviation increase in the *Mean Intensity* variable results in 2.71 *on-the-ground* expected additional violations recorded.

Finally, panel D of Table 2.3 shows estimates of the effect of a non-functioning camera, our indirect test of driver behavior. All point estimates are negative and statistically insignificant, with the point estimate in column (3) implying that a non-functioning camera results in 2.5 fewer daily traffic citations. Though statistically weak, the relatively large point estimate suggests that drivers are more careful in the presence of cameras.¹⁵ This proposition is supported by estimates presented later in the paper broken down into ticket and vehicle categories.

In summary, the data points to an increase in *on-the-ground* police tickets issued when the cameras are working, and suggests that drivers respond in predictable ways to CCTV cameras. Assuming that driver behavior does not differ between functioning and non-functioning cameras, these results are consistent with an increase in work effort and/or a decrease in bribe solicitation by officers. In the next section we attempt to disentangle these effects.

2.6.1 Mechanisms

In this section, we analyze potential mechanisms underlying this effect. In particular, in regards to corrupt behavior, the two hypotheses that we explore regarding the change in the number of violations recorded by traffic officers in response to the functionality of cameras

¹⁵We thank an anonymous referee for pointing out that fewer cases of illegal driving and more diligent enforcement may actually cancel out in the aggregate given our point estimates.

are:

(A) Shirking: some traffic officers are leisure-seeking individuals who are less likely to let violating drivers pass when monitored.

(B) Extortion: some traffic officers are rent-seeking individuals who are less likely to demand a bribe in exchange for not issuing a ticket when monitored.

Results in Table 2.3 indicate an overall increase in the number of tickets issues when a camera is functional, consistent with an incentive effect on officers when monitored. In panels A-C of Table 2.4 we explore whether the increase in number of tickets issued is being driven by any specific types of violations, with particular interest in violations that can be monitored by the CCTV cameras. Panels A-C of Table 2.4 show that common road rules citations drive the previously estimated increases in total tickets issued, with some additional increases in motorcycle helmet citations. Both are easily monitored remotely by the cameras.¹⁶ In contrast, violations that are not observable remotely and more likely to be correlated with corrupt behavior due to officer discretion, such as dangerous driving and issues with driver or vehicle documentation (“papers”), do not appear to be affected by the functionality of the cameras.

Panel D of Table 2.4 reports results for the non-functioning cameras and the nearby placebo locations with comparable ticketing patterns that did not have a camera installed.

¹⁶Examples of “road rules” violations include: not stopping before a pedestrian crossing, illegal turns, and running yellow and/or red lights. The “helmet” category mostly consists of helmet infractions but also includes more general unsafe riding citations to motorcycles. We choose the term “helmet” to avoid confusion with vehicle types in Table 2.5. See Data section for more details.

The table shows generally negative or statistically insignificant results. Here we find evidence that drivers appear to be more careful in regards to “dangerous” and CCTV observable violations, particularly motorcycle riders.

In Table 2.5 we further explore whether violations issued to vehicles according to their observable characteristics exhibit any changes. Anecdotal evidence and unofficial records of bribe reporting available online¹⁷ indicate that bribes are more commonly collected for bigger more expensive vehicles and vehicles with out-of-state license plates. We report results for out-of-state vehicles, big cars, cars and motorcycles. These categories are not always mutually exclusive, e.g. big car is a subset of cars, out-of-state vehicles could be cars or motorcycles.

Panel A of Table 2.5 shows aggregate increases in tickets issued to cars and motorcycles, but no changes to those more likely to be targeted for extortion, such as out-of-state vehicles and big cars. Panels B and C split these overall violations by whether they are observable or not by the CCTV cameras. Consistent with the aggregate results reported in panel A, for CCTV observable violations (panel B), cars and motorcycles show the highest increase in tickets. As these are the most common vehicles and officers know they are being monitored, increasing ticketing to these vehicles likely requires the lowest increase in effort and offers the greatest decrease in the probability of being punished for shirking. It is also worth noting that the results indicate a proportionally smaller (and statistically insignificant) increase in ticketing towards out-of-state and big cars. One interpretation of this result is that officers reduce shirking generally, but avoid ticketing big or out-of-state

¹⁷See <http://www.ipaidabribe.com/>

vehicles for fear of irritating a wealthy or powerful individual. They may also see little return to issuing such citations if camera monitoring reduces the potential payout of such interactions.

Panel C reports results for violations not observable by CCTV cameras. Intriguingly, both cars and big cars show a statistically significant increase in violations issued by *on-the-ground* officers after the onset of camera installation. It is possible that this is driven by a general change in norms prompted by the camera installations, inducing officers to increase effort even when unobserved. Another way such a pattern might arise is due to officers pocketing fewer bribes in exchange for not issuing challans. A reduction in active corruption seems more consistent with the results of Panel B, where officers appear to be reacting differently towards big cars when monitoring is introduced. Put another way - officers do not cite big cars more frequently for monitored violations, but do appear to cite big cars more frequently for discretionary violations. Ultimately, it is difficult to make definite statements about corruption given the available data. Still, it seems quite feasible that these ticketing patterns are consistent with a decrease in bribe-taking, at least within the immediate vicinity of the cameras.

Panel D of Table 2.5 reports results for the non-functioning cameras and the nearby placebo locations with comparable ticketing patterns that did not have a camera installed. Results show a decline in the number of tickets issued across all type of vehicles. This is consistent with drivers of all vehicles being more careful in the presence of the cameras, and encouragingly suggests that driver responses are relatively uniform across vehicle types

and driver responses are not driving our other estimates.

Taken together results in this section indicate that the observed increase in the number of *on-the-ground* tickets recorded after the installation of the CCTV cameras arise mostly from a decrease in shirking, in particular in regards to violations that are observable by CCTV cameras and issued to the most common types of vehicles. The fundamental problem of unobserved data prevents us from being able to say something definite about corruption. We view the results in Panels B and C of Table 2.5 as suggestive evidence that we should not rule out a modest decrease in rent-seeking behavior at CCTV intersections.

2.7 Robustness and Additional Results

In Tables 2.6 and 2.7 we present results that address our most important data quality and specification concerns. In panel A of Table 2.6, we limit the locations to only those that have a high number of tickets recorded on a regular basis (eliminating the low-functioning camera group used in the driver behavior test, plus one borderline case). We do this because our measures of camera functionality use information from the CCTV tickets issued by a camera on any given day. In particular, if no tickets are issued by a camera on a given location and day then the *Functional* indicator takes a value of zero. This seems to be a reasonable assumption given a pre-period mean of nearly 5.9 daily tickets. However, it is still possible that we are coding as “non-functional” cameras that were in fact working, but issued no tickets by chance. To check that our results still hold when we omit locations with an average lower number of tickets recorded on any given day, we re-run the regression just using places where having zero tickets is rare. Results in Panel A using as the key

explanatory variable the interaction between *Post* and *Mean Intensity* show an increase in the number *on-the-ground* tickets issued, consistent in magnitude with those reported in Table 2.3 with larger standard errors.

Panel B of Table 2.6 shows results using nighttime violations only, as a placebo check since cameras did not have night vision and were not functional during this time. We find that in two out of the three specifications, the results are insignificant, and much smaller in magnitude than those when we use daytime data.¹⁸

Figure 2.6 is used to ease potential concerns with statistical inference in the context of few clusters. We perform permutation inference by randomly re-allocating the “mean intensity” vector across locations 5000 times, and estimating how extreme the observed treatment effect is relative to these mean-zero synthetic placebo treatments. We note that the observed coefficient (corresponding to Panel C, Table 2.3, column (3)) is larger than 98% of permuted treatment effects.

Finally, in Table 2.7 we change how we define a CCTV “location” by varying the distance between a CCTV camera and the violations issued by *on-the-ground* officers considered to be in the “same” approximate location. We do this to test the robustness of our results to measurement error in our approximate violation coordinates, which are based on groups of written location descriptions in the police department’s administrative data. The result of these groupings are the two maps in Figures 2.1 and 2.2, the first (described in the Data section) showing approximate violation locations, and the second showing an

¹⁸We selected the functioning hours to be between 8 am and 7 pm. In robustness checks not shown, where we modify these hours slightly, our results remain consistent with the ones reported here.

example of buffer areas larger than 100 meters used in Table 2.3.

As we increase this distance to the closest camera, there could be two separate effects. First, it is likely that the traffic officer would be so far away from the CCTV that effectively, he/she would not be exposed to monitoring by the camera and the control room. This should lead to lowering of the impact of intensity of functionality of cameras on number of traffic violations. Second, it is also possible that by increasing the possible area of scope, we are reducing the measurement error in mapping the original violation. This could lead to an increase in the (potentially) downward biased estimates found in Table 2.3. Columns (1) to (6) in Table 2.7 show the estimates as distance increases from 100 meters to 800 meters. We find similar results with a slight increase in the magnitude of the estimate when distance of closest camera increases up to 300m, but then there is a sharp decline at 600m in the estimates of monitoring on violations. This may imply that after correction for the measurement error, extending the distance of the closest camera from the likely violation site leads to a dilution of the main effects.

2.8 Conclusion

Well-documented evidence of absenteeism and shirking by public sector workers in developing countries, and India in particular, has shown that this behavior translates into inefficiencies, decreased learning, and poor service delivery (Chaudhury and Hammer 2004, Chaudhury et al. 2006, Duflo et al. 2012). We add to this literature by studying a different group of public officials, traffic officers, and find that a surveillance technology meant mainly for drivers resulted in significant reductions in shirking among officers.

Our context is the Indian city of Chandigarh, where CCTV cameras were installed at twenty major intersections. To isolate the effect on traffic officers, we control for any potential changes in driver behavior due to the presence of cameras by exploiting the variation in how often the cameras worked at different locations. Regardless of the model used, we find consistent results indicating that the number of tickets issued by officers *on-the-ground* increased after the installation of the cameras. Although we cannot totally rule-out possible decreases in rent-seeking behavior in the vicinity of the cameras, our results appear to be driven by reductions in shirking. The increase in violations issued by *on-the-ground* officers are predominantly to cars and motorcycles with local registration and for violations that are easy-to-monitor for the CCTV cameras, such as failing to stop at a pedestrian crossings or failing to wear a helmet when on a motorcycle.

Banerjee et al. (2021) indicate that reforms to improve policing in India can be challenging due to the hierarchical nature of the organization. In a context of increases in worldwide adoption of CCTV cameras for traffic and crime monitoring, our results suggest that policing can be improved by exploiting the complementarities between *on-the-ground* enforcement and remote monitoring of both drivers and officers. This is the case even when, as in our context, camera operators' distinct role is not to supervise the behavior of on-the-ground officers. These complementarities are helpful not only in improving the efficiency, institutions, and delivery of public services, but they can also be leveraged to increase road safety. This is important as the majority of road traffic deaths occur overwhelmingly in developing countries (Duflo et al. 2019), and globally the World Health Organization (2018)

reports that road traffic injuries are now among the top 10 leading causes of deaths for all age groups.

2.9 Figures for Chapter 2




Figure 2.1: Map of Chandigarh and Location of CCTV Cameras. Gray dots indicate approximate locations of *on-the-ground* violations, black triangles indicate the location of CCTV cameras, and gray-filled triangles indicate placebo locations.



Figure 2.2: Location of CCTV Cameras and 300m Buffer Circles. Gray dots indicate approximate locations of *on-the-ground* violations, black triangles indicate the location of CCTV cameras, and circles indicate a 300 meter radius area from each CCTV camera.




Figure 2.3: **Daily Number of *On-the-Ground* Violations (left axis) and Cumulative Number of Days the Camera is Functional (right axis).** X-axis indicates number of days since the start of the sample. Figures by CCTV location using daily number of tickets recorded within 100 meters from each camera.




Figure 2.4: **Event Study Type Figure Using Months Pre/Post Camera Installation.** Figure reports coefficients from the regression specified in footnote 14 of section 2.5. Cameras were installed on May 7, 2010. Coefficients are normalized in relation to April 2010. Bars indicate 95% confidence intervals. Results show no pre-period trends and an increase in the number of *on-the-ground* violations in locations where cameras worked more often relative to locations where they worked less often, consistent with an incentive effect of monitoring.




Figure 2.5: **Event Study Type Figure Using Months Pre/Post Camera Installation on Driver Behavior.** Figure reports coefficients from the regression specified in footnote 14 of section 2.5 for the driver behavior test. Coefficients are normalized in relation to April 2010. Bars indicate 95% confidence intervals. Results show no pre-period trends and a decrease in the number of tickets of *on-the-ground* violations in locations with non-functioning cameras relative to nearby locations without cameras, consistent with drivers being more cautious in the presence of a camera regardless of whether it functions.




Figure 2.6: Permutation Inference Results. The figure presents a test of Fisher’s sharp null hypothesis of the “mean intensity” variable’s effect on the total number of traffic tickets issued. The “mean intensity” variable, a constant for each CCTV camera, was randomly re-assigned across locations 5000 times and the regression in equation 2.1 was estimated. Results correspond to Table 2.3, Panel C, column (3). The true observed coefficient exceeded 98% of these mean-zero permuted treatment effects (in absolute value), yielding the permutation test p-value.

2.10 Tables for Chapter 2

Table 2.1: Summary Statistics

VARIABLES	(1) N	(2) mean	(3) sd	(4) min	(5) max
Number of violations	23,660	5.343	9.016	0	193
Any violations	23,660	0.648	0.478	0	1
Functional	23,660	0.219	0.414	0	1
running	23,660	2.207	4.044	0	19.90
Mean intensity	23,660	0.413	0.342	0.00145	0.939
Post	23,660	0.585	0.493	0	1
Daytime	23,660	1	0	1	1
Out of state plate	23,660	2.496	4.096	0	81
Big vehicle	23,660	0.189	0.615	0	14
Car	23,660	2.189	4.701	0	123
Motorcycle	23,660	1.618	2.911	0	46
anycctv	23,660	4.213	7.841	0	193
Major road-rule (roadrules)	23,660	2.428	6.460	0	177
No or unauthorized helmet/unsafe motorcycle (helmet)	23,660	0.957	2.500	0	43
Other CCTV (other cctv)	23,660	0.828	1.656	0	41
nocctv	23,660	1.130	3.274	0	93
Parking tickets (parking)	23,660	0.177	0.784	0	20
License/insurance/registration (papers)	23,660	0.278	1.108	0	44
Dangerous driving (dangerous)	23,660	0.205	1.651	0	44
All other (other)	23,660	0.470	1.749	0	49

Note: The unit of observation is the CCTV location by date, using only daytime observations. *Functional* is a dummy variable equal to one when a CCTV camera recorded any violations. *Months' Running* is the cumulative number of days a CCTV camera was working after May 7, 2010, the day when the cameras were installed (divided by 30). *Mean Intensity* is defined as the percentage of time that the CCTV camera was functioning since the installation of the CCTV cameras. *Post* is a dummy variable equal to one after installation of the CCTV cameras. Panels B and C report the number of violations recorded according to each of the characteristics given.

Table 2.2: Effect of CCTV Cameras *Prior* to the Installation of the Cameras (Pre-Trends)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Effect of CCTV cameras on Officers Ticketing Behavior						
(Mean Intensity)*Post	-0.817 (1.696)	-0.233 (1.540)	-0.385 (1.480)	-0.218 (1.487)	-0.016 (1.444)	0.042 (1.403)
Observations	9,820	9,820	9,820	9,820	9,820	9,820
Panel B: Effect of Non-functional CCTV cameras on Drivers						
(Non-Functioning Camera)*Post	0.354 (0.572)	-0.044 (0.688)	-0.138 (0.575)	-0.135 (0.608)	-0.186 (0.660)	-0.266 (0.620)
Observations	6,383	6,383	6,383	6,383	6,383	6,383
Months before actual <i>Post</i>	2	4	6	8	10	12

Notes: Each coefficient comes from a separate regression. The unit of observation is the CCTV location by date, using only daytime observations. This table uses only data prior to installation of the cameras. Panel B includes only ‘Non-Functioning’ CCTV cameras, defined as the 8 locations where the cameras function less than 10% of the time, and compares them to 5 nearby placebo locations with comparable ticketing patterns that did not have a camera installed. All regressions include controls as described in equation 2.2, and report block bootstrapped standard errors re-sampled at the CCTV location level.

Table 2.3: Effect of CCTV Cameras on Officers *On-the-Ground* Ticketing Behavior

	Any violation	Violations	
	LPM	Poisson	Counts
	(1)	(2)	(3)
Panel A			
Functional*Post	0.118*** (0.045)	0.447*** (0.173)	2.469** (1.092)
Observations	23,660	23,660	23,660
R-squared	0.444		0.329
Panel B			
('Months' Running)*Post	0.022*** (0.008)	0.134*** (0.033)	0.614*** (0.188)
Observations	23,660	23,660	23,660
R-squared	0.450		0.347
Panel C			
(Mean Intensity)*Post	0.314** (0.126)	1.536*** (0.468)	7.921*** (2.944)
Observations	23,660	23,660	23,660
R-squared	0.451		0.344
Panel D: Driver Behavior			
(Non-Functioning Camera)*Post	-0.052 (0.103)	-0.392 (0.586)	-2.493 (1.541)
Observations	15,379	15,379	15,379
R-squared	0.513		0.492
Distance from CCTV	100m	100m	100m

Notes: Each coefficient comes from a separate regression. The unit of observation is the CCTV location by date, using only daytime observations. Panel D includes only 'Non-Functioning' CCTV cameras, defined as the 8 locations where the cameras function less than 10% of the time, and compares them to 5 nearby placebo locations with comparable ticketing patterns that did not have a camera installed. All regressions include controls as described in equation 2.1, and report block bootstrapped standard errors re-sampled at the CCTV location level. LPM = Linear Probability Model. *** p-value<0.01, ** p-value<0.05, * p-value<0.1.

Table 2.4: Violations by Type (Levels)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Observed by CCTV Cameras			Not Observed by CCTV Cameras			
Outcome variable:	Roadrules	Helmet	Other CCTV	Parking	Papers	Dangerous	Other
Panel A: Linear Regression							
Functional	1.761** (0.711)	0.518* (0.285)	0.145 (0.159)	0.020 (0.036)	-0.009 (0.050)	0.078 (0.090)	-0.044 (0.054)
Observations	23,660	23,660	23,660	23,660	23,660	23,660	23,660
Panel B: Linear Regression							
Months running	0.464*** (0.117)	0.081* (0.049)	0.035 (0.031)	0.015* (0.008)	0.002 (0.008)	0.024 (0.015)	-0.007 (0.017)
Observations	23,660	23,660	23,660	23,660	23,660	23,660	23,660
Panel C: Linear Regression							
(Mean Intensity)*Post	5.149*** (1.706)	1.475 (0.920)	0.786* (0.471)	0.181 (0.127)	0.070 (0.137)	0.181 (0.177)	0.078 (0.164)
Observations	23,660	23,660	23,660	23,660	23,660	23,660	23,660
Panel D: Linear Regression (Driver Behavior)							
(Non-Functioning Camera)*Post	-0.140 (0.431)	-1.437** (0.662)	-0.604* (0.344)	0.056 (0.074)	-0.081 (0.095)	-0.216* (0.118)	-0.071 (0.086)
Observations	15,379	15,379	15,379	15,379	15,379	15,379	15,379
Distance from CCTV	100m	100m	100m	100m	100m	100m	100m

Notes: Each coefficient comes from a separate regression. The unit of observation is the CCTV location by date, using only daytime observations. Panel D includes only ‘Non-Functioning’ CCTV cameras, defined as the 8 locations where the cameras function less than 10% of the time, and compares them to 5 nearby placebo locations with comparable ticketing patterns that did not have a camera installed. All regressions include controls as described in equation 2.1, and report block bootstrapped standard errors re-sampled at the CCTV location level. *** p-value<0.01, ** p-value<0.05, * p-value<0.1.

Table 2.5: Violations by Vehicle Type (Levels)

Outcome variable:	(1) Out of State	(2) Big Car	(3) Car	(4) Motorcycle
Panel A: All Violations				
(Mean Intensity)*Post	1.273 (1.298)	0.135 (0.091)	3.627*** (1.249)	2.028* (1.061)
Observations	23,660	23,660	23,660	23,660
Panel B: Observed by CCTV Cameras				
(Mean Intensity)*Post	1.010 (1.048)	0.072 (0.064)	3.306*** (1.186)	1.893* (0.980)
Observations	23,660	23,660	23,660	23,660
Panel C: Not Observed by CCTV Cameras				
(Mean Intensity)*Post	0.263 (0.271)	0.063** (0.030)	0.322** (0.152)	0.134 (0.095)
Observations	23,660	23,660	23,660	23,660
Panel D: Driver Behavior: All Violations				
(Non-Functioning Camera)*Post	-1.437 (0.880)	-0.111* (0.060)	-0.752* (0.456)	-1.169* (0.662)
Observations	15,379	15,379	15,379	15,379
Distance from CCTV	100m	100m	100m	100m

Notes: Each coefficient comes from a separate regression. The unit of observation is the CCTV location by date, using only daytime observations. Panel D includes only ‘Non-Functioning’ CCTV cameras, defined as the 8 locations where the cameras function less than 10% of the time, and compares them to 5 nearby placebo locations with comparable ticketing patterns that did not have a camera installed. All regressions include controls as described in equation 2.1, and report block bootstrapped standard errors re-sampled at the CCTV location level. *** p-value<0.01, ** p-value<0.05, * p-value<0.1.

Table 2.6: Main Results - Robustness

	LPM (1)	Poisson (2)	Counts (3)
<u>Panel A: High-Function Cameras</u>			
(Mean Intensity)*Post	0.518 (0.334)	2.640** (1.049)	11.378 (7.412)
Observations	14,196	14,196	14,196
R-squared	0.437		0.399
<u>Panel B: Nighttime</u>			
(Mean Intensity)*Post	0.118 (0.076)	0.578* (0.316)	0.445 (0.394)
Observations	23,660	23,660	23,660
R-squared	0.306		0.240
Distance from CCTV	100m	100m	100m

Notes: Each coefficient comes from a separate regression. The unit of observation is the CCTV location by date, using only daytime observations. All regressions include controls as described in equation 2.1, and report block bootstrapped standard errors re-sampled at the CCTV location level. LPM = Linear Probability Model. *** p-value<0.01, ** p-value<0.05, * p-value<0.1.

Table 2.7: Main Results by Distance to CCTV Camera

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome variable:	Count of Violations					
<u>Panel A: Poisson Model</u>						
(Mean Intensity)*Post	1.536*** (0.468)	1.532*** (0.468)	1.544*** (0.436)	1.470*** (0.369)	1.394*** (0.336)	0.519** (0.248)
Observations	23,660	23,660	23,660	23,660	23,660	23,660
<u>Panel B: Regression Model</u>						
(Mean Intensity)*Post	7.921*** (2.944)	7.927*** (2.943)	7.951*** (2.943)	8.451*** (2.827)	8.277*** (2.819)	6.213* (3.329)
Observations	23,660	23,660	23,660	23,660	23,660	23,660
Distance from CCTV	100m	200m	300m	400m	500m	800m

Notes: Each coefficient comes from a separate regression. The unit of observation is the CCTV location by date, using only daytime observations. All regressions include controls as described in equation 2.1, and report block bootstrapped standard errors re-sampled at the CCTV location level. *** p-value<0.01, ** p-value<0.05, * p-value<0.1.

Chapter 3: The Effect of Constrained Climate Policies: Evidence from Australia

3.1 Introduction

Economic theory suggests that carbon pricing is the most economically efficient policy instrument to mitigate climate change (Nordhaus, 1977; Poterba, 1991; Aldy et al., 2010). As of 2023, over 70 national or sub-national economies have introduced carbon pricing (World Bank, 2023), and hundreds of ex-ante studies have been carried out assessing the impacts of hypothetical carbon pricing schemes.¹ Yet a growing body of ex-post empirical evidence finds that carbon pricing has produced only modest reductions in emissions (e.g. Green, 2021).

Emissions reductions from carbon pricing policies may have been modest for a variety of reasons. It may be the case that prices levied on emissions have induced economically efficient, but small (potentially costly) reductions in emissions. On the other hand, emissions prices may be inefficiently low due to economic concerns or political concessions, such as exemptions built-in that limit policy coverage (Rafaty, Dolphin, and Pretis, 2020). Moreover, carbon pricing policies may be complex, leading to delays in emissions reductions, or emissions policies may have uncertain futures which lead firms to “wait-and-see”

¹Timilsina (2018) provides an introductory overview of different types of carbon pricing instruments, including carbon taxes, emissions trading system (ETS), and carbon offset mechanisms.

before changing input or output decisions.

This study provides insights on the efficacy of constrained carbon pricing instruments by focusing on the experience of Australia. We argue that Australia is a particularly interesting policy environment, exhibiting both complexity and significant uncertainty over a substantial national carbon pricing policy. In particular, over a backdrop of an expansion of subsidies for renewable energy generation in 2009, Australia announced a controversial carbon price in 2011, enacted in 2012, but subsequently repealed by 2015. To our knowledge, the focus of our study on a repealed policy is novel. Our study sheds light on the effects of political uncertainty on real-world environmental policies. Moreover, a mixed subsidy and “tax”² is a common feature of carbon pricing schemes proposed and in existence today.

Several interesting questions arise in the Australian context. What were the combined effects of the policies during the time of implementation? If emissions were reduced, were the reductions transitory, or did the policies shift emissions onto a long-run path below what would have occurred in the absence of the policy? Did firms sink costly investments that yielded future emissions reductions? Did they wait-and-see if uncertainty over the tax would be resolved? And finally, did the pricing policy lead to production leakage (was Australian coal shipped abroad during and after the tax was in place)?

We analyze these questions quantitatively. In the main part of the analysis, we apply synthetic control to Australia’s aggregate emissions from 1980-2018 in the spirit of Ander-

²For simplicity and clarity, we refer to the policy expansion of the existing policy as “subsidy”, and the new pricing policy as “tax” or “price” on carbon. We discuss the specific details of the policies in the following section.

sson (2019). There are several challenges unique to this context which we discuss in the following sections, most importantly for selection of a donor pool of economies to replicate the trajectory of Australia's emissions during the 2000s. Among other adjustments to the estimation, we limit the donor pool to countries with no national emissions policies during the sample period and focus on other high-emitting countries following the recommendations of Abadie (2021). We employ a battery of robustness and sensitivity checks including the analysis of placebo policies at different times, and leave-one-out analysis to ensure the results are robust to changes in the donor pool. We also analyze data that sheds further light on mechanisms, building intuition for the core results.

We estimate that Australia's increase in renewal subsidies in tandem with its short-lived carbon pricing scheme caused a 5-8% reduction in total emissions per capita from 2010 to 2018. We show that contemporaneous with policy implementation, power sector emissions declined as a result of declining coal-fired generation, while wind and especially solar generating capacity rapidly increased. We show that emissions declined rapidly while the pricing scheme was in place, and confirm qualitatively that emissions did not rebound to a level consistent with their likely path had the policy never been implemented. On the other hand, we present evidence that the rapid decline in emissions stalled after the repeal of the carbon price in 2015. Lastly, we show descriptive evidence suggesting that coal production and exports may have been accelerated by the implementation of the pricing scheme and the displacement of Australian coal.

3.2 Policy Details

Australia implemented a subsidy for renewable electricity generation in 2001 called the Renewable Energy Target (RET), which had modest objectives until 2009 when its ambition was substantially increased.³ The RET requires entities that acquire electricity in wholesale markets to also acquire green certificates in proportion to their market share. The green certificates can be created only by MWhs (MegaWatt-hours) from new renewable generation, effectively creating a tax and transfer system from load serving entities that sell retail electricity to new producers of renewable energy. Initially passed at the start of the 2000s, the target aimed to grow renewable generation to 2 percent of the national total by 2010. In 2009-2010 however, with broad legislative support, the RET increased its goal to an ambitious 20 percent of generation from renewable sources by 2020 (Cludius, Forrest, MacGill; 2014).

Australia also implemented a controversial short-lived carbon pricing scheme⁴ on the energy sector from 2012-2014, announced in 2011.⁵ As early as 2008, a series of referenda in the Australian congress strongly signalled the possibility of new climate policy. In 2010, the recently elected Prime Minister formed a multi-party commission to formally consider the details of a new national climate policy (Peel, 2014). The tax itself proposed to cover about 60% of Australia's total emissions by targeting the direct emissions (sometimes called Scope 1 emissions) of over 350 of the highest emitting firms in the country. Notably,

³<https://www.cleanenergyregulator.gov.au/RET/About-the-Renewable-Energy-Target/History-of-the-scheme>

⁴We interchangeably use "tax," for simplicity.

⁵<https://www.cleanenergyregulator.gov.au/Infohub/CPM/Pages/About-the-mechanism.aspx>

this designation did not cover emissions from exported fossil fuels. The overarching plan of the tax included a preliminary phase from 2012-2014, in which a fixed price fee of 23 to 25 AUD would be charged per ton of CO₂ equivalent emitted to covered firms, followed by a secondary phase in which the fee would float and Australian emissions allowances would integrate with world markets (specifically with the EU ETS). The plan was repealed in 2014, prior to integration.

Around the time that the carbon tax was implemented, there was considerable uncertainty about the future of Australian carbon pricing. In the summer of 2012, Jotzo, Jordan, and Fabian (2012) surveyed 76 experts from liable entities (including electricity generators), investors, advisors, and academics and found an unusual pattern in the responses. Among the surveyed experts, almost all expected the policy to be in place for at least the next two years, but 39% expected the policy to be repealed by 2016. Experts from liable entities were most “optimistic” in the sense that they were most likely to predict that the policy would be repealed. This suggests that firms may not have responded to the policy in meaningful ways because they accurately forecasted that it would not be binding in a few years. However, when asked whether there would be a carbon price in future years, 79% expected a price by 2020 and 81% by 2025. This suggests that over longer time horizons, Australian firms anticipated a carbon price and so may have responded to the policy with investment or other long-term decisions. Which effect dominated is ultimately an empirical question.

3.3 Methods: Synthetic Control and Extensions

To estimate the impacts of Australia's climate policies, we utilize the approach of the quantitative case study (Abadie and Gardazebal, 2003; Abadie et al, 2015). Recent developments in synthetic control and related estimation techniques have the advantage of minimal data requirements while providing a disciplined and transparent approach for the analysis of policies with aggregate data. Moreover, synthetic control and related techniques have been shown to be more adaptable than in its traditional form (Abadie, 2021), and to have desirable properties for the estimation of causal effects in many contexts.

In the basic setup for the case considered in the paper, we observe units $j = 1, \dots, J + 1$ for time periods $t = 1, \dots, T$. Without loss of generality let $j = 1$ be the treated unit, and consider an intervention that occurs at time $T_0 + 1 \leq T$. Denote observed outcomes (e.g. emissions per capita) by Y_{jt} and potential outcomes $Y_{jt}(1), Y_{jt}(0)$ with and without treatment, respectively. The treatment effect for unit j at time t is defined as

$$\tau_{jt} = Y_{jt}(1) - Y_{jt}(0). \quad (3.1)$$

Sometimes called the fundamental problem of causal inference, the object of interest $\tau_{1t}, t > T_0$ cannot be calculated because we only observe one of its two components. Synthetic control seeks to use a weighted sum of control units $2, \dots, J + 1$ to impute the missing counterfactual outcome of the first unit (i.e. annual emissions per capita in Australia had the tax not been implemented). The weights w_j^* are chosen subject to positivity and adding-up constraints $w_j \geq 0 \forall j$ and $\sum_{j=2}^{J+1} w_j = 1$. Thus the imputed counterfactual is a convex

combination of control units, and the estimated treatment effect is given by

$$\widehat{\tau}_{1t} = Y_{jt} - \sum_{j=2}^{J+1} w_j^* Y_{jt} \quad (3.2)$$

for periods $T_0 + 1, \dots, T$. We use the estimator implemented by Abadie et al. (2011), which defines a $(k \times 1)$ vector of characteristics of the treated unit X_1 , and a $(k \times J)$ matrix of the same characteristics for the control units X_0 . The matrices X_1, X_0 are typically composed of pre-period values of the variable of interest (i.e. emissions per capita) or other variables with significant predictive power. Stacking the weights into a length J vector $W = w_2, \dots, w_{J+1}$, the algorithm minimizes the objective function

$$\sqrt{(X_1 - X_0 W)' V (X_1 - X_0 W)} \quad (3.3)$$

where V is defined as some $(k \times k)$ symmetric and positive semidefinite matrix defining the relative importance of the k characteristics in the minimization. In practice, we rely on the default proposed by Abadie et al. (2011) which performs a two-step optimization, choosing both W and V jointly to minimize the objective function (in other words, to maximize pre-treatment fit).

There are many extensions to synthetic control that have been proposed to make the method more versatile and improve pre-treatment fit. In particular, in the case of Australia we implement the de-meaned estimator proposed both by Doudchenko and Imbens (2016) and Ferman and Pinto (2021), which we refer to as the DIFP estimator. We discuss reasons

why this estimator is a more suitable choice below. Intuitively, the DIFP estimator imputes the counterfactual $Y_{1t}(0), t > T_0$ with an included intercept, replacing equation (3.2) with

$$\widehat{\tau}_{1t} = Y_{jt} - \left(\mu + \sum_{j=2}^{J+1} w_j^* Y_{jt} \right) \quad (3.4)$$

where μ is also a parameter to be estimated. In practice, the DIFP estimator is implementable by demeaning observations for each country from a pre-period value and then estimating the synthetic control as normal. This poses no additional complications for computation or inference, and just requires adjusting expression (3.3) so that X_1, X_0 are replaced with their time-demeaned analogues $\check{X}_1, \check{X}_0$.

In future versions of the paper, we plan to consider the algorithm proposed in Ben-Michael et al. (2021) which uses regularization to allow for negative weights to improve pre-treatment fit, and consider more formal inference procedures as proposed by Cattaneo, Feng, and Titiunik (2021).

3.4 Data

We use annual country-level panel data over the period 1980-2018, with restrictions described below. In the analyses of per capita emissions, we form a control group that replicates the pre-existing trend in emissions prior to implementation of the new climate policies. As pre-period characteristics we consider some subset of the following variables that provide a good fit: emissions shares from different sectors, GDP per capita, population, and energy consumption of different emitting fuels (coal, natural gas, and petroleum

products).

The outcome variable for the main analysis is the emissions of CO₂ per capita. Emissions data come from the Emissions Database for Global Atmospheric Research (EDGAR). The EDGAR dataset uses a bottom-up methodology to calculate emissions in the power, industry, buildings, transport, and agricultural sectors. The methodology is applied consistently to all countries to facilitate comparison across countries.⁶

Two variables generally considered important predictors of emissions, and that are indeed correlated with emissions across countries, are population size and economic output. We use the harmonized long run estimates of GDP per capita and population in the MADISON database, which utilize similar methodology and are published alongside the Penn World Tables (PWT).⁷

Finally, to obtain a good fit and improve the credibility of the synthetic control estimates, we use fuel consumption of coal, natural gas, and petroleum products from the Energy Information Administration (EIA) International Energy Statistics.⁸

3.5 Estimation Details

The primary difficulty in applying synthetic control to Australia's climate policies is the country's status as an outlier in terms of emissions per capita. Australia's electricity sector (and exports) are highly coal-dependent and therefore emissions-intensive. Moreover, many natural comparison countries that might help form a plausible counterfactual

⁶Further information and downloads are available at https://edgar.jrc.ec.europa.eu/report_2022.

⁷Available at <https://www.rug.nl/ggdc/historicaldevelopment/maddison>. In practice, PWT and MADDISON GDP and population figures are almost identical for the periods in this study.

⁸<https://www.eia.gov/international/data/world>.

for Australia's emissions must be omitted from the sample because they implemented policies that would significantly affect emissions during the sample period. For instance, the implementation of the European Union's Emissions Trading System (EU ETS) in 2005 disqualifies all EU countries.

We form the donor pool for Australia in three main steps. First, we construct a global panel dataset of emissions, GDP per capita, and energy consumption using the data sources previously described. Second, we use data provided by the World Bank (primarily from the Carbon Pricing Dashboard)⁹ to identify every country that implemented a national-level carbon price or ETS from 1980-2017 and disqualify them from the donor pool. Third, to address Australia's position as an outlier in emissions per capita, we limit the donor pool to countries with the top 35 average highest emissions per capita over the period 1980-2000.¹⁰ The resulting donor pool is listed in Appendix C.1. This final restriction makes motivation of the DIFP estimator simpler (explained below), speeds computation, and helps maintain sparsity of the nonzero weights chosen by the algorithm. Including all countries with complete data does not meaningfully change the results.

Figure 3.1 motivates the use of the DIFP estimator. The solid line displays Australia's total emissions per capita from 1980-2018 while the dotted line displays the donor pool average over the same period. Both groups show a gradual upward trend for most of the sample period, while Australia's emissions declined starting in the mid- to late-2000s. The

⁹<https://carbonpricingdashboard.worldbank.org>. The Australian carbon tax, incidentally, does not appear in the World Bank's carbon tracker data.

¹⁰We then also exclude Brazil and Syria from the donor pool due to their economic crises and war, respectively, in the 2010s which caused their emissions to greatly decline.

broadly similar trend early in the sample (and the significant heterogeneity within the donor pool, not pictured) suggests that it may be difficult to match Australia’s *level* of emissions per capita, but it may be possible to match the trend, so this is the strategy we opt for. In practice, we implement the DIFP estimator by fitting a synthetic control model after “deameaning” the pre-period data by subtracting out its average emissions per capita from 1980-2000, as suggested by Doudchenko and Imbens (2016).

Given the year that the RET was updated (2009) and the ongoing controversy of a carbon pricing policy, we define T_0 , the last year before treatment, to be 2008. This allows us to simultaneously estimate the impact of the two policies by estimating the counterfactual and examining the dynamic treatment effect as uncertainty over the tax unfolded. This approach has the additional benefit, for analyzing the tax, of conforming with the back-dating recommendation of Abadie (2021) to incorporate possible anticipation effects.

3.6 Results

Figure 3.2 illustrates the fit of our preferred synthetic control, where again the solid line represents Australia’s true emissions profile over time. We use a mixture of lagged values of per capita emissions (prior to implementation of the tax) and covariates theoretically likely to be strong predictors of emissions. Figure 3.2 shows the model fit using the covariates displayed in Table 3.1 in addition to emissions per capita in 2001 and 2007-2008. The fit of the synthetic control has some error, but importantly is neither systematically high nor low, over-predicting in the mid-1990s and the mid-2000s, but under-predicting in the early 1990s and early 2000s. There is a significant under-prediction in the year 2009, but in 2010

the synthetic Australia and true Australia have almost exactly the same value. After 2010, Australia's emissions noticeably decline relative to its synthetic counterpart.

The estimated treatment effect is the gap between the synthetic Australia and true Australia after 2008 (T_0 , represented by the vertical dashed line). The estimated treatment effect increases from 2011-2014 when the tax was announced and in place (indicated by the shaded vertical bars), before immediately flattening in 2015 with the repeal of the policy. The treatment effect is most easily read from the axis on the right side of the figure, which adds back Australia's pre-period intercept to both series. The estimated average emissions reduction over the period 2011-2018 is just above -1.05 metric tons of CO₂ per capita, or approximately 5.4% of emissions per capita in 2008. The pre-period average (absolute) deviation is approximately 0.218 metric tons, or about 1.1% of emissions per capita in 2008.¹¹ Relative to the standard deviation of Australia's emissions per capita in the pre-period, the estimated effect size is about 0.61 standard deviations while the pre-period error is approximately 0.13 standard deviations.

Table 3.1 displays further details on covariates in the pre-period for Australia, the average of the donor pool and the synthetic unit. Each row displays the covariate average over the period 1998-2008 (the 10 years prior to T_0). As mentioned above, Australia (the "Treated" unit) has a high share of emissions from the power sector, is relatively rich, and has a very high coal consumption per capita relative to the donor pool average. The synthetic unit does not balance perfectly on pre-period characteristics - however, it bet-

¹¹The ratio of post-to-pre period mean squared prediction error is more than 21, but this quantity is best understood in relative terms, which we address after presenting the main results.

ter matches Australia on several important dimensions. In particular, it closely matches the share of power sector emissions, income per capita much more closely, and is much more similar in coal consumption per capita than the donor pool average. On the other hand, it has slightly worse fit in terms of transportation share of emissions and natural gas consumption per capita.

Table 3.2 shows the composition of the synthetic unit, displaying non-zero weights from the donor pool. The chosen weights are importantly quite sparse, with the synthetic unit composed mostly of Taiwan and the United States, with approximately 7% the combined economies of Israel and Palestine, with each of Hong Kong and South Africa receiving under 5%. We interpret Taiwan and Hong Kong as capturing regional (Pacific) economic and policy shocks. The United States and South Africa both have cultural similarities to Australia, particularly through British colonial history (as does Hong Kong). They also are significant producers, exporters, and consumers of coal, which helps to balance the synthetic unit with Australia in terms of power sector emissions share and coal consumption per capita prior to the implementation of Australia's new climate policies.

To build confidence in the empirical strategy, Figure 3.3 shows a leave-one-out sensitivity test of the result from Figure 3.2. To produce the figure, each of the five economies with nonzero weights in Table 3.2 was dropped from the donor pool and the synthetic control was re-estimated. The results show that the counterfactual path of Australian emissions does not qualitatively depend on the choice of the five economies listed in Table 3.2. In two cases the estimated treatment effect would have been qualitatively similar but smaller,

in two cases the estimated treatment effect would have been similar, and in one case the treatment effect would have been larger. Overall, the figure shows that the qualitative conclusions and approximate magnitudes described above are not idiosyncratically dependent on the selection of the five economies in the optimal synthetic unit.

To further build confidence in the empirical strategy, Figure 3.4 displays a placebo-in-time test of the methodology that produced Figure 3.2. To produce the figure, we backdated the pre-treatment period T_0 by 10 years, as well as all covariates from Table 3.1, and evaluated a fake policy affecting emissions in Australia in the year 1998 during the period 1998-2008.¹² Again, the synthetic control was able to replicate the trend in emissions per capita quite well for 20 years prior to the policy. In the case of Figure 3.4 however, there is no noticeable divergence between the synthetic unit and Australia's true emissions following the fake policy in 1998 or in 2001, as we should expect if the estimation strategy is valid. In Appendix C.2 we shift the placebo treatment forward and back by two years (1996 and 2000) to show that the placebo result is not dependent on the specific year chosen for "implementation" of the fake policy.

Finally, in Figures 3.5 and 3.6 we show that the estimated treatment effect in Australia is extreme relative to other countries in the donor pool using the permutation test proposed in Abadie et al (2010). The dark black line in Figure 3.5 shows the "gaps" between Australia's emissions per capita and its synthetic counterpart from 1980-2018. The light gray lines are the corresponding figure for each of the 33 donor pool countries, excluding those with

¹²Covariates are therefore averaged over the period 1988-1998, and lags of emissions per capita are from 1991 and 1997-1998.

much higher prediction error in the pre-period. Specifically, we exclude countries for which the mean squared prediction error was 5 or more times greater than Australia's pre-period MSPE.¹³

In Figure 3.6 we show the distribution of post-to-pre MSPE ratios for the same economies depicted in Figure 3.5. Australia's MSPE ratio between the post- and pre-periods was more than 20, and the 2nd most extreme of any of the other units included in the figure. Because permutation p-values can be highly sensitive to the number of sample countries considered in the donor pool, in future drafts we plan to explore uncertainty quantification using prediction intervals rather than relying on permutation procedures. Still, it is encouraging the Australia's estimated effect size is extreme regardless of the countries we include, as shown in Appendix C.4.

3.7 Mechanisms

Some descriptive analysis gives insight into what drives the emissions reductions estimated in the previous section. Figure 3.7 shows which sectors account for the estimated reduction in emissions per capita. From Figure 3.7 it is evident that total emissions declines are generally dominated by the power sector, and further that there is little or no evidence of declining emissions from other sectors corresponding to the timing of the policy changes. In Appendix C.5, we confirm that coal was the main fuel whose consumption meaningfully declined after implementation of the policies.

¹³ Appendix C.3 repeats this exercise with a less stringent MSPE threshold (20) and a more stringent MSPE threshold (2).

Figure 3.8 digs deeper into mechanisms, plotting electricity capacity by resource type according to the EIA. A few interesting facts appear. From Figure 3.8 it appears that the implementation of the pricing policy effectively halted fossil fuel capacity investments, evidence consistent with the idea that firms were unwilling to make irreversible investment decisions after the passage of the tax. Moreover, there appears to be a spike in fossil fuel capacity additions after its repeal. Additionally, in Figure 3.8 we see a substantial rise in renewable capacity, especially in solar generation. In Appendix C.5 we confirm that fossil fuel plants lost some market share (in terms of energy generated) to renewable, wind, solar, and hydro resources during this time. Rather than increasing with total electricity generation, fossil fuel generation declined from 2011-2014 and subsequently remained stable through 2018.

3.8 Exports

Some descriptive analysis also sheds light on the issue of leakage. This evidence should be taken with even more caution than the previous section, as it is an entirely different sector and in no way a decomposition of the effects estimated earlier in the paper. Still, the picture that emerges from Figure 3.9 is that coal production in Australia continued to increase during the policy period and may have even accelerated. Moreover, it seems likely that exports increased during the policy period given the apparent change in slope of the time series from 2011-2014.

In future versions, we plan to perform time series analysis of production and exports to corroborate the visual evidence in Figure 3.9. Preliminary tests indicate that both time se-

ries (production and exports) have unit roots, using the augmented Dickey-Fuller approach proposed in Elliott, Rothenberg, and Stock (1996). However, the empirically interesting question is whether or not these potentially non-stationary process exhibit a structural break around the times that new policies were implemented, and distinguishing between non-stationarity and structural breaks is not always a simple task.

We propose to use a traditional test for multiple structural breaks at unspecified times in the style of Bai and Perron (1998). Estimation will be performed in the context of two different time series models, with the dependent variable Y_t being coal exports in both cases. The first model is a regression only on the year variable (equation 3.5), and so would effectively test whether a linear trend or a linear trend with a break fits the export data better. The second (equation 3.6) would regress coal exports on coal production. The idea is that some portion of production is exported, and in theory the test could detect whether an increased share of production was exported when climate policy changed.

$$Y_t = \beta_2 t + \varepsilon_t \tag{3.5}$$

$$Y_t = \beta_1 X_t + \eta_t \tag{3.6}$$

3.9 Conclusion

This analysis explores the effect of real-world climate policies in Australia: a mixture of taxes and subsidies in an environment of major political uncertainty. Our analysis is novel in using synthetic control and its variants to estimate the effects of the Australian carbon

pricing policies on total CO2 emissions using aggregate data, and focusing on a policy that was ultimately repealed. If nothing else, the analysis shows that a patchwork of climate policies can be effective in reducing emissions, and that political constraints materially impact the outcomes of environmental policy, sometimes in ways that are predictable.

In the case of Australia, our analysis suggests that Australia's tax had a substantial impact on total emissions, increasing while the tax was in place and eventually levelling out over time. However, we highlight again the important caveat that what we are able to capture here is the effect of the suite of emissions-related policies implemented in Australia, most importantly to include the increase in renewable capacity subsidies at roughly the same time as the new pricing policy. Strikingly, in the 4 years after the policy was repealed, we don't see clear evidence that the emissions rebounded to a level consistent with expected emissions in the absence of a tax. Thus we cannot rule out that the tax has had a lasting effect that may continue today.

On the other hand, we find that the momentum of the decline in emissions per capita halted when the carbon price was repealed, and that the policy (which included no border adjustment and did not cover embodied emissions from produced fossil fuels) appears to have induced increased exports of Australian coal. While it seems clear then that Australia changed the trajectory of its emissions around 2011, our treatment effect may overstate the global impact on emissions if Australian coal was cheap enough to induce increased consumption by its trading partners. More to the point, we cannot rule out a complete reversal and rebound of emissions past our final sample year of 2018. We believe a total reversal of

this type to be unlikely given the long-lived nature of electric capacity investments and the clear and sustained upward trend in wind and solar power through the end of our sample.

3.10 Figures for Chapter 3




Figure 3.1: **Emissions per capita in Australia 1980-2018 versus donor pool countries**




Figure 3.2: **Path plot of emissions per capita in Australia 1980-2018 versus synthetic Australia**




Figure 3.3: **Leave-one-out path plots of emissions per capita versus synthetic Australia**




Figure 3.4: **Placebo-in-time test**




Figure 3.5: **Permutation tests.** Per capita emissions in Australia and placebo gaps for donor pool countries.




Figure 3.6: **Permutation tests.** Distribution of gaps and placebo gaps.




Figure 3.7: **Australia's emissions by sector**




Figure 3.8: Australia's electric generation capacity




Figure 3.9: Australia's coal production and exports

3.11 Tables for Chapter 3

Table 3.1: Emissions per Capita Predictors Before Policy Changes

Variables	Treated	Synthetic	Sample.Mean
Emissions share (power)	0.52	0.49	0.36
Emissions share (transportation)	0.21	0.20	0.20
Emissions share (other)	0.27	0.30	0.44
Log of gdp per capita (USD 000s)	3.72	3.57	2.68
Btus of coal per capita (Hundred MMs)	1.02	0.61	0.16
Btus of ng per capita (Hundred MMs)	0.53	0.43	0.47
Btus of petroleum per capita (Hundred MMs)	0.94	1.02	0.73
Emissions per capita in 2001	2.58	2.52	1.15
Emissions per capita in 2007-2008 (avg.)	3.00	2.96	2.26

Table 3.2: Nonzero Country Weights in Synthetic Australia

Weight	Country (Unit)
0.55	Taiwan
0.31	United States
0.07	Israel and Palestine, State of
0.04	South Africa
0.03	Hong Kong

Appendix A: Appendices for Chapter 1

A.1 Supplementary Tables

Table A.1: Supplemental Summary Statistics

	(1) N	(2) mean	(3) sd	(4) min	(5) max
Below poverty line (000s)	254	9.019	16.58	0.175	180.2
Median Home Value (000s)	254	98.09	49.85	30.50	417.1
Median Rent	254	380.9	105.8	176	889
Annual Inmigr.	254	2,212	5,416	66	76,116
Annual Outmigr.	254	2,213	5,661	86	78,034
Teachers	247	635.8	1,268	17.20	12,654
Students	254	9,648	20,903	221	251,213
Students/Teachers	254	15.16	2.112	9.600	22.20
GP Rev per capita (000s)	254	1.946	1.302	0.302	6.764
SD Rev per capita (000s)	254	1.898	0.869	0	5.298
IRS exemptions (000s)	254	57.47	137.2	1.862	1,673
IRS Returns (000s)	254	28.74	71.68	0.857	873.5

A.2 Data Appendix

This section describes the data sources and cleaning in more detail.

A.2.1 Energy Market Data

Electricity market data comes from the EPA's Continuous Emissions Monitoring Systems data. This contains hourly gross generation and emissions data for many of the fossil fuel power plants in the US. I downloaded these data directly from the EPA's FTP site.

Data on the **supply side of the coal market** comes from the Mine Safety and Health Administration (MSHA) and EIA Form 7A, which includes production in short tons and average employment for mines across the US. **Demand side data on the coal market** is derived from a combination of EIA (joint with FERC) 423 and EIA 923 data.

Combining the demand and supply side of the coal market, **fuel prices** are taken from the EIA's State Energy Data System (SEDS).

A.2.2 Outcomes

Wages, employment, unemployment, and labor force participation all come from the BLS. I corroborated income effects with data from the BEA and IRS, but ultimately cut those estimates from the paper.

Public finance data comes from the Census of Governments which I obtain from the Government Finance Database (<https://willamette.edu/mba/research-impact/public-datasets/index.html>)

Population data comes from various US Census datasets, most importantly the intercensal population estimates (both total and broken down by age and sex). **Immigration**

and outmigration are obtained from the IRS's county-to-county migration files, which they infer from changes in tax filings.

Median home values, rents, and educational attainment at the county level are derived from the American Community Survey (ACS) and collected via IPUMS NHGIS. To build the long differences, I use the decennial census in 2000 as an estimate of $t_1 = 2002$, and other values are inferred using the 5-year files (2005-2009, 2009-2013, 2014-2018). I considered other outcomes from this Census/ACS combination, but ultimately ruled them out because they did not appear to harmonize well.

Crime rates come from the FBI's Uniform Crime Reporting Data and **student and teacher counts** come from the National Center for Education Statistics.

A.2.3 Other (including controls)

Oil and gas production value come from Enverus (formerly DrillingInfo) through a subscription. These data provide gas and oil production at the well level for most areas of the country, which I aggregated up to the county-year level and valued using the West Texas Intermediate (WTI) and Henry Hub prices available from the EIA.

Appendix B: Appendices for Chapter 2

Table B.1: Effect of CCTV Cameras *Prior* to the Installation of the Cameras (Pre-Trends)

	(1)	(2)	(3)	(4)	(5)	(6)
Effect of CCTV cameras on Officers Ticketing Behavior						
<u>Panel A: LPM</u>						
(Mean Intensity)*Post	-0.011 (0.142)	-0.030 (0.132)	-0.031 (0.125)	-0.015 (0.122)	-0.026 (0.126)	-0.010 (0.123)
Observations	9,820	9,820	9,820	9,820	9,820	9,820
<u>Panel B: Poisson Model</u>						
(Mean Intensity)*Post	-0.289 (0.409)	-0.150 (0.399)	-0.188 (0.386)	-0.139 (0.377)	-0.089 (0.370)	-0.083 (0.378)
Observations	9,820	9,820	9,820	9,820	9,820	9,820
<u>Panel C: Linear Regression</u>						
(Mean Intensity)*Post	-0.817 (1.696)	-0.233 (1.540)	-0.385 (1.480)	-0.218 (1.487)	-0.016 (1.444)	0.042 (1.403)
Observations	9,820	9,820	9,820	9,820	9,820	9,820
Effect of Non-functional CCTV cameras on Drivers						
<u>Panel D: LPM</u>						
(Non-Functioning Camera)*Post	0.037 (0.084)	0.022 (0.102)	0.011 (0.100)	0.009 (0.093)	0.015 (0.097)	0.016 (0.097)
Observations	6,383	6,383	6,383	6,383	6,383	6,383
<u>Panel E: Poisson Model</u>						
(Non-Functioning Camera)*Post	-0.006 (0.539)	-0.026 (0.739)	-0.111 (0.712)	-0.066 (0.683)	-0.034 (0.713)	-0.039 (0.666)
Observations	6,383	6,383	6,383	6,383	6,383	6,383
<u>Panel F: Linear Regression</u>						
(Non-Functioning Camera)*Post	0.354 (0.572)	-0.044 (0.688)	-0.138 (0.575)	-0.135 (0.608)	-0.186 (0.660)	-0.266 (0.620)
Observations	6,383	6,383	6,383	6,383	6,383	6,383
Months before actual Post	2	4	6	8	10	12

Note: The unit of observation is the CCTV location by date, using only daytime observations. The camera variable in panels D-F takes a value of one for non-functioning CCTV cameras, defined as locations where the cameras function less than 10% of the time, and a value of zero for locations without CCTV cameras. This table uses only data prior to installation of the cameras. All regressions include controls as described in equation 2.2, and report block bootstrapped standard errors clustered at the CCTV location level. LPM = Linear Probability Model.

Appendix C: Appendices for Chapter 3

C.1 Donor Pool

Table C.1: Donor Pool for Australia

(1)	(2)	(3)
Kuwait	South Korea	Jamaica
United States	Taiwan	Algeria
Canada	Hong Kong	Lebanon
Australia	Venezuela	Turkey
Saudi Arabia	Mongolia	China
Trinidad and Tobago	Iraq	Botswana
Oman	Iran	Tunisia
Singapore	Malaysia	Thailand
Libya	Jordan	Ecuador
South Africa	Argentina	Angola
Israel and Palestine, State of	Cuba	Egypt

C.2 Alternate Placebo-in-time

Figure C.1: Placebo-in-time (2000)



Figure C.2: Placebo-in-time (1996)



C.3 Alternate Permutation Gaps Plots

Figure C.3: Placebo-in-space Plot (MSPE ratio threshold: 20)





Figure C.4: Placebo-in-space Plot (MSPE ratio threshold: 2)



C.4 Alternate Permutation Distributions

Figure C.5: Placebo-in-space Distribution (MSPE ratio threshold: 20)





Figure C.6: Placebo-in-space Distribution (MSPE ratio threshold: 2)



C.5 More Detail on Mechanisms

Figure C.7: Australia's Energy Consumption by Source





Figure C.8: Australia's Electricity Generation



Bibliography

- Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2):391–425.
- Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program. *Journal of the American statistical Association*, 105(490):493–505.
- Abadie, A., Diamond, A., and Hainmueller, J. (2011). Synth: An R package for synthetic control methods in comparative case studies. *Journal of Statistical Software*, 42(13).
- Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2):495–510.
- Abadie, A. and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132.
- Aldy, J. E., Krupnick, A. J., Newell, R. G., Parry, I. W. H., and Pizer, W. A. (2010). Designing climate mitigation policy. *Journal of Economic Literature*, 48(4):903–934.
- Allcott, H. and Keniston, D. (2018). Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. *The Review of Economic Studies*, 85(2):695–731.
- Andersson, J. J. (2019). Carbon taxes and CO2 emissions: Sweden as a case study. *American Economic Journal: Economic Policy*, 11(4):1–30.
- Ariel, B., Farrar, W. A., and Sutherland, A. (2015). The effect of police body-worn cameras on use of force and citizens’ complaints against the police: A randomized controlled trial. *Journal of Quantitative Criminology*, 31(3):509–535.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.

- Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American economic review*, 103(6):2121–2168.
- Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, pages 47–78.
- Bandiera, O., Prat, A., and Valletti, T. (2009). Active and passive waste in government spending: Evidence from a policy experiment. *American Economic Review*, 99(4):1278–1308.
- Banerjee, A., Chattopadhyay, R., Duflo, E., Keniston, D., and Singh, N. (2021). Improving police performance in Rajasthan, India: Experimental evidence on incentives, managerial autonomy and training. *American Economic Journal: Economic Policy*, 13(1):36–66.
- Bartik, A. W., Currie, J., Greenstone, M., and Knittel, C. R. (2019). The local economic and welfare consequences of hydraulic fracturing. *American Economic Journal: Applied Economics*, 11(4):105–55.
- Ben-Michael, E., Feller, A., and Rothstein, J. (2021). The augmented synthetic control method. *Journal of the American Statistical Association*, 116(536):1789–1803.
- Bertrand, M., Djankov, S., Hanna, R., and Mullainathan, S. (2007). Obtaining a driver’s license in India: An experimental approach to studying corruption. *The Quarterly Journal of Economics*, 122(4):1639–1676.
- Beuermann, D. (2011). Telecommunication technologies, agricultural profitability and child labor in rural Peru. Working paper, Central Bank of Peru, Working Paper Series.
- Black, D., McKinnish, T., and Sanders, S. (2005a). The economic impact of the coal boom and bust. *The Economic Journal*, 115(503):449–476.
- Black, D. A., McKinnish, T. G., and Sanders, S. G. (2005b). Tight labor markets and the demand for education: Evidence from the coal boom and bust. *ILR Review*, 59(1):3–16.
- Bobonis, G. J., Cámara Fuertes, L. R., and Schwabe, R. (2016). Monitoring corruptible politicians. *American Economic Review*, 106(8):2371–2405.
- Borcan, O., Lindahl, M., and Mitrut, A. (2017). Fighting corruption in education: What works and who benefits? *American Economic Journal: Economic Policy*, 9(1):180–209.
- Borenstein, S., Bushnell, J. B., and Wolak, F. A. (2002). Measuring market inefficiencies in California’s restructured wholesale electricity market. *American Economic Review*, 92(5):1376–1405.

- Brown, M. and Heywood, J. S. (2005). Performance appraisal systems: Determinants and change. *British Journal of Industrial Relations*, 43(4):659–679.
- Burgess, S. and Ratto, M. (2003). The role of incentives in the public sector: Issues and evidence. *Oxford Review of Economic Policy*, 19(2):285–300.
- Busso, M., Gregory, J., and Kline, P. (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, 103(2):897–947.
- Callaway, B. and Sant’Anna, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.
- Cattaneo, M. D., Feng, Y., and Titiunik, R. (2021). Prediction intervals for synthetic control methods. *Journal of the American Statistical Association*, 116(536):1865–1880.
- Chang, Juin-Jen and Lai, Ching-Chong (1999). Carrots or sticks? A social custom viewpoint on worker effort. *European Journal of Political Economy*, 15(2):297–310.
- Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K., and Rogers, F. H. (2006). Missing in action: Teacher and health worker absence in developing countries. *Journal of Economic Perspectives*, 20(1):91–116.
- Chaudhury, N. and Hammer, J. S. (2004). Ghost doctors: Absenteeism in rural Bangladeshi health facilities. *The World Bank Economic Review*, 18(3):423–441.
- Chyi, N. (2021). The workplace-surveillance technology boom. Accessed 9 August 2021.
- Cicala, S. (2015). When does regulation distort costs? lessons from fuel procurement in us electricity generation. *American Economic Review*, 105(1):411–444.
- Cilliers, J., Kasirye, I., Leaver, C., Serneels, P., and Zeitlin, A. (2018). Pay for locally monitored performance? A welfare analysis for teacher attendance in Ugandan primary schools. *Journal of Public Economics*, 167:69–90.
- Cludius, J., Forrest, S., and MacGill, I. (2014). Distributional effects of the Australian Renewable Energy Target (RET) through wholesale and retail electricity price impacts. *Energy Policy*, 71:40–51.
- Çubukçu, S., Sahin, N. M., Tekin, E., and Topalli, V. (2021). Body-worn cameras and adjudication of citizen complaints of police misconduct. Working paper no. w29019, National Bureau of Economic Research.
- Das, J., Hammer, J., and Leonard, K. (2008). The quality of medical advice in low-income countries. *Journal of Economic perspectives*, 22(2):93–114.
- Davis, R. J., Holladay, J. S., and Sims, C. (2022). Coal-fired power plant retirements in the United States. *Environmental and Energy Policy and the Economy*, 3(1):4–36.

- De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96.
- Diamond, R. (2016). The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *American Economic Review*, 106(3):479–524.
- Dickinson, D. and Villeval, M. C. (2008). Does monitoring decrease work effort? *Games and Economic Behavior*, 63(1):pp–56.
- Donovan, K. (2017). *Anytime, Anywhere: Mobile Devices and Services and their Impact on Agriculture and Rural Development*. The World Bank.
- Doudchenko, N. and Imbens, G. W. (2016). Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. *NBER Working Paper*, (w22791).
- Duflo, E., Banerjee, A., and Keniston, D. (2019). The efficient deployment of police resources: Theory and new evidence from a randomized drunk driving crackdown in India. Working paper no. w26224, National Bureau of Economic Research.
- Duflo, E., Hanna, R., and Ryan, S. P. (2012). Incentives work: Getting teachers to come to school. *American Economic Review*, 102(4):1241–78.
- Elliot, B., Rothenberg, T. J., and Stock, J. H. (1996). Efficient tests of the unit root hypothesis. *Econometrica*, 64(8):13–36.
- Falk, A. and Kosfeld, M. (2004). Distrust-the hidden cost of control. Working paper, IZA, Working Paper Series.
- Fama, E. F. and Jensen, M. C. (1983). Separation of ownership and control. *The Journal of Law and Economics*, 26(2):301–325.
- Feler, L. and Senses, M. Z. (2017). Trade shocks and the provision of local public goods. *American Economic Journal: Economic Policy*, 9(4):101–143.
- Ferman, B. and Pinto, C. (2021). Synthetic controls with imperfect pretreatment fit. *Quantitative Economics*, 12(4):1197–1221.
- Frey, B. S. (1993). Does monitoring increase work effort? The rivalry with trust and loyalty. *Economic Inquiry*, 31(4):663–670.
- Furnaro, A., Herpich, P., Brauers, H., Oei, P.-Y., Kemfert, C., and Look, W. (2021). German just transition: A review of public policies to assist german coal communities in transition. *Resources for the Future Report*, pages 21–13.
- Gallagher, J. and Fisher, P. (2020). Criminal deterrence when there are offsetting risks: Traffic cameras, vehicular accidents, and public safety. *American Economic Journal: Economic Policy*, 12(3):202–237.

- Giang, L., Nguyen, C., and Tran, A. (2017). Does PAPI monitoring improve local governance? Evidence from a natural experiment in Vietnam. *International Journal of Development Issues*.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.
- Green, J. F. (2021). Does carbon pricing reduce emissions? A review of ex-post analyses. *Environmental Research Letters*, 16(4):043004.
- Grossman, G., Platas, M. R., and Rodden, J. (2018). Crowdsourcing accountability: ICT for service delivery. *World Development*, 112:74–87.
- Hubbard, T. N. (2000). The demand for monitoring technologies: The case of trucking. *The Quarterly Journal of Economics*, 115(2):533–560.
- Janssens-Maenhout, G., Crippa, M., Guizzardi, D., Muntean, M., Schaaf, E., Dentener, F., Bergamaschi, P., Pagliari, V., Olivier, J. G., Peters, J. A., et al. (2019). Edgar v4. 3.2 Global Atlas of the three major greenhouse gas emissions for the period 1970–2012. *Earth System Science Data*, 11(3):959–1002.
- Jauregui, B. (2016). *Provisional Authority: Police, Order and Security in India*. University of Chicago Press.
- Johnsen, R., LaRiviere, J., and Wolff, H. (2019). Fracking, coal, and air quality. *Journal of the Association of Environmental and Resource Economists*, 6(5):1001–1037.
- Jotzo, F., Jordan, T., and Fabian, N. (2012). Policy uncertainty about Australia’s carbon price: Expert survey results and implications for investment. *Australian Economic Review*, 45(4):395–409.
- Kline, P. and Moretti, E. (2014). People, places, and public policy: Some simple welfare economics of local economic development programs. *Annu. Rev. Econ.*, 6(1):629–662.
- Linn, J. and McCormack, K. (2019). The roles of energy markets and environmental regulation in reducing coal-fired plant profits and electricity sector emissions. *The RAND journal of economics*, 50(4):733–767.
- Lohr, S. (2014). Unblinking eyes track employees. *The New York Times*.
- Look, W., James, J., Fedorko, E., Pesek, S., Mazzone, D., Barone, A., and Shelton, R. (2022). POWER for transition: Investment in coal communities through the partnerships for opportunity and workforce and economic revitalization (POWER) initiative. Technical report, Resources for the Future.
- Makofske, M. (2020). Pretextual traffic stops and racial disparities in their use. Technical report.

- Manson, S. M. (2020). IPUMS national historical geographic information system: version 17.0.
- Matheis, M. (2016). Local economic impacts of coal mining in the United States 1870 to 1970. *The Journal of Economic History*, 76(4):1152–1181.
- Metcalf, G. E. and Wang, Q. (2019). Abandoned by coal, swallowed by opioids? *NBER Working Paper*, (w26551).
- Morris, A. (2016). Build a better future for coal workers and their communities. *Climate and Energy Economics Discussion Paper*, Brookings Institution, Washington DC.
- Morris, A. C., Kaufman, N., and Doshi, S. (2021). Revenue at risk in coal-reliant counties. *Environmental and energy policy and the economy*, 2(1):83–116.
- Muralidharan, K., Niehaus, P., Sukhtankar, S., and Weaver, J. (2021). Improving last-mile service delivery using phone-based monitoring. *American Economic Journal: Applied Economics*, 13(2):52–82.
- Nagavarapu, S. and Sekhri, S. (2016). Informal monitoring and enforcement mechanisms in public service delivery: Evidence from the public distribution system in india. *Journal of Development Economics*, 121:63–78.
- Nordhaus, W. D. (1977). Economic growth and climate: The carbon dioxide problem. *American Economic Review*, 67(1):341–346.
- Olken, B. A. (2007). Monitoring corruption: Evidence from a field experiment in indonesia. *Journal of Political Economy*, 115(2):200–249.
- Peel, J. (2014). The Australian carbon pricing mechanism: Promise and pitfalls on the pathway to a clean energy future. *Minn. JL Sci. & Tech.*, 15:429.
- Pierce, L., Snow, D. C., and McAfee, A. (2015). Cleaning house: The impact of information technology monitoring on employee theft and productivity. *Management Science*, 61(10):2299–2319.
- Poterba, J. (1991). Tax policy to combat global warming: On designing a carbon tax.
- Prashar, S. (2012). Cop caught on facebook for graft. Last accessed 6 June 2019.
- Rafaty, R., Dolphin, G., and Pretis, F. (2020). Carbon pricing and the elasticity of CO2 emissions.
- Roback, J. (1982). Wages, rents, and the quality of life. *Journal of political Economy*, 90(6):1257–1278.
- Roemer, K., Raimi, D., Glaser, R., et al. (2021). Coal communities in transition: A case study of Colstrip, Montana. *Washington, DC: Resources for the Future Report*.

- Schulze, G. G. and Frank, B. (2003). Deterrence versus intrinsic motivation: Experimental evidence on the determinants of corruptibility. *Economics of Governance*, 4(2):143–160.
- Scott, A., Balthrop, A., and Miller, J. W. (2021). Unintended responses to it-enabled monitoring: The case of the electronic logging device mandate. *Journal of Operations Management*, 67(2):152–181.
- Stoker, T. M., Berndt, E. R., Ellerman, A. D., and Schennach, S. M. (2005). Panel data analysis of US coal productivity. *Journal of Econometrics*, 127(2):131–164.
- Tellez, J. F., Wibbels, E., and Krishna, A. (2020). Local order, policing, and bribes: Evidence from India. *World Politics*, 72(3):377–410.
- Timilsina, G. R. (2022). Carbon taxes. *Journal of Economic Literature*, 60(4):1456–1502.
- Tolbert, C. M. and Sizer, M. (1996). US commuting zones and labor market areas: A 1990 update. *Rural Economy Division, Economic Research Service, US Department of Agriculture*, (ERS Staff Paper Number 9614).
- Watson, B., Lange, I., and Linn, J. (2023). Coal demand, market forces, and US coal mine closures. *Economic Inquiry*, 61(1):35–57.
- Weber, J. G. (2020). How should we think about environmental policy and jobs? An analogy with trade policy and an illustration from US coal mining. *Review of Environmental Economics and Policy*.
- West, D. M. (2021). How employers use technology to surveil employees. Last accessed 9 August 2021.
- World Bank (2023). State and trends of carbon pricing 2023. *World Bank, Washington, DC*.
- World Health Organization (2018). Global status report on road safety. Technical report, Geneva.
- Zamboni, Y. and Litschig, S. (2018). Audit risk and rent extraction: Evidence from a randomized evaluation in Brazil. *Journal of Development Economics*, 134:133–149.