2016-04-28

Essays in Environmental Economics

Luke Fitzpatrick
University of Miami, l.fitzpatrick@umiami.edu

UNIVERSITY OF MIAMI

ESSAYS IN ENVIRONMENTAL ECONOMICS

By

Luke Fitzpatrick

A DISSERTATION

Submitted to the Faculty
of the University of Miami
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Coral Gables, Florida

May 2016

UNIVERSITY OF MIAMI

A dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

ESSAYS IN ENVIRONMENTAL ECONOMICS

Luke Fitzpatrick

Approved:

David Kelly, Ph.D.
Professor of Economics

Christopher Parmeter, Ph.D.
Associate Professor of Economics

Laura Giuliano, Ph.D.
Associate Professor of Economics

Guillermo Prado, Ph.D.
Dean of the Graduate School

Christian Traeger, Ph.D.
Assistant Professor of Economics
University of California - Berkeley

FITZPATRICK, LUKE (Ph.D., Economics)
Essays in Environmental Economics (May 2016)

Abstract of a dissertation at the University of Miami.

Dissertation supervised by Professor David Kelly.
No. of pages in text. (134)

Efficiently addressing climate change and waning our dependence on fossil fuels will prove central to ensuring a viable habitat for humanity in the coming centuries. Chapter 1 analyzes stabilization targets - widely employed and well-intentioned environmental policies that ultimately reduce global welfare. Chapter 2 quantifies a health-care externality stemming from surface mining of coal. Chapter 3 uses a new, novel approach to estimating how the coal mining industry affects markets for residential real estate in three Appalachian counties.

To my family

Acknowledgments

I thank Dave Kelly, Chris Parmeter, Laura Giuliano, and Christian Traeger for their invaluable guidance on this project.

LUKE FITZPATRICK

University of Miami

May 2016

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Addressing climate change by waning our dependence on fossil fuel is central to ensuring the viability of our planet for mankind. A common feature of international climate agreements are stabilization targets: the maximum allowable change in an environmental variable tied to greenhouse gas (GHG) emissions. The most well-known stabilization target, and a central component of the 2015 Paris Climate Accord, calls for limiting the increase in global mean surface temperature to 2°C above its pre-industrial level. Staying below the 2°C target is widely advocated to prevent irreversible and catastrophic climate damages. With a target in hand, researchers assume certainty equivalence and compute the least cost emissions path that ensures remaining below the stipulated threshold.

However, uncertain parameters and random weather shocks affect climate dynamics. A sequence of large shocks, or an unexpectedly sensitive relationship between temperature and GHG concentrations (the climate sensitivity), can cause the temperature to exceed the target, irrespective of the emissions policy.[1] *Pure* stabilization targets are thus not feasible in a dynamic setting with uncertainty. An alternative is *probabilistic*

[1] We show that even if GHG emissions are immediately and permanently reduced to zero, enough inertia exists in the climate that the global mean temperature will exceed 2°C with 15% probability.

stabilization targets (see for example, Held et al. (2009)) which instead require that the environmental variable stay beneath the target with a maximum allowable probability.[2] A probabilistic target is a feasible policy if the maximum probability of exceeding the target is sufficiently high. We analyze stabilization targets in an integrated assessment model of the climate and economy with uncertainty, Bayesian learning, and random weather shocks.[3] We show that a probabilistic target in an environment with uncertainty induces two welfare losses previously undocumented in the literature.

First, the arrival of new information changes beliefs about the climate sensitivity over time, changing the optimal temperature. For example, an unexpectedly high climate sensitivity raises the optimal temperature, since achieving a given temperature requires more abatement expenditures, while the damages from that temperature are unchanged. By definition, however, the target temperature remains unchanged. Therefore, a welfare cost ensues because the target is *inflexible*. Second, we show that stabilization targets force overly stringent policy responses to transient weather shocks. Weather shocks cause the temperature to exceed the target with positive probability. If so, abatement expenditures must rise to bring the temperature back to the target. The abatement cost is excessive, since the temperature naturally returns to the target as the shock dies out regardless. Therefore, a second welfare cost occurs because stabilization targets cause an *overreaction* to transient shocks.

Both welfare losses are new to this paper, and are present only when the temperature is a function of uncertain and/or random variables. In models which assume certainty, a welfare cost only occurs if the target is set below the temperature resulting from the optimal emissions policy. For example, Nordhaus (2007) calculates the welfare loss of

[2]A pure stabilization target is a special case of a probabilistic target where the probability of exceeding the target is zero. Therefore, without loss of generality, we consider only probabilistic targets.

[3]See Kelly & Kolstad (1999*b*) for a survey of integrated assessment models.

restricting the temperature to 2°C under certainty, relative to an optimal temperature change of 3.5°C. This result is sensitive to assumptions about the damage function, the discount rate, the lack of tipping points, and other assumptions in the DICE model. Indeed, we show that when uncertainty about the climate sensitivity is added to the DICE model, the optimal temperature without a target falls to 3.26°C, which reduces the welfare loss associated with the 2°C target.[4] In contrast, welfare losses associated with inflexibility and overreaction depend only on the interaction of uncertainty and a pure target and are therefore robust to alternative modeling assumptions.

We show that introducing a stabilization target into the baseline model with uncertainty causes a welfare loss of 4.7%, which is 66% higher than the welfare cost of introducing a stabilization target into the model with certainty.[5] We find that most of the increase in welfare loss comes from the inflexibility of the target, but overreaction accounts for 8.2% of the increase in welfare loss. Further, the welfare loss is relatively insensitive to the policy parameter which specifies the maximum probability of exceeding the target. The welfare loss increases, however, with the unknown true value of the climate sensitivity. If the climate sensitivity is higher than expected, the abatement cost of keeping the temperature at the target increases, which increases the optimal temperature. We show that the welfare loss can increase to 14% or more, depending on the true value of the climate sensitivity.

We also study the dynamics of learning in the presence of stabilization targets. Since the target reduces emissions, the planner learns the climate sensitivity more slowly. We find the standard deviation of the prior distribution of the climate sensitivity is approx-

[4] With uncertainty, the risk averse planner becomes more conservative, reducing emissions which are more damaging if the climate sensitivity turns out to be higher than expected.

[5] We follow Neubersch, Held & Otto (2014) and calibrate a maximum probability of exceeding the target of 0.33, based on an interpretation of IPCC statements calling for policies for which achieving the 2°C target is likely.

imately 10% higher in 2050 than without the stabilization target, if the true climate sensitivity is equal to the mean of the prior. If the true climate sensitivity is unexpectedly high, then learning slows further. In turn, slower learning makes the temperature target more difficult to achieve, as the planner needs more time to increase abatement if the climate sensitivity is high.

Asthma is a long-term illness characterized by debilitating short-term symptoms. Annually, asthma places a substantial burden on the U.S. economy. Barnett & Nurmagambetov (2010) estimate that lost work and school days amount to annual economic of $3.8 billion from asthma morbidity, and $2.1 billion from asthma mortality. Although the causes of asthma are unknown, air pollution is widely accepted as a trigger of asthma symptoms (Centers for Disease Control and Prevention 2013). In West Virginia, where the incidence of asthma is over 2% higher than the national average, environmental advocates finger coal mining and the environmental pollution it generates as an underlying cause of the health disparities seen there.[6] If true, then the industry's decline could foretell an improvement in health in the West Virginia, where coal mining for many years contributed to the economy and cultural fabric. Although there is significant research on Appalachian health disparities, none has conclusively linked coal mining to worsened public health outcomes. This paper aims to bridge that gap.

I find that county populations respond to nearby surface mining of coal with increased demand for asthma-related healthcare. A standard deviation of surface-mined coal within 36 kilometers of a county's population centroid leads to 21% more asthma hospitalizations in a given quarter. According to my estimates, around 11% of all hospitalizations for asthma in West Virginia between 2005 and 2010 were caused by surface mining of coal. Combined with cost estimates from the Center for Medicare and Med-

[6] http://www.cdc.gov/asthma/most_recent_data_states.htm

icaid Services 2012 Provider Utilization and Payment Data Inpatient Public Use File, this suggests that surface mining led to an additional 9.13 million dollars of healthcare expenditures over this six-year period. In general, women are more sensitive to surface mining than men are. Youths and the elderly are more sensitive than those between the age of 18 and 64, who demonstrate no sensitivity to surface mining. The single most sensitive subgroup are females below the age of 18; a standard deviation of surface tonnage within 34 kilometers of a county's population centroid increase asthma hospitalization among that group by 42%. Among the entire population, the same quantity of coal increases asthma hospitalizations by 20%. Among the elderly, female demonstrate a slightly higher sensitivity than men.

This paper builds on an earlier literature that clearly identifies health disparities between mining and non-mining Appalachian counties. These cross-sectional studies find that populations in mining counties are more likely to suffer from hypertension, kidney disease, and respiratory disorders than those in non-mining counties (Brink et al. 2014, Hendryx & Ahern 2008, Hendryx & Ahern 2009, Hendryx, Ahern & Nurkiewicz 2007). A sign of one problem with a cross-sectional approach in this context is that populations in coal counties are also hospitalized more frequently for some ailments unrelated to coal mining. Therefore, to test internal validity, I estimate the same econometric model using hospitalizations for bone fractures or hernias as dependent variables. If the frequency with which populations seek care for those ailments similarly covaries with surface mining, then my findings are likely spurious. These regressions do not yield similar results. Neither surface nor underground mining consistent affect either alternative ailment. This is evidence that my method reduces at least some of the bias present in earlier epidemiological studies on the topic.

There are two primary contributions of the present work with respect to identifying a health externality from coal mining.[7] First, I use a panel data framework. I combine hospitalization records and mine-level production data to explore how quarterly variation in mining activity influences quarterly demand for asthma-related healthcare. My findings match Brink et al. (2014), who find surface mining regions have worse respiratory health, while eliminating the cross-sectional bias of Hendryx & Ahern (2008), and Hendryx, Ahern & Nurkiewicz (2007) by addressing unobservable sources of heterogeneity. County fixed effects address the unobserved heterogeneity between counties that might contribute to consistently higher or lower rates of hospital treatments. Asthmatics may cluster in coal counties because those counties have more hospitals or better ambulatory services. Populations in non-mining counties might be better educated on how to manage their asthma symptoms, meaning they are less likely to seek hospital care in the event of an attack. To the extent that these relative differences are unchanged over time, county fixed effects will absorb these differences. Further, I use quarterly dummy variables to absorb the exogenous seasonality of asthma symptoms.

The second important contribution is a more flexible definition of exposure. Previous work on this topic uses data that lacks the geographic precision in my data. As a result, every previous measure of exposure to mining is based on the tons of coal mined within a county's borders during some time-period, thereby ignoring the possibility that mining beyond a county's borders shaped that population's health. Figure 2.14 illustrates the problem with the county-border definition of exposure: of the three mines shown, only the mine inside of Logan County's borders affects the health of Logan County residents. Considering the proximity of the two mines nearest to each other, it is more likely that both affect the health of Logan County residents, if either does.

[7]Coal mining's affect on miners' health is well documented: time spent working underground leads to miner's pneumonconiosis, otherwise known as black lung. To the extent that the risk of developing black lung are priced into underground miners' wages, that is not an externality.

The mining data I use contains the geographic coordinates of every mine in the United States, and quarterly production statistics since 1983. This allows me to model exposure as the amount of coal mined within a circular exposure buffer of any size, centered upon a county's population centroid. This maintains the spirit of earlier work, which implicitly assumed an exposure buffer shaped like a county, while making one less assumption about the size of the buffer - namely, that an exposure buffer is limited to the size of a county. I test exposure buffers with radii up to 100 kilometers, and find that the most accurate (ranked by mean squared prediction error) buffer is between 34 and 36 kilometers. Tables 2.10 and 2.11 provides evidence of the attenuation bias caused by mis-measurement.

Section 2.6 discusses sources of air pollution at surface mines, and earlier attempts at identifying a health externality. Section 3.3 details the data. I present the econometric model and address threats to identification in section 2.8. Section 2.9 presents response estimates for the general population and subpopulations, and section 2.10 concludes.

The final chapter attempts to better understand the housing market impacts of the coal industry. The environmental and ecological damages of the coal mining are well documented, and occur via emissions of air pollution (Ghose & Majee 2007, Aneja, Isherwood & Morgan 2012, Knuckels et al. 2012) and water pollution (McAuley & Kozar 2006, Pond et al. 2008, Lindberg et al. 2011). These papers conclude that coal mining pollutes both the air above and the groundwater below mining areas with elevated levels of dangerous toxins. Further, public health researchers document clear health disparities in between mining and non-mining (Hendryx, Ahern & Nurkiewicz 2007, Hendryx & Ahern 2008).

Broadly, this paper considers the claim in Randall et al. (1978) that residents living near coal mining operations and in coal mining regions are most likely to bear these externalized costs. I consider the relationship between coal mining operations and the real estate markets with which they are interwoven. I perform a hedonic analysis using seven years of data on sales of single-family homes from three northern Appalachian counties and mine level quarterly production data to understand how coal operations affect house prices. Our results show that in two out of three markets analyzed, homes near standalone surface mines sell at steep and significant discounts; between 12% in Fayette, Pennsylvania; and 15% in Monongalia, West Virginia. These estimates come from regressions that control for observed structural characteristics, as well as unobserved and spatially correlated characteristics.

There is surprisingly little hedonic research on coal mines. (Williams, A.M 2011) performs a cross sectional analysis of median house values using county-level Census data, and the number of surface mines within each county. They conclude surface mines within a county's borders reduce median home values there by less than 2%. (Williamson, Thurston & Heberling 2008) consider sales in two West Virginia counties that took place near a stream impacted my acid mine drainage. They find that homes located within 0.25 miles of such a stream sell for $4,783 less.

This paper is more closely related to a literature on the housing market impacts of different environmental disamenities. Researchers have used hedonics to study Superfund sites and their cleanups (Gamper-Rabindran & Timmins 2011, Kiel & Williams 2007, Greenstone & Gallagher 2008), power plants (Davis 2011), point sources of water pollution (Leggett & Bockstael 2000), wind farms (Heintzelman & Tuttle 2012), and, more recently, shale gas development (Gopolakrishnan & Klaiber 2013, Muehlenbachs, Spiller & Timmins 2015).

These studies model a property's exposure to a disamenity with the distance between it and the nearest disamenity of interest. Models typically include the linear distance or logged distance (for which a positive coefficient indicates a net disamenity); inverse linear distance or logarithm of inverse distance (for which a negative coefficient indicates a net disamenity); or an indicator variable equal to one if the property falls within a pre-specified distance buffer of a disamenity. The overwhelming majority of work on sources of environmental externalities find highly localized impacts. For example, a review of panel data studies of Superfund sites notes that the maximum distance for which price impacts are felt are generally 2-2.5 miles (Jenkins, Kopits & Simpson 2006).

The present paper builds on this LULU literature in two ways. First, we broaden its scope by focusing on the coal industry, a previously unstudied source of environmental externalities. We combine seven years of single family home sales from three northern Appalachian counties with mine-level quarterly production data from the Department of Labor's Mine Safety and Health Administration (MSHA.) These counties provide a natural setting for our research due to the relatively large number of sales that took place there, and their exposure to the mining industry (over 90% of sales in each county in our dataset take place within 10 kilometers of coal related operations.)

Second, we take a new approach to defining which homes are exposed to the industry. Using every sale from a given county, we construct the distribution of exposure (as measured by the distance between a parcel and the nearest coal site) and test 99 different possible treatment buffers. We perform a leave-one-observation-out prediction procedure for every possible threshold, and select as our buffer whichever minimizes the sum of squared prediction errors from this exercise. In this way, we can abstain from imposing any *ad hoc* assumptions on how markets view coal operation. Two out of the three counties we consider have similar exposure buffers, and the impacts of coal mines (discounts of 10 and 15 percent) are driven by surface mines. In the third county, we find even steeper discounts, but these are for homes sold near preparation plants.

That county's exposure buffer is also considerably smaller than the other two, suggesting that buyers demonstrate different sensitivities to the environmental and health risks associated with the mining industry, just as they demonstrate different preferences for structural attributes. We also test several measures of exposure commonly employed in earlier work: linear distance, logged distance, inverse logged distance, and two *ad hoc* discrete measures of exposure.

Chapter 2

Probabilistic Stabilization Targets

2.1 Stabilization Targets

Stabilization targets are ubiquitous in climate change policy. Policy makers recommending a 2°C stabilization target include the European Commission (2007), the Copenhagen Accord (2009), and the German Advisory Council on Global Change (Schubert et al. 2006). Many atmospheric scientists (Hansen 2005, O'Neill & Oppenheimer 2002) also advocate for warming targets. Other stabilization differ according to the choice of target. For example, the German Advisory council recommends limiting sea level rise to at most 1 meter and ocean acidification to at most 0.2 units of pH below its preindustrial level. Some policy groups such as 350.org favor a GHG concentration target.[8] Replacing a temperature target with an alternative target variable would not change our qualitative conclusions. However, the overreaction welfare loss would change quantitatively, as alternative target variables may be more or less variable than temperature.

A large literature computes the least-cost emissions path which stabilizes the climate at 2°C under certainty (Nordhaus 2007, Richels, Manne & Wigley 2004, Lemoine & Rudik 2014). However, it is well known that parameters of the climate system are

[8] den Elzen, Meinshausen & van Vuuren (2007) and Lemoine & Rudik (2014) study GHG targets.

uncertain. For example, the climate sensitivity, which measures the elasticity of the global mean temperature with respect to GHG concentrations, is notoriously uncertain (Intergovernmental Panel on Climate Change 2007, Kelly & Kolstad 1999*a*). Therefore, following the least cost pathway calculated under certainty can, under uncertainty, result in warming that exceeds the target by a considerable margin.[9] Indeed, a branch of the literature focuses on the likelihood of meeting stabilization targets for various emissions scenarios proposed by policy makers, or what emissions paths satisfy the target for various values of the climate sensitivity. For example, Hare & Meinshausen (2006) and Keppo, O'Neill & Riahi (2007) compute temperature changes for various emissions scenarios; Harvey (2007) proposes allowable CO_2 emissions paths for different ranges of the climate sensitivity. This research provides an important first step in estimating the probability of exceeding the target. Here we take two important next steps by introducing learning and fat-tailed uncertainty. Learning allows the emissions policy to adjust as new information arrives, while fat tailed uncertainty implies a relatively large value of the climate sensitivity is expected to occur more often than if the uncertainty was normally distributed. Learning and fat-tailed uncertainty have competing effects on the ability of the planner to stabilize the climate.

Learning allows the planner to more easily meet a target by quickly adjusting emissions if new information indicates the climate sensitivity is greater than previously thought. Further, learning allows the planner to move closer to the target, while still remaining below the target with the same probability. However, Kelly & Kolstad (1999*a*), Leach (2007), and Roe & Baker (2007) show that learning about the climate sensitivity is a slow process, because noisy weather fluctuations obscure the climate change signal.

[9]Paradoxically, stabilization targets evolved as method of dealing with uncertainty in integrated assessment models. The idea was to propose limits on temperature changes which, if exceeded, would cause damages high enough so that uncertainties in the cost of abatement and other parameters are less relevant.

Further, small uncertainties about the climate sensitivity ultimately have large effects on the temperature through feedback effects (Kelly & Tan 2015). Therefore, controlling the temperature requires a precise estimate of the climate sensitivity, which takes time. We find the optimal near term policy with Bayesian learning is similar to the case without learning. We find that learning eventually allows the temperature to move closer to the target, but the effect is marginal since learning is slow.

It is well known that the prior distribution of the climate sensitivity has a fat tail, since small, uncertain feedback effects eventually have large effects on the temperature (Kelly & Tan 2015, Weitzman 2009). We show that stabilizing the climate is more difficult when the climate sensitivity is high, as the climate has more inertia and learning slows. The welfare loss from inflexibility of the target also increases with the climate sensitivity.

We analyze targets when the climate sensitivity is unknown. Rudik (2014) considers targets with damage uncertainty and learning. He shows that if the learning model is misspecified, targets become attractive since they prevent the planner from increasing the temperature when believing incorrectly that damages are not very convex. Here we show stabilization targets more problematic when the climate sensitivity is uncertain. The relationship between the policy variable, emissions, and the target, temperature, is unknown and nonlinear, which makes the target difficult to achieve.

Other authors compute optimal emissions paths under certainty which keep temperatures below a threshold, beyond which specific irreversible and disastrous consequences occur. Keller et al. (2005) propose emissions paths which prevent coral bleaching or the disintegration of ice sheets. Kvale et al. (2012) propose emission paths which limit ocean level rise and acidification. Additionally, Bruckner & Zickfield (2009) compute emission paths that reduce the likelihood of a collapse of the Atlantic thermohaline circulation. The aforementioned studies employ the tolerable windows approach, an inverse modeling method that asks: in order to limit GHG concentrations or warming below a threshold at all future dates, how should emissions be controlled in every period moving forward? These studies have some similarities to stabilization targets with certainty in that the optimal policy stabilizes the climate below the threshold. Lemoine & Traeger (2014) compute the optimal emissions path in a model with various climate tipping points in an environment with uncertainty and Bayesian learning, but do not study stabilization targets. While we do not specifically model irreversibilities, our model combines convex damages with uncertainty. Therefore, the planner insures against very high damages by pursuing a conservative emissions policy. Thus, the literature with irreversibilities and thresholds consider either certainty, or do not consider stabilization targets. We show that the stabilization targets are fundamentally different with uncertainty: control of the climate becomes difficult and additional welfare losses occur.

2.2 Model

We consider an infinite horizon version of the Nordhaus DICE model (Nordhaus 2007). In the DICE model, economic production causes GHG emissions, which raise the global mean temperature. Higher temperatures reduce total factor productivity (TFP). The social planner chooses capital investment and an emissions abatement rate to maximize welfare. Our model has four differences from the DICE model. First, we use an annual time step rather than the 10 year step in DICE. Second, we use the simplified model of the atmosphere/climate due to Traeger (2014), in which the ocean temperature changes exogenously and only two reservoirs exist for carbon (atmosphere and ocean/biosphere). Third, the model is stochastic, with an uncertain climate sensitivity and random weather shocks that obscure the effect of GHGs on temperature. The planner learns about the uncertain climate sensitivity over time by observing temperature changes. Fourth, we impose stabilization targets to ascertain their effects on welfare, temperature, and economic growth. Sections 2.2.1-2.2.2 describe the economic and climate models briefly (for a detailed discussion, refer to Traeger (2014).)

2.2.1 Economic System

The global economy produces gross output, Q, from capital K and labor L according to:

$$Q_t = A\left(t\right) K_t^{\gamma} L\left(t\right)^{1-\gamma}. \tag{2.1}$$

Here variables denoted as a function of t, such as $L\left(t\right)$ and TFP, $A\left(t\right)$, grow exogenously. Appendix A.1 gives the growth rates for all variables which change exogenously over time. Variables with a t subscript are endogenous.

An emissions abatement technology exists which can reduce emissions by a fraction x_t at a cost of $\Lambda\left(x_t\right) = \Psi\left(t\right) x_t^{a_2}$ fraction of gross output. Here Λ is the cost function and $\Psi\left(t\right)$ is the exogenously declining cost of a backstop technology which reduces emissions to zero. Further, increases in global mean temperatures above preindustrial levels, T_t, reduce TFP by a factor $1/\left(1 + D\left(T_t\right)\right)$, where $D\left(T_t\right) = b_1 T^{b_2}$ is the damage function. Therefore, output net of abatement spending and climate damages, Y_t, is:

$$Y_t = \frac{1 - \Psi\left(t\right) x_t^{a_2}}{1 + b_1 T_t^{b_2}} A\left(t\right) K_t^{\gamma} L\left(t\right)^{1-\gamma}. \tag{2.2}$$

Let C_t be consumption and let capital depreciate at rate δ_k. Then the resource constraint is:

$$Y_t = C_t + K_{t+1} - \left(1 - \delta_k\right) K_t. \tag{2.3}$$

Period utility is constant relative risk aversion:

$$u = \frac{\left(C_t/L_t\right)^{1-\phi} - 1}{1 - \phi}. \tag{2.4}$$

We follow Costello et al. (2010) and assume $T_t \leq T^{max} << \infty$ to prevent expected utility from being unbounded from below.

The discount factor for future utility is $\exp(-\delta_u)$, where δ_u is the pure rate of time preference.

2.2.2 Climate System

Current period GHG emissions, E_t, from production depend on the planner's choice of emissions abatement rate x_t, the emissions intensity of output $\sigma(t)$, exogenous emissions from land use changes, $B(t)$, and gross global output:

$$E_t = (1 - x_t)\,\sigma(t)\,Q_t + B(t)\,. \tag{2.5}$$

The stock of GHG equivalents, M_t, depends on current period emissions and the natural decay rate of GHGs into the biosphere and ocean. Let $\delta_m(t)$ denote the decay rate (which changes exogenously) and M_B denote the stock of GHGs during pre-industrial times. Then M_t accumulates according to:

$$M_{t+1} - M_B = (1 - \delta_m(t))\,(M_t - M_B) + E_t. \tag{2.6}$$

We normalize GHG stocks relative to pre-industrial. Let $m_t \equiv \frac{M_t}{M_B}$, then:

$$m_{t+1} - 1 = (1 - \delta_m(t))\,(m_t - 1) + \frac{E_t}{M_B}. \tag{2.7}$$

Radiative forcing of GHGs, F_t, increases the temperature:

$$F_{t+1} = \Omega \log_2(m_{t+1}) + EF(t)\,. \tag{2.8}$$

Here $EF(t)$ is forcing from other sources, which grows exogenously.

The global mean temperature evolves according to:

$$\hat{T}_{t+1} = \hat{T}_t + \frac{1}{\alpha}\left(F_{t+1} - \frac{\hat{T}_t - \Gamma}{\tilde{\lambda}} + \xi\left(\hat{O} - \hat{T}\right)(t)\right) + \tilde{\nu}_{t+1} \tag{2.9}$$

Here $\hat{T}$ and $\hat{O}$ denote the absolute global atmospheric and oceanic temperatures in °C, respectively; α is the thermal capacity of the upper oceans; Γ is the pre-industrial atmospheric temperature; ξ is a coefficient of heat transfer from the upper oceans to the atmosphere; $\tilde{\nu}_t \sim N(0, 1/\rho)$ is the random weather shock; $\tilde{\lambda}$ is the uncertain climate sensitivity. We assume the ocean-atmosphere temperature differential changes exogenously. The climate sensitivity $\tilde{\lambda}$ describes how sensitive the atmospheric temperature is to GHG concentrations, and is the subject of great uncertainty.

Let $\Delta T_{2\times}$ be the steady state atmospheric temperature deviation from pre-industrial time resulting from a doubling of GHG concentrations, also relative to pre-industrial levels. Then:

$$\Delta T_{2\times} = \Omega\tilde{\lambda}. \tag{2.10}$$

Since the climate sensitivity is uncertain, $\Delta T_{2\times}$ is also uncertain. Stocker, Dahe & Plattner (2013) estimate that $\Delta T_{2\times}$ is most likely to lie somewhere between 1.5°C and 4.5°C. The initial mean of the prior distribution is 3.08, taken from the mean of estimates in the atmospheric science literature (Roe & Baker 2007).[10]

Let $T_t = \hat{T}_t - \Gamma$ and $O_t = \hat{O}_t - \Gamma$ be the current deviations from pre-industrial temperatures, $\tilde{\beta}_1 = 1 - 1/\tilde{\lambda}\alpha$ denote the climate feedback parameter, $\beta_2 = \frac{1}{\alpha}$, and $\beta_3 = \xi/\alpha$. The climate system simplifies to:

$$T_{t+1} = \tilde{\beta}_1 T_t + \beta_2 F_{t+1} + \beta_3 \left(O - T\right)(t) + \tilde{\nu}_{t+1}. \tag{2.11}$$

Since $\tilde{\lambda}$ is uncertain, the climate feedback parameter is also uncertain. The climate feedback parameter is increasing in $\tilde{\lambda}$. For example, if GHG induced warming reduces ice cover, which reduces the amount of sunlight reflected back into space (the albedo

[10]More precisely, the initial mean of the climate feedback parameter, defined below, is identical to Roe & Baker (2007). Our implied $\Delta T_{2\times}$ is slightly lower because of small differences in the other parameters of the model.

effect), causing still higher temperatures, we have a positive feedback (higher $\tilde{\beta}_1$ and $\tilde{\lambda}$).[11] Uncertain climate feedbacks cause uncertainty in the climate sensitivity (Stocker, Dahe & Plattner 2013).

2.2.3 Learning

Following the literature (Roe & Baker 2007, Kelly & Tan 2015), we assume the planner has prior beliefs that the climate feedback parameter is drawn from a normal distribution with mean μ_t and precision η_t. However, values of the feedback parameter greater than or equal to one imply infinite steady state temperature changes for any positive concentration of GHGs. Therefore, the computational solution requires truncating the prior distribution at a value less than one.[12]

When prior beliefs about the climate sensitivity are normal, the prior distributions of the climate sensitivity and $\Delta T_{2\times}$ have fat upper tails (Kelly & Tan 2015). Thus, the probability of relatively large temperature changes from a doubling of GHGs is large relative to the normal distribution.

The weather shock ν occurs at the beginning of each period, before the abatement rate is chosen. We combine the two uncertain terms in equation (2.11) and denote the sum $\tilde{H}$:

$$\tilde{H}_{t+1} = \tilde{\beta}_1 T_t + \tilde{\nu}_{t+1}. \tag{2.12}$$

Since $\tilde{H}_{t+1}$ is the sum of two normally distributed random variables, it is also normally distributed with mean $\mu_t T_t$ and variance:

$$\sigma^2_{\tilde{H}} = T_t^2/\eta_t + 1/\rho. \tag{2.13}$$

[11] Other feedbacks include changes in cloud cover and water vapor.

[12] See Costello et al. (2010) for a justification of truncation. We assume numerically that beliefs evolve ignoring the truncation. See appendices for more details of the truncation.

The planner observes $H_{t+1} = T_{t+1} - \beta_2 F_{t+1} - \beta_3 (O - T)(t)$ at the beginning of $t+1$ and updates beliefs of $\tilde{\beta}_1$. Bayes' Rule implies that the posterior distribution of $\tilde{\beta}_1$ is also normal, with mean and precision:

$$\mu_{t+1} = \frac{\eta_t \mu_t + \rho H_{t+1} T_t}{\eta_t + \rho T_t^2}, \tag{2.14}$$

$$\eta_{t+1} = \eta_t + \rho T_t^2. \tag{2.15}$$

Perfect information implies that $\mu = \tilde{\beta}_1$ and $\eta = \infty$. The information set used by the planner to select x_t includes μ_t.

2.2.4 Recursive Problem

The planner chooses emission abatement rate, x_t, and capital investment each period to maximize social welfare.

Let $k \equiv K / \left(L A^{\frac{1}{1-\gamma}} \right)$ denote normalized capital per productivity adjusted person (Kelly & Kolstad 2001, Traeger 2014) and the same for y and c, and $s = [k, T, m, t, \mu, \eta]$. The recursive version of the social planning problem is:

$$V(s) = \max_{k', x \in [0,1]} \left\{ u(c) + \beta(t) \int_{-\infty}^{\infty} V[s'] N\left(\mu T, \frac{T^2}{\eta_t} + \frac{1}{\rho} \right) d\tilde{H} \right\}, \tag{2.16}$$

subject to:

$$C = y - \exp(-g_A(t) - g_L(t)) k' + (1 - \delta_k) k. \tag{2.17}$$

$$T' = \beta_2 F' + \beta_3 (O - T)(t) + \tilde{H}', \tag{2.18}$$

$$F' = \Omega \log_2(m') + EF(t), \tag{2.19}$$

$$m' = 1 + (1 - \delta_m(t))(m - 1) + \frac{E}{M_B}, \tag{2.20}$$

$$E = (1 - x) \sigma(t) A(t)^{\frac{1}{1-\gamma}} L(t) k^{\gamma} + B(t). \tag{2.21}$$

$$\mu' = \frac{\eta\mu + \rho\tilde{H}'T}{\eta + \rho T^2}, \tag{2.22}$$

$$\eta' = \eta + \rho T^2. \tag{2.23}$$

$$t' = t + 1. \tag{2.24}$$

Equation (2.16) condenses the double expectation over $\tilde{\beta}_1$ and $\tilde{\nu}_{t+1}$ into one expectation over the random variable H_{t+1}. Appendix A.1 gives the equations which govern the evolution of variables that change exogenously over time, including the growth rates of productivity and population, g_A and g_L. Time, t, is a state variable which determines all exogenous variables. The discount factor accounts for growth in population and productivity. Because the growth rates change over time, the normalized discount factor $\beta(t)$ is not constant, but is exogenous.

In the model, two state variables, t and η, are unbounded. Therefore, the computational solution replaces the precision η with the variance $1/\eta$, and replaces time with a bounded, monotonic increasing function.[13] Table A.1 gives parameter values and definitions for the above problem.

2.2.5 Stabilization Targets

As discussed in Section 2.1, stabilization policies differ primarily by the choice of target variable. Here we analyze a temperature target, which is the most common. However, the qualitative channels we find apply regardless of the target variable: all targets are by definition inflexible and force policy variables to respond to transient shocks. However, the magnitude of the effects differ depending on, for example, the volatility of any shocks to the target variable.

Other differences in stabilization policies include the probability of exceeding the target and possibly the timing (when the constraint becomes active). The model treats

[13] See for example, Kelly & Kolstad (1999*a*) or Traeger (2014).

the probability of exceeding the constraint, ω, as a policy variable. The probabilistic stabilization target is a sequence of constraints:

$$Pr\left(T_{t+1} \geq T^*\right) \leq \omega,\ \forall t = 0, \ldots. \tag{2.25}$$

A pure stabilization target is a special case of equation (2.25) with $\omega = 0$. The constraint is always satisfied if $\omega = 1$, so $\omega = 1$ corresponds to the unconstrained case.

A constraint exists for each period, $t = 1, \ldots$, which restricts the probability that $T_{t+1} \geq T^*$. The constraints limit the temperature change in every period. Other alternatives are more difficult to reconcile with the intent of the policy. One alternative requires that the temperature limit holds only in the steady state. However, in the DICE the backstop cost $\Psi\left(t\right)$ eventually decreases to the point where zero emissions becomes optimal. Thus, under the optimal solution, the climate eventually returns back to prehistoric temperature levels. Therefore, a steady state temperature limit of 2°C is not binding and the unconstrained solution results. Further, one can show that a steady state constraint is non-binding more generally. Since the transversality condition only requires T_t remain bounded, a constraint limiting T_t to T^* in the infinite future can be ignored: discounting implies that the shadow cost of such a constraint is zero. Intuitively, the temperature converges to the unconstrained solution, with a plan to satisfy the constraint at $t = \infty$, which never arrives.

A second alternative is a peak temperature target, which requires only the maximum temperature change not exceed 2°C. However, uncertainty may cause the peak warming target to move. Suppose the temperature target is accidentally exceeded in period t, so that say $T_t = 2.5$°C. Now in periods after t, the temperature cannot exceed 2.5°C, otherwise a new peak temperature exists which is above 2°C, violating the constraint again. However, the temperature may remain at any level below 2.5°C, since such temperatures do not create a new peak temperature and the constraint only applies to the peak. Hence, under uncertainty, accidental violations of the constraint effectively weaken the constraint over time.[14]

The stabilization target is effectively a sequence of constraints on the abatement rate x, since x_t affects T_{t+1}. Appendix A.1.1 shows that condition (2.25) may be written as:

$$x_t \geq PC\left(s_t, T^*, \omega\right). \tag{2.26}$$

Mathematically, PC equals the abatement rate which reduces emissions and therefore GHG concentrations and radiative forcing by just enough to cause the probability of T_{t+1} exceeding T^* to be exactly ω. A stabilization target is therefore equivalent to a minimum abatement rate. For t close to the initial period, $PC \approx 0$. The temperature is very unlikely to exceed 2°C regardless of the emissions level, so the constraint is non-binding. Instead, in the early periods the planner chooses x_t anticipating that constraints in the future will bind with high probability.

Let θ denote the Lagrange multiplier on the probabilistic constraint. The recursive version of the problem, which includes the probabilistic constraint, is then:

$$V(s) = \max_{k', x \in [0,1]} \left\{ u\left(c\right) + \theta \left[x - PC\left(s, T^*, \omega\right) \right] + \beta\left(t\right) \int_{-\infty}^{\infty} V\left[s'\right] N\left(\mu T, \sigma_H\right) d\tilde{H} \right\}, \tag{2.27}$$

subject to equations (2.17)-(2.24) and the Kuhn-Tucker conditions. In period t, the planner anticipates facing constraints in periods $t+i$, which restrict the probability that $T_{t+i+1} \geq T^*$ for all $i = 1, 2, \ldots$. Therefore, constraints on the probability that $T_{t+i+1} \geq T^*$ are implicit in V, since the planner anticipates in period t facing a probabilistic constraint in $t+i$.

[14] Note that this difference only occurs with uncertainty. In a certainty model, the peak warming target and the period-by-period target are identical. Thus the alternative illustrates one way in which temperature targets work very differently under uncertainty.

2.3 Feasibility

2.3.1 Tightness of the Constraint

Given current information and an emissions path $X_{t-1+i} \equiv \{x_j\}_{j=t...t-1+i}$, an expected probability $\omega\left(s_t; X_{t-1+i}\right)$ exists of exceeding the target. Appendix A.1.1 calculates the probabilities $\omega^{\min}_{t+i} = \omega\left(s_t; X_{t-1+i} = 1\right)$ and $\omega^{\max}_{t+i} = \omega\left(s_t; X_{t-1+i} = 0\right)$, which are the expected probabilities of exceeding the target when the abatement rate is set to 1 and 0, respectively, for all $t, \ldots, t-1+i$, conditional on current information. For $i = 1$, the probabilities satisfy:

$$1 = PC\left(s_t, T^*, \omega^{\min}_{t+1}\right) \tag{2.1}$$

$$0 = PC\left(s_t, T^*, \omega^{\max}_{t+1}\right) \tag{2.2}$$

Values of ω greater than $\omega^{\min}_{t+i}$ are expected to be feasible policies in $t+i$, conditional on current information. For $\omega > \omega^{\max}_{t+i}$, the constraint is expected not to bind in period $t+i$, since the planner can set $X_{t+i} = 0$ for all periods up to $t+i$ and still expect to satisfy the constraint.[15] Conversely, any value of $\omega < \omega^{\min}_{t+i}$ implies in expectation $PC\left(s_{t+i}\right) > 1$, even if the abatement rate is set to one immediately for all periods up to $t+i$. By doing this exercise we glean intuition as to how tight the constraint is as a function of ω, which helps to explain the results in the next section. The tightness of the constraint may equivalently be controlled by altering T^*, but we assume here T^* is a given policy.

The set of probabilities achievable with an interior abatement rate expands over time, since the planner can lower future temperatures via a sustained reduction in emis-

[15]Note that the planner may optimally set $x_t > 0$, even if the constraint does not bind in $t+i$, anticipating that the constraint will bind at some point after $t+i$. Here a low value of $\omega^{\max}_{t+i}$ indicates the constraint is relatively slack in $t+i$, since even a zero abatement policy will in expectation satisfy the constraint in $t+i$ for most values of ω.

sions. On the other hand, the uncertain climate feedback parameter has a multiplicative effect over time. Therefore, future temperatures are more uncertain and therefore are more difficult to control. Figure 2.1 plots the mean of 10,000 simulated values of $\omega_t^{\min}$ and $\omega_t^{\max}$ for the given parameter values.

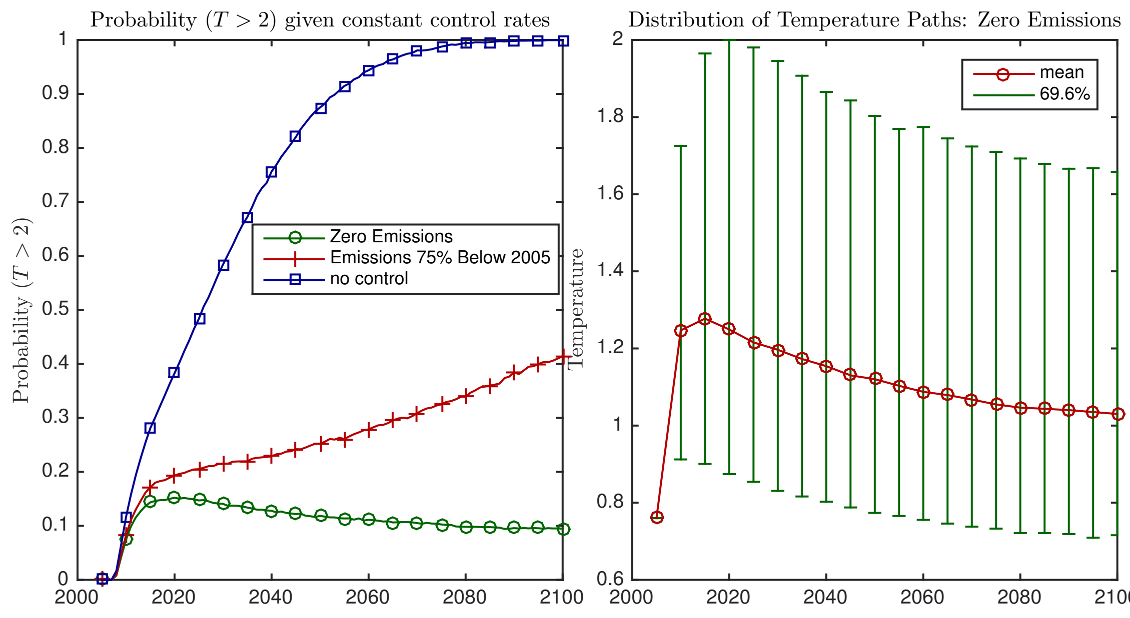

Figure 2.1: Probability of exceeding $T^* = 2$°C, given various emissions policies.

For the first period, the probability of exceeding 2°C is nearly zero. Hence the current probabilistic constraint is non-binding for almost any ω. Given current information, however, there exists an approximately 15.2% chance that the 2°C target will eventually be exceeded, even with an immediate, permanent drop to zero emissions. Therefore, values of $\omega < 0.152$ are expected to be infeasible given today's information.

The fat upper tailed uncertainty, calibrated from Roe & Baker (2007), is visible in the right panel of Figure 2.1, where the error bar extends further above the mean than below. Fat tailed uncertainty implies very high values of the climate sensitivity occur relatively more often than under the normal distribution. Therefore, the temperature change exceeds the 2°C target with surprisingly high probability, regardless of the emissions policy.[16]

An immediate, permanent, 75% drop in endogenous emissions below 2005 levels is expected to violate the constraints for any $\omega < 0.41$ by 2100.[17]

Figure 2.1 also shows that a zero emissions policy can eventually overcome the inertia in the climate and reduce the probability of exceeding 2°C to near zero. A policy of zero emissions over time will slowly return the GHG concentrations to preindustrial levels (see equation 2.7).

The planner chooses a current emissions policy, anticipating future emissions policies over a period of decades, which is expected to satisfy future constraints with an interior abatement rate, given current information. However, learning and inertia in the climate may cause the temperature to enter the infeasible range, even along a path where all long run probabilistic constraints were expected to be satisfied.

2.3.2 A Feasible Constraint

Assuming emissions cannot be negative, constraint (2.26) may not be feasible; the set of $x \in [0, 1]$ which also satisfy (2.26) may be empty, violating a necessary condition for the existence of a maximum. The problem is not feasible when $PC(s, T^*, \omega) > 1$, since the maximum abatement rate is one. This occurs when the temperature rises close to or above T^*. In this case, given the inertia of the climate, even an abatement rate of 1 cannot reduce the temperature enough to satisfy the constraint.

[16] Relatively high values of the feedback parameter cause both GHG concentrations and weather shocks to have more persistent effects on temperature, both of which amplify the variance of temperature over time.

[17] An immediate 75% drop below 2005 emissions is a far stricter policy than, for example, a 350 ppm GHG target or the Kyoto agreement of 7% below 1990 levels.

Feasibility is also affected by ω: as $\omega \to 1$, the planner is allowed to exceed the target with high probability, and so the model has a feasible solution even if the temperature is relatively high.

To solve the infeasibility problem while keeping with the spirit of a stabilization target, we assume that $x = 1$ whenever $PC > 1$. That is, whenever even zero emissions results in the target being exceeded with probability greater than ω, the planner must return to the feasible range as quickly as possible by setting $x = 1$.[18] Let θ_t denote the Lagrange multiplier for the probabilistic constraint in t, then the constrained problem becomes:

$$V(s) = \max_{k', x \in [0,1]} \left\{ u(c) + \theta \left[x - \min \left\{ PC(s, T^*, \omega), 1 \right\} \right] + \beta(t) \int_{-\infty}^{\infty} V[s'] \, N(\mu T, \sigma_H) \, d\tilde{H} \right\}, \tag{2.3}$$

subject to equations (2.17)-(2.24) and the Kuhn-Tucker conditions.

Problem (2.3) always has a feasible solution. Further, when the probabilistic constraint is exceeded due to a large random weather shock or an unexpectedly high realization of β_1, the planner must return as quickly as possible (by setting $x = 1$) to the range where $PC < 1$. This is in keeping with the idea that exceeding the target causes damages and should be avoided if possible. The planner anticipates future constraints in

[18]Other options are possible, but less attractive for welfare analysis. Ignoring infeasible constraints is not attractive because the planner would have an incentive to push the climate change as close as possible to T^* in the hope that the climate will go over the target so that the constraint can be ignored. The other option would be to include a penalty function for going over T^* (Neubersch, Held & Otto 2014). However, a penalty for high temperatures already exists (the damage function), so it is unclear from a welfare perspective what real world phenomena to calibrate the penalty to. In contrast, here exceeding the constraint causes two penalties: first damages increase, and second the planner must incur the cost of returning as quickly as possible back to the target. Both penalties are therefore implicitly consistent with the welfare costs and damages in the model. Finally, note that our solution is not equivalent to a sufficiently strong penalty function. Both a very strong penalty function and our assumption imply $x = 1$ when the target is exceeded. But they have different effects on current emissions: if the penalty function is sufficiently strong, current emissions are less as the planner has a stronger incentive to reduce the risk of exceeding the target.

(2.3), through the future value $V[s']$. For example, a future value s' that has a relatively high T' has a lower value in the constrained problem, because the planner anticipates having to set $x = 1$ in the future to satisfy the constraint. The planner faces a tradeoff: choose a relatively low emissions abatement rate and take the risk that the temperature exceeds the target and be forced to set $x = 1$, or raise the current abatement rate as a precaution. Since abatement costs are convex, the planner prefers to smooth abatement costs over time and therefore raises current abatement as a precaution. However, the planner cannot eliminate all risk (even with zero emissions), and so the computational question is how much extra risk is optimal.

The next section analyzes the how the constraints affect optimal abatement and temperature paths, determines how the planner adjusts abatement today in response to future constraints, and determines how learning affects the constrained optimal policy.

2.4 Results

2.4.1 Optimal Policy and Uncertainty

Appendix A.1.3 details the solution method. First, we analyze how the optimal abatement policy varies with the probabilistic target and the uncertainty. Figure 2.2 plots the optimal abatement policy for two different true values of β_1: the mean of the prior, for which a doubling of GHGs causes a steady state temperature change of $\Delta T_{2\times} = 3.08$°C, and a relatively high value for which $\Delta T_{2\times} = 4$°C.

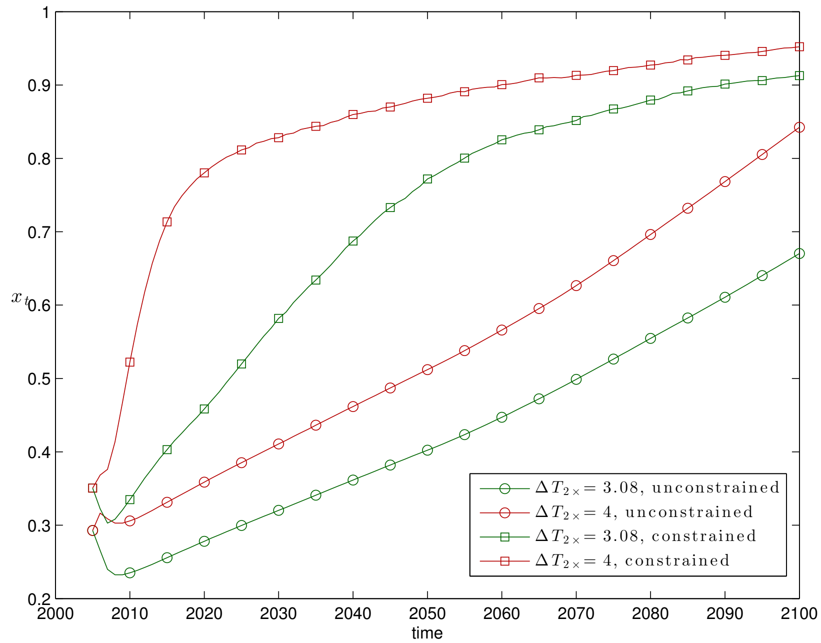

Figure 2.2: Emissions abatement rate for $\Delta T_{2\times} = \{3.08, 4\}$, unconstrained and constrained with $\omega = 0.5$.

In both cases, the probabilistic constraint increases the abatement rate. The constraint increases the abatement rate initially as the planner must begin reducing emissions immediately to prevent the target from eventually being exceeded. The initial abatement rate with the probabilistic constraint is 35%, in contrast to the unconstrained initial abatement rate of 29%. As new information arrives, when $\Delta T_{2\times} = 4$, the planner updates the prior and therefore must increase the abatement rate to meet the target. When $\Delta T_{2\times} = 4$, the constrained planner must dramatically increase the abatement rate to 71% by 2015. In contrast, when $\Delta T_{2\times} = 3.08$, the constrained planner has more time to slow the climate inertia, and the abatement rate is only 33% in 2015.

The unconstrained abatement rate drops over the period 2005-2010. Kelly & Tan (2015) show that, while overall learning of the climate sensitivity is a slow process, the planner can rule out extreme values of $\Delta T_{2\times}$ relatively quickly if the true climate sensitivity is close to the mean of the prior. Therefore, when $\Delta T_{2\times} = 3.08$, learning reduces one motivation for early abatement- to insure against potentially extreme values of the climate sensitivity. The abatement rate then rises over time as abatement becomes less expensive, unabated emissions rise due to economic growth, and wealthier future households are more willing to purchase abatement.

The difference between the constrained and unconstrained abatement rates in 2010-2060 is much greater when $\Delta T_{2\times} = 4$. When $\Delta T_{2\times} = 4$, the optimal temperature rises. The planner learns the climate is more sensitive to GHG concentrations, and so the abatement expenditure required to keep the temperature at a given level increases, but the benefits are unchanged. In contrast, by definition, the probabilistic constraint employs a fixed target. Therefore, the difference between the unconstrained and constrained policies rises with $\Delta T_{2\times}$ because the constraint is inflexible: T^* and ω cannot adjust as new information arrives.

Figure 2.3 plots the average temperature changes for the above two cases.

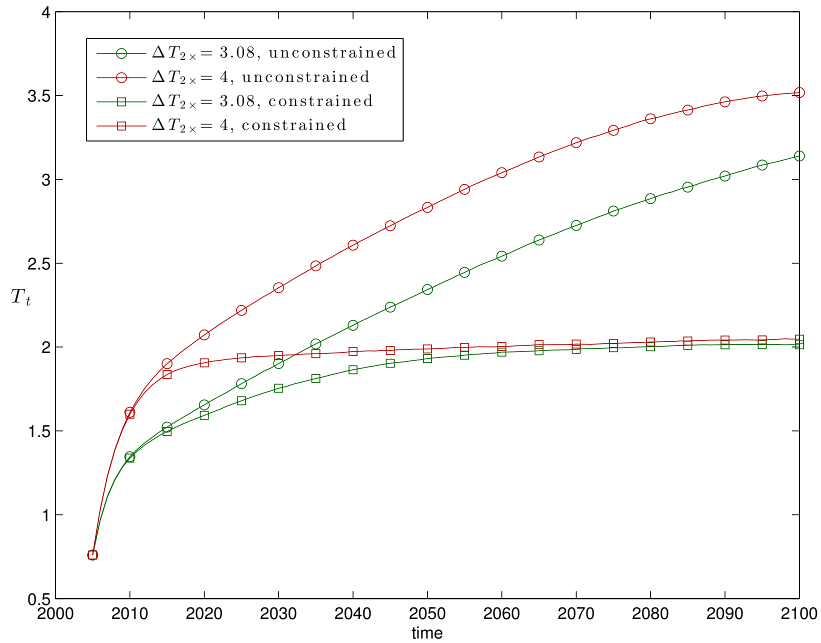

Figure 2.3: Temperature change for $\Delta T_{2\times} = \{3.08, 4\}$, unconstrained and constrained with $\omega = 0.5$.

When $\Delta T_{2\times} = 3.08$, the unconstrained optimal temperatures cross the target in 31 years. Therefore, the planner has more time to slow climate change and can spread out the increase in the abatement rate. In contrast, when $\Delta T_{2\times} = 4$, the unconstrained temperature crosses the target in only 14 years. Therefore, in the constrained case, the abatement rate must rise more quickly to keep the temperature below the target. Notice the unconstrained optimal temperature rises when $\Delta T_{2\times}$ is higher as the planner responds to the higher required expenditure to keep the temperature at a given level by letting the temperature rise more. However, the target stays fixed at 2°C.

Next, we examine how abatement policy and the temperature respond to changes in the probability of exceeding the target, ω. We solve the model (2.3) for various values of ω, and simulate each solution 10,000 times. Figure 2.4 reports the mean optimal abatement rate for each ω, assuming the true climate sensitivity equals the mean of the prior.

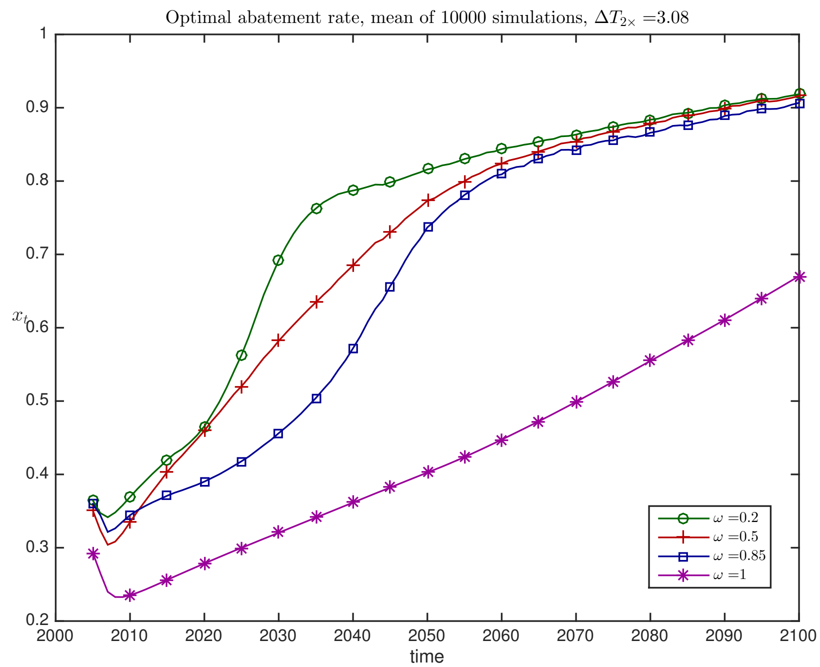

Figure 2.4: Optimal constrained abatement policy.

In Figure 2.4, $\omega = 1$ corresponds to the unconstrained optimum. As the maximum allowable probability of exceeding the target (ω) decreases, the constraint becomes more strict, causing the optimal abatement to rise. Non-monotonic regions arise in the first few periods, as the optimal abatement rate responds to a complex set of changing variables. Uncertainty is decreasing, causing the planner to become less concerned about the risk that the climate sensitivity and damages will be high. The cost of abatement and emissions intensity are falling exogenously, while unconstrained emissions and wealth are rising due to growth in productivity and capital.

The planner must ramp up abatement spending relatively early when $\omega = 0.2$, since in fact the temperature must converge to a level in which random weather shocks cause the temperature to exceed 2°C with probability of at most 0.2. This temperature is below 2°. In contrast, when $\omega = 0.85$ the planner can wait a few more years before ramping up abatement, since the constraint allows the temperature to converge to slightly above 2°C, requiring only that the weather shock causes the temperature to fall below 2°C with probability 0.15.

The unconstrained optimal abatement is far less than the constrained optimal abatement, even for $\omega = 0.85$. Regardless of ω, the ultimate temperature must stay relatively close to 2°, since the standard deviation of the weather shocks is only $0.11°$. In contrast, the unconstrained optimal temperature rises to over 3°C, which requires much less abatement.

After 2050, abatement rates are similar for $\omega = 0.2$ and $\omega = 0.85$. At this point, learning has resolved much of the uncertainty and the climate no longer has an upward trend.[19] The difference in abatement rates therefore reflects only that the mean steady state temperature differs with ω. However, the mean steady state temperature is only about 0.2° higher when $\omega = 0.85$ versus $\omega = 0.2$. Therefore, the abatement rates are relatively close in later years. Figure 2.5 shows the temperature change for the same simulations as Figure 2.4.

[19] Our results are sensitive to the endogenous rate of learning. Kelly & Tan (2015) perform a detailed analysis of the rate of learning in a model with fat tailed uncertainty about the climate sensitivity, and consider a number of extensions, including adding additional uncertainties. Adding additional uncertainties would amplify the effect of targets on emissions and welfare. Learning would slow and the variance of the temperature would increase, both causing the target to be more difficult to adhere to.

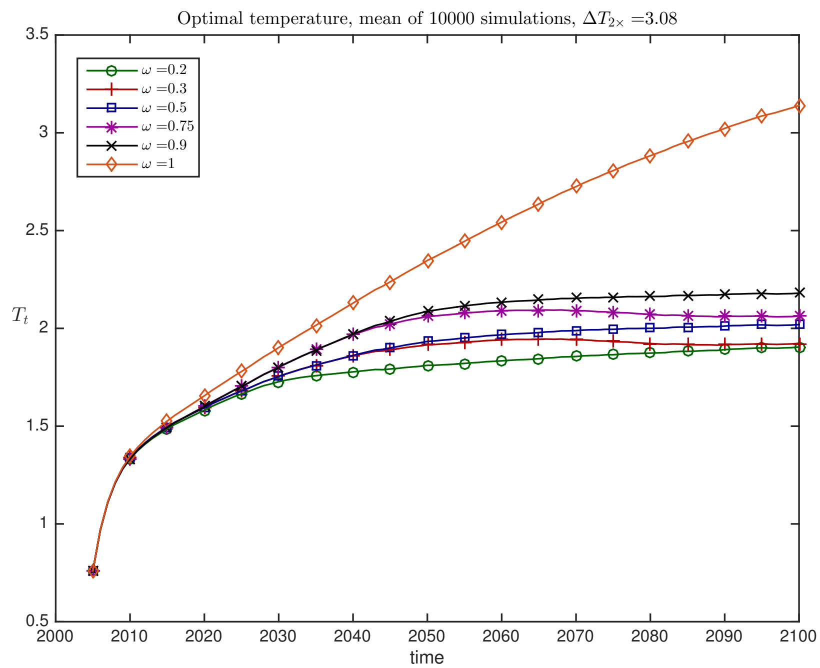

Figure 2.5: Optimal constrained temperature, with true $\Delta T_{2\times} = 3.08$.

The unconstrained optimal maximum mean temperature change is 3.26°C, so the 2° target is about 1.26 degrees too stringent, if the true climate sensitivity equals the mean of the prior.[20] Smaller values of ω imply a smaller maximum mean temperature change. For $\omega = 0.5$, the maximum mean temperature is very slightly below 2°C by 2100. The weather shock is normally distributed with mean zero. Hence, starting at 2°C, the temperature next period exceeds the constraint with probability 0.5. Similarly, since the standard deviation of the weather shock is 0.11°, with probability 0.2 the weather shock is greater than 0.09°. Hence, the planner sets the maximum temperature change when $\omega = 0.2$ to slightly below 1.91 (actually 1.89). A high temperature shock implies costs rise next period, as the planner must increase the abatement level, whereas reducing the

[20]Our model is based on the Nordhaus DICE model, which has no tipping points, irreversibilities, etc. Other models with these features may feature smaller optimal maximum temperature changes.

temperature below 1.91 increases abatement costs today, but reduces the risk of higher abatement costs later. Abatement costs are convex, however, which implies the planner will reduce the temperature below 1.91 as a precaution. This effect is small, however.

Overall, the maximum mean temperature is not very sensitive to ω. At $\omega = 0.9$, the maximum temperature change is 2.16°C (the level at which a weather shock causes the temperature to drop below 2°C with probability 0.1). The difference in maximum temperatures between $\omega = 0.2$ and $\omega = 0.9$ is only 0.25°C. The weather shock does not have enough variance to cause changes in ω to generate large differences in temperature. Although the maximum temperature must approach the unconstrained maximum temperature as $\omega \to 1$, the convergence is highly nonlinear. Suppose for example that $\omega = 0.999997$. Then, the planner can set the abatement rate so that the temperature converges to 2.5°C, since at 2.5° the probability of a large negative weather shock that causes the temperature to decrease to 2° is exactly $1 - \omega$. However, even with $\omega = 0.999997$, the mean temperature of 2.5° is still far from the unconstrained temperature of 3.26°. Therefore, the imposition of a probabilistic constraint causes a significant drop in temperature and requires significantly more abatement, for almost any value of ω.

The above analysis is for the case where the true climate sensitivity equals the mean of the prior. Sufficient time exists to slow the inertia in the climate and stabilize the temperature, both because the true climate sensitivity is moderate and because the planner is not very surprised by the true climate sensitivity. However, if the climate sensitivity is sufficiently high, the temperature exceeds the target, regardless of ω. For example, Figure 2.6 plots average temperature change over time as a function of ω when $\Delta T_{2\times} = 5$.[21]

[21] Weitzman (2009) using IPCC data assigns prior probability that $\Delta T_{2\times} \geq 5 = 0.07$.

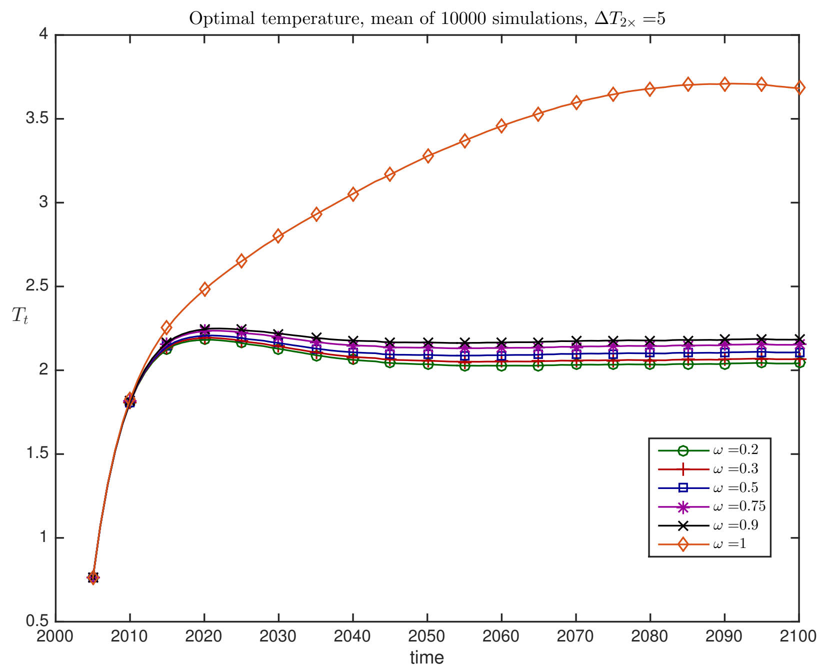

Figure 2.6: Optimal constrained temperature with true $\Delta T_{2\times} = 5$

The mean optimal unconstrained temperature change rises to 3.71°C, indicating the optimal unconstrained temperature is sensitive to climate sensitivity. The mean optimal unconstrained temperature change increases because when the climate is more sensitive to GHG concentrations, the abatement expenditures required to keep the temperature at a given level rises, but the benefits are unchanged. In contrast, the target is by definition inflexible and does not vary with the resolution of uncertainty. The mean maximum temperature change falls with ω as in Figure 2.5, but is not very sensitive to ω. When $\Delta T_{2\times} = 5$ and $\omega = 0.2$, on average the temperature exceeds the target for about 104 years. Since $\Delta T_{2\times} = 5$ is unexpectedly high, the climate has unexpected inertia, and crosses the target after just a few years. The problem is exacerbated by the fact that the planner believes $\Delta T_{2,\times} < 5$ for many years, and underestimates the effect of emissions on the climate. Once the temperature crosses the target, the planner sets emissions to zero, but the climate continues to rise on average because the inertia is unexpectedly strong. Note that Figure 2.5 is an average of 10,000 simulations. In 18% of the simulations, the weather shocks are low enough such that the temperature falls below 2°C in less than 40 years. Conversely, 34% of the simulations require 150 years or more to reduce the temperature back to the target.

2.4.2 Optimal Reaction to Future Constraints

Given a sequence of emissions levels, the planner knows that the temperature may exceed the target in the future. If so, the planner endures a cost of setting emissions equal to zero (and high damages). If abatement is inexpensive, the planner may elect to increase abatement, paying today to reduce the risk of additional costs in the future. The planner may even elect to increase abatement to the point where the probability of exceeding the target in the future is less than ω, if the cost of abatement is inexpensive relative to the cost of exceeding the target in the future. Conversely, if abatement is expensive, the planner may take on more risk, perhaps even to the point where today the planner expects the target to be exceeded in the future with probability greater than ω.

The probability today of exceeding the target in the future is therefore a constrained optimal policy, trading off these costs. In contrast, ω acts in the future to lower the threshold temperature beyond which the planner incurs the extra costs. Thus reducing ω creates an extra incentive to hedge by reducing emissions today, but is different from optimal probability of exceeding the target.

Table A.2 gives the percentage of 10,000 simulations for which the temperature exceeds the target. In the unconstrained case the temperature exceeds the target for 72.1% of the simulations in 2050. As ω falls below one, the planner increases current abatement. Since abatement is low and abatement cost is convex, small increases in abatement are not very costly and by increasing abatement the planner reduces the probability of exceeding the target in the future and incurring the cost of setting emissions equal to zero. Thus for $\omega > 0.3$, the planner sets abatement high enough so that the probability of exceeding the target in 2050 is strictly less than ω.

However, given the convexity of abatement costs, eventually the cost of increasing abatement today rises to the point where further reductions in the probability of exceeding the constraint in the future are no longer optimal. For $\omega = 0.2$, the target is still exceeded approximately 35% of the time in 2050 (from Figure 2.1, an immediate 75% reduction in emissions cause the target to be exceeded with probability 0.25 in 2050). The planner is unwilling to pay the cost of more severe emissions reductions today and accepts the penalty of occasionally exceeding 2° and having to set future emissions equal to zero. The probability of exceeding the target is therefore constrained optimal in the sense that the planner chooses a future probability of exceeding the target today given the costs and benefits of reducing that probability.

Finally, the last column of Table A.2 shows the autocorrelation of periods in which the target is exceeded. The correlation is very positive: once the target is exceeded, inertia typically causes the the temperature to exceed the target in subsequent periods.

2.4.3 Learning

Learning plays an important role in the presence of a stabilization target. If learning reduces uncertainty relatively quickly, the planner can adjust abatement and the temperature is less likely to exceed the target. Further, learning directly affects the stringency of the probabilistic constraint.

With regard to the speed of learning, from equation (2.23), the precision of the feedback parameter grows approximately linearly, which implies the variance shrinks according to a power law (t^{-1}). Nonetheless, the rate of learning about the climate sensitivity differs from the rate of learning about the feedback parameter. From equation (2.10):

$$\Delta T_{2\times} = \Omega\tilde{\lambda} = \frac{\Omega}{\alpha\left(1-\tilde{\beta}_1\right)}, \tag{2.1}$$

so small differences in β_1 can imply relatively large differences in λ and $\Delta T_{2\times}$. Since the sensitivity of the climate to GHG concentrations is ultimately more important for determining abatement policy, especially with a stabilization target, we focus on the rate of learning about $\Delta T_{2\times}$ rather than the feedback parameter.

Kelly & Tan (2015) show that overall learning about the climate sensitivity is a slow process, although the planner can effectively rule out values of the climate sensitivity much higher than the true value relatively quickly. Figures 2.7 and 2.8 confirm these results. The left panel of Figure 2.7 plots the evolution of the mean of the prior for 200 simulations where the true $\Delta T_{2\times}$ equals the prior. The percentage of simulations for which the estimate of $\Delta T_{2\times}$ is above 3.5° falls from 28.5% initially to 2% in only 14 years. However, even in 2050 significant uncertainty remains. The difference between the largest and smallest value of the estimate of $\Delta T_{2\times}$ is 0.42° in 2050. Further, learning slows as the true value increases. The right panel shows that if the $\Delta T_{2\times} = 5°$, in 2050 the difference between the largest and smallest estimate of $\Delta T_{2\times}$ is 0.89°.

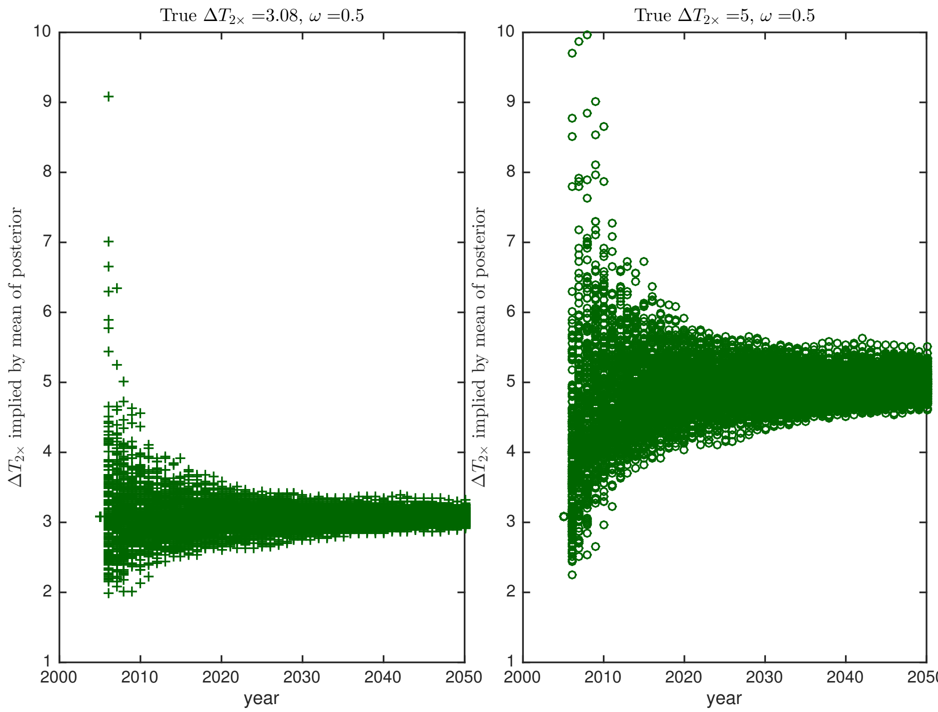

Figure 2.7: Evolution of the estimate of $\Delta T_{2\times} = 5°$ when $\omega = 0.5$

Alternatively, Figure 2.8 shows that the standard deviation of the estimate of $\Delta T_{2\times}$ falls initially as very high values are ruled out,[22] but then falls at a much slower rate after.

[22]However, if the true $\Delta T_{2\times}$ is greater than the prior, the standard deviation of $\Delta T_{2\times}$ does not decrease monotonically. The standard deviation of the feedback parameter does decrease monotonically (equation 2.23), but increases in the mean estimate of the feedback parameter implies small changes in the feedback parameter have larger changes in $\Delta T_{2\times}$. Therefore, increases in the mean of the feedback parameter estimate, which occur when the true value exceeds the prior, increase the standard deviation of the estimate of $\Delta T_{2\times}$ (equation 2.1).

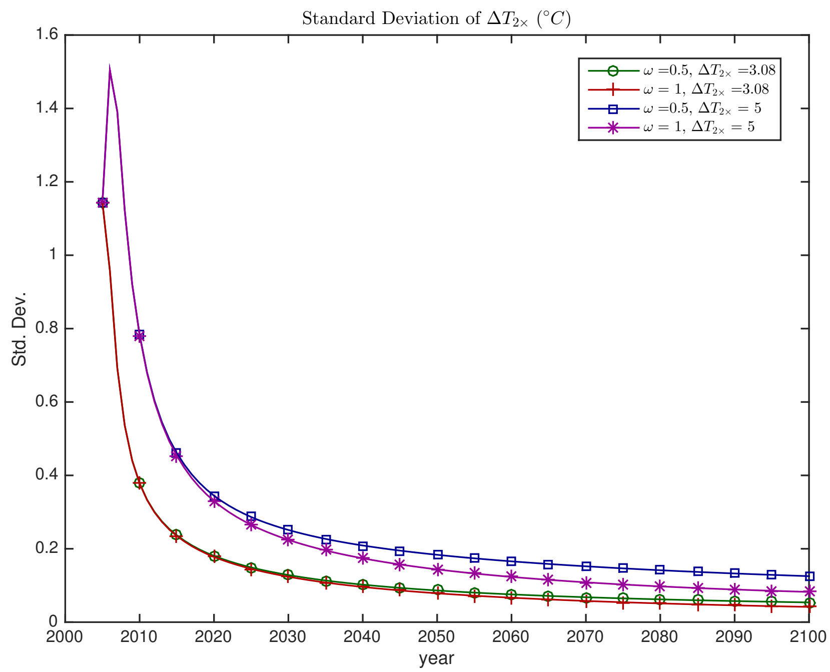

Figure 2.8: Evolution of the standard deviation of the prior belief distribution of $\Delta T_{2\times}$.

Figure 2.8 also shows that the stabilization target slows the rate of learning. Because the probabilistic constraint reduces emissions, the climate change signal is less visible in the noisy weather shocks, which slows learning. When the true $\Delta T_{2\times} = 3.08$°C, the standard deviation of the estimate of $\Delta T_{2\times}$ is 9.8% higher in 2050 when the stabilization target is present. When the true $\Delta T_{2\times} = 5$°C, the standard deviation of the estimate of $\Delta T_{2\times}$ is 27.7% higher in 2050 when the stabilization target is present. Thus the stabilization target slows learning most when learning is most helpful for keeping the temperature below the target. Thus, slow learning and the inertia in the climate explain why the temperature exceeds the target when the climate sensitivity is relatively high (Figure 2.6).[23]

[23] Although consistent with the previous literature, learning on a scale of decades may still appear

The primary channel through which learning affects abatement policy is the effect of reducing the variance of the feedback parameter on the probabilistic constraint. Since the constraint binds frequently, changes in the constraint induced by a reduction in the variance of uncertainty directly affect the abatement rate. It is straightforward to verify that:

$$\frac{\partial PC\left(s_t, T^*, \omega\right)}{\partial 1/\eta_t} > (<, =)0 \;\Leftrightarrow\; \omega < (>, =)0.5. \tag{2.2}$$

Therefore, how learning (reducing the variance, $1/\eta$, of the feedback parameter) affects abatement depends on ω.

Suppose, for example, the constraint binds so that $x_t = PC\left(s_t, T^*, \omega\right)$, or equivalently $Pr\left(T_{t+1} \geq T^*\right) = \omega$. Suppose now learning reduces the variance of the estimate of the feedback parameter. Reducing the variance means extreme realizations of the climate sensitivity, which imply $T_{t+1} > T^*$, are less likely. Therefore, if $\omega < 0.5$, reducing the variance causes $Pr\left(T_{t+1} \geq T^*\right) < \omega$. The constraint relaxes (PC decreases), and x_t falls and T_t rises closer to T^*. Therefore, for $\omega < 0.5$, we expect the temperature with learning to move closer to T^* over time as the variance falls, relative to the case where no learning takes place.

Surprisingly, equation (2.2) shows this result reverses when $\omega > 0.5$. Suppose again that $x_t = PC\left(s_t, T^*, \omega\right)$, or equivalently $Pr\left(T_{t+1} \geq T^*\right) = \omega$. Recall from equation and magenta lines of Figure 2.5. Now learning makes extremely high realizations of the climate sensitivity less unaffected by reducing the probability of extremely high realizations. However, extremely low realizations of the climate sensitivity are also less likely, and realizations slightly less than the mean are more likely. increase and T_t must

relatively rapid given that many climatic processes evolve over a period of centuries. Indeed, our two equation climate model with a single uncertainty likely overstates the rate of learning in complex climate models with many uncertainties. Introducing additional uncertainties which slow learning would magnify the results here, since additional uncertainties would make a stabilization target more difficult to adhere to.

decrease. Therefore, for $\omega > 0.5$, we expect the temperature with learning to move closer (from above) to T^* over time as the variance falls, relative to the case where no learning takes place. For the case of $\omega = 0.5$, reducing the variance has no effect on PC, since the normal distribution is symmetric.

Figure 2.9 confirms the above intuition. For $\omega = 0.25 < 0.5$, the model with learning eventually has a higher temperature, as learning reduces the likelihood that the climate sensitivity is high, which allows the planner to get closer to T^* without violating the constraint, relative to the case without learning.[24] Conversely, for $\omega = 0.75 > .5$, the model with learning eventually has a lower temperature, as learning reduces the likelihood that the climate sensitivity is low, which increases the probability of exceeding the target. Therefore, learning tightens the constraint and abatement rises and the temperature falls, relative to the case without learning. For $\omega = 0.5$ the constraint is unaffected by learning and, since the constraint binds both with and without learning, the temperatures are almost identical. Note the true climate sensitivity equals the prior in Figure 2.9, which allows us to focus on the effect of reducing the variance of uncertainty in isolation.[25]

[24]Note that for the case without learning, the model is re-solved so that the planner anticipates no further information about the climate sensitivity is forthcoming.

[25]In addition, the difference in temperatures with and without learning widens as $\Delta T_{2\times}$ increases.

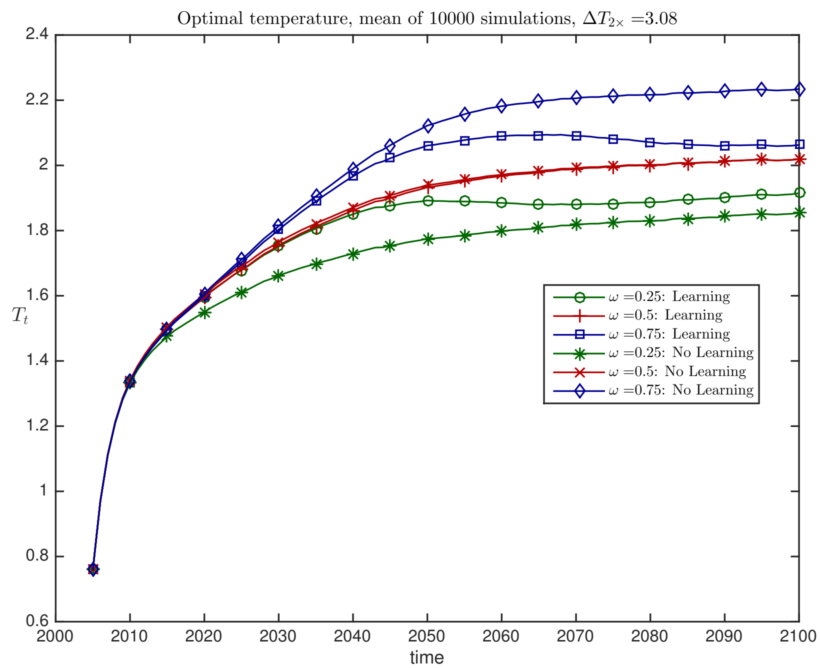

Figure 2.9: Optimal constrained temperature with true $\Delta T_{2\times} = 3.08$

2.4.4 Welfare Loss

The probabilistic constraint at least weakly restricts the choice set for the planner, and therefore must result in a welfare loss, relative to the unconstrained problem.[26] Our interest is in how the welfare loss varies with ω and how uncertainty affects the welfare loss.

Figure 2.10 graphs the welfare loss as a percentage of the welfare of the unconstrained problem, $\omega = 1$.

[26] We are following, for example, Nordhaus (2007), who imposes a 2°C constraint a version of the model with no uncertainty, and then calculates the welfare loss. Other authors use cost effective analysis (CEA) or cost risk analysis (CRA), which replace the damage function with the probabilistic constraint (for an excellent discussion of CEA and CRA, see (Neubersch, Held & Otto 2014).) A damage function, despite being uncertain, allows for a transparent interpretation of the welfare costs of temperature changes, which is the focus of this paper.

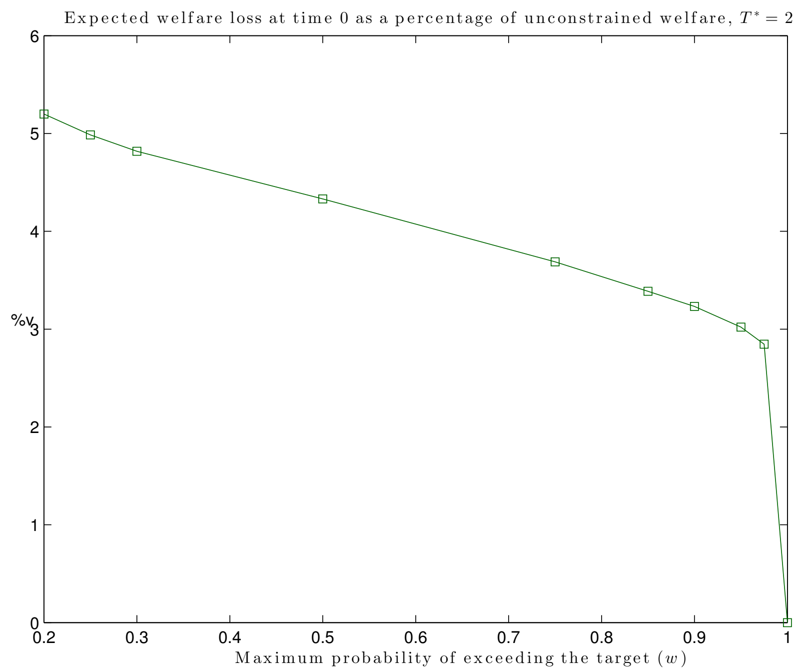

Figure 2.10: Welfare loss as a function of ω

From the graph, the unconstrained problem has no welfare loss, and welfare loss is decreasing in ω since relaxing the constraint reduces the welfare loss. The slope is relatively flat for $\omega < 0.975$, since even relatively large values of ω force the maximum climate change to be close to two degrees, rather the optimum which is 3.26°C and rises with the true value of the climate sensitivity (see Figure 2.5). Welfare loss is about 3-5% of the unconstrained policy, depending on ω (Table A.3 reports the exact welfare losses).

Mastrandrea et al. (2010) and Neubersch, Held & Otto (2014) calibrate a value of $\omega = 0.33$ based on an interpretation IPCC statements calling for policies for which achieving the 2° target is likely. We therefore let $\omega = 0.33$ be the baseline stabilization policy recommendation.[27] Table A.3 then indicates that the baseline stabilization policy

[27]One may alternatively interpret the IPCC statement as the optimal probability of exceeding the target

results in a welfare loss of approximately 4.7%.[28]

Figure 2.11 plots the welfare loss for various true values of $\Delta T_{2\times}$.

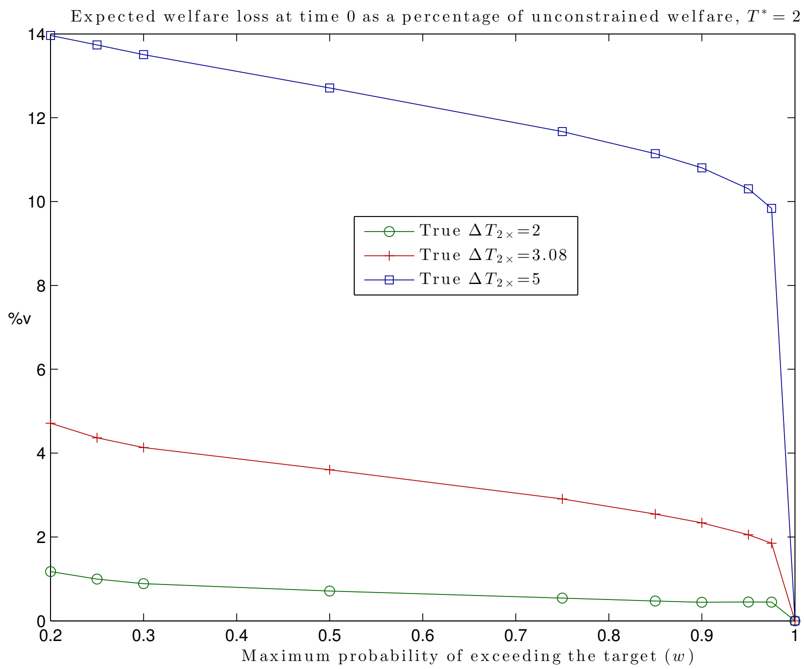

Figure 2.11: Welfare loss as a function of ω and true $\Delta T_{2\times}$.

The welfare loss increases with the true value of $\Delta T_{2\times}$, rising to 14% when $\Delta T_{2\times} = 5$ and $\omega = 0.2$.

When the arrival of new information indicates the climate sensitivity is higher than expected, the unconstrained planner learns that the cost of keeping the temperature at a given level has risen. The planner increases abatement in response to the higher expected damages. Nonetheless, the high climate sensitivity implies temperature is on a higher trajectory. In contrast, the constrained planner must increase abatement

at some point in the future is 0.33. If so, from Table A.2, the appropriate calibration is $\omega \approx 0.2$.

[28]The welfare loss in consumption equivalents is approximately 1%. That is, the household would require an annual consumption increase of 1% in perpetuity to be willing to accept the stabilization constraint (see Table A.3).

still further, to keep the temperature on the same trajectory despite the high climate sensitivity, because the target does not change. Therefore, higher values of $\Delta T_{2\times}$ cause greater welfare losses, because the constraint is inflexible.

Traeger (2014) shows that, for the DICE model under certainty, the maximum temperature change is approximately 3.6°C above preindustrial.[29] Therefore, in the certainty version of the model, the target is set 1.6° too low. After imposing a pure ($\omega = 0$) stabilization target on the certainty version of the model,[30] we find the welfare loss of a 2°C target is 2.84%. This result depends on the specification of the damage function, assumptions about the discount rate, the lack of tipping points, the lack of intrinsic preferences for low temperatures, etc. In the certainty version of the model, more pessimistic parameter assumptions would lower the maximum optimal temperature change to 2° or lower, in which case imposing the 2° target would create no welfare loss.

In contrast, a welfare loss exists in the model with uncertainty, irrespective of parameter assumptions, because the target is inflexible and abatement policy overreacts to transient shocks. Table A.3 shows that the welfare loss in the model with uncertainty ranges between 5.2% for $\omega = 0.2$ to 2.85% for $\omega = 0.975$. The welfare loss with uncertainty differs from the welfare loss with certainty for three reasons. First, the optimal maximum temperature changes with the resolution of uncertainty, but the target is inflexible, creating a welfare loss. In contrast, adjusting the target is not necessary under certainty. Second, with uncertainty the constrained planner must respond to transient weather shocks. These shocks are absent in the model with certainty. Third, the difference between the maximum unconstrained mean temperature and the maximum

[29] In our model with uncertainty, the maximum of the mean temperature change is 3.26°C. The lower temperature occurs because the risk averse planner responds to the risk of damages from a potentially high climate sensitivity by increasing abatement.

[30] A probabilistic target cannot be imposed on the certainty version of the model, since the planner can always choose a path which satisfies the target with probability one, and any path which exceeds the target also does so with probability one.

constrained mean temperature when the true climate sensitivity equals the mean of the prior is smaller under uncertainty than the difference between the maximum unconstrained temperature and the maximum constrained temperature under certainty. The risk of GHG emissions eventually causing large temperature changes makes the planner more cautious under uncertainty. Therefore, the 2° target is not as restrictive when the true climate sensitivity equals the mean of the prior with uncertainty relative to certainty.

The first two channels cause a larger welfare loss relative to certainty, while the third channel causes a smaller welfare loss relative to certainty. Overall the first two channels dominate for all values of ω. The welfare loss of the baseline model ($\omega = 0.33$) under uncertainty is 66% greater than the model with certainty. Table A.3 indicates the welfare loss under uncertainty is 82.8% greater than the model with certainty for $\omega = 0.2$, although for $\omega = 0.975$ the third effect becomes more important, and the welfare loss is only slightly higher under uncertainty.

Unfortunately, the three channels by which uncertainty affects the welfare loss of a stabilization target are difficult to isolate. For example, eliminating the weather shocks from the model removes the need for the optimal constrained abatement policy to respond to transient shocks. However, without weather shocks, the planner learns the climate sensitivity after one observation, and the model reduces to the certainty case.

However, some channels can be approximately isolated. For example, consider an exercise designed to isolate the welfare loss caused by the response of abatement policy to transient shocks. We perform a simulation which keeps two sets of temperature data. The first set of temperature data includes the weather shocks and the second set of data sets all shocks equal to zero. We use the temperature data with shocks to compute the updates to the mean and variance of the prior uncertainty distribution, and use the temperature data without shocks to compute damages, the optimal decisions, and the future temperature. In this way, the speed of learning is nearly identical to the previous case, but abatement policy does not directly adjust to transient shocks, which are absent.[31] Table A.3 shows that eliminating weather shocks (except for updating priors) reduces the welfare loss by 6.5-9.7%.

The third effect is that the probabilistic target used under uncertainty is less restrictive than the pure stabilization target under certainty. Under certainty, the maximum unconstrained temperature is 3.6°, whereas the constrained target is 2°C, a difference of 1.6°C. In contrast, Figure 2.5 shows that the unconstrained maximum mean temperature is 3.26°, and that the maximum constrained mean temperature is 1.89-2.24°, depending on ω. When $\omega = 0.2$, the temperature difference is 1.37° versus 1.6° under certainty. Therefore, for $\omega = 0.2$, the third effect, the difference between optimal and constrained temperature, is similar for uncertainty and certainty. The welfare loss for $\omega = 0.2$ is 82.8% higher under uncertainty, indicating that removing the third effect causes the welfare loss to increase.

Overall, then the majority of the added welfare loss in the model with uncertainty results from the inflexibility of the target. The climate sensitivity is uncertain, and the inability of the constrained planner to adjust policy as new information arrives reduces the value of learning in the model.

2.5 Conclusions

In this paper, we evaluated the common policy recommendation of a 2°C temperature target in the context of an integrated assessment model with uncertainty and learning about the climate sensitivity. Uncertainty fundamentally changes the feasibility and

[31] That is, we compute $x\left(S\left(T\left(\nu=0\right),\mu\left(\nu\right),\eta\left(\nu\right)\right)\right)$, which differs from the abatement policy when weather shocks affect the temperature, eventually slightly affecting the speed of learning.

welfare implications of a stabilization target. We show that even with an immediate reduction of all GHG emissions to zero, the temperature eventually exceeds 2°C with probability 0.15. Given that the temperature is already 0.8°C above pre-industrial and climatic processes are highly uncertain and subject to inertia, it is difficult to envision controlling the climate to the precise degree implied by the 2°C target in the short run. Nonetheless, we show that as learning resolves the planner can eventually achieve the target with sustained emissions reductions over time. An immediate policy implication is that implementing a stringent emissions reduction policy is imperative if the temperature change is to be kept under 2°C: even with an immediate 75% reduction in emissions, the 2°C target is exceeded with probability 0.41 by 2100.[32]

Further, we show that adding uncertainty changes the welfare effect of stabilization targets in three ways: first, with uncertainty, as new information arrives the optimal temperature path adjusts. Because the stabilization target is by definition inflexible, adhering to a stabilization target causes welfare loss. Second, the temperature randomly evolves over time, occasionally exceeding the target. In this case, the planner must set an excessively high abatement rate to immediately return the temperature back to the target. Third, uncertainty makes the risk averse planner more cautious: the average unconstrained temperature change is lower under uncertainty. Thus, the welfare loss associated with setting the target lower than is optimal is less under uncertainty. We show that the total welfare loss for the baseline policy is 4.7%, most of which is due to inflexibility of the stabilization target.

An important policy implication of our results is that introducing some flexibility into the target can reduce the welfare loss. For example, many regulations and international agreements have provisions whereby the regulation comes up for renewal after a

[32] Indeed, some scientists now argue that exceeding the 2°C target is inevitable (Sanford et al. 2014).

prespecified number of years, at which point the regulation may be changed to reflect new information. Another policy implication is that a less variable target, such as GHG concentrations, is more attractive than temperature. The welfare loss of the target increases with the variance of the shock, through the overreaction channel. Since GHGs are less variable than temperature, less overreaction would occur.

Our model may be extended in a number of ways. We could consider other targets such as limit GHG concentrations to 350 ppm or limiting sea level rise or ocean acidification. Our approach would also likely extend to regulations other than climate change. For example, some fisheries regulations try to achieve a particular stock of fish, even though fish stocks are affected by many uncertain processes other than the size of the catch. One may envision stabilization targets as providing some welfare benefits outside the current model. For example, they could serve as a commitment device to induce firms to invest in irreversible abatement capital.

A 2°C stabilization targets is easier to convey to the public than, for example, a particular carbon tax. Since damages are a function of temperature, it is also easier to understand the impacts of 2°C temperature limit. However, one must use caution in that a stabilization target may convey the false impression that we have precise control over an uncertain climate and that our understanding of the optimal temperature change will not change as new information arises.

2.6 Coal Mining and Health in Appalachia

As a result of the large-scale mechanization in the mining industry over the last 30 years, modern day surface mines generate prolific amounts of air pollution.[33] Although

[33]Scientists agree about other ecological impacts of coal mining. Biologists find mutated species of wildlife in stream beds near surface mining sites (Pond et al. 2008), and hydrologists document elevated levels of heavy metal and salinity in watersheds surrounding both active and reclaimed surface mines (Lindberg et al. 2011). Water samples from drinking wells near reclaimed surface mining sites contains

the U.S. Environmental Protection Agency only detailss sources air pollution at surface mines in Western states, the same activities they analyze take place at surface mines around the world. They list over a dozen sources at surface mines (International Energy Agency 2012). These include drilling, blasting, bulldozing, draglining, loading and unloading of coal and overburden, heavy vehicles operating on unpaved roads, crushing and processing, wind erosion of coal and overburden stockpiles, and exhaust from heavy machinery.

To uncover coal seams efficiently, mining firms make wide use of nitrate-based explosives. The drilling used to places explosives underground generates up to 10.2 kilograms of coal dust per meter of depth for a hole 10 centimeters wide (Nair & Sinha 1987). The dust clouds that linger post-blast contain nitrogen oxides, ammonia, and respirable particulates which scientists show induce cardiovascular failure among laboratory rats (Ghose & Majee 2000, Knuckels et al. 2012). After a seam is uncovered, bucket-wheel excavators or dragline excavators scrape chunks of coal out of seams. More fugitive dust comes from loading and unloading haul trucks, and from trucks idling at mines (Siami-Irdemoosa & Dindarloo 2015). Once a surface mine's useful life ends, bulldozers use discarded topsoil and overburden to mitigate the aesthetic and ecological consequences of mining, paradoxically generating more fugitive dust. Coal mines also produce carbon monoxide and volatile organic compounds, which are known to form ozone in the presence of sunlight.

Vehicular traffic around surface mines agitates deposits of waste material, which become breathable once airborne. Samples of air quality along coal roads contain - in addition to respirable particulates - arsenic, lead, and mercury (Aneja, Isherwood & Morgan 2012). Finally, wind blown fugitive dust can escape aboveground waste piles.

dangerous levels of dissolved heavy metals (Palmer et al. 2010).

Since blasting and drilling are more common at surface mines Ghose & Majee (2007) find worse ambient air around surface mines than underground mines.

Asthma is a chronic lung disease.[34] Air pollution causes asthma attacks by inflaming airways upon inhalation, thereby inhibiting proper respiratory functioning (Laumbach & Kipen 2012). When people with asthma are exposed to particulates, ozone, nitrogen dioxide, and other pollutants, their airways become inflamed and narrow. This leads to wheezing, discomfort, and difficulty breathing. Most asthmatics use long-term medications to prevent frequent inflammation of the airways, and short-term medication (bronchodilators) to open the airways in the event of a sudden attack. In the event of a serious attack that cannot be controlled with a bronchodilator, an asthmatic may be forced to visit an emergency room or seek hospital care. In addition to short term risks, research suggests air pollution might cause new-onset asthma cases (Guarnieri & Balmes 2014).[35]

Many studies either provide evidence for the link between pollution and respiratory health, identify sub-populations who are particularly susceptible to air pollution, or rely on it and quasi-experimental research designs to draw causal conclusions about the effect of air pollution on respiratory health. Schwartz & Marcus (1990) find that mortality in London from 1958 to 1972 strongly and positively correlates with daily particulate concentrations, which were elevated during this period due to the London Fog. Choudhury, Gordian & Morris (1997) use another quasi-experimental approach to show that organic particulate concentrations positively affect asthma hospitalization rates. They use individual-level Alaskan insurance data to show that, following a volcanic eruption which caused particulate concentrations to spike, visits to local hospitals for asthma

[34] National Heart, Lung, and Blood Institute: http://www.nhlbi.nih.gov/health/health-topics/topics/asthma

[35] See Currie et al. (2014) for a summary of the long-term effects of early-life exposure to air pollution, which include slowed human capital formation.

jumped. More recently, Schlenker & Walker (2011) use an instrumental variables approach to demonstrate the short-term effects of air pollution from airplane taxiing on hospitalizations.

Unfortunately, the dearth of air quality monitors in West Virginia prevents the same type of approach as Schlenker & Walker (2011), who link variation in zip-code level pollution measures with airplane taxiing times at nearby airports.[36] Moreover, since disentangling the separate effects of highly correlated airborne particulates on health is difficult (Environmental Protection Agency 2004), I refrain from attempting to identify the different effects of each component of coal dust on asthma, and focus instead on quantifying a total effect of coal mining. This approach acknowledges that deteriorated respiratory functioning results from many of the harmful airborne byproducts of coal mining. These include particles known to penetrate the lungs (those less than 10 micrometers in diameter, known as "PM10"), particles that penetrate deep into the lungs (PM2.5), and other mining byproducts like nitrogen dioxide and carbon monoxide - known triggers of asthma among children (Neidell 2004). I seek a total effect of mining on respiratory health because all of the aforementioned pollutants induce asthma attacks, and because this approach yields more straightforward policy recommendations.

2.6.1 Regional Health Disparities in Appalachia

Anti-coal advocacy groups frequently cite a strand of public health research that studies the differences in health outcomes between populations exposed to mining, and popula-

[36] It would be difficult to use variation in quarterly mining to demonstrate a precise link between mining activity and the short term spikes in particulates that cause asthma attacks. Further, any relationship found using a single air quality monitor, of which there are few in the state, would not necessarily describe an entire county's exposure. In the state, PM10 is monitored at four locations: three in the northern panhandle, and one in Charleston. PM2.5 is monitored at one location in the northern panhandle. Figure 2.12 provides some evidence of surface mining's effect on short-term spikes in air pollution: there is small positive correlation ($\rho = 0.14$) between surface tonnage within 20 kilometers of a monitoring station and days with an hourly observation of PM10 above $50\mu g/m^3$. Figure 2.13 shows a negative correlation between underground mining and PM10 spikes.

tions that are not.[37] Hendryx & Ahern (2009) investigate human mortality rates across four different county groups: heavily mined Appalachian coal counties, lightly mined Appalachian coal counties, non-mining Appalachian counties, and non-mining counties outside of Appalachia.[38] They estimate that, between 1979 and 2005, coal mining counties experienced between 1,736 and 2,889 more deaths annually than non-coal counties, after adjusting for covariates like socio-demographic traits, smoking rates, and the supply of physicians. Hendryx & Ahern (2008) use survey data and the same county stratification to find that residents of coal counties experience higher rates of lung disease, kidney disease, and hypertension. They also find that living in a coal county is associated with overall worse health, measured by a hierarchical health variable - residents of coal counties feel sicker than residents of non-coal counties. In addition to coal sensitive ailments, Hendryx, Ahern & Nurkiewicz (2007) also analyze conditions that they deem to be coal-insensitive, and find that diabetes, muskuloskelatal disorders, and psychosis occur independently of exposure to coal. Pless-Mulloli et al. (2000) uses matching techniques and finds no association between living near mines and the prevalence of respiratory illness or asthma severity, but that children near mines sought more doctor consultations for those illnesses, relative to children living far from mines. They also find that children living near surface mines are exposed to higher levels of respirable particulates.

[37] Appalachian Voices: http://appvoices.org/end-mountaintop-removal/health-impacts/

[38] They separate counties based on median coal production over their study period.

Hendryx, Ahern & Nurkiewicz (2007) perform a cross-sectional analysis of Appalachian counties' coal production and hospitalization records from 2001 and find that an additional 1,462 tons in an individual's county of residence leads to a 1 percent increase in their likelihood of hospitalization for chronic obstructive pulmonary disorder (COPD) and that an additional 1,873 tons led to a 1 percent increase in their likelihood of hospitalization for hypertension. This is a surprisingly large result considering that some counties produce tens of millions of tons of coal per year.

In the aforementioned studies, unobserved differences between counties that correlate with mining activity would confound conclusions. For example, if coal counties have more hospitals or their populations are more likely to be covered by insurance, then a researcher may mistakenly attribute inflated hospital visits to coal mining, when in reality the cost (either implicit or explicit) of visiting a hospital is lower in those counties. Further, these studies use as a treatment variable the tons of coal mined within a county's borders, by either method.[39] Given the transboundary nature of air pollution, such an approach implicitly ignores potential spillover effects. This is another dimension upon which I improve.

2.7 Data

The Department of Labor Mine Safety and Health Administration maintains data on short tons mined per quarter, extraction techniques used, and geographic coordinates for every coal mine in the United States. I classify surface production as coal mined via strip, quarry, open pit, and auger mining. In the MSHA data, underground mining is classified as underground. This high level of detail allows me to paint a geographically precise picture of Appalachian coal extraction for the six years 2005-2010. Tables 2.4,

[39] The exception here is Brink et al. (2014), who link respiratory problems to surface mining.

2.5, and 2.6 broadly summarize the Appalachian mining industry.[40] Unlike the Western states, where surface techniques dominate production, underground mines produce the majority of Appalachian. Over the six-year period, Appalachian states produced 1.8 billion tons of coal. Underground mining accounted to 63% of this total. On average, underground operations produce more coal, employ more workers, and use more labor hours than surface mines. Table 2.7 describes quarterly tons produced by method within exposure buffers of four different shapes: three circular buffers of radius 25, 35, and 75 kilometers, and the irregularly shaped county buffer. Table 2.7 indicates that underground mining accounts for the majority of mining activity taking place in a county's exposure buffer, no matter the shape of the buffer. Figure 2.16 plots total underground and total surface mining tonnage during the study period. Surface methods account for between 33 and 47 percent of quarterly production. The correlation between the two methods is 0.6. Among the 55 counties in the state, 10 experienced zero nearby mining over the study period.[41] Low-exposure counties experienced an average of 189 thousands tons of surface mining per quarter, and high-exposure counties experienced nearly 2 million tons mined per quarter.

My primary dependent variable is the aggregate number of hospitalizations for asthma among each county's population in a give quarter. To compute it, I obtained hospitalization records from the West Virginia Health Care authority for three maladies: asthma, broken bones, and hernias (the latter two were obtained for falsification tests.) These include the first eighteen diagnosis codes issued upon admission, as well as gender, age group, county of residence, week of admittance, length of stay, and type of

[40] Statistics shown are for underground mines and surface mines in Ohio, Pennsylvania, Maryland, Virginia, West Virginia, and Kentucky. Table 2.4 describes production for observations with positive tonnage. Underground mines were idle (produced zero tons) 17% of quarters, compared to 24% of quarters for surface mines.

[41] I define exposure as the sum of coal mined within 35 kilometers of a county's population centroid over the six years of interest. High-exposure counties had greater than median production.

admission (emergency, elective, etc.) Due to WVHCA policy, my outcome is missing individuals who seek emergency care and are discharged directly from the emergency room. To the extent that surface mining also impacts asthmatics who can manage their care and do not require admission into a hospital, the magnitude of the health impact found here is understated. Table 2.2 presents summary statistics by age, gender, and types of admission for each ailment. Adults between the age of 18 and 65 make up the bulk of asthma and hernia admissions. The elderly are hospitalized most for fractures. Women are more likely than men to be hospitalized for each ailment. Among hospitalizations for asthma and fractures, the most common type of admissions are emergencies.

Table 2.1 reports summary statistics for the county groups described in Hendryx & Ahern (2009). Just as Hendryx & Ahern (2008), Hendryx, Ahern & Nurkiewicz (2007) and Hendryx & Ahern (2009) find, coal counties have more hospitalizations for respiratory problems: low-coal counties have 57% more asthma hospitalizations than non-coal counties, and high-coal counties experience over twice the rate of non-coal counties. However, the same pattern exists for broken bones and hernias. To the extent that the elevated rates of respiratory problems are solely due to the air pollution, higher fractures and hernias weakens the argument that mining causes poor respiratory health. Although it is possible that residents of high-coal mining counties experience more broken bones due to mining related accidents, it is improbable that mining accidents generate 88 percent more fractures per quarter than coal counties. Therefore, table 2.1 offers evidence of cross-sectional bias: coal counties have more hospitalizations for coal-sensitive *and* coal-insensitive ailments.

Table 2.1 also indicates that counties with more coal exposure are more populous, have lower median incomes, and are more insured. The proportion of populations with high school diplomas is moderately less in mining counties.

Demographic data came from the 2010 Census Summary File 1, West Virginia. Data on health insurance coverage is available at the annual level and comes from the U.S. Census Small Area Health Insurance Estimates. Quarterly estimates of income for other pollution intensive industries whose output may covary with coal mining, and the number of workers in the natural resource and mineral extraction industry industry come from the BLS Quarterly Census of Employment and Wages. The final dataset contains 1,320 observations - 55 West Virginian counties observed quarterly from 2005 through 2010.

2.8 Methodology

Cross-sectional research on this topic conclusively finds health disparities between Appalachian counties after controlling for observable differences. A panel data approach allows me to move beyond highlighting health disparities and demonstrate that populations respond to coal mining activity with increased demand for healthcare, after controlling for observed and unobserved differences between counties. Identification therefore comes from the variation in quarterly surface mining activity that populations are subjected to in that quarter. This leads to similar but refined evidence of a healthcare externality by showing that intertemporal variation in coal surface mining leads to increased demand for healthcare.

County-level unobserved effects address the unobserved heterogeneity between counties that contribute to consistently higher or lower rates of hospital treatments. For instance, asthmatics may cluster in coal counties because those counties have more hospitals or better ambulatory services. Populations in non-mining counties might be better educated on how to manage their asthma symptoms, making them less likely to seek hospital care in the event of an attack. Hospitals in coal counties might be more likely than others to discharge directly, rather than recommend admission to a hospital. For these reasons and more, populations vary in their proclivities to be hospitalized for asthma. To the extent that these differences remain unchanged over time, county fixed effects help to deliver unbiased evidence that surface mining drives poor health outcomes. I also use quarterly dummy variables to absorb the exogenous seasonality of asthma symptoms. Importantly, tables 2.10 and 2.11 demonstrate how conclusions change once either of these effects, as well as other covariates, are included in my analysis.

Since there is no clear definition of a population's exposure to coal mining, any analysis of aggregate health and mining requires assumptions about constitutes a population's exposure. Prior work used the tons of coal extracted over some period of time within a county's borders as a proxy. The lack of air quality monitors in the state prevents any study from precisely characterizing and quantifying the air pollution that each population is exposed to. Still, the spikes in air pollution resulting from blasting, scraping, bulldozing, and machinery exhaust at mines within a county contribute to the health of individuals distributed non-uniformly throughout that county. One interpretation of the county-border approach is that a population exists at a single point within a county (the population centroid), and is exposed to mining activity taking place within an area (the exposure buffer) shaped like their county. Since the mining data I use here includes the geographic coordinates of each mine, I can relax the second part by allowing for the possibility that mining activity outside a county's buffers contribute to a population's health. I test 100 different circular buffers and objectively select (for each age group) an optimal buffer size.[42] This is a more reasonable approach since air pollution - the channel through which coal mines impact respiratory health - is transboundary, and a population's health could be shaped by the air pollution from neighboring counties.

Count data applications rely on Poisson regression models. These use the Poisson distribution, which features an equal mean and variance, represented by the parameter λ_i. In these models, this conditional expectation is a non-linear function of the covariates:

$$\mathrm{E}\left(y_i \mid x_i\right)=\lambda_i=\exp \left(\mathbf{x}_i \beta\right) \tag{2.1}$$

In this study, the outcome of interest is an overdispersed: table 2.3 makes clear that quarterly hospitalizations in most counties violate the equal mean and variance assumption of the Poisson distribution. Therefore, I use negative binomial regressions, which relax the assumption of equidispersion: assuming a Poisson distribution in the presence of over-dispersion leads to undersized standard errors, and therefore incorrect inference. In a panel data setting such as mine, one can include individual effects to address unobserved heterogeneity. I include county-level effects, as well as quarter dummies to absorb the seasonality of asthma symptoms. I estimate:

$$\log y_{it}=\beta_0+\mathbf{x}_{it} \beta+\boldsymbol{surface}_{it} \gamma+c_i+q_t+\epsilon_{it}, \tag{2.2}$$

where $i=[1,55]$ denotes county, $t=[1,24]$ denotes time period. The left hand side variable y_{it} is the number of hospital admissions with a primary diagnosis of asthma experienced by residents of county i in quarter t. The variable $\boldsymbol{surface}_{it}$ is my improved treatment measure: short tons of coal from surface mines within an exposure buffer centered upon county i's population centroid in quarter t. When estimating model 2.2, I standardize surface mining for ease of interpretation. The coefficient of primary interest

[42] This process is detailed in the appendix.

is the semi-elasticity γ: the *ceteris paribus* percent increase in asthma hospitalizations per quarter resulting from an additional standard deviation surface coal mining. County and quarter fixed effects are c_i and q_t respectively. The set $\mathbf{x}_{it}$ contains county-level covariates which may explain some of the variation in asthma hospital visits, and ϵ_{it} is an idiosyncratic error. A coefficient on exposure that is greater than zero and statistically significant is evidence that county populations respond to surface mining activity by increasing their consumption of healthcare services. I cluster standard errors by county to address any remaining within-county correlation of residuals. In an alternative specification, I use ordinary least squares to estimate models whose dependent variables are quarterly hospitalizations per 100 thousand residents.

2.8.1 Threats to Identification

In this section, I lay out possible sources of bias in estimation of the parameter γ, and justify the covariates in equation 2.2. I consider the possibility that county-level effects are time-invariant, and discuss possible omitted variables that would confound the effect of mining on hospitalizations.

Coal mines, like other environmental disamenities, may be sited in low-value neighborhoods with high rates of poverty. Given the links between poverty and asthma (Katz, Kling & Liebman 2001), differences in physical well-being may reflect differences in economic well-being. Therefore, the explanation for why coal counties are sicker on average is a consequence of mines siting themselves in low-income counties where more people already suffer from asthma.[43] Similarly, residents of high-income counties could be better prepared to deal with an asthma attack, reducing the likelihood that any one resident requires a hospitalization for asthma. If differences in income remain unchanged over time, then the county effects in equation 2.2 would eliminate this prob-

[43] See some reference for a discussion of research on the links between income and asthma.

lem. If income levels change over time in a way that is correlated with surface mining activity, however, then failing to account could bias estimates of γ if incomes covary with surface mining.

To that end, I use estimates of quarterly county income from the BLS Quarterly Census of Employment and Wages. This data covers 98% of all U.S. businesses, and reports number of employees and and wages paid for thirteen different industries. If including quarterly income substantively changed estimates of γ or its standard error, then the positive correlation I find could merely result from a relationship between asthma and aggregate economic activity.[44] I also use this data to control for the possibility that the effect I find is driven by some other pollution intensive industry: the vector $\mathbf{x}$ in model 2.2 includes estimates of the quarterly income paid to the manufacturing, construction, and utilities/trade/transportation industries in each county.

Another possible source of cross-sectional bias is access to care: how easy is it for someone to receive medical treatment they need, and how does this vary between populations over time throughout the state? A person living in a county with few clinics or hospitals faces higher implicit costs of care than someone in a city where hospitals are numerous. Counties with more hospitals and more insured populations should then experience higher rates of asthma hospital visits, regardless of exposure to mining. Table 2.1 confirms this fact. Therefore I include covariates that proxy for a population's access to care: the proportion of its population that is covered by some form of health insurance (from the Census Bureau's Small Area Health Insurance Estimates), as well as the hospitals and total hospital beds in each county (from the West Virginia Office of Health Facility Licensure and Certification).

[44]There is limited evidence that respiratory morbidity or mortality is pro-cyclical. **?** finds no evidence that mortality from respiratory illnesses increases with economic activity, proxied for with state-level annual unemployment rates.

A key assumption for the validity of my analysis is the time-invariant unobserved heterogeneity between counties. I rely upon this assumption when using county fixed effects to absorb the differences between populations that lead to varying rates of asthma hospitalizations. If this heterogeneity changes in a way that is correlated with surface mining of coal, then my estimates would be biased. To illustrate this, consider two extreme examples. First, suppose that during the third quarter of every year, asthmatics in all counties changed their residency to surface mining counties.[45] Surface mining counties would experience temporary increases in mean sensitivity, thereby temporarily increasing the population's likelihood of hospitalization for asthma. In this case of short-term Tiebout sorting (Tiebout 1956), in which asthmatics briefly vote with their feet and receive more intense exposure, estimates of γ in equation 2.2 would be biased upwards. If the opposite were true, and asthmatics flee coal counties in the heavy summer months, then my estimates would be biased downward. My identification strategy relies on the implicit assumption that Tiebout sorting over characteristics correlated with surface mining occurs slowly relative to changes in surface mining exposure.

A more likely source of bias is within-quarter avoidance behavior. If asthmatics across the state systematically avoid exposure to air pollution by spending time indoors, then my estimates of the healthcare impact of surface mining would be biased downward (Neidell 2008). If asthmatics maintained their regular habits, exposing themselves to air pollution from surface mines, then the quarterly spikes in hospitalizations would be higher than reported here. To the extent that asthmatics across the state are more likely to demonstrate short-term avoidance behavior rather than overexpose themselves to air pollution during the summer months, my estimates therefore reflect a lower bound on the true sensitivity of hospitalizations with respect to mining activity.

[45]Treated counties are subjected to the highest amount of surface mining in the third quarter of every year, though the difference is small.

Table 2.16 extends the categorization of table 2.1 by separating coal-exposed counties according to the proportion of total exposure that comes from surface mining. Among coal-exposed counties, those counties where surface methods account for greater than 27 percent of total exposure have fewer hospital beds, and roughly the same number of asthma visits per thousand residents per quarter. The median number of hospitals per county for all groups is 1. This is evidence that if asthmatics sort according to their preferences for access to care, then they are not moving to surface mining counties. If access to care was noticeably larger among high-exposure, high-surface counties, then heightened asthma visits in those counties could be a consequence of overly sensitive individuals moving to those counties. However, it does not appear that surface mining counties offer compelling healthcare reasons to move to them.

Finally, excess asthma hospitalizations could be attributed to miners themselves. If miners are at a higher risk of asthma attacks since they are constantly in the presence of respirable pollutants, and the size of counties' mining workforces vary over time, then mining mining pollution may only affect a very small subset of the population. There are several reasons this is unlikely to bias my estimate of γ.. First, all regressions include the number of workers employed in natural resource extraction (from the BLS Quarterly Census of Employment and Wages.) Second, miners who seek care for respiratory illnesses are typically diagnosed with miner's pneumoniosis, ("miner's lung"), a distinct illness from asthma. Third, males between the ages of 18 and 65, the demographic group that comprises the overwhelming majority of coal miners, demonstrates no sensitivity to surface mining. Fourth, considering that miners working underground are exposed to far more concentrated levels of coal dust, underground mining activity would demonstrate a larger and more consistent relationship with asthma.

2.9 Results

2.9.1 Surface Mining Matters

Table 2.10 presents the main results. Table 2.11 presents similar estimates from linear models that take hospitalizations per 100 thousand residents as outcome variables. All estimates in these tables are computed holding other covariates fixed at their means, and have been exponentiated.

The first important fact to take away from table 2.10 is that coefficients on a surface mining within an area larger than a county are positive and significant at the 1% level. Column (1) is a pooled model of without county or quarter effects: without accounting for unobservable differences between counties or seasonality, a standard deviation of surface mining (1.8 million tons of coal) increases asthma hospitalizations by 18.7%. As expected, population growth increases quarterly rates of hospitalizations for asthma: an additional thousand residents increases hospitalizations by a little more than 3%. A more powerful effect comes from miners: an additional thousand workers in the natural resources extraction industry increase asthma hospitalizations by 10%. The estimate on quarterly income somewhat supports the hypothesis above, in that higher incomes reduce asthma hospitalizations. Accounting for seasonality by adding quarterly fixed effects has little impact on these estimates.

Next, including county fixed effects leads to an increase in the estimate of surface mining's impact: the estimate increases from 18.5% to 21.3%. This is evidence that failing to account for the unobservable differences across counties leads to understimating the health impact of coal. Finally, a larger source of bias is mis-specifying populations' exposure to surface mining. Columns (4) and (5) contain estimates from models that use within-border surface mining and within-border total mining, as in Brink et al.

(2014), Hendryx & Ahern (2008), Hendryx & Ahern (2009), and Hendryx, Ahern & Nurkiewicz (2007). Coefficients on these measures are less than half the magnitude of the estimates from models that account for spillover effects, and are statistically insignificant. This is evidence that restricting attention to within-borders mining activity, and failing to account for mining activity in neighboring counties severely underestimates the relationship between coal mining and health. Once county and quarter fixed effects are included in models, income, insurance coverage, and even population are insignificant determinants of hospitalizations for asthma. Curiously, an additional hospital reduces asthma hospitalizations by nearly 17% in my preferred specification. Once county fixed effects are included, population, income, hospital beds, and miners in the labor force lose significance.

Table 2.11 presents similar estimates from linear models with dependent variables of asthma hospitalizations per 100 residents. Controlling for unobserved heterogeneity across counties, and seasonality leads to larger estimates of surface mining's affect on asthma. Using larger exposure buffers also leads to larger estimates of coal's impact. According to column (3) of table 2.11, an additional standard deviation increase in surface tonnage leads to nearly 10 more hospitalizations for asthma per 100,000 residents per quarter. Neither within-borders surface mining nor total coal mined have any significant effect on this outcome.

2.9.2 A Healthcare Cost of Surface Mining

I use the coefficients in table 2.10 to estimate how many hospitalizations are due to surface mining by predicting each county-quarter's asthma hospitalizations in the absence of surface mining. This exercise indicates that of the 16,935 hospitalizations in West Virginia between 2005 and 2010, over 1,800 resulted from surface mining. To calculate the monetary cost of these visits, I use the mean cost of a West Virginia asthma hospitalization found in the Centers for Medicare and Medicaid Services (CMS) 2012 Provider Utilization and Payment Data Inpatient Public Use File. The mean total payments for an inpatient hospitalization for bronchitis or asthma are $5,000, which suggests that the cost of these extra hospitalizations is in the neighborhood of 9.2 million dollars. While this amount is relatively small (the cost per ton of surface coal is between 2 and 3 cents), my analysis is unable to capture the cost of emergency room visits that do not result in hospital admissions, extra spending on defensive medicine like bronchodilators, reduced productivity of asthmatic laborers, lost school days of asthmatic children, or other costly avoidance behaviors. Therefore, my results are less useful for estimating the total magnitude of the negative health externality of surface mining, and more as concrete evidence that variation in surface mining activity induces a healthcare response among Appalachian populations.

2.9.3 Heterogeneous Effects

Given the level of detail in my hospitalization records, I can disaggregate the total number of asthma hospitalizations by subpopulations. I compute male and female specific hospitalizations by broad age groups: children under the age of 18, adults between 18 and 65, and the elderly. Estimating equation 2.2 for those groups produces the gender and age specific responses to exposure to surface mining.

Tables 2.14 and 2.15 present coefficient estimates on surface mining exposure from models that use the same set of covariates as the models in table 2.10. Table 2.14 are from models in which the effects of other covariates are constrained to be equal across genders, and table 2.15 allows those estimate to vary across genders. In both approaches, individuals between 18 and 65 demonstrate small and insignificant sensitivities. Among the elderly, females demonstrate a higher sensitivity: a standard deviation of surface mining exposure causes a 32% spike in hospitalizations for asthma. The single most sensitive age group (in the unconstrained models) are young females: a standard deviation of surface mining exposure leads to a 42% spike in asthma hospitalizations.

2.9.4 Falsification Tests

One shortcoming of a cross-sectional approach is that, when applied to this data, certain coal insensitive ailments follow the same pattern as coal sensitive ones. Table 2.1 indicates that high-exposure counties experience higher rates of asthma *and* higher rates of broken bones and hernias. Since my framework, by construction, assumes that residents respond to the airborne pollution generated at mines, its application to fractures and hernias should yield effects not significantly different from zero. To test this, I estimate equation 2.2 using hernias and broken bones as dependent variables. Results are presented in tables **??** and **??**, respectively. When pooling observations, surface mining appears to have a positive impact. However, after including county effects, the relationship disappears. Once these models include county effects, surface mining activity has no impact on hospitalization rates for fractures or hernias. This supports the hypothesis that the air pollution from surface mines is driving a respiratory health externality, and that panel data approaches largely eliminate spurious correlations between mining and other ailments.

2.10 Conclusions

I present new evidence of a public health externality from surface mining of coal by exploiting variation in counties' exposure to the industry, and use CMS cost estimates to place a lower bound on the externalized healthcare cost of coal. With a simple fixed effects framework and a more flexible measure of exposure to mining, I take the first steps in identifying the causal effect of coal mining on respiratory functioning and show that populations respond to the the quantity of coal mined nearby by increasing consumption of healthcare services for asthma. I deal with previously unaccounted for observable and unobservable heterogeneity, such as populations' access to care, and therefore rule out some potential sources of confoundedness. I estimate that surface mining of coal required an additional 9 million dollars of spending on asthma hospitalizations.

These results refine the findings of Hendryx & Ahern (2008) and Hendryx & Ahern (2009) who document differences in health outcomes between Appalachian county groups. It offers stronger evidence of the air pollution channel they hypothesize: spikes in surface mining activity induce increases in demand for healthcare services. The fact that coefficients on underground mined coal are never significant is consistent with the findings of Ghose (2009), Ghose & Majee (2007), and Brink et al. (2014) who show that surface mining is more harmful to both ambient air quality and health than underground mining. I find that an additional standard deviation of coal from surface mines within 35 kilometers of a population centroid increases total asthma hospitalizations by 21% percent. I provide evidence that failing to account for mining activity beyond a county's borders severely underestimates the true impact of mining on heath.

I disaggregate the main response variable and show that youths and the elderly are most sensitive to nearby surface mining activity. Tests of internal validity revealed that my approach is free of the bias that earlier cross-sectional works suffer from. Regressions of hernia and broken bone visits on mining activity do not generate results that match those from the asthma regressions. As with most empirical studies, threats to identification remain. Undoubtedly, the ideal study which identifies causal effects would employ data on air quality to make firm the link between mining activity, pollution, and human health, but data limitations prevent such a study.

One feasible extension of this research would be to account for quarterly prevailing wind patterns each county experiences. However, I am limited to modeling populations as existing at a population centroid. Wind patterns can vary greatly over a county during a quarter, so it is unclear how much information would be gained by including prevailing wind patterns at a single point. Further, the exposure method I use could be extended to a number of settings where point sources of industrial pollution may affect the health of local populations. Although studying coal mines in developing countries would be difficult due to lack of sufficient data, a similar study could be performed for other coal-producing states such as Kentucky, Tennessee, or Wyoming.

As the largest domestic users of coal - thermal power plants - transition to burning natural gas, demand for Appalachian coal will continue to fall. My findings suggest that one bright side of this regional challenge, besides lower national carbon dioxide emissions, will be fewer childhood cases of asthma, and therefore fewer life-long asthmatics.

2.11 Figures and Tables

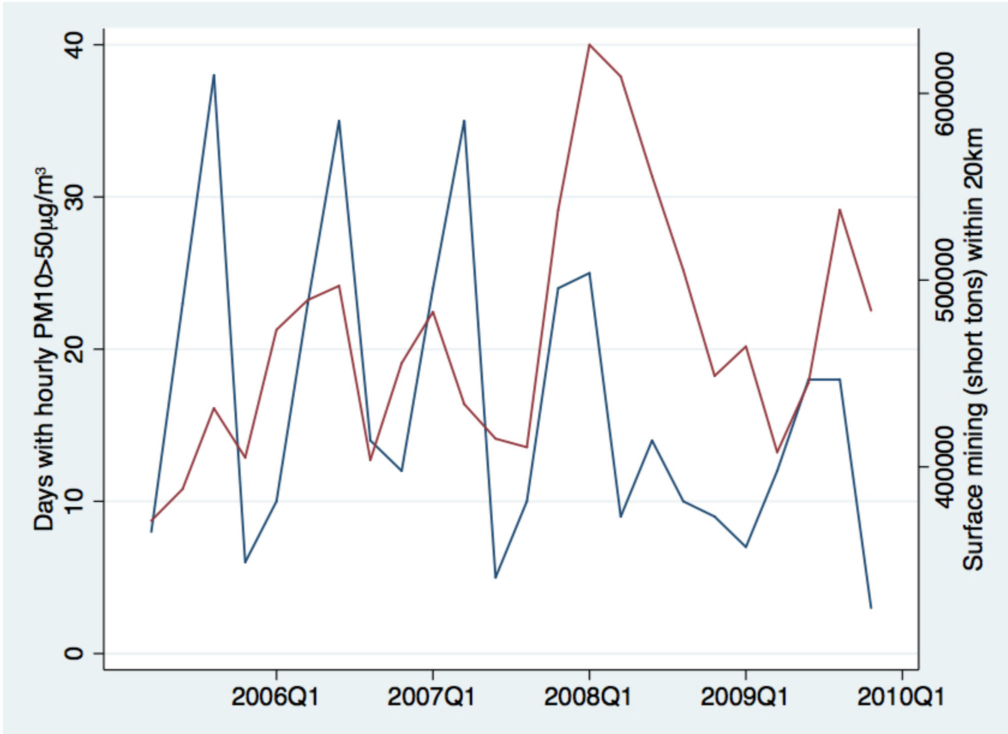

Figure 2.12: Kanawha County: quarterly spikes in PM10 (blue) and surface coal mining (red): $\rho = 0.14$

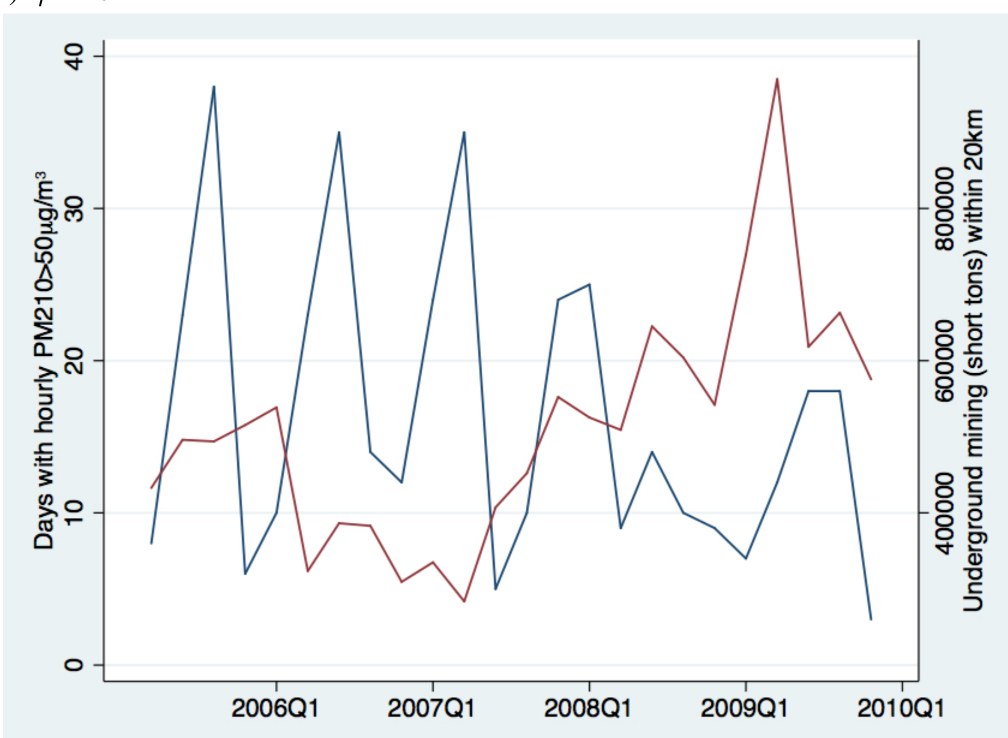

Figure 2.13: Kanawha County: quarterly spikes in PM10(blue) and underground coal mining (red): $\rho = -0.33$

Figure 2.14: Within-border definition of exposure: Logan County

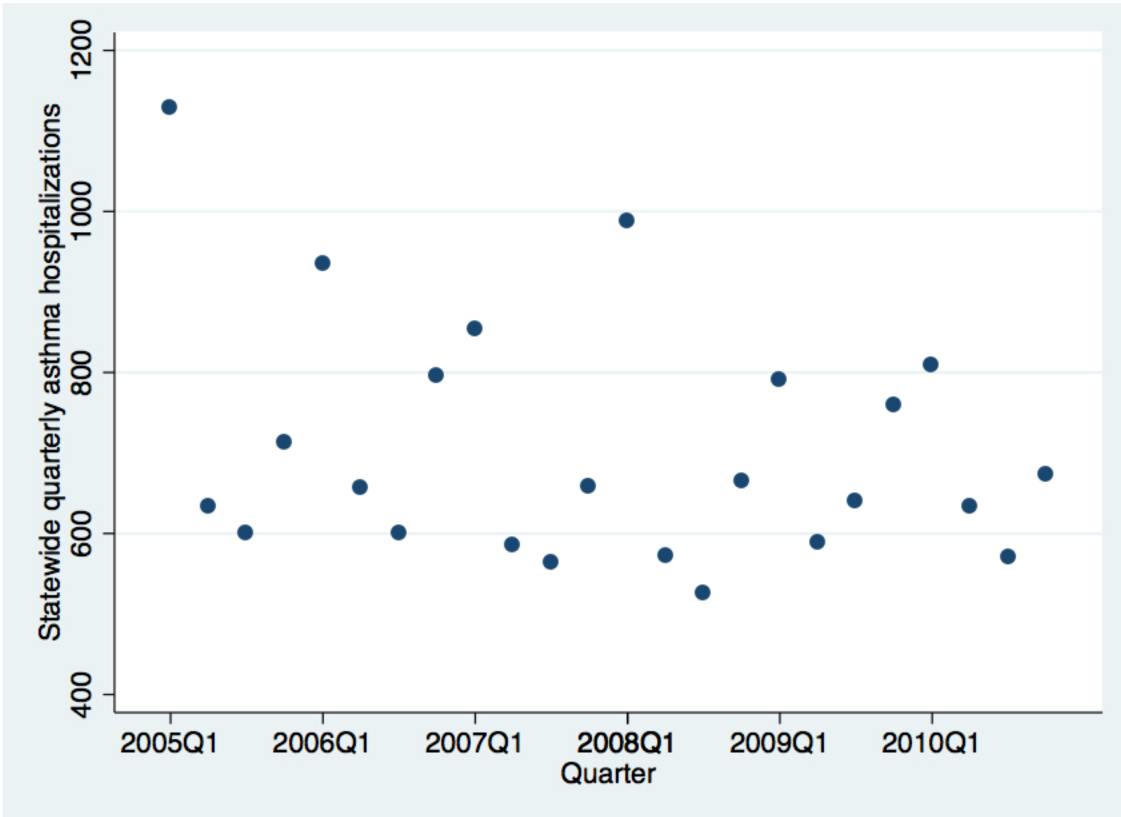

Figure 2.15: Statewide asthma hospital visits

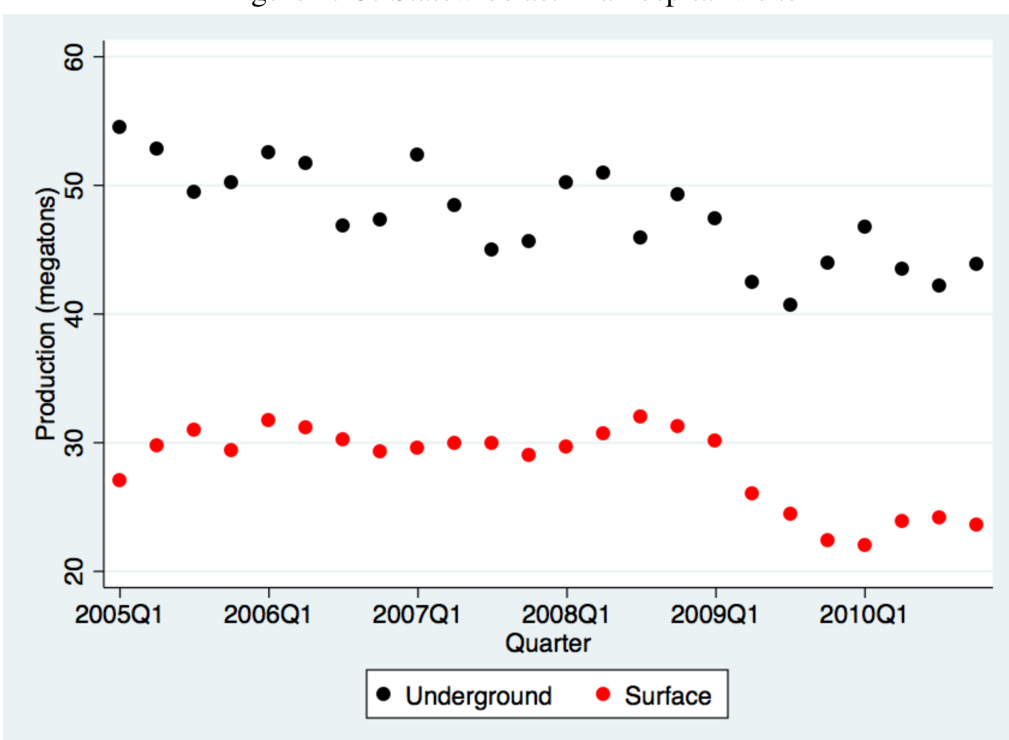

Figure 2.16: Quarterly tons mined by method: Appalachia 2005-2010

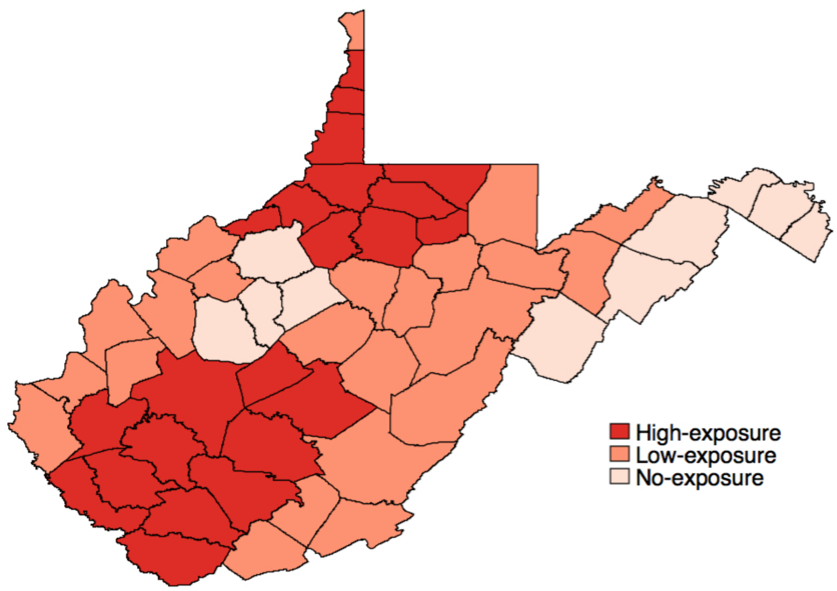

Figure 2.17: West Virginia counties and exposure to coal mining

Table 2.1: Health care and sociodemographic statistics by exposure category

	Exposure Category			
	No coal	Low-coal	High-coal	All
Asthma hospitalizations	21.4	33.7	44.0	35.6
(per 100,000 residents)	(20.5)	(22.6)	(27.4)	(25.6)
Fracture hospitalizations	58.5	97.7	115.7	97.8
	(41.7)	(37.7)	(36.7)	(43.1
Hernia hospitalizations	10.6	14.6	17.4	15.0
	(12.7)	(11.4)	(9.3)	(11.1)
Number of hospitals	0.6	0.9	1.2	1.0
	(.5)	(.7)	(1.4)	(.9)
Population (1000s)	24.4	30.2	40.5	33.5
	(27.7)	(23.6)	(40.2)	(32.4)
% White	93. 6	96.9	95.9	95.9
	(6.4)	(2.1)	(3.4)	(3.9)
% Black	3.3	1.5	2.4	2.2
	(4.5)	(1.7)	(2.7)	(2.9)
% Hispanic	1.9	0.7	0.6	0.9
	(1.7)	(0.4)	(0.5)	(1.0)
Median income (1000s)	38.2	34.4	34.1	35.0
	(9.7)	(5.2)	(4.9)	(6.4)
% of population with some	79.6	81.3	82.3	81.4
form of health insurance	(3.3)	(2.8)	(2.6)	(2.9)
% H.S. diploma or higher	45.0	44.5	43.5	44.1
	(6.2)	(4.7)	(4.8)	(5.1)
% Bachelor's or higher	8.3	9.1	8.4	8.7
	(3.5)	(2.3)	(3.8)	(3.2)
	$n = 10$	$n = 23$	$n = 22$	$n = 55$

Counties separated based on total coal mined within 35 kilometers of their population centroids over the six-year period of interest. Standard deviations in parentheses.

Table 2.2: Summary Statistics: Hospitalizations

		Asthma (ICD-9 493.0)	Hernias (ICD-9 550.0-553.0)	Fractures (ICD-9 800.0-829.0)
Age group	0-17	3,822	169	2,467
	18-65	8,371	3,963	17,346
	65+	4,772	2,894	25,799
Gender	Men	5,554	2,806	18,195
	Women	11,411	4,220	27,417
Type of Admission	Emergency	9,957	2,352	30,051
	Urgent	3,913	1,156	5,036
	Elective	3,000	3,473	7,728
	Trauma	210	115	3,431
	Newborn	0	1	3
	Not classified	544	253	308
	Total	16,965	7,026	45,612

Table 2.3: Quarterly Asthma Hospitalizations by County: 2005-2010

County	Mean	Var.	Min.	Max
Barbour	5.7	5.7	3	11
Berkeley	17.3	33.3	8	31
Boone	7.8	13.5	2	15
Braxton	3.7	7.1	0	10
Brooke	7.4	15.5	2	19
Cabell	29.5	95.0	12	47
Calhoun	2.1	3.1	0	6
Clay	2.5	4.4	0	10
Doddridge	2.9	4.6	0	8
Fayette	20.4	51.7	8	37
Gilmer	3.4	5.8	0	8
Grant	2.6	1.6	0	5
Greenbrier	10.5	23.4	1	22
Hampshire	2.3	2.8	0	6
Hancock	15.5	31.2	4	27
Hardy	1.7	2.8	0	6
Harrison	58.5	199.4	42	95
Jackson	15.5	31.1	6	26
Jefferson	5.7	5.9	1	10
Kanawha	66.9	217.7	49	102
Lewis	8.9	17.9	3	20
Lincoln	5.3	3.6	2	10
Logan	20.6	43.2	8	36
McDowell	10.0	22.2	4	21
Marion	25.8	78.4	11	43
Marshall	11.0	16.6	5	21
Mason	9.8	13.2	5	17
Mercer	46.8	117.7	30	67
Mineral	3.7	6.5	0	10
Mingo	25.1	117.7	10	49
Monongalia	24.3	57.4	11	42
Monroe	3.9	8.4	0	13
Morgan	2.9	8.2	0	11
Nicholas	7.4	16.6	3	20
Ohio	16.1	40.8	5	30
Pendleton	0.8	0.9	0	3
Pleasants	3.6	3.8	0	8
Pocahontas	2.3	7.0	0	12
Preston	11.6	14.4	4	20
Putnam	11.6	15.0	4	20
Raleigh	55.4	219.7	34	87
Randolph	14.0	73.5	5	34
Ritchie	4.9	5.2	1	9
Roane	3.4	3.7	1	8
Summers	3.4	4.3	0	8
Taylor	7.1	17.3	1	16
Tucker	1.5	2.9	0	6
Tyler	2.2	1.9	0	5
Upshur	5.3	12.0	0	13
Wayne	12.1	16.3	4	20
Webster	1.6	1.5	0	4
Wetzel	10.0	31.7	4	26
Wirt	2.0	2.1	0	5
Wood	36.3	68.0	21	56
Wyoming	13.5	28.7	6	25
Total	12.8	250.2	0	102

Table 2.4: Tons extracted per quarter at active Appalachian Mines, by type of mine

Type	**Mean**	**Median**	**Std. Dev.**	**Max.**
Underground	145,477	52,470	343,081	3,089,612
Surface	87,706	33,288	152,836	1,352,654

Table 2.5: Hours worked per quarter, by type of mine

Type	**Mean**	**Median**	**Std. Dev.**	**Max.**
Underground	38,999	20,094	61,668	534,159
Surface	20,207	9,654	29,753	270,725

Table 2.6: Miners employed per quarter, by type of mine

Type	**Mean**	**Median**	**Std. Dev.**	**Max.**
Underground	68	35	103	894
Surface	33	17	46	361

Table 2.7: Mean tonnage by method within three distance buffers

Method	**25km**	**35km**	**50km**	**Within borders**
Surface	412,650	850,028	1,602,863	283,242
	(992,873)	(1,802,121)	(2,811,137)	(711,157)
Underground	758,907	1,645,279	3,403,032	412,240
	(1,200,453)	(2,212,485)	(3,919,896)	(978,646)
Total	1,171,557	2,495,307	5,005,895	695,483
	(1,887,164)	(3,387,091)	(5,718,368)	(1,342,854)

Standard deviations in parentheses.

Table 2.8: Negative binomial models: response of asthma hospitalizations to 100k tons of coal from mines within x kilometers of county population centoids

x (kilometers)	Surface tonnage	Underground tonnage	Total tonnage
90	-0.0019	-0.0003	-0.001
	(0.0012)	(0.0012)	(0.0008)
80	0.0002	-0.0010	-0.0004
	(0.0013)	(0.0013)	(0.0008)
70	0.0006	-0.0010	-0.0002
	(0.0015)	(0.0014)	(0.0009)
60	0.0029*	-0.0020	0.0004
	(0.0017)	(0.0016)	(0.0011)
50	0.0042**	-0.0021	0.0010
	(0.0021)	(0.0021)	(0.0013)
40	0.0076***	-0.0041	0.0013
	(0.0028)	(0.0026)	(0.0017)
30	0.0114***	-0.0033	0.0035
	(0.0036)	(0.0033)	(0.0022)
20	-0.0004	0.0021	0.0013
	(0.0084)	(0.0052)	(0.0040)

Standard errors (in parentheses) clustered at county level.

Table 2.9: Linear models: response of asthma hospitalizations to 100k tons of coal from mines within x kilometers of county population centoids

x (kilometers)	Surface tonnage	Underground tonnage	Total tonnage
90	-0.090	0.096	0.006
	(0.061)	(0.059)	(0.037)
80	-0.017	0.070	0.029
	(0.069)	(0.064)	(0.041)
70	0.020	0.073	0.049
	(0.079)	(0.071)	(0.047)
60	0.097	0.056	0.075
	(0.088)	(0.084)	(0.055)
50	0.141	0.116	0.128*
	(0.111)	(0.109)	(0.068)
40	0.300**	0.003	0.140
	(0.148)	(0.134)	(0.089)
30	0.479**	0.112	0.287**
	(0.191)	(0.181)	(0.121)
20	0.013	0.210	0.141
	(0.404)	(0.283)	(0.214)

Standard errors (in parentheses) clustered at county level.

Table 2.10: Negative binomial models: asthma hospitalizations per quarter

	Dependent variable:				
	Asthma hospitalizations per quarter				
	(1)	(2)	(3)	(4)	(5)
Std. dev. surface tonnage (36km)	0.187***	0.185***	0.213***		
	(0.022)	(0.020)	(0.060)		
Std. dev. surface tonnage (borders)				0.083	
				(0.055)	
Std. Dev. total tonnage (borders)					0.037
					(0.063)
Population (thousands)	0.032***	0.032***	0.001	0.001	0.002
	(0.002)	(0.002)	(0.011)	(0.011)	(0.011)
Miners in labor force (thousands)	0.101***	0.098***	−0.020	−0.019	−0.015
	(0.016)	(0.015)	(0.020)	(0.021)	(0.021)
Total Quarterly Income (millions)	−0.008***	−0.007***	0.001	0.001*	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Construction income (millions)	−0.008*	−0.006	0.0001	−0.001	−0.001
	(0.005)	(0.005)	(0.003)	(0.003)	(0.003)
Manufacturing income (millions)	0.023***	0.021***	0.001	−0.001	−0.0003
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)
Utilities income (millions)	0.032***	0.031***	−0.001	−0.002	−0.001
	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Hospitals	−0.182***	−0.201***	−0.166**	−0.134**	−0.129*
	(0.038)	(0.036)	(0.066)	(0.069)	(0.069)
Total beds	0.003***	0.003***	0.001	0.0002	0.0002
	(0.0002)	(0.0002)	(0.001)	(0.001)	(0.001)
Proportion with insurance	0.027***	0.036***	−0.003	−0.0003	0.001
	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)
Constant	−1.114*	−1.265**	2.588***	2.397***	2.249***
	(0.601)	(0.595)	(0.569)	(0.572)	(0.562)
County effects	No	No	Yes	Yes	Yes
Quarter effects	No	Yes	Yes	Yes	Yes
Observations	1,320	1,320	1,320	1,320	1,320
Log Likelihood	−4,045.398	−3,976.768	−3,319.204	−3,324.605	−3,325.802
Akaike Inf. Crit.	8,112.796	8,021.536	6,814.408	6,825.210	6,827.604
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 2.11: Linear models: asthma hospitalizations per capita

	Dependent variable:				
	Asthma hospitalizations per 100k residents per quarter				
	(1)	(2)	(3)	(4)	(5)
Std. dev. surface tonnage (34km)	4.19***	3.97***	9.85***		
	(1.00)	(0.90)	(3.14)		
Std. dev. surface tonnage (borders)				3.93	
				(2.91)	
Std. Dev. total tonnage (borders)					2.87
					(3.43)
Miners in labor force (thousands)	3.34***	3.07***	−1.45	−1.36	−1.29
	(0.59)	(0.58)	(0.99)	(1.00)	(1.03)
Total Quarterly Income (millions)	−0.21***	−0.18***	0.07**	0.07**	0.07**
	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Construction income (millions)	−0.26**	−0.12	−0.09	−0.10	−0.09
	(0.13)	(0.13)	(0.12)	(0.13)	(0.13)
Manufacturing income (millions)	0.46***	0.33***	−0.13	−0.18	−0.17
	(0.09)	(0.09)	(0.16)	(0.16)	(0.16)
Utilities income (millions)	1.31***	1.14***	−0.48*	−0.50*	−0.44*
	(0.14)	(0.14)	(0.26)	(0.27)	(0.26)
Hospitals	−5.69***	−6.30***	−6.48**	−4.06	−3.53
	(1.26)	(1.19)	(3.03)	(2.99)	(3.00)
Total beds	0.08***	0.07***	0.02	−0.0001	−0.003
	(0.01)	(0.01)	(0.03)	(0.04)	(0.04)
Proportion with insurance	−0.02	0.40	−0.19	−0.05	−0.001
	(0.28)	(0.29)	(0.33)	(0.32)	(0.32)
Constant	31.23	26.08	85.10***	73.71***	69.08***
	(22.17)	(22.62)	(26.40)	(26.09)	(25.84)
County effects	No	No	Yes	Yes	Yes
Quarter effects	No	Yes	Yes	Yes	Yes
Observations	1,320	1,320	1,320	1,320	1,320
R^2	0.18	0.29	0.62	0.62	0.62
Adjusted R^2	0.17	0.27	0.60	0.59	0.59
Note:			*p<0.1; **p<0.05; ***p<0.01		

Table 2.12: Negative binomial models: hernia hospitalizations per quarter

	Dependent variable:				
	Hernia hospitalizations per quarter				
	(1)	(2)	(3)	(4)	(5)
Std. dev. surface tonnage (36km)	0.13***	0.13***	−0.02		
	(0.03)	(0.03)	(0.09)		
Std. dev. surface tonnage (borders)				−0.07	
				(0.06)	
Std. Dev. total tonnage (borders)					−0.12*
					(0.07)
Population (thousands)	0.02***	0.02***	−0.002	−0.0005	0.0002
	(0.002)	(0.002)	(0.01)	(0.01)	(0.01)
Miners in labor force (thousands)	0.05***	0.05***	−0.02	−0.02	−0.02
	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Total Quarterly Income (millions)	−0.01***	−0.01***	0.0001	−0.0001	−0.0001
	(0.0005)	(0.0005)	(0.001)	(0.001)	(0.001)
Construction income (millions)	0.01***	0.01***	−0.01	−0.005	−0.01
	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)
Manufacturing income (millions)	0.01***	0.01***	0.001	0.002	0.002
	(0.002)	(0.002)	(0.004)	(0.004)	(0.004)
Utilities income (millions)	0.02***	0.02***	0.01	0.01*	0.01
	(0.003)	(0.003)	(0.01)	(0.01)	(0.01)
Hospitals	−0.16***	−0.17***	0.06	0.05	0.02
	(0.04)	(0.04)	(0.09)	(0.09)	(0.09)
Total beds	0.002***	0.002***	−0.001	−0.001	−0.001
	(0.0002)	(0.0002)	(0.001)	(0.001)	(0.001)
Proportion with insurance	0.03***	0.04***	−0.002	−0.001	−0.003
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Constant	−2.13***	−2.37***	1.41**	1.33*	1.50**
	(0.63)	(0.70)	(0.69)	(0.69)	(0.67)
County effects	No	No	Yes	Yes	Yes
Quarter effects	No	Yes	Yes	Yes	Yes
Observations	1,320	1,320	1,320	1,320	1,320
Log Likelihood	−3,023.64	−3,015.65	−2,529.75	−2,529.20	−2,528.24
Akaike Inf. Crit.	6,069.28	6,099.30	5,235.51	5,234.40	5,232.48
Note:			*p<0.1; **p<0.05; ***p<0.01		

Table 2.13: Negative binomial models: fracture hospitalizations per quarter

	Dependent variable:				
	Fracture hospitalizations per quarter				
	(1)	(2)	(3)	(4)	(5)
Std. dev. surface tonnage (36km)	0.18***	0.17***	−0.02		
	(0.02)	(0.02)	(0.04)		
Std. dev. surface tonnage (borders)				−0.02	
				(0.03)	
Std. Dev. total tonnage (borders)					−0.03
					(0.03)
Population (thousands)	0.03***	0.03***	0.002	0.003	0.003
	(0.002)	(0.002)	(0.01)	(0.01)	(0.01)
Miners in labor force (thousands)	0.07***	0.06***	−0.02	−0.02	−0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Total Quarterly Income (millions)	−0.01***	−0.01***	0.0003	0.0002	0.0002
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Construction income (millions)	0.001	−0.0001	−0.001	−0.0004	−0.0005
	(0.004)	(0.005)	(0.002)	(0.002)	(0.002)
Manufacturing income (millions)	0.01***	0.01***	−0.003	−0.002	−0.002
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Utilities income (millions)	0.02***	0.02***	0.01	0.01*	0.01
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Hospitals	−0.12***	−0.12***	−0.09**	−0.09**	−0.10**
	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
Total beds	0.002***	0.002***	0.001**	0.001***	0.001***
	(0.0002)	(0.0002)	(0.001)	(0.001)	(0.001)
Proportion with insurance	0.03***	0.04***	0.001	0.001	0.0003
	(0.01)	(0.01)	(0.004)	(0.004)	(0.004)
Constant	−0.24	−0.86*	2.80***	2.81***	2.84***
	(0.44)	(0.46)	(0.33)	(0.33)	(0.32)
County effects	No	No	Yes	Yes	Yes
Quarter effects	No	Yes	Yes	Yes	Yes
Observations	1,320	1,320	1,320	1,320	1,320
Log Likelihood	−5,136.83	−5,121.27	−4,056.40	−4,056.45	−4,056.30
Akaike Inf. Crit.	10,295.67	10,310.54	8,288.80	8,288.91	8,288.61

Note: *p<0.1; **p<0.05; ***p<0.01

Table 2.14: Semi-elasticities of hospitalizations with respect to SD of surface tonnage: age group and genders, other coefficients constrained across groups

	All ages	Youth (0-18years)	Adults (18-65)	Elderly(65+)
Male	0.23***	0.28***	0.02	0.22*
	(0.06)	(0.06)	(0.08)	(0.12)
Female	0.20***	0.29***	0.04	0.32***
	(0.06)	(0.06)	(0.06)	(0.10)

Table 2.15: Semi-elasticities of hospitalizations with respect to SD of surface tonnage: age group and genders, other coefficients unconstrained across groups

	All ages	Youth (0-18yrs)	Adults (18-64)	Elderly (65+)
Both	0.21***	0.28***	0.08	0.33***
	(0.06)	(0.07)	(0.07)	(0.10)
Male	0.18**	0.23**	0.11	0.23
	(0.08)	(0.09)	(0.12)	(0.15)
Female	0.23***	0.39***	0.06	0.34***
	(0.06)	(0.09)	(0.06)	(0.07)

Table 2.16: Access to care, by exposure and method type

Exposure Category	**Surface Proportion**	**Hospitals**	**Beds per 1000**	**% Insured**	**Asthma per 1000**
None	-	1	1.8	79.6	1.1
		(0.5)	(1.9)	(3.3)	(0.7)
Low ($n = 23$)	$\leq 27\%$	1	3.2	82.6	1.7
		(0.9)	(2.9)	(2.5)	(0.5)
	$> 27\%$	1	2.6	80.1	1.8
		(0.6)	(2.4)	(2.4)	(0.9)
High ($n = 22$)	$\leq 27\%$	1	5.3	81.9	2.3
		(0.6)	(4.5)	(2.9)	(0.7)
	$> 27\%$	1	2.8	82.6	2.3
		(1.9)	(2.6)	(2.1)	(0.7)

Chapter 3

Mines in the Housing Market

3.1 Mining externalities

Coal mines generate numerous environmental externalities, and have well documented impacts on other natural resources. Randall et al. (1978) categorize five broad categories of environmental damage from surface mines: water pollution, degradation of life-support systems for wildlife, increased frequency of flooding, damage to land and structures, and aesthetic damages. Underground mines and preparation plants have similar impacts: acid mine runoff occurs when water seeps into underground mines, collects toxins, and enters the water table. Preparation plants, which require large volumes of water to clean coal, similarly require special permits allowing them to discharge harmful waste water.

Air pollution, acidic runoff, damaged ecology, badly worn road surfaces, health risks, altered topography, and risk of soil subsidence are some. As noted in Randall et al. (1978), these externalized costs are borne largely by nearby landowners, residents in the region, and downstream water users.

3.1.1 Particulate Matter

Any part of the extraction process that require moving soil or exposing erodible surfaces generates fugitive particulate matter (Environmental Protection Agency 1995).

First, large scrapers remove topsoil and subsoil at the surface. Bulldozers then move this dirt in temporary stockpiles, or use it to cover previously mined areas - a process known as reclamation. Next, the overburden - material between topsoil and coal seam - is drilled and blasted to fully uncover the mineral below. Overburden is similarly bulldozed into stockpile or used in reclamation. Finally, the coal seam itself is drilled and blasted. Machines load the broken coal trunks into haul trucks, which take the mineral out of the mine to processing facilities. During surface mine reclamation, bulldozers smooth and contour piles of spoil. Dump trucks account for 30% of the energy used at surface mines (Siami-Irdemoosa & Dindarloo 2015).

At coal preparation plants, haul trucks unload their coal into crushers and screeners. The coal is screened into coarse and fine varieties, which each undergo their own cleaning process.[46] After cleaning, the coal is either placed in stockpile or loaded into railcars for transportation to power plants or seaports.

In general, surface mining is more harmful to air quality than underground mining . This is true because of the continual aboveground activity at surface mines that generates large quantities of fugitive dust (Ghose & Majee 2007, Papagiannis et al. 2014).

3.1.2 Water Pollution

Coal mines also produce effluent. Surface mines in Appalachian states must be permitted according to two separate sections of the Clean Water Act: one for their discharge of fill into streams, and another for discharging wastewater. As evidence of their im-

[46] The differences between cleaning fine and coarse coal lay in the severity of the mechanical processes.

pacts on local water quality, the EPA's 2011 decision to veto the Spruce Mine (a surface mine in southern West Virginia) centered on its anticipated violations of the Clean Water Act. The EPA also regulates effluent from coal preparation plants, which require vast amounts of water to clean coal and equipment; active underground coal mines, which pump wastewater deep underground; and abandoned mines, which produce acid drainage when rainwater enters mines and mixes with waste products (Code of Federal Regulations 2014).

Further, the water quality impacts of surface mining persist even after a mine's useful life ceases. McAuley & Kozar (2006) find that in the high-sulfur coal regions studied in this paper, ground water in mined area has higher median concentrations of mine-drainage constituents than non-mined areas, and that this water also had higher levels of specific conductance. They cite acid mine drainage as the source of these differences. These discrepancies fade when comparing locations beyond three hundred meters from reclaimed mines.

Figure 3.1: Northern Appalachia Coal Industry: Tons Produced

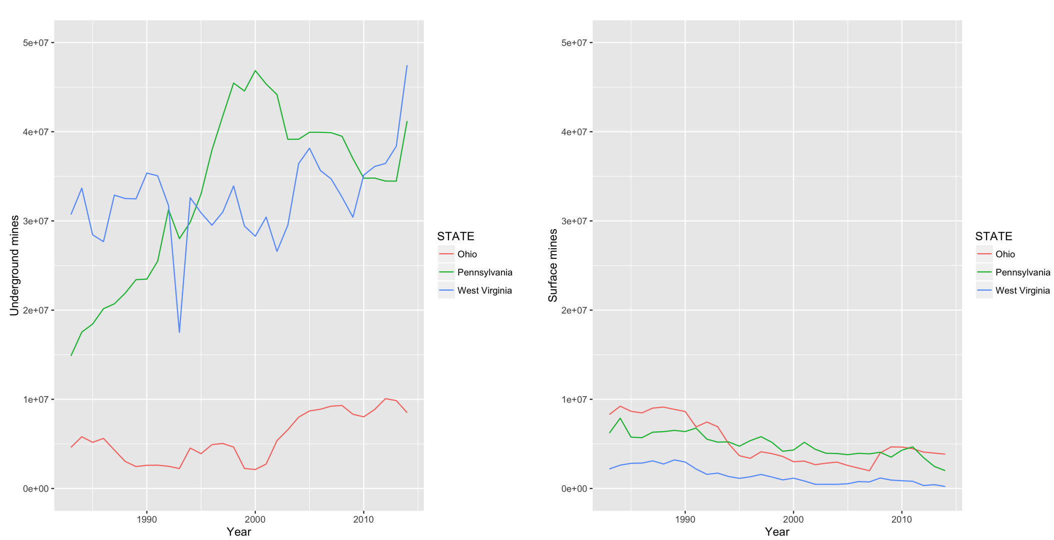

3.1.3 Ecological Impacts

Mining firms have converted 1.1 million hectares of forests into mines and buried more than 2,000 kilometers of streams in Appalachia (Bernhardt & Palmer 2011). These reduce the supply of ecosystem services in these areas, leading to lower use and non-use value in impacted regions (Mazzotta et al. 2015)

3.1.4 Soil Subsidence

Any geology-altering activity will place nearby buildings at risk of soil subsidence. This occurs in mining areas when underground caverns collapse, leading to a sudden drop at every soil level between cavern and surface. For a structure located near underground mining activity, this poses a serious risk. At least one state publishes information regarding soil subsidence for homeowners near mines.[47]

3.1.5 Positive Impacts

Local mining operations could boost house prices through three channels. First, miners who work at mining operations might prefer to live near the coal operation where they work. If miner households display sufficiently high demand, then prices near mines would reflect a higher price from this demand driven price differential. ed.

Second, homeowners might witness coal operations and then expect firms to pay them royalties for the rights to mine on their property. One possible way to isolate this channel would use information on whether or not a parcel sits atop a coal bed. According to recent data from the U.S. Geological Survey, virtually every home in our dataset sits atop a coal bed. A similar possibility is that coal firms are willing to purchase entire parcels in order to gain access to a particular seam. To the extent that this poses

[47] What You Need to Know About: Living Near Indiana Coal Mines. www.in.gov/dnr/reclamation/files/whatyouneedtoknow.pdf

a threat from clearly identifying the environmental externality driven impacts of coal mines, we focus on these counties because they contain a relatively low proportion of land owned by private firms.[48]

3.2 Literature

There is a large body of research that uses the hedonic approach to value environmental disamenities. These studies vary by time frame, geographic scope, type of disamenity, and identification strategy. The key concern in this literature that seeks to understand the market impacts of disamenities is the potential for omitted variable bias (Parmeter & Pope 2011). For example, if neighborhoods where power plants locate also feature high crime rates, then simply regressing house prices on proximity to a power plant would generate coefficient estimates that conflate the negative impacts of the power plant with those of crime. The gold standard approach for overcoming this problem is with quasi experiments, which exploit seemingly exogenous variation in the location or intensity of the amenity of interest. The second best approach explicitly controls for spatially correlated unobservables. These geographic fixed effects can mollify concerns that spatially correlated factors are obscuring the true price impacts of disamenities. Below, I review some well known studies that use as a measure of exposure some function of the distance between a parcel or tract centroid and the nearest disamenity.

Recently, much attention has been paid to the housing market impacts of shale gas development. Muehlenbachs, Spiller & Timmins (2015) find that groundwater-dependent homes located near fracking wells experience declines in sale prices of nearly 25%, reflecting perceived risk of contaminated water. They select a 2 kilometer buffer by comparing the difference in house prices for homes located near gas wells that only

[48] Energy firms own over 50% of land in some West Virginia counties.

vary in the timing of sale relative to drilling. They find that at 2 kilometers, the difference in before and after drilling prices becomes insignificant. Homes in public water service areas appear to benefit from fracking, but this positive impact (which they hypothesize comes from anticipated mineral leases) go away when moving too close to a well. Gopolakrishnan & Klaiber (2013) study the same topic using a smaller dataset, and find that for homes already located near wells, an additional well reduces prices by over $3,000. That study finds significant impacts for buffers of size 0.75 mile, 1 mile, and 2 miles. They similarly find evidence of perceived risk of groundwater contamination.

A topic that receives particular attention is Superfund sites, and whether spending on their cleanups is justified by expected or realized gains in values. Greenstone & Gallagher (2008) compare property values around the first 400 Superfund sites designated for cleanups to those that narrowly missed designation. They find small and insignificant benefits accrue to homes in Census tracts within 3 miles of a cleanup. Gamper-Rabindran & Timmins (2011) broaden the scope of Greenstone & Gallagher (2008). They use the entire distribution of house values in census tracts within 3 miles of each site, and find that houses at the low end of the distribution increase in value more than those at the top, following nearby Superfund sites' deletion from the National Priority List. Kohlhase (1991) use house prices from Houston near toxic waste sites before and after they are designated a Superfund site, and find evidence that the distance from a property to the nearest toxic waste site is a significant determinant of prices following that site's official EPA designation.

That study pools all superfund sites - which vary by size and pollutant - in their analysis and find that, post-EPA designation, increasing an property's distance to a superfund site by a mile boosts prices by 5.5%.

Kiel & Williams (2007) take a finer approach to valuing Superfund sites. They consider different types of sites and surrounding demographics, and find that larger sites that are located in areas with fewer blue-collar workers have larger negative impacts. They also estimate models for different time periods (related to the cleanup process) and find that during the cleanup process, distance increases sale prices, suggesting that mere information about being near a site harms prices.

Davis (2011) focuses on the housing market impacts of new electric power plants. They use restricted access Census data to estimate modest decreases in median housing values and rents within 2 miles of a newly opened power plant. Economists have also looked at links between clean energy generation and house prices.Heintzelman & Tuttle (2012) use sale prices in three counties in upstate New York, and find negative impacts associated with homes near wind turbines. These negative effects are felt up to 5 miles away.

Some authors construct indices of exposure that are functions of the intensity of a disamenity and a home's distance to it. Palmquist, Roka & Vokina (1997) construct a measure based on manure produced within three buffers around a property. The coefficient on this measure is negative and significant, so more poop lowers house prices. Leggett & Bockstael (2000) analyze the market impacts of water pollution using the three point sources nearest to each property sold on Chesapeake Bay.

Other authors use the hedonic method to understand perceived risk of violent crime. Pope (2008) uses the move-in and move-out dates of sex offenders in Hillsborough County, Florida. He finds that prices within 0.1 miles of a sex offender sell at a roughly

2% discount after a sex offender moves in, and a slight price rebound following a sex offender moving out. Linden & Rockoff (2008) find large declines (12%) among homes directly adjacent to a sex offender's residence, and small declines (4%) up to a tenth of a mile away.

3.3 Data

3.3.1 Home Sales

Data on property sales come from Corelogic, a private firm that compiles real estate transaction data from county tax assessors and property appraisers. Data here come three northern Appalachian counties: Belmont (Ohio), Monongalia (West Virginia), and Fayette (Pennsylvania). I focus on these counties because they provide a large number of sales, and largely exposed to the coal industry, and are located geographically near each other. To ensure commodity consistency, I focus only on single-family homes. To minimize the risk of price impacts that are conflated with another recent energy boom, I restrict the analysis to sales that took place prior to 2007 - widespread shale exploration began in this region in 2008 (Gopolakrishnan & Klaiber 2013). To filter out non-arm's length sales and reduce the impact of outliers, I drop observations from the top and bottom percentile of the distribution of prices in each county. The final dataset used here contains 2,923 records from Belmont County, Ohio; 3,688 sales from Fayette County, Pennsylvania; and 1,477 sales from Monongalia County, West Virginia.

In addition to sale price, the samples used here contain a full set of physical attributes such as the number of bedrooms and bathrooms; the logged values of living space and plot size; whether the structure features a garage, fireplace, or swimming pool, or septic tank; and whether the property is served by public water or has its own well. I use the U.S. Census geocoding tool and each property's mailing address to assign Census tracts and geographic coordinates.

Tables 3.1, 3.2, and 3.3 the sales data used in this analysis of transacted single family homes in each county from 2000 to 2006. In addition to the physical characteristics of each county's sales, the tables include statistics that broadly describe exposure to the coal extraction industry: these tables contain statistics that describe nearby underground mines, surface mines, preparation plants, and combinations of those.

Fayette and Belmont counties are relatively similar. The median home in these counties sold for around 50 thousand dollars, had three bedrooms and one bathroom, and sat on plot of roughly 1.2 acres. Further, median home age is around 63 years old. Homes in Monongalia County, West Virginia are more expensive, larger, newer, and have more bathrooms. The median home sold for almost twice as much as those in either Fayette or Belmont counties.

These tables also contain demographic, economic, and housing statistics from the 2000 Census. The three counties are largely similar in terms of race and income: the population in each county is roughly 95% white, and median income is around $28,000. The poverty rates are similar (around 12%). In terms of total population, Fayette is nearly twice as populous as the other two (it also has nearly twice as many housing units as the other two.) Another point worth noting is that the percentage of 25-year olds with a high school diploma is much lower in Monongalia than the other two counties. This could reflect the lower quality of public education in West Virginia. Also, between 2 and 4% of the labor force in each county works in mining.

3.3.2 Mining Sites

Coal operations in northern Appalachia consist of some combinations of three subunits. Extraction occurs either at underground mines, if coal forms more than 100 feet belowground, and surface mines, for coal forming close to ground level. Preparation receive recently mined coal for cleaning, screening, and processing before it's shipped to domestic power plants for or shipped abroad. The Miners' Health and Safety Administration assigns each mine a unique identifier, and maintains statistics for each type of activity at every location. We classify each mine in the MSHA dataset based on what combination of subunits operate in a given quarter. The quarterly production and employment statistics in tables 10, 11, and 12 are based on the most frequently occurring combination at a given location.

These locations are: standalone underground mines, surface mines, and preparation plants; two-part combination underground mine and surface mines, and surface mines with preparation plants; and three-part combination underground mine, surface mine, and preparation plant.[49] Statistics in these tables are for all mining operations in the counties we study here, plus surrounding counties that are contiguous to study counties.[50]

The large three-way combination operations account for the overwhelming majority of coal production in each region. Mean quarterly tonnage produced exceeds one million tons for these sites. They also employ the most miners and office workers. The least common type of operation is a standalone underground: only four sites across the study study region are classified as such. This is not to say that underground mining is uncommon, but that mining firms tend to operate both underground and surface mines together instead of an underground mine alone.

The most common type of operation in each study region is the standalone surface mine. These produce between 10 and 50 thousand tons per quarter in each state, and employ a relatively low number of workers per site. Given the large number of these operations in each state, it is therefore unsurprising that in each dataset of sales, the closest type of coal site is a surface mining operation.

3.4 Analytic Framework

(Rosen 1974) was the first to formalize the functional relationship between the price of a differentiated product and its attributes. A hedonic equilibrium is a schedule of prices resulting from the interactions between buyers and sellers of the differentiated

[49] There were no combination underground mine/preparation plant locations operating in the region during the period of study

[50] This allows for the possibility that the price of a property in our dataset might be affected by mining activity in an adjacent county.

Table 3.1: Belmont: n=2,923

Statistic	Mean	St. Dev.	Min	Median	Max
Sale amount (2000 dollars)	64,254.59	50,776.23	4,762.19	51,331.84	351,037.80
Sq. feet	1,409.69	552.53	451	1,293.5	8,727
Lot size (acres)	1.99	5.37	1.00	1.20	156.09
Bedrooms	2.75	0.72	1	3	6
Bathrooms	1.54	0.73	1	1	7
Age	63.33	35.11	0	64	200
Fireplace	0.22	0.41	0	0	1
Swimming pool	0.02	0.14	0	0	1
Garage	0.49	0.50	0	0	1
Well water	0.00	0.00	0	0	0
Septic tank	0.00	0.00	0	0	0
Nearest underground (km)	16.46	3.78	2.16	17.62	28.91
Nearest surface (km)	8.11	5.30	0.03	7.37	21.79
Nearest prep plant (km)	8.21	4.50	0.15	7.46	21.20

Table 3.2: Fayette: n=3,688

Statistic	Mean	St. Dev.	Min	Median	Max
Sale amount (2000 dollars)	65,525	48,952	4,676	54,283	389,720
Sq. feet	1,470	587	240	1,380	6,416
Lot size (acres)	1.66	1.11	1.00	1.29	11.00
Bedrooms	2.77	0.77	1	3	6
Bathrooms	1.51	0.69	1	1	6
Age	61.74	31.70	0	63	247
Fireplace	0.23	0.42	0	0	1
Swimming pool	0.02	0.15	0	0	1
Garage	0.65	0.48	0	1	1
Well water	0.09	0.28	0	0	1
Septic tank	0.35	0.48	0	0	1
Nearest underground (km)	23.95	8.66	3.07	23.57	47.56
Nearest surface (km)	6.82	3.71	0.10	6.50	26.42
Nearest prep plant (km)	9.31	5.65	0.18	7.86	31.89
Census stats					
housing units	66,490				
Population	148,644				
mean commute	26.5				
% white	95.3				
% black	3.9				
% hispanic	0.4				
% labor force in mining	3.5				
Median hh income	27,451				
% below poverty	13.8				
% hs grad	47.9				

Table 3.3: Monongalia: n=1,477

Statistic	Mean	St. Dev.	Min	Median	Max
Sale amount (2000 dollars)	118,621	77,316	4,762	101,776	403,426
Sq. feet	1,813	814	364	1,620	6,794
Lot size (acres)	1.54	1.73	1.00	1.22	36.74
Bedrooms	2.99	0.78	1	3	6
Bathrooms	2.12	0.94	1	2	6
Age	37.09	30.82	0	29	178
Fireplace	0.39	0.49	0	0	1
Swimming pool	0.00	0.00	0	0	0
Garage	0.35	0.48	0	0	1
Well water	0.03	0.16	0	0	1
Septic tank	0.15	0.36	0	0	1
Nearest underground (km)	7.93	4.20	0.47	7.28	21.00
Nearest surface (km)	5.25	2.67	0.09	5.07	11.76
Nearest prep plant (km)	7.00	3.34	0.26	7.07	16.75
Census stats					
housing units	36,695				
Population	81,866				
mean commute	20.6				
% white	93.5				
% black	3.8				
% hispanic	1.0				
% labor force in mining	2.7				
Median hh income	28,625				
% below poverty	11.3				
% hs grad	30.5				

Figure 3.2: Nearest coal sites: three counties

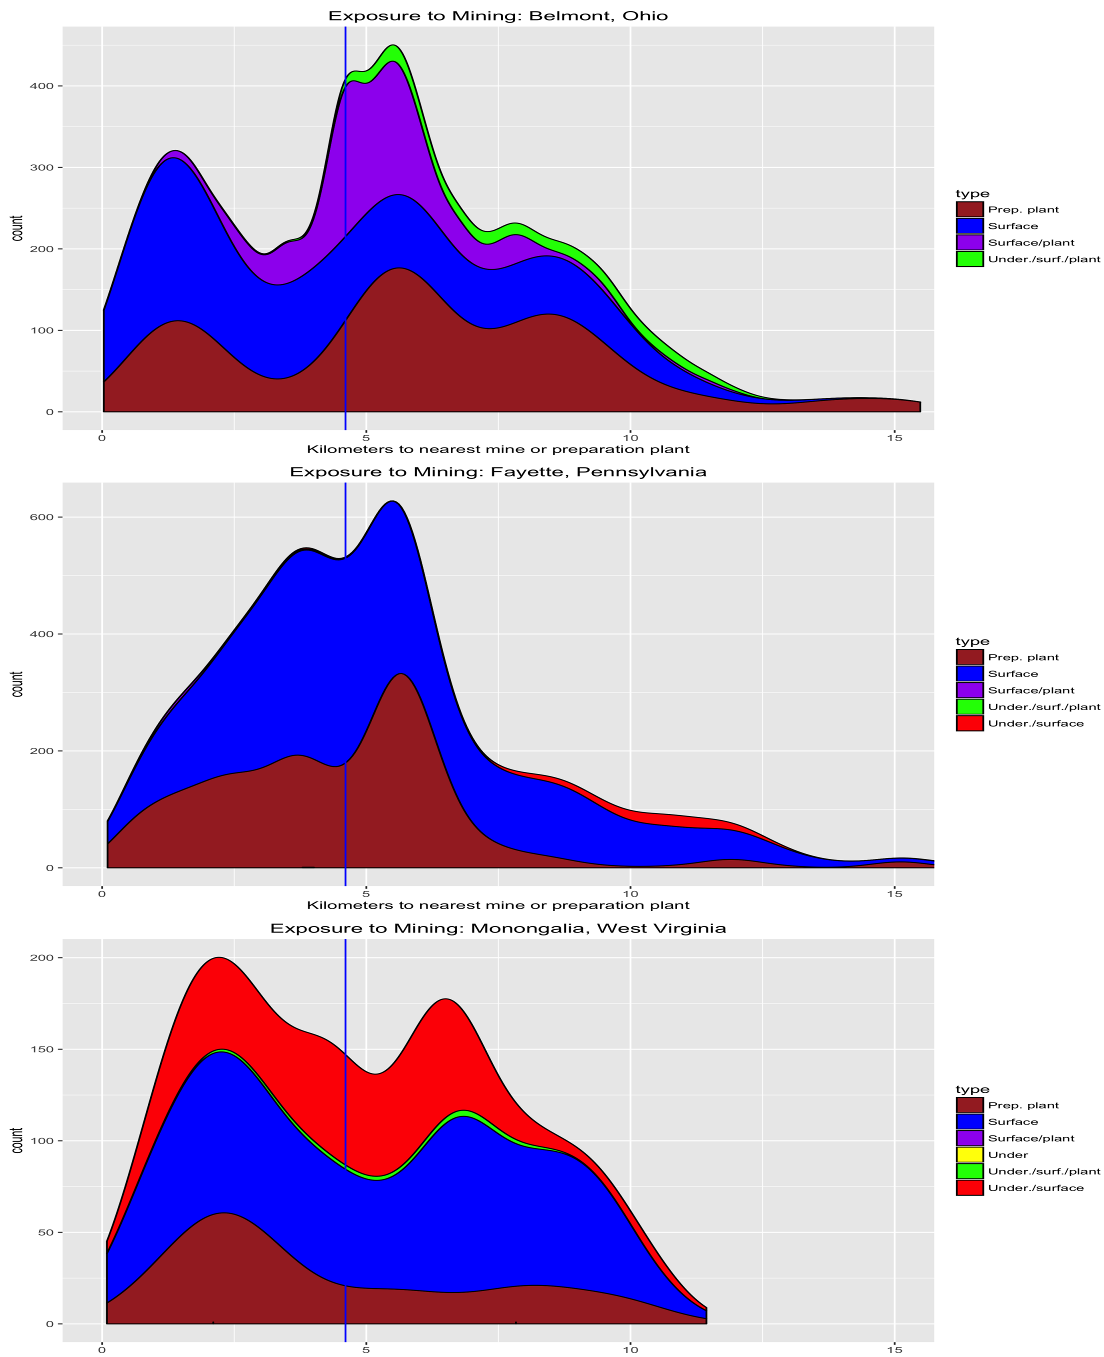

good. The hedonic prices schedule represents the tangency points between buyers' bid functions and sellers' offer functions:

$$P = P(x_1, \ldots, x_n, \ldots, z) \tag{3.1}$$

In the case of homes, the hedonic price schedule represents the relationship between a home's price and its structural and neighborhood characteristics x, and environmental characteristics z. In a hedonic equilibrium where individual utility curves lie tangent to the hedonic price schedule, the slope of the utility curve (and of the HPS) represents the increase in consumption goods households would accept in exchange for a reduction in environmental quality.

Suppose z reflects each property's insulation from the coal mining industry, and buyers value living far away from coal sites. If two homes differ only in their exposure to coal mining and households prefer to be insulated from mining, then the home closer to a coal site would command a lower price in order to convince buyers to take that home. At any point on the hedonic price schedule, the marginal price of a characteristic represents the marginal willingness to pay of households at that point (Rosen 1974). Valuation studies proceed by specifying a functional form (most often log-log or semi-log) of $P(x, z)$ and estimating its parameters. The coefficient on environmental quality would therefore be either an elasticity (i.e. the elasticity of price with respect to distance from a coal site) or semi-elasticity (i.e. the percent change in price resulting from an additional kilometer between a property and the nearest site.)

A simple hedonic model (of the semi-log form) one could estimate is:

$$\ln P = \beta_0 + X\beta_1 + Z\gamma + \epsilon. \tag{3.2}$$

According to model 3.2, the log of house prices are a linear function of structural characteristics X and environmental quality Z. Suppose that Z measure insulation from the mining industry - an attribute which households value. Then $\gamma > 0$ would confirm that households view coal sites as a disamenity, and would be wiling to pay for more distance between themselves and the coal industry. If the location of coal mines or preparation plants are spatially correlated with unobservable characteristics that households also care about, then regressing logged sale prices on structural characteristics and coal exposure would generate biased estimates of γ (Kuminoff, Parmeter & Pope 2010). To address this concern, we estimate models that include census tract fixed effects. Census tracts are designed to contain roughly 4,000 people, and can therefore vary in size, depending on local population density. Since we use seven years of data, we also include sale year fixed effects to dampen the impact of market-wide shocks. The main specification we use is:

$$\ln P_{ijt} = \beta_0 + X_{ijt}\beta_1 + Z_{ijt}\beta_2 + tract_j + year_t + \epsilon_{ijt} \tag{3.3}$$

where X_{ijt} is a vector containing the structural characteristics of home i sold in census tract j during year t, and Z is some function of distance to the nearest operational mining site.[51] In all of our specifications, our proxy for exposure to the coal industry is some fuction of linear distance to the nearest coal mine or preparation plant d_{ijt}. We first test three continuous measures of exposure: distance to nearest site ($Z(d) = d$), the logarithm of that distance to nearest site ($Z(d) = \ln d$), and the inverse of the logarithm of nearest distance $Z(d) = \frac{1}{\ln d}$. According to a linear measure of exposure, γ represents the expected percent change in price for an extra kilometer between a home and the nearest coal site. For the log-log specification, γ is a standard elasticity of price with respect to distance to site. For both, a positive value indicates coal sites are net disamenities. Also in the spirit of earlier work, we test two *ad hod* exposure buffers of

[51] According to these measures, "operational" means non-zero hours worked at a site in the same or previous quarter of sale. Need to test robustness here by perhaps only using same quarter, or previous quarter.

1 and 5 kilometers. According to these measures, a home is exposed to the industry if it falls within 1 or 5 kilometers of a site. For the last three, a negative value indicates mines are net disamenties. In the case of inverse logged distance, a value of β_m greater than zero would mean mines are net amenities: as inverse logged distance grows (and so logged distance falls) home prices rise. In this study, identification arises from variation in exposure among homes in the same census tract that are sold in the same year.

To advance earlier work on undesirable land uses, we take aim at the *ad hoc* nature of earlier treatment buffers by letting the sales data dictate the thresholds markets implicitly use to define exposure. Figure 3.2 presents the distribution of nearest coal site for each county. In each, over 90% of home sales take place within 10 kilometers of some kind of mining site. We separate the distribution into 99 roughly equal segments. For threshold b, we estimate the following model for a county while omitting observation k:

$$\ln P_{ijt} = \beta_0 + X_{ijt}\beta_1 + 1\left(d_{ijt} < w_b\right)\gamma + tract_j + year_t + \epsilon_{ijt}. \tag{3.4}$$

We use the estimates from equation 3.4 to predict the the outcome for observation k. This is repeated for every sale in a given county. We can then rank each potential buffer in order of how successfully they predict house prices, as measured by the sum of squared prediction errors. Note that by testing the entire distribution in roughly 1% increments, we avoid imposing any *a priori* assumptions about how housing markets view coal operations. Results from these exercises are in the next section.

3.5 Results

Tables 3.4, 3.5, and 3.6 present coefficient estimates from six different models, each of which uses a separate measure of exposure. Several facts are worth noting. First, the implicit prices of structural characteristics are stable across specifications of exposure. Moving across columns (1) to (6) in each county demonstrates that markets' valuation of structural attributes is consistent across models that use linear distance, logged linear distance, inverse logged linear distance, and discrete measures of exposure. Second, implicit prices are not stable across counties. Square footage is twice as valuable in Monongalia as it is in the other two counties, but the price of acreage in Monongalia is not statistically different from zero. This is most likely due to the relatively more urban nature of real estate in Monongalia. Bathrooms in Belmont are half as valuable as in the other two counties. Drinking water wells have a large, negative impact on prices in Monongalia, and a small and positive (though insignificant) impact in Fayette. Only in Belmont does an additional bedroom increase sale prices (by 6%) though this may be due to the fact that, *ceteris paribus*, an additional bedroom implies smaller rooms. Homes with fireplaces command a roughly 20% premium over homes without one in each county. While no homes with pools were transacted in Monongalia, pools in Belmont commanded twice the price premium as in Fayette. In short, these distinct markets feature buyers who have varied preferences for structural characteristics. With this in mind, it's unsurprising that these same markets would value environmental quality, in the form of exposure to mining, differently.

Of primary interest in tables 3.4, 3.5, and 3.6 are the coefficients on exposure to coal sites. The only consistent results using the continuous measures of exposure comes from Monongalia. There, the first two continuous measures of exposure indicate modest negative price impacts: an additional kilometer of distance between a parcel and the nearest mining site is associated with an additional three percent in sale price. Column (2) indicates that a one percent increase in this distance would lead to a 0.081 percent increase in sale price.

Both specifications indicated that buyers in homes farther from coal sites sell for more. Coefficients on inverse logged distance is near zero for each county.

The discrete measures yield more consistent results. In each model that uses a 1 kilometer buffer, Belmont and Fayette generate negative estimates of price impacts. Monongalia, where only 77 sales took place within that buffer, is slightly greater than zero. For the 5 kilometer buffer, Monongalia has a large, significant, and negative impact (nearly 13%) less. The negative impact in Fayette goes away. The impact in Belmont switches from insignficant and negative to insignificant and positive. For the preferred buffer, selected from the leave-one-out prediction exercise, the price impact of being exposed to the industry is negative. In Fayette, the 81% of sales exposed to the coal industry sold at a 12 percent discount. In Monongalia, the 80% of sales exposed sold at a 15% discount. These two impacts are significnt at the 1% level. In Belmont, the negative impact is insignificant.

Table **??** contains the empirically determined exposure thresholds, the proportion of sales in each county that fall within that distance of a site, the estimates of price impacts, and the standard errors, t-statistics, and p-values of each estimate. In all three counties, markets view coal operations as net disamenities. Price impacts are highly significant in Fayette and Monongalia. Both of those counties also have exposure buffers that are considerably larger than that of Belmont county, which indicate that buyers in those counties are more sensitive to or are more aware of the environmental impacts of mining. A larger buffer indicates that more homes are considered to be affected. These differences could be due to topography, information about mining, county building regulations, or some other factor.

Table 3.4: Exposure to mining: Belmont

	logpricedef					
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Sq. feet)	0.339***	0.338***	0.337***	0.339***	0.338***	0.340***
	(0.043)	(0.043)	(0.043)	(0.043)	(0.043)	(0.043)
ln(acres)	0.042*	0.042*	0.043*	0.042	0.044*	0.040
	(0.025)	(0.025)	(0.025)	(0.025)	(0.025)	(0.025)
Age	−0.007***	−0.007***	−0.007***	−0.007***	−0.007***	−0.007***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Beds	0.058***	0.058***	0.059***	0.058***	0.059***	0.058***
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)
Baths	0.043**	0.043**	0.043**	0.043**	0.043**	0.043**
	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)
Swimming pool	0.397***	0.398***	0.397***	0.398***	0.396***	0.400***
	(0.081)	(0.081)	(0.080)	(0.081)	(0.080)	(0.080)
Garage	0.134***	0.134***	0.134***	0.134***	0.134***	0.134***
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Fireplace	0.248***	0.248***	0.248***	0.248***	0.248***	0.248***
	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)
Linear Distance	0.002					
	(0.005)					
Logged Distance		0.004				
		(0.018)				
Inverse logged distance			0.000			
			(0.000)			
Within 1km				−0.031		
				(0.048)		
Within 5km					0.035	
					(0.030)	
Within my 1.24km						−0.065
						(0.044)
Observations	2,916	2,916	2,916	2,916	2,916	2,916
R^2	0.454	0.454	0.454	0.454	0.454	0.454

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 3.5: Exposure to mining: Fayette

	logpricedef					
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Sq. feet)	0.343***	0.343***	0.342***	0.344***	0.342***	0.343***
	(0.035)	(0.035)	(0.035)	(0.035)	(0.035)	(0.035)
ln(acres)	0.329***	0.328***	0.328***	0.327***	0.328***	0.335***
	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)
Age	−0.006***	−0.006***	−0.006***	−0.006***	−0.006***	−0.006***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Beds	0.021	0.021	0.021	0.021	0.021	0.021
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Baths	0.098***	0.099***	0.098***	0.099***	0.098***	0.098***
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Swimming pool	0.222***	0.223***	0.221***	0.219***	0.221***	0.221***
	(0.062)	(0.062)	(0.062)	(0.062)	(0.062)	(0.062)
Garage	0.310***	0.310***	0.310***	0.308***	0.310***	0.310***
	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)
Fireplace	0.235***	0.236***	0.237***	0.237***	0.237***	0.231***
	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
Linear Distance	0.005					
	(0.004)					
Logged Distance		0.031				
		(0.020)				
Inverse logged distance			−0.000			
			(0.001)			
Within 1km				−0.127**		
				(0.052)		
Within 5km					−0.001	
					(0.029)	
Within 7.47 km						−0.115***
						(0.035)
Observations	3,692	3,692	3,692	3,692	3,692	3,692
R^2	0.477	0.477	0.477	0.478	0.477	0.478

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Table 3.6: Exposure to mining: Monongalia

	logpricedef					
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Sq. feet)	0.668***	0.667***	0.665***	0.664***	0.667***	0.670***
	(0.049)	(0.049)	(0.050)	(0.050)	(0.049)	(0.049)
ln(acres)	0.017	0.013	0.012	0.012	0.011	0.022
	(0.050)	(0.050)	(0.050)	(0.050)	(0.050)	(0.050)
Age	−0.004***	−0.004***	−0.004***	−0.004***	−0.004***	−0.004***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Beds	0.039*	0.037*	0.038*	0.038*	0.038*	0.037*
	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Baths	0.093***	0.094***	0.093***	0.093***	0.094***	0.090***
	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Garage	0.095***	0.095***	0.094***	0.093***	0.095***	0.094***
	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)
Fireplace	0.176***	0.175***	0.176***	0.177***	0.180***	0.176***
	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)	(0.031)
Linear Distance	0.032***					
	(0.011)					
Logged Distance		0.081**				
		(0.035)				
Inverse logged distance			−0.001			
			(0.001)			
Within 1km				0.005		
				(0.069)		
Within 5km					−0.129***	
					(0.044)	
Within 6.42 km						−0.148***
						(0.048)
Observations	1,471	1,471	1,471	1,471	1,471	1,471
R^2	0.544	0.542	0.541	0.541	0.544	0.544

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

3.5.1 Impacts of different types of sites

Having identified these thresholds, we can turn to identifying which types of coal sites buyers are taking note of. For each county, we estimate the model

$$\ln P_{ijt} = \beta_0 + X_{ijt}\beta_1 + 1\left(d_{ijtm} < w_b\right)\gamma_m + tract_j + year_t + \epsilon_{ijt} \tag{3.5}$$

where d_{ijtm} is the distance from home i in tract j and year t to mining site of type m. This includes: standalone underground mines, surface mines, and preparation plants; combination underground and surface mines, and surface mines with preparation plants; and sites that feature an underground mine, a surface mine, and a preparation plant. These results are in tables impacts. In Monongalia and Fayette counties, surface mines are associated with large negative price impacts. In Monongalia, a home exposed to a standalone surface mine sells for 15% less on average, after controlling for structural characteristics, unobserved neighborhood characteristics, and temporal variation. In Fayette, homes exposed to surface mines sell for 10% less. Homes near preparation plants in Belmont sell at an even larger discount of 32%.

In both of these counties, homes located near combination underground mine, surface mine, preparation plants sell at a premium, though this result is statistically insignificant. The weak positive impact could reflect the fact that these locations (in all three states studied) employ by far the largest number of workers of any type of location.

Table 3.7: Exposure Thresholds and Price Impacts: Three Appalachian Housing Markets

County	Buffer (km)	Prop. treated	Price Impact	Std. Err.	t value	Pr($>\|t\|$)
Belmont, OH	1.2	0.12	-0.07	0.04	-1.48	0.14
Fayette, PA	7.5	0.81	-0.12	0.04	-3.27	0.00
Monongalia, WV	6.4	0.67	-0.15	0.05	-3.09	0.00

Table 3.8: Price Impact of Mining Operations

Belmont, OH: $n = 2,916$	# treated	Price Impact	Std. Error
Any Mine	349	-0.07	0.04
Surface mine	243	0	0.07
Prep. plant	103	-0.32	0.11
S. mine & plant	3	-0.84	1.51

Fayette, PA: $n = 3,692$	Proportion treated	Price Impact	Std. Error
Any operation	2,990	-0.12	0.04
U. mine	2	-0.64	0.87
S. mine	2,290	-0.10	0.04
Prep. plant	1,678	-0.02	0.08
U. & S. mine	22	-0.13	0.16
S. mine & plant	27	0.13	0.08
U. & S. mine & plant	45	0.07	0.07

Monongalia, WV	Proportion treated	Price Impact	Std. Error
Any operation	985	-0.15	0.05
U. mine	4	-0.03	0.3
S. mine	701	-0.15	0.04
Prep. plant	618	0.05	0.05
U. & S. mine	617	0.02	0.05
S. mine & plant	189	-0.05	0.05
U. & S. mine & plant	12	0.15	0.34

Table 3.9: Impact of a Single Operation within Buffer

Belmont, OH	# Treated	Price Impact	Std. Error
Surface mine	243	0	0.05
Prep. plant	103	−0.32	0.11
S. mine & plant	3	−0.84	0.35

Fayette, PA	Proportion treated	Price Impact	Std. Error
S. mine	1,288	−0.08	0.03
Prep. plant	698	−0.02	0.04

Monongalia, WV	Proportion treated	Price Impact	Std. Error
S. mine	150	−0.15	0.04
Prep. plant	60	0.00	0.00
U. & S. mine	93	0.00	0.07
U. & S. mine & plant	9	0.25	0.24

Table 3.10: Quarterly Production Statistics: Ohio

Site	# in study area	Tons mined	Miners	Office workers
Underground	2	69,942	16	0
		(29,541)	(6)	(0)
Surface	44	47,043	20	3
		(59,047)	(19)	(2)
Under/surface	5	185,412	44	3
		(157,684)	(31)	(3)
Prep plant	8		10	1
			(7)	(0)
Surface/plant	3	19,869	10	13
		(29,635)	(7)	(3)
Under/surface/plant	2	1,178,876	130	19
		(388,913)	(20)	(12)

Table 3.11: Quarterly Production Statistics: Pennsylvania

Site	# in study area	Tons mined	Miners	Office workers
Underground	1	19	3	2
		(40)	(0)	(0)
Surface	86	22,264	9	3
		(52,546)	(17)	(2)
Under/surface	20	317,675	50	4
		(711,753)	(77)	(4)
Prep plant	17		21	2
			(24)	(1)
Surface/plant	5	24,590	9	3
		(15,423)	(3)	(1)
Under/surface/plant	6	1,562,219	167	19
		(730,630)	(34)	(10)

Table 3.12: Quarterly Production Statistics: West Virginia

Site	# in study area	Tons mined	Miners	Office workers
Underground	1	0	1	0
		(0)	(0)	(0)
Surface	23	13,329	7	2
		(34,198)	(7)	(0)
Under/surface	8	163,716	37	1
		(179,943)	(40)	(0)
Prep plant	9		11	3
			(14)	(3)
Surface/plant	2	8,347	2	0
		(10,850)	(1)	(0)
Under/surface/plant	6	1,228,869	148	5
		(609,607)	(44)	(3)

Appendix A

A.1 Exogenous variables

The DICE model includes many variables which change exogenously over time. Further, unlike DICE, we assume the ocean temperature also changes exogenously to reduce the state space. Reducing the state space from seven to six variables significantly reduces the computation time. Traeger (2014) presents a deterministic DICE model with exogenous ocean temperature. Therefore, we take the evolution of the exogenous variables directly from that study. For completeness, they are listed below. We refer the reader to Traeger (2014) for details.

Population growth:

$$L(t) = L_0 + \left(\bar{L} - L_0\right)\left(1 - \exp\left[-g_L t\right]\right). \tag{A.1}$$

TFP growth:

$$A(t) = A_0 \exp\left[g_{A,0}\frac{1 - \exp\left[-\delta_A t\right]}{\delta_A}\right]. \tag{A.2}$$

Backstop cost:

$$\Psi(t) = \frac{\sigma(t)}{a_2} a_0 \left(1 - \frac{1 - \exp\left[g_\Psi t\right]}{a_1}\right). \tag{A.3}$$

Emissions intensity of output:

$$\sigma(t) = \sigma_0 \exp\left[g_{\sigma,0}\frac{1-\exp\left[-\delta_\sigma t\right]}{\delta_\sigma}\right]. \tag{A.4}$$

Exogenous emissions:

$$B(t) = B_0 \exp\left[-\delta_B t\right]. \tag{A.5}$$

Decay rate of GHGs:

$$\delta_m(t) = \bar{\delta}_m + \left(\delta_{m,0} - \bar{\delta}_m\right)\exp\left[-\delta_m^* t\right]. \tag{A.6}$$

Exogenous forcing:

$$EF(t) = EF_0 + 0.01\left(EF(100) - EF_0\right)\cdot \min\left\{t, 100\right\}. \tag{A.7}$$

Heat uptake from ocean:

$$O(t) = \max\left\{\Delta_{T1} + \Delta_{T2}t + \Delta_{T3}t^2, 0\right\}. \tag{A.8}$$

Discount factor:

$$\beta(t) = \exp\left[-\delta_u + (1-\eta)\log\left(\frac{A(t+1)}{A(t)}\right) + \log\left(\frac{L(t+1)}{L(t)}\right)\right]. \tag{A.9}$$

A.1.1 Derivation of Equation (2.26)

Rewriting the left hand side of (2.25) for $i = 1$ gives:

$$Pr\left(T_{t+1} \geq T^*\right) = Pr\left(\beta_2 F_{t+1} + \beta_3 (O-T)(t) + \tilde{H}_{t+1} \geq T^*\right), \tag{A.10}$$

$$= Pr\left(\tilde{H}_{t+1} \geq T^* - \beta_2 F_{t+1} - \beta_3 (O-T)(t)\right), \tag{A.11}$$

$$= 1 - \text{NCDF}\left[\frac{T^* - \beta_2 F_{t+1} - \beta_3 (O-T)(t) - \mu_t T_t}{\sqrt{\frac{T_t^2}{\eta_t} + \frac{1}{\rho}}}\right]. \tag{A.12}$$

Here NCDF is the cumulative distribution function of the standard normal distribution. Let:

$$P_{t+1} \equiv \beta_2 F_{t+1} + \beta_3 \left(O - T\right)(t) + \mu_t T_t \tag{A.13}$$

be the expected (predicted) value of T_{t+1}, then using equation (2.13):

$$Pr\left(T_{t+1} \geq T^*\right) = 1 - \text{NCDF}\left[\frac{T^* - P_{t+1}}{\sigma_{H,t}}\right]. \tag{A.14}$$

Constraint (2.25) is therefore equivalent to:

$$\text{NCDF}\left[\frac{T^* - P_{t+1}}{\sigma_{H,t}}\right] \geq 1 - \omega, \tag{A.15}$$

$$\frac{T^* - P_{t+1}}{\sigma_{H,t}} \geq \text{NCDF}^{-1}\left[1 - \omega\right], \tag{A.16}$$

$$F_{t+1} \leq \frac{1}{\beta_2}\left(T^* - \beta_3\left(O - T\right)(t) - \mu_t T_t - \sigma_{H,t} \cdot \text{NCDF}^{-1}\left[1 - \omega\right]\right). \tag{A.17}$$

Equations (2.19)-(2.21) then imply (A.17) reduces to:

$$\begin{aligned} x_t \geq & 1 + \frac{M_B}{\sigma(t) L(t) A(t)^{\frac{1}{1-\gamma}} k_t^{\gamma}} \Bigg(1 + \left(1 - \delta_m(t)\right)\left(m_t - 1\right) - \\ & \exp\left\{\frac{\log(2)}{\Omega \beta_2}\left[T^* - \beta_3\left(O - T\right)(t) - \mu_t T_t - \sigma_{H,t} \cdot \text{NCDF}^{-1}\left[1 - \omega\right]\right]\right\}\Bigg), \end{aligned} \tag{A.18}$$

$$x_t \geq PC\left(s_t, T^*, \omega\right). \tag{A.19}$$

Here PC is the right hand side of (A.18).

We now calculate the highest possible probability for which $T_{t+1} \geq T^*$, which occurs with a zero abatement rate. In this case, we have from (2.20):

$$E_t^{\max} = \sigma(t) A(t)^{\frac{1}{1-\gamma}} L(t) k_t^{\gamma} + B(t), \tag{A.20}$$

$$m_{t+1}^{\max} - 1 = \left(1 - \delta_m(t)\right)\left(m_t - 1\right) + \frac{E_t^{\max}}{M_B}, \tag{A.21}$$

$$F_{t+1}^{\max} = \Omega \log_2 \left(m_{t+1}^{\max} \right) + EF\left(t\right). \tag{A.22}$$

Next, from (2.18):

$$T_{t+1}^{\max} = \tilde{H}_{t+1} + \beta_2 F_{t+1}^{\max} + \beta_3 \left(O - T\right)\left(t\right), \tag{A.23}$$

We can then calculate $\omega^{\max}$ as:

$$\omega_{t+1}^{\max} = Pr\left(T_{t+1}^{\max} \geq T^*\right), \tag{A.24}$$

$$\omega_{t+1}^{\max} = Pr\left\{\tilde{H}_{t+1} \geq T^* - \beta_2 F_{t+1}^{\max} - \beta_3 \left(O - T\right)\left(t\right)\right\}, \tag{A.25}$$

$$\omega_{t+1}^{\max} = 1 - \text{NCDF}\left(\frac{T^* - P_{t+1}^{\max}}{\sigma_{H,t}}\right), \text{ where} \tag{A.26}$$

$$P_{t+1}^{\max} \equiv \mu_t T_t + \beta_2 F_{t+1}^{\max} + \beta_3 \left(O - T\right)\left(t\right). \tag{A.27}$$

Since the difference between the target and the expected temperature in period $t + 1$ under maximum emissions is not infinite, we have immediately from (A.26) that $\omega^{\max} < 1$.

Analogously, the minimum probability of exceeding the constraint occurs when the abatement rate is one.

$$E_t^{\min} = B\left(t\right), \tag{A.28}$$

$$m_{t+1}^{\min} - 1 = \left(1 - \delta_m\left(t\right)\right)\left(m_t - 1\right) + \frac{E_t^{\min}}{M_B}, \tag{A.29}$$

$$F_{t+1}^{\min} = \Omega \log_2 \left(m_{t+1}^{\min} \right) + EF\left(t\right), \tag{A.30}$$

$$T_{t+1}^{\min} = \tilde{H}_{t+1} + \beta_2 F_{t+1}^{\min} + \beta_3 \left(O - T\right)\left(t\right), \tag{A.31}$$

$$\omega_{t+1}^{\min} = 1 - \text{NCDF}\left(\frac{T^* - P_{t+1}^{\min}}{\sigma_{H,t}}\right), \text{ where} \tag{A.32}$$

$$P_{t+1}^{\min} \equiv \mu_t T_t + \beta_2 F_{t+1}^{\min} + \beta_3 \left(O - T\right)\left(t\right). \tag{A.33}$$

A.1.2 Tables

Parameter	Description	Value
L_0	Initial population	6514
$\bar{L}$	Steady state population	8600
g_L	decline rate in population growth	0.035
A_0	Initial TFP	0.0058
g_{A0}	Initial TFP growth rate	0.0131
δ_A	Decline rate in TFP growth rate	0.001
γ	Capital share	0.3
δ_k	Depreciation rate of capital	0.1
ϕ	Coefficient of risk aversion	2
δ_u	Pure rate of time preference	0.015
g_Ψ	Backstop cost growth rate	-0.005
a_0	Initial backstop cost	1.17
a_1	Backstop cost parameter	2
a_2	Cost convexity	2
b_1	Damage function parameter	0.00284
b_2	Damage function convexity	2
σ_0	Initial emissions intensity	0.1342
$g_{\sigma,0}$	Initial growth rate in σ	-0.0073
δ_σ	Decline rate in emissions intensity growth	0.003
B_0	Initial exogenous emissions	1.1
δ_B	decay rate in exogenous emissions	0.0105
M_B	Pre-industrial GHG stock (gigatons)	590
δ_m^*	decay rate in GHG decay rate	0.0083
$\bar{\delta}_m$	steady state GHG decay rate	0.01
$\delta_{m,0}$	initial decay rate	0.014
EF_0	Initial exogenous forcing	-0.06
EF_{100}	Exogenous forcing at $t = 100$	0.3
Ω	Forcing coefficient	3.8
α	Ocean heat uptake	0.2837^{-1}
ξ	Heat transfer coefficient from the ocean	0.07
Δ_{T1}	Ocean Temperature Parameter	0.7
Δ_{T2}	Ocean Temperature Parameter	0.02
Δ_{T3}	Ocean Temperature Parameter	-0.00007
k_0	Initial capital per TFP adjusted person	3.6261
T_0	Current air warming above pre-industrial	0.76
m_0	Current GHG stock, relative to pre-industrial	1.371
μ_0	Mean of climate feedback prior distribution	0.65
η_0	Precision of climate feedback prior distribution	0.13^{-2}
ρ	Precision of weather shock	0.11^{-2}

Table A.1: Description and values of model parameters.

		Year			Auto-
		2005	2050	2075	correlation
	0.2	0	25.3	34.8	0.67
	0.25	0	30.4	37.3	0.67
	0.3	0	32.6	41.4	0.66
	0.5	0	35.5	50.6	0.67
ω	0.75	0	49.5	62.7	0.70
	0.85	0	51.8	69.1	0.73
	0.9	0	52.7	71.0	0.74
	0.95	0	53.8	72.8	0.76
	0.975	0	54.9	74.5	0.77
	1	0	72.1	93.6	0.88

Table A.2: Columns 2-4 give the percent of 10,000 simulations for which $T_t > T^*$. The last column is the first order autocorrelation of exceeding the constraint. That is, let $q_{it} = 1$ if and only if $T_{it} > T^*$ for time t and simulation i. The last column is the correlation between q_{it} and $q_{i,t-1}$.

	Total	No weather shocks	
ω	Loss (%)	Loss (%)	Difference (%)
0.20	5.20	4.86	6.49
0.25	4.99	4.57	8.29
0.30	4.82	4.42	8.16
0.33	4.73	4.34	8.23
0.50	4.33	3.95	8.84
0.75	3.69	3.39	8.12
0.85	3.39	3.07	9.47
0.90	3.23	2.92	9.73
0.95	3.02	2.73	9.49
0.97	2.85	2.60	8.78
1.00	0	0	NA
		Certainty	
0	2.84		

Table A.3: Decomposition of the welfare loss. The baseline case is $\omega = 0.33$. The second column reports $1 - V\left(s_0; \omega\right) / V\left(s_0; 1\right)$. The third column is computed in the same way as column two, except that only the evolution of the mean and variance of the prior are affected by the weather shocks. That is $s' = \left[k', m', T'\left(\nu = 0\right), t', \mu'\left(\nu\right), \eta'\left(\nu\right)\right]$.

A.1.3 Solution

We solve the model by forming a grid of values for the state space, and then assume v takes the form of a cubic spline with continuous first and second derivatives. The model can then be solved by choosing an initial spline, optimizing at each grid point, and then using the solution to update the parameters of the spline. Kelly & Tan (2015) give a detailed explanation of this solution method. The solution method requires specifying a maximum climate sensitivity, which is discussed in Section **??**. The constraint, $x \geq PC\left(s\right)$, is linear in the controls and is therefore straightforward to add to the optimization routine in the solution method (e.g. Matlab's optimization routines automatically generate the Kuhn-Tucker conditions). Potential points of non-differentiability in the value function caused by the constraint are discussed in Section **??**.

Once the model converges, we obtain the optimal decision rules, $x\left(s\right)$ and $k'\left(s\right)$. We then simulate the model using the decision rules and the transition equations (2.17)-(2.24). The algorithm is:

1. Draw a true value of the climate feedback parameter β_1^* from $N[\mu_0, \frac{1}{\eta_0}]$.

2. Given s_0 compute $x(s_0) = x_0$.

3. Given x_0, s_0, β_1^*, and a randomly drawn weather shock ν_0, compute s_1 from transition equations.

4. Repeat steps (2)-(3) for n_p years.

5. Repeat steps (1)-(4) n_s times with different draws for ν and β_1.

6. Compute means over all simulations to get the expected value of each variable in each time period.

The above algorithm gives the expected value of each variable conditional on the prior information for β_1. In some cases, we fix a value for β_1 and vary only ν across simulations. In this case, we obtain the expected value of each variable conditional on β_1.

The solution algorithm requires a maximum value of the climate sensitivity. As discussed in Kelly & Tan (2015) appendix B.2, if climate sensitivity is unbounded, then the temperature can always exceed the largest temperature grid point. Interpolation of the value function outside the grid is likely to be inaccurate. We therefore set the largest value of the climate sensitivity equal to 15°C. Values greater than 15°C are collected into a mass point at 15°C.[52]

The results, however, are not sensitive to the maximum allowable climate sensitivity. First, the simulations aggregate probability mass greater than the maximum climate sensitivity into a mass point at the maximum climate sensitivity. Therefore, reducing the maximum $\Delta T_{2\times}$ essentially makes a small range of low probability events less harmful.

Further, unconstrained abatement is increasing in the climate sensitivity. Therefore, beyond a certain critical climate sensitivity, the 2°C target is almost certainly exceeded, so constrained emissions quickly fall to zero and remain at zero for a relatively long period of time. But unconstrained emissions also drop to near zero. The peak temperature increases for both the constrained and unconstrained models. Therefore, beyond a certain climate sensitivity, the only welfare difference between high and very high values of the climate sensitivity is in the far future, where the constrained model still has zero emissions in an attempt to return to 2°, while the unconstrained model remains at an optimal temperature which is increasing in the climate sensitivity, with more emissions.

For example, if all weather shocks are zero and emissions are zero, the target is exceeded for any $\Delta T_{2\times} \geq 5.36$. Therefore, the distribution can be truncated anywhere

[52]The literature provides some expert opinion the upper bound of the climate sensitivity. Weitzman (2009) aggregates the densities of 22 studies and estimates the probability that $\Delta T_{2\times} > 10$°C is approximately 1%. However, he also suggests that because of feedback uncertainty, the probability that $\Delta T_{2\times} > 20$ is 1%. Roe & Baker (2007) plot the densities of several studies, none of which have any mass above 10°C. Similarly, Intergovernmental Panel on Climate Change (2007) normalize a number of papers so that the probability that $\Delta T_{2\times} > 10°$ is zero. Our calibration matches the fitted distribution of Roe & Baker (2007), for which the probability of $\Delta T_{2\times} > 15$ is 1.8%. Therefore, a mass point exists at $\Delta T_{2\times} = 15$ of 1.8%.

between 8-15°C with little change in the near term emissions path. Differences in the far future exist, but the low probability of such events and discounting combine to limit the effect on welfare.

Table A.4 shows how the welfare loss changes with the maximum value of $\Delta T_{2\times}$. The maximum difference in the table is about 12%. Note that welfare losses are lower as a percentage of the unconstrained welfare when the maximum $\Delta T_{2\times}$ is higher. This occurs because both constrained and unconstrained welfare fall as the maximum $\Delta T_{2\times}$ increases, but the unconstrained welfare falls more. Total welfare (not as a percentage of unconstrained) is monotonically decreasing as a function of the maximum $\Delta T_{2\times}$ for all cases.

		Maximum $\Delta T_{2\times}$ (°C)		
		8	10	15
	0.20	5.54	5.35	5.20
	0.25	5.32	5.14	4.99
	0.30	5.16	4.97	4.82
	0.50	4.67	4.49	4.33
ω	0.75	4.04	3.85	3.69
	0.85	3.73	3.55	3.39
	0.90	3.58	3.39	3.23
	0.95	3.37	3.19	3.02
	0.98	3.19	3.01	2.85
	1	0	0	0
Minimum Feasible ω		0.152	0.152	0.152

Table A.4: Welfare loss as a function of the maximum $\Delta T_{2\times}$. The baseline maximum is 15°C above preindustrial (column 3). The last row gives the probability that the temperature exceeds the target at some time in the future with zero emissions (the peak in Figure 2.1).

The presence of the min operator in the constraint implies the value function is continuous, but not differentiable (kinked), at a particular point. Therefore, error will exist at the kink point as the spline has continuous first and second derivatives. Here we examine the error in the kink point.

The first step is to calculate the kink point. Below is a heuristic argument, for a more formal treatment, see Rincón-Zapatero & Santos (2009). Problem (2.3) has potentially three regions: (1) the constraint does not bind ($\theta = 0$), (2) the constraint binds and $PC(s) < 1$, and (3) the constraint binds with $PC(s) \geq 1$.

(2),

$$\theta(s_0) \frac{\partial}{\partial s} \min\{PC(s_0), 1\} = \theta(s_0) PC_s(s_0). \tag{A.34}$$

Therefore, one can again apply standard arguments since the constraint is continuously differentiable in this region. At the point where the constraint just binds, $\theta = 0$ and $x^*(s_0) = PC(s_0)$. Further, V is differentiable at s_0 if the left and right derivatives are equal. Since from the binding region $\theta \to 0$ as $s \to s_0$ and from the non-binding region $\theta = 0$, the left and right derivatives are equal.

Next, in the binding region where $PC(s_0) > 1$: Hence any kink point s_0 in the value function satisfies $PC(s_0) = 1$. Further, the left derivative is smaller than the right derivative. For states such that the derivative of V is negative (increases in T, m, and μ reduces lifetime utility), then the left slope is steeper than the right slope.

Figure **??** graphs a sample kink point. The value function switches from concave to almost linear at the kink point, and the right slope is flatter than the left slope at the kink point, as expected. The green line is the (differentiable) spline approximation, which smooths the kink point. The magenta line is composed of two spline approximations, one on either side of the kink point. Hence, the magenta line need not be differentiable at the kink point. From the graph, the error associated with assuming the value function is differentiable is small.[53]

[53] We do not have a proof which bounds the approximation error. Nonetheless, a check of kink points in many areas of the state space revealed very small errors.

A.2 Determining Exposure Zones

Earlier research on the non-miner health impacts of coal define a county population's exposure as the quantity of coal (in short tons) extracted from mines within a county's borders. As described, this definition ignores nearby mining activity that very likely contributes to health. To improve on this, I maintain the implicit assumptions of earlier research by explicitly defining a county population's exposure as the quantity of coal extracted from mines within different sized exposure buffers.

First, I use Census estimates of the population centroid according to the 2010 Diennial Census. With these coordinates and the locations of every mine in the MSHA dataset, I sum the tons of coal mined per quarter (by technique) from mines within $x \in [1 - 99]$ kilometers of each centroid. An exposure buffer is therefore a circular area of radius x centered on a population's geographic centroid, and a population's exposure to either type of mining is the quantity of coal produced at mines of that type within an exposure buffer.

Next, I regress asthma hospital visits on two measures of exposure (surface and underground production within x kilometers), population, and county and quarter fixed effects. I estimate the following model:

$$\ln y_{it} = \beta_0 + \mathbf{pop}'_{it}\beta + \mathbf{surface}'_{xit}\gamma_x + \mathbf{underground}'_{xit}\delta_x + c_i + q_t + \epsilon_{it}, \quad \text{(A.35)}$$

where γ_x and δ_x are the semi-elasticities of asthma hospital visits with respect to an additional one-hundred thousand tons of coal extracted using surface techniques and underground techniques, respectively, within x kilometers of a population centroid. Here c_i and q_t denote county-specific and quarter-specific fixed effects, and ϵ_{it} is an idiosyncratic error. Table 2.8 contains estimates of γ_x and δ_x at 10 kilometer increments. Surface exposure yields positive and significant coefficients when defining a buffer of up to 60 kilometers. Given the lack of theoretical justification for one buffer over another, I select for each age group the exposure measure that delivers the minimum mean squared prediction error. These are 36 kilometers, for the entire population; 34 kilometers, for youths and adults; and 39 kilometers for the elderly. Given the lack of association between any measure of underground mining exposure and asthma hospitalizations, results present focus solely on measures of surface mining. I also repeat the above exercise for linear models that use as dependent variables the number of asthma hospitalizations per 100,000 residents. These results are in 2.9.

Bibliography

Aneja, V., A. Isherwood, and P. Morgan. 2012. "Characterization of particulate matter (PM10) related to surface coal mining operations in Appalachia." *Atmospheric Environment*, 54: 496–501.

Barnett, S., and T. Nurmagambetov. 2010. "Costs of asthma in the United States: 2002-2007." *Journal of Allergy and Clinical Immunology*, 127: 145–152.

Bernhardt, E., and M. Palmer. 2011. "The environmental costs of mountaintop mining valley fill operations for aquatic ecosystems of the central Appalachians." *Annals of the New York Academy of Sciences*, 1223: 39–57.

Brink, L., et al. 2014. "The association of respiratory hospitalization rates in WV counties, total, underground, and surface coal production and sociodemographic covariates." *Journal of Ocupational and Environmental Medicine*, 56: 1179–1188.

Bruckner, T., and K. Zickfield. 2009. "Emissions corridors for reducing the risk of a collapse of the Atlantic thermohaline circulation." *Mitigation and Adaptation Strategies for Global Change*, 14: 61–83.

Centers for Disease Control and Prevention. 2013. "Asthma facts - CDC's national asthma control program grantees." Department of Health and Human Services, Centers for Disease Control and Prevention.

Choudhury, A., M. Gordian, and S. Morris. 1997. "Associations between respiratory illness and PM10 air pollution." *Archives of Environmental Health*, 52: 113–117.

Code of Federal Regulations. 2014. "Coal mining point source category BPT, BAT, BCT limitations and new source performance standards." U.S. Federal Government 434.

Copenhagen Accord. 2009. "Report of the Conference of the Parties on its fifteenth session, held in Copenhagen from 7 to 19 December 2009." United Nations Framework Convention on Climate Change. `http://unfccc.int/resource/docs/2009/cop15/eng/11a01.pdf`.

Costello, C., M. Neubert, S. Polasky, and A. Solow. 2010. "Bounded Uncertainty and Climate Change Economics." *Proceedings of the National Academy of Sciences*, 107(18): 8108–10.

Currie, J., J. Zivin, J. Mullins, and M. Neidell. 2014. "What do we know about short- and long-term effects of early-life exposure to pollution?" *Annual Review of Resource Economics*, 6: 217–247.

Davis, L. 2011. "The effect of power plants on local housing values and rents." *The Review of Economics and Statistics*, 93: 1391–1402.

den Elzen, M. G. J., M. Meinshausen, and D. van Vuuren. 2007. "Multi-gas emissions envelopes to meet greenhouse gas concentration targets: cost versus certainty of timiting temperature increase." *Global Environmental Change-Human and Policy Dimensions*, 17: 260–280.

Environmental Protection Agency. 1995. "Emissions factors and AP42: Fifth edition." Environmental Protection Agency Technology Transfer Network, Washington, DC.

Environmental Protection Agency. 2004. "Air quality criteria for particulate matter." Environmental Protection Agency, Brussels, Belgium.

European Commission. 2007. "Communication from the commission to the council, the European parliament, the European economic and social committee and the committee of the regions - limiting global climate change to 2 degrees celsius - the way ahead for 2020 and beyond." European Commission, Brussels, Belgium. `http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:52007DC0002:EN:NOT`.

Gamper-Rabindran, S., and C. Timmins. 2011. "Does cleanup of hazardous waste site raise housing values? Evidence of spatially localized benefits." *Journal of Environmental Economics and Management*, 65: 345–360.

Ghose, M.K. 2009. "Pollution due to opencast coal mining and the characteristics of air-borne dust - an Indian scenario." *Public Health Reports*, 124: 541–550.

Ghose, M.K., and S. Majee. 2000. "Sources of air pollution due to coal mining and their impacts in Jharia coalfield." *Environment International*, 26: 81–85.

Ghose, M.K., and S. Majee. 2007. "Characteristics of hazardous airborne dust around an Indian surface coal mining area." *Environmental Monitoring and Assessment*, 130: 17–25.

Gopolakrishnan, S., and A. Klaiber. 2013. "Is the shale energy boom a bust for nearby residents? Evidence from housing values in Pennsylvania." *American Journal of Agricultural Economics*, 96: 43–66.

Greenstone, M., and J. Gallagher. 2008. "Does hazardous waste matter? Evidence from the housing market and the superfund program." *Quarterly Journal of Economics*, 123: 951–1003.

Guarnieri, M., and J. Balmes. 2014. "Outdoor air pollution and asthma." *The Lancet*, 63: 498–516.

Hansen, J. 2005. "A slippery slope: how much global warming constitutes 'dangerous anthropogenic interference'?" *Climatic Change*, 68(3): 269–79.

Hare, B., and M. Meinshausen. 2006. "How much warming are we committed to and how much ban be avoided?" *Climatic Change*, 75: 111–149.

Harvey, L. 2007. "Allowable CO2 concentrations under the United Nations framework convention on climate change as a function of the climate sensitivity probability distribution function." *Environmental Research Letters*, 2: 10.

Heintzelman, D., and C. Tuttle. 2012. "Values in the wind: a hedonic analysis of wind power facilities." *Land Economics*, 88: 571–588.

Held, H., E. Kriegler, K. Lessmann, and O. Edenhofer. 2009. "Efficient climate policies under technology and climate ukelncertainty." *Energy Economics*, 31: S50–S61.

Hendryx, M., and M. Ahern. 2008. "Relations between health indicators and residential proximity to coal mining in West Virginia." *American Journal of Public Health*, 98: 669–671.

Hendryx, M., and M. Ahern. 2009. "Mortality in Appalachian coal mining regions: the value of statistical life lost." *International Journal of Environmental Studies*, 124: 541–550.

Hendryx, M., M. Ahern, and T. Nurkiewicz. 2007. "Hospitalization patterns associated with Appalachian coal mining." *Journal of Toxicology and Environmental Health*, 70: 2064–2068.

Intergovernmental Panel on Climate Change. 2007. "Climate Change 2007: The Physical Science Basis: Summary for Policy Makers, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change." IPCC Secretariat, World Meteorological Organization, Geneva, Switzerland. `http://www.ipcc.ch/ipccreports/ar4-wg1.htm`.

International Energy Agency. 2012. "Key World Energy Statistics." International Energy Agency, Paris, France.

Jenkins, R., E. Kopits, and D. Simpson. 2006. "Measuring the social benefits of EPA land cleanup and reuse programs." National Center for Environmental Economics, Washington, DC.

Katz, L., J. Kling, and J. Liebman. 2001. "Moving to opportunity in Boston: early results of a randomized mobility experiment." *Quarterly Journal of Economics*, 116: 607–654.

Keller, K., et al. 2005. "Avoiding Dangerous Anthropogenic Interference with the Climate System." *Climatic Change*, 73: 227–238.

Kelly, David L., and Charles D. Kolstad. 1999*a*. "Bayesian learning, pollution, and growth." *Journal of Economic Dynamics and Control*, 23(4): 491–518.

Kelly, David L., and Charles D. Kolstad. 2001. "Solving infinite horizon growth models with an environmental sector." *Journal of Computational Economics*, 18: 217–31.

Kelly, David L., and Zhuo Tan. 2015. "Learning and climate feedbacks: optimal climate insurance and fat tails." *Journal of Environmental Economics and Management*, 72: 98–122.

Kelly, D. L., and C. Kolstad. 1999*b*. "Integrated assessment models for climate change control." In *International Yearbook of Environmental and Resource Economics 1999/2000: A Survey of Current Issues*. , ed. T. Tietenberg and H. Folmer, Chapter 5, 171–97. Cheltenham, UK:Edward Elgar.

Keppo, I., B. O'Neill, and K. Riahi. 2007. "Probabilistic temperature change projections and energy system implications of greenhouse gas emissions scenarios." *Technological Forecasting and Social Change*, 74: 936–961.

Kiel, K., and M. Williams. 2007. "The impact of Superfund sites on local property values: Are all sites the same?" *Journal of Urban Economics*, 61: 170–192.

Knuckels, T., P. Stapleton, V. Minarchick, L. Esch, M. McCawley, M. Hendryx, and T. Nurkiewicz. 2012. "Air pollution particulate matter collected from an Appalachian mountaintop mining site induces microvascular dysfunction." *Microcirculation*, 20: 158–169.

Kohlhase, J. 1991. "The impact of toxic waste sites on housing values." *Journal of Urban Economics*, 30: 1–26.

Kuminoff, N., C. Parmeter, and J. Pope. 2010. "Which hedonic models can we trust to recover the marginal willingness to pay for environmental amenities." *Journal of Environmental Economics and Management*, 20: 145–160.

Kvale, K., et al. 2012. “Carbon dioxide emission pathways avoiding dangerous ocean impacts.” University of Victoria Working Paper.

Laumbach, R.J., and H.M. Kipen. 2012. “Respiratory health effects of air pollution: update on biomass smoke and traffic pollution.” *Journal of Allergy, Asthma, and Immunology*, 129: 3–11.

Leach, A. J. 2007. “The climate change learning curve.” *Journal of Economic Dynamics and Control*, 31: 1728–52.

Leggett, C., and N. Bockstael. 2000. “Evidence of the effects of water quality on residential land prices.” *Journal of Environmental Economics and Management*, 39: 121–144.

Lemoine, D., and I. Rudik. 2014. “Steering the climate system: using inertia to lower the cost of policy.” University of Arizona, Department of Economics 14-03.

Lemoine, D. M., and C. P. Traeger. 2014. “Watch your step: optimal policy in a tipping climate.” *American Economic Journal: Economic Policy*, 6(1): 1–31.

Lindberg, T., E. Bernhardt, R. Bier, A. Helton, R. Merola, A. Vengosh, and R. Di Giulio. 2011. “Cumulative impacts of mountaintop mining on an Appalachian watershed.” *Proceedings of the National Academy of Sciences*, 108: 20,929–20,934.

Linden, L., and J. Rockoff. 2008. “Estimates of the impact of crime risk on property values from Megan’s Laws.” *American Economic Review*, 98: 1103–1127.

Mastrandrea, M., et al. 2010. “Guidance note for lead authors of the IPCC Fifth Assessment Report on consistent treatment of uncertainties Intergovernmental Panel on Climate Change IPCC.” IPCC.

Mazzotta, M., L. Wainger, S. Sifleet, J. Petty, and B. Rashleigh. 2015. “Benefit transfer with limited data: an application to recreational fishing losses from surface mining.” *Ecological Economics*, 119: 384–398.

McAuley, S., and M. Kozar. 2006. “Ground-water quality in unmined areas and near reclaimed surface coal mines in the northern and central Appalachian coal regions, Pennsylvania and West Virginia.” USGS Scientific Investigations Report Report 20065059.

Muehlenbachs, L., E. Spiller, and C. Timmins. 2015. “The housing market impacts of shale gas development.” *American Economic Review*, 102: 3633–3659.

Nair, P.K., and B. Sinha. 1987. “Dust control consolidation in opencast mines- a new approach.” *Journal of Mines, Metals, and Fuels*, 35: 360–364.

Neidell, M. 2004. “Air pollution, health, and socio-economic status: the effect of outdoor air quality on childhood asthma.” *Journal of Health Economics*, 23: 1209–1236.

Neidell, M. 2008. “Information, avoidance behavior, and health: the effect of ozone on asthma hospitalizations.” NBER Working Paper 14209.

Neubersch, D., H. Held, and A. Otto. 2014. “Operationalizing climate targets under learning: An application of cost-risk analysis.” *Climatic Change*, 126: 305–18.

Nordhaus, W. 2007. “DICE Model Version as of May 22, 2007.” `http://www.econ.yale.edu/nordhaus/DICEGRAMS/DICE2007.htm`.

O’Neill, B., and M. Oppenheimer. 2002. “Climate change - dangerous climate impacts and the Kyoto Protocol.” *Science*, 296(5575): 1971–2.

Palmer, M., et al. 2010. “Mountaintop mining consequences.” *Science*, 327: 148–149.

Palmquist, R., F. Roka, and T. Vokina. 1997. “Hog operations, environmental effects, and residential property values.” *Land Economics*, 73: 114–124.

Papagiannis, A., D. Roussos, M. Menegaki, and D. Damigos. 2014. “Externalities from lignite mining-related dust emissions.” *Energy Policy*, 74: 414–424.

Parmeter, C., and J. Pope. 2011. “Handbook on Experimental Economics and the Environment.” , ed. J. List and M. Price, Chapter Quasi-experiments and hedonic property value methods, 69–71. Cheltenham, UK:Edward Elgar Publishing.

Pless-Mulloli, T., et al. 2000. “Living near opencast coal mining sites and children’s respiratory health.” *Occupational and Environmental Medicine*, 57: 145–151.

Pond, G., M. Passmore, F. Borsuk, L. Reynolds, and C. Rose. 2008. “Downstream effects of mountaintop coal mining: comparing biological conditions using family- and genus-level macroinvertebrate bioassessment tools.” *Journal of the North American Benthological Society*, 27: 717–737.

Pope, J. 2008. “Buyer information and the hedonic: the impact of a seller disclosure on the simplicity price for airport noise.” *Journal of Urban Economics*, 63: 498–516.

Randall, A., O. Grunewald, S. Johnson, R. Ausness, and A. Pagoulatos. 1978. “Reclaiming coal surface mines in central Appalachia: a case study of the benefits and costs.” *Land Economics*, 54: 472–489.

Richels, R.G., A.S. Manne, and T.M.L. Wigley. 2004. “Moving beyond concentrations: the challenge of limiting temperature change.” AEI-Brookings Joint Center for Regulatory Studies 04-11. `http://ssrn.com/abstract=545742`.

Rincón-Zapatero, J. P., and Manuel S. Santos. 2009. "Differentiability of the value function without interiority assumptions." *Journal of Economic Theory*, 144(5): 19481964.

Roe, G.H., and M.B. Baker. 2007. "Why is climate sensitivity so unpredictable?" *Science*, 318(5850): 629–32.

Rosen, S. 1974. "Hedonic prices and implicit markets: product differentiation in pure competition." *Journal of Political Economy*, 82: 34–55.

Rudik, I. 2014. "Targets, taxes, and learning: optimizing climate policy under knightian damages." University of Arizona, Department of Economics, `http://dx.doi.org/10.2139/ssrn.2516632`.

Sanford, T., P. Frumhoff, A. Luers, and Jay Gulledge. 2014. "The climate policy narrative for a dangerously warming world." *Nature Climate Change*, 4: 164–166.

Schlenker, W., and R. Walker. 2011. "Airports, air pollution, and contemporaneous health." NBER Working Paper 17684.

Schubert, R., et al. 2006. "The Future Oceans - Warming Up, Rising High, Turning Sour." German Advisory Council on Global Change. Special report, `http://www.wbgu.de`.

Schwartz, J., and A. Marcus. 1990. "Mortality and air pollution in London: a time series analysis." *American Journal of Epidemiology*, 131: 185–194.

Siami-Irdemoosa, E., and S. Dindarloo. 2015. "Prediction of fuel consumption of mine dump trucks: a neural networks approach." *Applied Energy*, 151: 77–84.

Stocker, T., Q. Dahe, and G. Plattner. 2013. "Technical Summary." In *Climate Change 2013: The Physical Science Basis*. 1–79. IPCC Secretariat.

Tiebout, C. 1956. "A pure theory of local expenditures." *The Journal of Political Economy*, 64: 416–424.

Traeger, C. 2014. "A 4-stated DICE: quantitatively addressing uncertainty effects in climate change." *Environmental and Resource Economics*, 59: 1–37.

Weitzman, M. 2009. "On modeling and interpreting the economics of catastrophic climate change." *Review of Economics and Statistics*, 91: 1–19.

Williams, A.M. 2011. "The impact of surface coal mining on residential property values: a hedonic price analysis." Tennessee Research and Creative Commons 5.

Williamson, J., H. Thurston, and M. Heberling. 2008. "Valuing acid mine drainage remediation in West Virginia: a hedonic modeling approach." *Annals of Regional Science*, 42: 987–999.