

Climate Policy and Innovation: A Quantitative Macroeconomic Analysis

Stephie Fried[†]

Carleton College

December 2, 2016

Abstract

A carbon tax can induce innovation in green technologies. I evaluate the quantitative impact of this channel in a dynamic, general equilibrium model with endogenous innovation in fossil and green energy inputs. I discipline the parameters using evidence from historical oil shocks, in which both energy prices and energy innovation increased substantially. I find that a carbon tax induces large movements in innovation that have considerable effects on energy prices, production, and other macroeconomic aggregates. Moreover, analyses that omit endogenous innovation result in a substantial overestimation of the carbon tax necessary to attain a given reduction in emissions.

[†]Email: sfried@carleton.edu

I thank seminar participants at UAlbany, Arizona State University, Cal Poly San Luis Obispo, Carleton College, University of California San Diego, and University of Stockholm IIES for their helpful insights. In particular, I thank Thomas Baranga, Betty Daniel, Jim Hamilton, Mark Jacobsen, David Lagakos, Kevin Novan, William Peterman, Valerie Ramey, Irina Telyukova, and three anonymous referees for their many suggestions.

1 Introduction

A carbon tax can induce innovation in green technologies. Over time, these technological advances lower the cost of reducing carbon emissions. However, how much innovation responds and the magnitudes of the accompanying effects on energy prices, production, and carbon emissions remain open questions. A quantitative understanding of the consequences of endogenous innovation is important, since government agencies often evaluate climate mitigation projects based partly on climate-economy models which abstract from endogenous innovation.¹

Much of the macroeconomic literature studying climate policy assumes that innovation is exogenous (e.g., Golosov et al. (2014); Krusell and Smith (2009); Hassler and Krusell (2012); and Nordhaus (2008)) while much of the environmental literature has concentrated on endogenous innovation in partial equilibrium (see Popp et al. (2009) for an overview). This paper combines these two approaches by studying a carbon tax in a general equilibrium model with endogenous innovation. I use this model to analyze the dynamic effects of a carbon tax and to quantify the importance of endogenous innovation for climate policy outcomes. I find that the carbon tax induces large movements in innovation that have considerable effects on energy prices, production, and other macroeconomic aggregates. Moreover, abstracting from endogenous innovation –and modeling technological progress as exogenous– results in a substantial overestimation of the carbon tax necessary to attain a given reduction in emissions.

The central contribution of this paper is to quantify the interaction between endogenous innovation and climate policy in a dynamic, general equilibrium framework that explicitly models innovation in fossil energy, green energy, and non-energy sectors. The model builds on the macroeconomic literature on directed technical change and climate (e.g., Acemoglu, Aghion, Burzysten, and Hemous (2012) (AABH), Acemoglu, Akcigit, Hanley, and Kerr (2014); Hart (2012); Hassler, Krusell and Olovsson (2012); Hemous (2014); Smulders and de Nooij (2003). For an overview, see Heutel and Fischer (2013)). This earlier work is mainly theoretical, and the models are generally not designed for quantitative analysis.²

¹For example, the social cost of carbon that the EPA uses to evaluate climate policies is, in part, based on climate-economy models which do not incorporate endogenous innovation.

²For example, AABH state that their “objective is not to provide a comprehensive quantitative evaluation” (AABH, p. 154). One exception is Acemoglu et al. (2014), which is a quantitative paper focused on the

I complement this earlier work by quantifying the importance of endogenous innovation for climate policy outcomes. In many of the existing models, such as AABH, innovation occurs in only one energy sector (i.e., fossil or green) on the long-run balanced growth path. However, US data on fossil and green innovation show positive and substantial amounts of innovation in both of these sectors since the 1970s. To match this empirical fact, I incorporate technology spillovers across the different sectors. The spillovers imply that technology developed for one sector increases the productivity of innovation in the other sectors. One example of these spillovers between the fossil and green energy sectors is that the first mass commercialization of solar cells was driven by demand from oil companies to power the lights on their offshore rigs (Perlin (2000)). If spillovers such as these are sufficiently strong, then the balanced growth path is an interior solution in which innovation occurs in both the fossil and green energy sectors.³

I develop a novel calibration strategy using the energy price increases triggered by oil shocks and the accompanying changes in energy production and innovation. It is important for the model to capture the empirical relationships among energy prices, production, and innovation. These are key links because many climate policies, including a carbon tax and a cap and trade system, create incentives to reduce fossil energy consumption through changes in energy prices. The oil shocks provide empirical evidence of the response of energy innovation and production to an aggregate increase in the energy price. This variation is particularly useful for disciplining the parameter values since economy-wide historical examples of climate policies are scarce.

I perform two exercises to fully explore the interactions between endogenous innovation and climate policy. First, to evaluate the dynamic effects of climate policy with endogenous innovation, I introduce a constant carbon tax into my benchmark model with endogenous innovation. I compare the movements in technology, relative prices, and other macroeconomic aggregates in the model with the tax to their values in the model without the tax. Next, to quantify the importance of endogenous innovation for climate policy evaluation, I introduce a carbon tax into an alternative model with the endogenous innovation channel shut down. I refer to this model as the exogenous-innovation model because innovation cannot respond

relative roles of carbon taxes and subsidies to green energy research in the structure of optimal climate policy.

³Acemoglu (2002) and Hart (2012) show that the strength of cross-sector technology spillovers can determine stability of an interior long-run balanced growth path in models of directed technical change.

to the tax. Comparing the effects of the tax in the endogenous- and exogenous-innovation models allows me to quantify the interaction between endogenous innovation and climate policy. In both models, I choose the size of the carbon tax to achieve a 30-percent reduction in emissions by 2030, one version of the emissions target that the US government set with the announcement of the Clean Power Plan.⁴

There are two main findings. First, comparing the endogenous-innovation model with and without the tax, I find that the tax induces substantial movements in innovation, energy prices, and other macroeconomic aggregates. For example, by 2030, the tax causes green innovation to be 50 percent higher and fossil innovation to be 60 percent lower than what they would have been without the tax. These movements in innovation are accompanied by substantial changes in relative prices. In the model with the tax, the relative price of green compared to fossil energy is 7 percent lower in 2030 and 17 percent lower on the new balanced growth path than in the model without the tax. There is little change in non-energy innovation in response to the tax, suggesting that the increased green innovation crowds out fossil energy innovation instead of non-energy innovation.

Second, comparing the results from the tax in the exogenous- and endogenous-innovation models, I find that endogenous innovation has substantial implications for the effectiveness of the carbon tax and for the relative price of green energy. The carbon taxes required to achieve the Power Plan target in the exogenous- and endogenous-innovation models are 30.3 and 24.5 in 2013 dollars per ton of CO₂, respectively. Endogenous innovation reduces the carbon tax by 19.2 percent because it increases incentives for carbon abatement. The intuition for this result is that regardless of whether innovation is endogenous, the carbon tax operates through prices to shift demand from fossil to green energy, reducing emissions. However, when innovation is endogenous, this shift in demand spurs green innovation. Over time, the increase in green innovation reduces the marginal cost of producing green energy, lowering its equilibrium price and creating stronger incentives for agents to switch from fossil

⁴The primary focus of this paper is on quantifying the dynamic relationship between a carbon tax and fossil, green, and non-energy innovation. I choose to quantify these mechanisms for a realistic climate policy and not for the optimal policy. Specifically, I analyze a carbon tax that achieves the US Clean Power Plan's emissions targets. I choose this route both because it is more realistic in the current environment and also because calculating optimal policy requires additional assumptions which can make the underlying mechanisms governing the relationship between endogenous innovation and climate policy less transparent. For example, a rigorous calculation of the optimal policy requires a damage function, a realistic depiction of the carbon cycle, a plausible time frame for the analysis, a reasonable rate of time preference, and assumptions about carbon emissions from other, non-modeled, countries.

to green. Thus, endogenous innovation amplifies the price incentives created by the carbon tax, implying that the emissions target can be achieved with a 19.2 percent smaller tax.

The standard equivalence between carbon taxes and carbon permits holds in this model. An alternative interpretation of these results is that endogenous innovation reduces the predicted carbon permit price from an equivalent cap and trade system by 19.2 percent. This interpretation is consistent with the United States' experience using tradable permits to reduce acid rain in the 1990s. Initial forecasts of the permit price were orders of magnitude higher than the realized prices, partly because of technological advances in low-sulfur coal mining, fuel mixing, and scrubber installation and performance (Sandor et al. (2014)).

Additionally, I find that endogenous innovation has offsetting effects on the gross welfare costs of attaining a given abatement target.⁵ The carbon tax is smaller when innovation is endogenous, and, hence, the accompanying gross distortionary cost is smaller. However, the shift in innovation from fossil to green energy in response to the tax reduces the aggregate growth rate along the transition path to a new long-run equilibrium, raising the gross welfare cost of the policy. As a result, the overall effect of endogenous innovation on the gross welfare costs of the carbon tax is small. In particular, the consumption equivalent variation (CEV) of the tax is -0.5 percent in the endogenous-innovation model and -0.6 percent in the exogenous-innovation model.

There is a substantial environmental literature on the effects of endogenous innovation in integrated assessment climate-economy models (see for example, Grubb et al. (1994); Goulder and Schneider (1999); van der Zwaan et al. (2002); Nordhaus (2002); Popp (2004); Gerlagh (2008); again Popp et al. (2009) provides a nice overview). Of these earlier models, the most closely related to the present paper are Goulder and Schneider (1999), Popp (2004), and Gerlagh (2008).

In their seminal paper, Goulder and Schneider develop both analytical and numerical climate-economy models with endogenous innovation. While their models are largely qualitative, they find that the inclusion of endogenous innovation increases the amount of abatement from a given sized carbon tax, consistent with the present paper.

Popp (2004) modifies the DICE model of climate change (Nordhaus and Boyer (2000)) to

⁵The carbon tax is designed to correct the externality from carbon emissions. Therefore, relative to the social optimum, the carbon tax should not generate first-order distortionary or welfare costs net of the climate damages. Here, I use the term “gross” to denote changes relative to the baseline outcome which does not include the consequences of climate change.

include endogenous innovation in a single energy sector and quantifies its effects on climate policy outcomes. Relative to the present paper, Popp (2004) finds that including endogenous energy innovation has very small implications for the size of the carbon tax necessary to achieve a given emissions target, but considerably larger effects on welfare.⁶ Gerlagh (2008) develops a model which allows for endogenous innovation in multiple sectors. Relative to the present paper, he finds larger effects of endogenous innovation on the size of the carbon tax necessary to achieve a given emissions target.⁷

These differences largely arise because the models in the earlier work are more complex in some ways but are reduced form in other ways. For example, Popp (2004) focuses on the effects of innovation on the optimal time path of the carbon tax. Solving for this optimal time path comes at the expense of incorporating general equilibrium features such as the production of fossil energy, endogenous energy prices, a green energy sector, or endogenous innovation in areas besides improvements in fossil energy efficiency. Similarly, Gerlagh (2008) does not directly model a green energy sector and his calibration procedure assumes that the economy is in steady state with respect to energy (and other variables) from 1970-1990. Such an assumption is at odds with the data on energy prices and innovation and Gerlagh stresses that a more robust calibration procedure is essential for future work.

Unlike much of the previous environmental literature, the present paper specifically models the general equilibrium effects from endogenous innovation in each of two energy sectors, a green and a fossil sector, and in a third sector comprising the rest of the economy. These features influence the effects of endogenous innovation on climate policy outcomes along three important dimensions.

First, the potential for innovation in fossil, green, and non-energy sectors is important for obtaining a plausible calibration that applies to the whole economy. This three-sector design facilitates a direct mapping to the data on fossil, green, and non-energy R&D which can be difficult to obtain otherwise. Moreover, the model framework allows for the distinction between the innovation incentives offered by the carbon tax versus those offered by higher

⁶Specifically, Popp (2004) compares the size of the carbon tax necessary to reduce emissions to their 1995 levels in his model with and without endogenous energy innovation. The difference in the size of the tax is less than 1 percent. Additionally, he finds that including endogenous energy innovation increases welfare under the optimal policy by 9.4 percent.

⁷Specifically, Gerlagh (2008) finds that endogenous innovation reduces the size of the carbon tax necessary to achieve a given emissions target by a factor of two.

energy prices due to non-tax induced changes in energy supply or demand. This distinction is crucial for obtaining a realistic calibration based on historical data in which higher energy prices occurred because of supply or demand changes instead of from a carbon tax.

Second, the general equilibrium, three-sector framework fully endogenizes the relative price of green to fossil energy. This relative price is the primary determinant of firms' energy choices and, hence, of aggregate emissions. The relative price depends on the levels of innovation in both the fossil and the green energy sectors. If the policy causes green innovation to increase above its baseline level, then the marginal cost of producing green energy falls, reducing the relative price of green to fossil energy. If the policy also causes fossil innovation to decrease below its baseline level, then the marginal cost of producing fossil energy rises (relative to the baseline), causing the relative price of green to fossil energy to fall further. The quantitative impact of the reduced fossil energy innovation on this relative price is considerable and clearly captured within this three-sector, directed technical change framework.

Third, the three sectors imply that increased green innovation can crowd out fossil innovation and/or non-energy innovation.⁸ These two dimensions for crowd-out have substantially different implications for both the effectiveness and gross welfare cost of the carbon tax. Increased green innovation at the expense of fossil innovation amplifies the impact of green innovation on the relative price of green to fossil energy, increasing the emissions reduction from the carbon tax. In contrast, increased green innovation at the expense of non-energy innovation could result in a larger reduction in economic growth, amplifying the gross welfare costs of the policy.⁹

The paper proceeds as follows: Sections 2 and 3 describe the model. Section 4 discusses the oil shocks and the calibration strategy. Sections 5 and 6 present the results and robustness analysis, respectively. Section 7 concludes.

⁸It is also possible that the carbon tax will increase the overall level of innovation in the economy, and, thus, neither fossil nor non-energy innovation will be crowded out by the tax. However, I do not find evidence to support this outcome; the relative return to supplying labor as an innovator (scientist) compared to as a worker is almost unchanged by the carbon tax, suggesting that the tax would not lead to a substantial change in the aggregate level of innovation in the economy.

⁹While both types of crowd-out are possible, I find that green energy innovation almost exclusively crowds-out fossil innovation, and not non-energy innovation.

2 Model

I adapt the standard directed technical change framework (Acemoglu (2002)) to a setting with fossil, green, and non-energy intermediate inputs and oil price shocks. Fossil energy refers to energy derived from coal, oil, or natural gas. Green energy refers to energy derived from any non-carbon energy source. This category includes renewable energy, such as wind and solar, as well as nuclear energy and energy savings from improved fossil energy efficiency, such as better insulation, higher fuel economy, etc.

2.1 Final good

The unique final consumption good, Y , is produced competitively from energy, E , and non-energy inputs, N , according to the CES production function

$$Y_t = \left(\delta_y E_t^{\frac{\varepsilon_y - 1}{\varepsilon_y}} + (1 - \delta_y) N_t^{\frac{\varepsilon_y - 1}{\varepsilon_y}} \right)^{\frac{\varepsilon_y}{\varepsilon_y - 1}}, \quad (1)$$

where $\varepsilon_y < 1$ is the elasticity of substitution between the energy and non-energy inputs. Energy is a nested CES composite of fossil energy, green energy, and oil imports,

$$E_t = \left(\tilde{F}_t^{\frac{\varepsilon_e - 1}{\varepsilon_e}} + G_t^{\frac{\varepsilon_e - 1}{\varepsilon_e}} \right)^{\frac{\varepsilon_e}{\varepsilon_e - 1}} \quad \text{and} \quad \tilde{F}_t = \left(\delta_{\tilde{F}} F_t^{\frac{\varepsilon_f - 1}{\varepsilon_f}} + (1 - \delta_{\tilde{F}}) (O_t^*)^{\frac{\varepsilon_f - 1}{\varepsilon_f}} \right)^{\frac{\varepsilon_f}{\varepsilon_f - 1}}. \quad (2)$$

Parameter $1 < \varepsilon_f < \infty$ denotes the elasticity of substitution between fossil energy (produced domestically), F , and oil imports, O^* . Since fossil energy is a mixture of coal, oil, and natural gas, oil imports and fossil energy are not perfect substitutes. Parameter $\varepsilon_e > 1$ is the elasticity of substitution between green energy, G , and the CES composite comprised of fossil energy and oil imports.¹⁰ The final good is the numeraire.

¹⁰Following AABH and Hemous (2014), I do not include a distribution parameter between green energy, G , and the CES composite comprised of fossil energy and oil imports, \tilde{F} . Differences in the quantities of \tilde{F} and G result exclusively from differences in their relative prices and not from an underlying asymmetry in the production function. Both \tilde{F} and G contribute equally at the margin to the energy composite, E , when relative prices are the same. For example, a boiler that uses one less ton of coal (higher G) is equivalent to additional coal (higher \tilde{F}). However, the finite elasticity of substitution implies that there is some heterogeneity in the production process, so agents do not substitute indefinitely into either \tilde{F} or G .

2.2 Fossil, green, and non-energy intermediate inputs

Fossil, green, and non-energy intermediate inputs are produced competitively and sold at market prices to the final good producer.¹¹ The production functions are constant returns to scale in labor, L_j , and a unit mass of sector-specific machines, each indexed by i , X_{ji} where $j \in \{f, g, n\}$,

$$\begin{aligned} F_t &= L_{ft}^{1-\alpha_f} \int_0^1 X_{fit}^{\alpha_f} A_{fit}^{1-\alpha_f} di, & G_t &= L_{gt}^{1-\alpha_g} \int_0^1 X_{git}^{\alpha_g} A_{git}^{1-\alpha_g} di \\ N_t &= L_{nt}^{1-\alpha_n} \int_0^1 X_{nit}^{\alpha_n} A_{nit}^{1-\alpha_n} di. \end{aligned} \quad (3)$$

Variable A_{ji} denotes the factor-augmenting technology embodied in machine X_{ji} , and α_j is the factor share of machines in sector J . A representative intermediate-goods producer chooses machines and labor to maximize profits, taking prices as given. Labor market clearing requires that $L_{ft} + L_{gt} + L_{nt} \leq L$, where L is the fixed exogenous supply of workers in the economy.

2.3 Machines

There is a unit mass of machine producers in each of the three sectors. The machine producers sell their machines to the intermediate-goods producers in their specific sectors. Each machine embodies technology. A machine producer can hire scientists to innovate on the embodied technology. A machine costs one unit of the final good to produce, regardless of the sector or the level of technology embodied in the machine. The market for scientists is competitive, and the machine producer must pay the scientists he hires the market wage, w_{sj} , where $j \in \{f, g, n\}$. However, the market for machines is monopolistically competitive, and the machine producers earn positive profits from the sale of their machines.

¹¹This model of fossil energy production abstracts from resource scarcity. While the supplies of fossil energy are finite, historically, fossil energy prices have not followed the predictions from the standard Hotelling model of an exhaustible resource (Hamilton (2009)). Moreover, given the presence of climate change, scarcity constraints on fossil energy extraction are less likely to bind. For example, the IEA estimates that if the world is to remain below the two degree target, then no more than one third of the proven reserves of fossil energy can be consumed prior to 2050 (World Energy Outlook (2012)).

The evolution of technology for machine producer i in each of the sectors F , G , and N is

$$\begin{aligned} A_{fit} &= A_{ft-1} \left(1 + \gamma \left(\frac{S_{fit}}{\rho_f} \right)^\eta \left(\frac{A_{t-1}}{A_{ft-1}} \right)^\phi \right), & A_{git} &= A_{gt-1} \left(1 + \gamma \left(\frac{S_{git}}{\rho_g} \right)^\eta \left(\frac{A_{t-1}}{A_{gt-1}} \right)^\phi \right), \\ A_{nit} &= A_{nt-1} \left(1 + \gamma \left(\frac{S_{nit}}{\rho_n} \right)^\eta \left(\frac{A_{t-1}}{A_{nt-1}} \right)^\phi \right), \end{aligned} \quad (4)$$

where S_{ji} denotes the number of scientists working for machine producer i in sector $j \in \{f, g, n\}$. Scientists affect the growth rate of the machine producer's technology. Hence, there is path dependence in innovation; higher existing technology in a sector increases the marginal product of research in that sector.¹²

Parameter $\eta \in (0, 1)$ implies that there are diminishing returns to scientific research within a given period. This modeling choice captures the “stepping on toes” effect discussed in the endogenous growth literature, where scientists are more likely to duplicate discoveries within a given period (Jones and Williams (1998)). Parameter γ measures the efficiency with which scientists produce new ideas.

Parameters (ρ_f, ρ_g, ρ_n) adjust for differences in sector diversity. Specifically, ρ_f is the number of processes on which a scientist can innovate in fossil energy. Fossil energy scientists divide their time equally among all available processes (and likewise for green and non-energy scientists). Accounting for differences in sector diversity is particularly important because there are diminishing returns to innovation in each sector. Without a diversity adjustment, the marginal product of a non-energy scientist is much lower than that of an energy scientist simply because there are more non-energy scientists.

Variable A_j , $j \in \{f, g, n\}$ denotes the aggregate (average) level of technology in sector J :

$$A_{ft} = \int_0^1 A_{fit} di, \quad A_{gt} = \int_0^1 A_{git} di, \quad A_{nt} = \int_0^1 A_{nit} di. \quad (5)$$

I define aggregate technology (A) as the average of the technologies in each sector weighted

¹²The main differences between the specification in equation (4) and the specification used in AABH are the TFP catchup term $\left(\frac{A_{t-1}}{A_{gt-1}} \right)^\phi$ and the diminishing returns to innovation, η .

by the number of processes:

$$A_t = \frac{\rho_f A_{ft} + \rho_g A_{gt} + \rho_n A_{nt}}{\rho_f + \rho_g + \rho_n}. \quad (6)$$

The TFP catchup ratios $\left(\frac{A_{t-1}}{A_{ft-1}}\right)^\phi$, $\left(\frac{A_{t-1}}{A_{gt-1}}\right)^\phi$, and $\left(\frac{A_{t-1}}{A_{nt-1}}\right)^\phi$ incorporate technology spillovers across the different sectors. Parameter $\phi \in (0, 1)$ determines the strength of these spillovers. The intuition for the spillovers and their implications for the long-run behavior of the model are discussed in Section 3.

In addition to the across-sector technology spillovers, the technology accumulation process also incorporates technology spillovers within a sector after one period. The technology of machine producer i in sector J tomorrow depends on the level of knowledge in sector J today and on any new ideas that machine producer i accumulates from hiring scientists. Hence, a given machine producer's discoveries are secret for one period. After the period is over, other machine producers in his sector observe his discoveries and can incorporate them into their own innovation processes. This modeling choice is empirically reasonable, provided that the period is sufficiently long and is in line with similar assumptions made in the literature (e.g., AABH; Hemous (2014)). I discuss evidence of these within-sector spillovers in fossil and green energy and an appropriate period length in Appendix B.

Each machine producer chooses the quantity of machines, the machine price, and the number of scientists, to his maximize profits. He takes the existing levels of technology as given. Scientist market clearing requires that $S_{ft} + S_{gt} + S_{nt} \leq S$, where S is the fixed exogenous supply of scientists in the economy and S_{jt} is the number of scientists in sector J .

2.4 An oil shock and a carbon tax

Carbon emissions, \mathcal{E} , accumulate from the use of fossil energy and oil imports,

$$\mathcal{E}_t = \omega_f F_t + \omega_o O_t^*.$$

Parameters ω_f and ω_o convert fossil energy and oil imports into carbon emissions.

The supply of oil imports is perfectly elastic at exogenous price, P_{ot}^* . An oil shock is an

exogenous increase in P_{ot}^* . I choose to model the price of oil imports as exogenous because this is a simple way to model the oil price shocks, which I use for calibration. Note that all other prices are endogenous and respond through general equilibrium channels to oil price shocks as well as to other economic shocks (such as the introduction of a carbon tax).

A carbon tax places a price on the externality, carbon. Thus, the tax, τ , is a tax per unit of carbon consumed, which is independent of the price. The tax increases the price of fossil energy from P_{ft} to $P_{ft} + \tau_f$ and the price of oil imports from P_{ot}^* to $P_{ot}^* + \tau_o$, where $\tau_f = \tau \times (\text{carbon content of fossil energy})$ and $\tau_o = \tau \times (\text{carbon content of oil imports})$.

2.5 Household

The representative household is inhabited by a unit mass of machine producers in each sector, L workers, and S scientists. The relative supplies of workers and scientists are fixed.¹³ Additionally, I assume that both workers and scientists are mobile across sectors so that they can switch sectors without incurring adjustment costs. Zero adjustment costs are reasonable provided the time period is sufficiently long. Such an assumption can be further justified with a sufficiently broad view of scientists and workers. For example, the skills of a chemist (scientist) and a construction worker (worker) are needed in all three sectors, suggesting that these types of scientists and workers would not incur substantial adjustment costs from switching sectors in the long run.

The utility function is $U(C) = \frac{C^{1-\theta}}{1-\theta}$, where $\frac{1}{\theta}$ is the intertemporal elasticity of substitution. There is no mechanism through which the household can save, and, thus, it consumes its income.¹⁴ The budget constraint is

$$C_t = w_{lft}L_{ft} + w_{lgt}L_{gt} + w_{lnt}L_{nt} + w_{sft}S_{ft} + w_{sgt}S_{gt} + w_{snt}S_{nt} + \int_0^1 (\pi_{fit} + \pi_{git} + \pi_{nit}) di + T_t. \quad (7)$$

¹³This assumption implies that the overall level of innovation in the economy cannot change in response to the carbon tax. However, as discussed in Section 5, the relative wage of a scientist compared to a worker is almost unchanged by the carbon tax, suggesting that the tax would not lead to substantial changes in the relative numbers of workers and scientists.

¹⁴This is a standard simplification in the directed technical literature, which vastly simplifies the solution (e.g., Acemoglu (2002); AABH). Empirically, consumption volatility is relatively small so abstracting from this volatility by not allowing agents to smooth is unlikely have huge welfare implications (e.g., Lucas (2003)).

Variable π_{ji} denotes profits to machine producer i in sector $j \in \{f, g, n\}$, and T_t is lump sum transfers from a carbon tax.

The aggregate resource constraint implies that the final good can be consumed, converted to machines, or used to purchase oil imports:

$$Y_t = C_t + \int_0^1 (X_{fit} + X_{git} + X_{nit}) di + P_{ot}^* O_t^*. \quad (8)$$

2.6 Equilibrium

A *decentralized equilibrium* consists of sequences of wages $(w_{lft}, w_{lgt}, w_{lnt}, w_{sft}, w_{sgt}, w_{snt})$, prices for machines $(P_{fit}^x, P_{git}^x, P_{nit}^x)$, prices for intermediates (P_{ft}, P_{gt}, P_{nt}) , demands for machines $(X_{fit}^d, X_{git}^d, X_{nit}^d)$, demands for intermediates (F_t^d, G_t^d, N_t^d) , demands for labor $(L_{ft}^d, L_{gt}^d, L_{nt}^d)$, demands for scientists $(S_{ft}^d, S_{gt}^d, S_{nt}^d)$, supplies of machines $(X_{fit}^s, X_{git}^s, X_{nit}^s)$, supplies of intermediates, (F_t^s, G_t^s, N_t^s) , supplies of labor $(L_{ft}^s, L_{gt}^s, L_{nt}^s)$, and supplies of scientists $(S_{ft}^s, S_{gt}^s, S_{nt}^s)$ such that given a sequence of oil import prices (P_{ot}^*) ,

1. Agents optimize: $(P_{jit}^x, S_{jit}^d, X_{jit}^s)$ maximize the machine producers' profits, $j \in \{f, g, n\}$; $(X_{fit}^d, X_{git}^d, X_{nit}^d, L_{ft}^d, L_{gt}^d, L_{nt}^d)$ maximize intermediate-goods producers' profits; $(F_t^d, G_t^d, N_t^d, (O_t^*)^d)$ maximize final-good producer's profits; $(L_{ft}^s, L_{gt}^s, L_{nt}^s, S_{ft}^s, S_{gt}^s, S_{nt}^s)$ maximize the household's utility.
2. Markets clear: $(P_{fit}^x, P_{git}^x, P_{nit}^x)$ clear the machine producer markets; (P_{ft}, P_{gt}, P_{nt}) clear the intermediate input markets; $(w_{lft}, w_{lgt}, w_{lnt}, w_{sft}, w_{sgt}, w_{snt})$ clear the labor and scientist markets.

3 Discussion

The model is designed to endogenize the innovation response to energy price increases triggered by carbon taxes and oil shocks. Both oil shocks and carbon taxes enter the model through the final-good producer's demand for energy inputs. The optimization problem of the representative final-good producer is

$$\max_{F_t, G_t, N_t, O_t^*} \{Y_t - (P_{ft} + \tau_f)F_t - P_{gt}G_t - (P_{ot}^* + \tau_o)O_t^* - P_{nt}N_t\} \quad (9)$$

subject to the production technology defined in equations (1) and (2). The crucial difference between a carbon tax and an oil shock is that the direct effect of the carbon tax is to increase both the fossil energy and the oil import prices, while the direct effect of the oil shock is to increase only the oil import price. Thus, the carbon tax increases only the demand for green energy while an oil shock increases the demand for both fossil and green energy.¹⁵

The first-order conditions for the machine producer imply the wages to scientists in each sector are given by (see Appendix A for the full derivation):

$$\begin{aligned}
 w_{sft} &= \frac{\eta\gamma\alpha_f A_{ft-1} \left(\frac{S_{ft}}{\rho_f}\right)^\eta \left(\frac{A_{t-1}}{A_{ft-1}}\right)^\phi P_{ft} F_t}{\left(\frac{1}{1-\alpha_f}\right) S_{ft} A_{ft}}, & w_{sgt} &= \frac{\eta\gamma\alpha_g A_{gt-1} \left(\frac{S_{gt}}{\rho_g}\right)^\eta \left(\frac{A_{t-1}}{A_{gt-1}}\right)^\phi P_{gt} G_t}{\left(\frac{1}{1-\alpha_g}\right) S_{gt} A_{gt}} \\
 w_{snt} &= \frac{\eta\gamma\alpha_n A_{nt-1} \left(\frac{S_{nt}}{\rho_n}\right)^\eta \left(\frac{A_{t-1}}{A_{nt-1}}\right)^\phi P_{nt} N_t}{\left(\frac{1}{1-\alpha_n}\right) S_{nt} A_{nt}} & & (10)
 \end{aligned}$$

Since the market for scientists is perfectly competitive, the wage of a scientist in a given sector equals the marginal return to innovation in that sector. Thus, equation (10) shows that the marginal return to fossil innovation is increasing in the value of fossil energy production, $P_{ft}F_t$. This relationship implies that the product of price and quantity, $P_{ft}F_t$, as opposed to each individual component, (P_{ft} , and F_t) is what matters for innovation incentives, and thus, it is important that the calibrated model match the product as opposed to a single component for the quantitative analysis. Similarly, the marginal return to green innovation is increasing in the value of green energy production, $P_{gt}G_t$.

Oil shocks and carbon taxes have opposite effects on fossil energy innovation incentives. An oil shock increases fossil energy demand, raising the equilibrium value of $P_{ft}F_t$ and the accompanying innovation incentives. A carbon tax decreases fossil energy demand, reducing the equilibrium value of $P_{ft}F_t$, and the accompanying innovation incentives. In contrast to their opposite effects on fossil innovation, both oil shocks and carbon taxes increase green innovation incentives. Each of these shocks increases demand for green energy, raising the equilibrium value of $P_{gt}G_t$ and the accompanying innovation incentives.¹⁶

¹⁵Importantly, while the direct effect of the oil shock is the increase the price of oil imports, the resulting increase in fossil energy demand raises the equilibrium fossil energy price as well, consistent with the empirical observations of oil price shocks and domestic fossil energy prices.

¹⁶This model abstracts from leakage from the carbon tax; fossil energy producers are not able to sell their

The technology accumulation process incorporates both path dependence and across-sector technology spillovers. These two drivers of innovation are captured in the marginal return by the term, A_{ft-1} and the catchup ratio, $\left(\frac{A_{t-1}}{A_{ft-1}}\right)^\phi$. Path dependence implies that higher existing technology increases the marginal return to innovation. However, the catchup effect raises the marginal productivity of innovation in sectors that are behind the frontier. The intuition is that if sector J is relatively backward, then there are many ideas from other sectors that have not yet been applied in sector J . This “low-hanging fruit” increases the productivity of research in sector J .

Parameter ϕ measures the strength of the productivity catchup effect. If $\phi = 0$, there are no across-sector spillovers and there is full path dependence, as in AABH and Hemous (2014). Since fossil and green energy are gross substitutes (i.e., $\varepsilon_e > 1$), this strong path dependence implies that innovation in one energy sector raises the relative marginal product of innovation in that sector by so much that the only stable balanced growth paths are corner solutions in which innovation occurs in only one form of energy.¹⁷ In contrast, if $\phi = 1$, the marginal return to innovation in a sector is independent of the previous level of technology in that sector, and, hence, there is no path dependence. In this case, there exists a stable interior balanced growth path in which innovation occurs in both forms of energy. The value of ϕ determines the relative strengths of the path dependence and across-sector spillover channels and, thus, governs the stability of the interior balanced growth path.¹⁸

4 Calibration

I begin with a discussion of the time period. Next I describe the data I use to calibrate the model parameters. Finally, following standard procedure (e.g., Gourinchas and Parker (2002)), I calibrate the production and innovation components of the model in two steps. In the first step, I calibrate a group of parameters directly from the data series. In the second

energy abroad. If fossil energy producers were allowed to sell their energy abroad and they did not face a similar carbon tax abroad, then the carbon tax would lead to a smaller change in domestic fossil energy demand.

¹⁷Innovation will also occur in the non-energy sector since the non-energy and energy sectors are gross complements. The diminishing returns to innovation imply that the corner solution balanced growth paths only exist asymptotically.

¹⁸Acemoglu (2002) and Hart (2012) show that the strength of cross-sector technology spillovers can determine stability of an interior long-run balanced growth path in models of directed technical change.

step, I use historical oil price shocks and the accompanying data on energy production and innovation to jointly calibrate the remaining parameters. A growing empirical literature that finds a causal relationship between a change in energy prices and energy innovation supports this approach.¹⁹

4.1 Time period

The time period in the model is five years. This choice implies that technology spillovers within a sector occur in five years. To determine this time period, I examine the rate of technology spillovers experienced in a green and in a fossil industry. In particular, I focus on solar power and offshore drilling. In both cases, within-sector technology spillovers frequently occur in less than five years. For a full discussion of the spillovers in these two industries, see Appendix B.

4.2 Data

The National Science Foundation’s (NSF) Survey of Industrial Research and Development reports innovation expenditures by US companies from 1953-2007. The data include both company- and government-funded R&D expenditures.^{20,21} From 1972-2007, the survey reports energy specific R&D expenditures. I split the R&D expenditures into fossil, green, and non-energy categories.²² Fossil innovation corresponds to any R&D expenditures on coal, oil, or natural gas. Green innovation corresponds to any energy R&D expenditures that are not in coal, oil, or natural gas. This category includes renewables and nuclear, as well

¹⁹See, for example, Aghion et al. (2015); Crabb and Johnson (2007); Hassler et al. (2012); Lanzi and Sue Wing (2010); Newell et al. (1999); Popp (2002).

²⁰Government-funded research expenditures are defined as “the cost of R&D performed within the company under federal R&D contracts or subcontracts, and R&D portions of federal procurement contracts and subcontracts” (Documentation for NSF Industrial Research and Development Information System available at: <http://www.nsf.gov/statistics/iris/glossary.cfm>)

²¹I include both government- and company-funded research expenditures because government-funded R&D in the early 1970s arguably responded to market based incentives. Prior to President Reagan taking office in 1981, a specific goal of federal energy policy was to accelerate the development of new marketable technologies, making federally-funded R&D a potential substitute for company funded R&D (see Popp (2002) for further discussion). Additionally Lichtenberg (1987) finds a substantial response of government funded R&D to changes in the relative price of energy.

²²The 1972 data only include energy and non-energy R&D; the split between fossil and green is not available during this year. The data after 1972 does include the split between fossil and green. Therefore, I assume that the relative split between fossil and green in 1972 is the same as it is in 1973.

as energy conservation and efficiency.²³ This mapping reflects the broad definition of green energy to encompass both non-carbon sources of energy and improvements in conservation and efficiency discussed in Section 2. Finally, I measure non-energy R&D expenditures as the difference between total and energy R&D expenditures.

Data on fossil energy prices, fossil energy production, and oil import prices and quantities are from the US Energy Information Administration. Data on labor, fixed assets, output, and employee compensation are from the US Bureau of Economic Analysis (BEA) industry accounts. Following Mork (1989), I use the refiner acquisition cost of imported crude oil to measure the price of oil imports. This measure captures differences in the foreign and domestic prices of crude oil due to price controls and other policies.

4.3 Direct calibration

Table 1 reports the parameter values. I calibrate the following six parameters directly from the data series: $\{\alpha_f, \alpha_n, \rho_f, \rho_g, S, \omega\}$, where $\omega = \frac{\omega_o}{\omega_f}$ measures the carbon content of oil imports relative to that of domestic fossil energy.

I calibrate the labor share in fossil energy, $1 - \alpha_f$, as the cost share of labor in value added in the fossil energy sector. Fossil energy corresponds to coal, oil, and natural gas extraction, as well as to the production of petroleum and coal products (such as gasoline). I map fossil energy to the mining and the petroleum and coal products industries (NAICS codes 21 and 324) in the BEA accounts. The average labor share over the past twenty-five years (1987-2012) in these two industries combined is 0.28. I use the standard value for labor share in GDP, 0.64 for non-energy labor share, $1 - \alpha_n$, since the non-energy sector comprises most of the economy.

I normalize the workforce to unity, $L = 1$. Approximately one percent of workers are engaged in R&D in the US (Jones and Vollrath (2013)), and, so I set the number of scientists $S = 0.01$. I also normalize ρ_n to unity. Thus, parameters (ρ_f, ρ_g) capture the number of

²³Conceptually, it is important to include nuclear energy R&D as part of green R&D because this was seen as the main viable alternative to fossil energy in the early 1970s. Major disasters such as the meltdowns at Chernobyl and three-mile island had not yet occurred. Excluding nuclear energy would therefore understate the innovation in non-fossil energy. Specifically, if agents did not see nuclear as a viable alternative to fossil energy, then there would have presumably been more investment in other green energy forms of R&D. However, quantitatively, re-calibrating and re-solving the model excluding nuclear R&D does not make a substantial difference in the effects of endogenous innovation on the size of the carbon tax required to achieve the Power Plan target.

processes in the fossil and green energy sectors relative to the number of processes in non-energy. I measure these relative levels of diversity by the long-run average fractions of fossil R&D to non-energy R&D and green R&D to non-energy R&D. This measure assumes that average R&D is equal across all processes in the long run.

Table 1: Parameter Values

Parameter	Model Value	Source
Final good production		
Output elasticity of substitution: ε_y	0.05	-
Energy elasticity of substitution: ε_e	1.50	-
Fossil elasticity of substitution: ε_f	6.24	Method of moments
Distribution parameter: [†] δ_y	1.44e-38	Method of moments
Distribution parameter: $\delta_{\tilde{F}}$	0.47	Method of moments
Intermediates production		
Labor share in fossil energy: $1 - \alpha_f$	0.28	Data
Labor share in green energy: $1 - \alpha_g$	0.09	Method of moments
Labor share in non-energy: $1 - \alpha_n$	0.64	Data
Number of workers: L	1	Normalization
1971-1975 productivity shock: ν	0.64	Method of moments
Research		
Across-sector technology spillovers: ϕ	0.50	-
Diminishing returns: η	0.79	Method of moments
Scientist efficiency: γ	3.96	Method of moments
Sector size: ρ_f	0.01	Data
Sector size: ρ_g	0.01	Data
Sector size: ρ_n	1	Normalization
Number of scientists: S	0.01	Data
Climate		
Emissions Conversion: ω	1.03	Data

[†]The value of the economically relevant quantity is $\left(\frac{\delta_y}{1-\delta_y}\right)^{\varepsilon_y} = 0.01$.

Additionally, I design the model so that the elasticity of substitution between energy and non-energy in the production of output, ε_y , is close to zero. As this elasticity reaches zero, the specification becomes Leontief. The Leontief condition implies that non-energy inputs and the CES composite comprised of the energy inputs (F , G , and O^*) are required in fixed proportions to produce output. Even with this Leontief condition, the amount of the composite comprised of oil imports and fossil energy, \tilde{F} , used to produce a unit of output can

vary since agents can substitute green energy for \tilde{F} . Empirically, this substitution occurs through increases in renewable energy, nuclear, and/or energy efficiency. As discussed in Section 2, green energy includes all of these channels. Thus, any reduction in \tilde{F} requires an increase in green energy to produce the same quantity of output. Note that when the elasticity of substitution is exactly zero, there are kinks in the equilibrium conditions that are difficult to handle numerically. To avoid these numerical difficulties, I set the elasticity of substitution slightly greater than zero, $\varepsilon_y = 0.05$.

I use a conservative value for the elasticity of substitution between green energy and the composite comprised of fossil energy and oil imports, $\varepsilon_e = 1.5$. This parameter is particularly difficult to pin down because of the lack of aggregate data on green energy prices and quantities. Values of similar parameters used in integrated assessment and macroeconomic models typically range from around unity to ten (Lanzi and Sue Wing (2010); AABH) while empirical estimates range from 1.6-3 (Lanzi and Sue Wing (2010); Papageorgiou et al. (2013)). Section 6 considers robustness analysis for different values of ε_e .

Finally, to calibrate ω , I measure the carbon content of fossil energy as the weighted average of the carbon content of coal, oil, and natural gas, where the weights are determined by the average quantities produced in the US in 2012. Parameter ω is the ratio of the carbon content of oil to the carbon content of fossil energy.

4.4 A method of moments

I jointly calibrate the remaining parameters $\{\alpha_g, \varepsilon_f, \delta_{\tilde{F}}, \delta_y, \eta, \gamma\}$ to capture the relationships between energy prices, production, and innovation. To obtain empirical evidence of these relationships, I analyze the energy price increases triggered by historical oil shocks and the accompanying changes in energy production and innovation. Empirically, these oil shocks led to large increases in the prices of substitute fossil fuels (such as coal and natural gas) in addition to the increases in the price of oil. Thus, like a carbon tax, the oil shocks created a substantial increase in the price of fossil energy.

In an ideal setting, to calibrate the model parameters I would use data on energy price increases triggered by climate policy instead of by oil shocks. However, there are not many economy-wide historical examples of climate policies. The closest example is the Emissions Trading System (ETS) in the European Union. However the ETS carbon permit price

has been very unstable. In both the pilot period (2005-2007) and the first trading period (2008-2012), the EU over-allocated carbon permits and the price effectively fell to zero. Another alternative to using oil shocks is to use the variation in gas taxes (or other energy taxes) across countries. However, these taxes are usually specific to a single sector, such as transportation, and, thus, are not necessarily representative of how the aggregate economy would respond to a carbon tax that applies to all carbon-emitting fuels. The oil shocks and the accompanying data on energy production and innovation are a rare historical example of the economic response to an aggregate increase in fossil energy prices. As a robustness check, I compare empirical, sector-specific estimates of the price elasticity of green patents with the corresponding aggregate model value. The results, reported in Section 4.5, show that the model estimates are within the range spanned by the empirical studies for different green energy sectors.

I focus on the oil shocks triggered by the rise of OPEC in the first half of the 1970s. I use the oil shocks of the early 1970s instead of more recent oil shocks for two reasons. First, because a carbon tax will likely be permanent, it is important to calibrate to an aggregate increase in energy prices that agents at least believe to be permanent. After the rise of OPEC in the early 1970s, there was a sense that the economy had permanently switched from a low-energy-price regime to a high-energy-price regime. Energy price forecasts during the 1970s and early 1980s generally do not predict falling energy prices, suggesting that agents believed that the oil shocks were very long-lived.²⁴ However, after oil prices began to fall in the mid-1980s, agents potentially learned that this regime switch was not permanent and that oil shocks could be temporary. The model implicitly assumes that oil price changes are expected to be permanent. This makes using later oil price shocks inappropriate since expectations likely violated this assumption.

Second, a convenient way to introduce an oil shock is to model the economy on a balanced growth path (in which energy prices are constant) and then shock it with an oil shock. The 1970s is the most recent time period that matches these dynamics –that is, a long period of price stability followed by an unexpected jump in the oil price. Real energy prices were relatively constant for the 20 years prior to the 1970s. To summarize, I have calibrated to the early 1970s oil shocks because it is the only historical episode that arguably matches the

²⁴See, for example, Annual Energy Outlook (EIA, 1979); World Oil (EMF, 1981); Levy, “Beyond the Oil Bonanza: When the Wells Run Down” (NY Times January 4, 1979).

model's assumptions of being on a balanced growth path when there is a large and exogenous change in the price of oil imports that agents perceive to be permanent.

One limitation with using the early 1970s oil shocks to pin down the model parameters is that they happened forty years ago. It's possible that some of the parameter values could have changed over time. Even so, any meaningful inference from a calibrated growth model requires the assumption of parameter constancy. And the parameters can be constant at values calibrated from any episode along the equilibrium path, whether the episode is early or late. As a check on the assumption of parameter constancy in Appendix C, I analyze the effects of the 2003 oil shock in both the model and the data.²⁵ In particular, I calculate the responses of fossil and green innovation to a change in the price of oil imports. The responses are similar in the model and the data, suggesting that the parameter values that govern these responses have not changed substantially over time.

The early 1970s oil shocks coincided with a decline in the capacity of US oil fields (Hamilton (2009)) and with changes in energy and environmental policies. These events likely affected energy innovation incentives and so it is important to account for them in the calibration strategy. In particular, the Environmental Protection Agency (EPA) was initiated on December, 2, 1970, and with it came the authority for the federal government to implement and enforce environmental regulation. This was a major regulatory change which launched the US into a new era of environmental stewardship (Berman and Bui (2001)). Examples of influential environmental regulation from the early 1970s include the Clean Air Act, which limited emissions from coal power plants and oil refineries, and the Clean Water Act and Safe Drinking Water Act, which placed restrictions on fossil energy companies' hazardous waste. Congress also passed a set of long overdue health and safety regulations in underground coal mines that reduced labor productivity (Bohi and Russell (1978)).

In addition to this new era of environmental protection, the government implemented a series of oil price controls and windfall profits taxes on oil companies which lasted from 1971 until 1982, when President Reagan deregulated the industry. These price distortions drove a wedge between the prices of imported and domestic oil and led to energy shortages. Furthermore, oil import restrictions were relaxed considerably in 1973 (Bohi and Russell

²⁵As discussed earlier, I do not calibrate to the 2003 oil shock because energy prices and energy innovation are not stable for a sustained period preceding the 2003 oil shock, suggesting that the assumption that economy was on a long-run balanced growth path prior to the shock is imperfect. Moreover, after the 1970s, agents learned that energy prices are uncertain, and they formed expectations over future energy prices.

(1978)). The share of oil imports increased throughout most of the 1970s despite their rising cost.

All of these policies likely reduced the profitability of fossil energy extraction and the accompanying innovation incentives. To account for these coincident events in the calibration, I model effects of the policy changes together with the decline in the capacity of US oil fields as a negative productivity shock, ν , to fossil energy production:

$$F_t = \nu_t L_{ft}^{1-\alpha_f} \int_0^1 X_{fit}^{\alpha_f} A_{fit}^{1-\alpha_f} di. \quad (11)$$

Since the model is not sufficiently detailed to accurately incorporate each individual regulation change, I use the reduced-form productivity shock to capture the overall effects of the new regulation and the decline in oil capacity.

I jointly calibrate the parameters to match the data generated by the following experiment in the model with the data generated by the oil and productivity shocks of the early 1970s in the US economy.

Initial balanced growth path (1961-1970): The economy is on a balanced growth path with respect to the price of oil imports and environmental and energy policies.

Shock period (1971-1975): Two unexpected shocks realize: (1) the price of oil imports increases from its value on the balanced growth path; and (2) a negative productivity shock affects domestic fossil energy production.

Environmental policy and the price of oil imports were relatively constant prior to the 1970s, allowing me to begin the experiment on a balanced growth path. I match this balanced growth path to data from 1961-1970. I begin the shock period in 1971 because the EPA was created in December of 1970, launching the US into a new era of environmental regulation. Since this regulation was a major turning point in US environmental policy, it arguably knocked the US off its balanced growth path and stimulated green energy investment and innovation. I measure the size of the oil shock by the observed percentage change in the average price of oil imports from 1971-1975 relative to its average value from 1961-1970. Both shocks are unexpected by the agents on the balanced growth path since they were

unprecedented in the data. Machine production decisions are made prior to the realization of the shocks, while scientist and labor decisions are made after the shocks realize.²⁶

I construct moments from this experiment so that the model matches the innovation incentives that coincided with the oil price shocks and regulatory changes. Four important moments are the values of fossil energy production and oil imports (as shares of GDP) in both the balanced growth path and in the shock period. As shown in equation (10) and discussed in Section 3, the values of fossil and green energy production ($P_{ft}F_t$ and $P_{gt}G_t$) are primary determinants of the innovation incentives in each of these sectors. However, data on the value of green energy production is not available. Data is available for the value of imported oil and the first order conditions for the final good producer imply that one important component of the value of green energy production is the value of imported oil. Specifically, equation (12) (derived from the first order conditions for the final good producer, see Appendix A) shows that P_gG is directly proportional to the value of the composite comprised of oil imports and fossil energy, $P_{\tilde{F}}\tilde{F}$,

$$P_gG = P_{\tilde{F}}\tilde{F} \left(\frac{P_{\tilde{F}}}{P_g} \right)^{\varepsilon_e - 1}. \quad (12)$$

The CES properties of the production functions imply that the value of this composite equals the sum of the values of fossil energy production and oil imports: $P_{\tilde{F}}\tilde{F} = P_{ft}F_t + P_{ot}^*O_t^*$. Therefore, the value of oil imports is also important for capturing innovation incentives.

Two more relevant moments are the research expenditures on fossil and green energy as a fraction of total research expenditures. The energy research data are not available until 1972, so I construct the empirical averages from 1972-1975. Research expenditures in the data correspond to the wage multiplied by the number of scientists in the model. Scientist market clearing implies that the scientists' wages are equated across all sectors. Therefore, the fraction of research expenditures in fossil energy in the data corresponds to the fraction of scientists in fossil energy in the model (and likewise for green research). Table 2 reports the empirical values of the moments in both the balanced growth path and the shock period. Additionally, I target the annualized long-run growth rate of GDP per capita of 2 percent.

²⁶The empirical evidence supports these timing assumptions. The change in the fraction of fixed assets in the fossil energy sector (relative to total fixed assets) is very small from 1971-1975. In contrast, the fraction of energy research expenditures relative to total research expenditures almost doubles from 1972 to 1975.

Table 2: Data in the BGP and Shock Period

	BGP	Shock Period
Fossil energy production share	1.9	2.1
Oil imports share	0.2	0.8
Percent of R&D expenditures in fossil	-	2.1
Percent of R&D expenditures in green	-	3.4

This process yields seven moments (those listed in Table 2 plus the long-run growth rate of per-capita GDP) for the six parameters and the productivity shock, ν . For each set of parameters, I solve the model, compute the moments, and compare them with the moments in the data. I use the Nelder-Meade simplex algorithm (Nelder and Mead (1965)) to minimize the sum of the square of the residuals between the empirical and model values of the moments. The model fits the data very closely; the minimized distance is 2.2×10^{-21} . Appendix C evaluates the fit of the model against five non-targeted moments. The values of these moments are relatively similar in the model and the data, suggesting that the model's fit is reasonably good. Appendix D reports bootstrapped standard errors for the parameters calibrated from the method-of-moments procedure. The standard error estimates suggest a reasonable degree of precision for most of the parameter estimates.

While all the parameters are jointly determined, the shares of fossil energy production and oil imports on the initial balanced growth path are pinned down primarily by the CES distribution parameters, $\delta_{\tilde{F}}$ and δ_y . The movements in these shares are largely governed by the productivity shock, ν , and the elasticity of substitution between fossil energy and oil imports, ε_f . For example, if fossil energy and oil imports are more substitutable, then the oil shock leads to a larger increase in demand for fossil energy, which leads to a bigger increase in the fossil energy price, quantity, or both. Hence, increases in this substitution elasticity result in a larger increase in fossil energy production (as a share of GDP) in response to the oil shock.

The research expenditure moments primarily pin down the level of diminishing returns η and the labor share in green energy, $1 - \alpha_g$. The price elasticity of demand for green machines is $\frac{1}{1 - \alpha_g}$. All else constant, increases in labor share reduce the price elasticity of demand. Less elastic demand increases the machine producer's optimally chosen machine price, raising the returns to innovation (see Appendix A for the derivation and further discussion). Parameter

γ determines the long-run growth rate.

The calibration strategy does not pin down the strength of the cross-sector technology spillovers, ϕ . This parameter is governed by the relative levels of technology on a long-run balanced growth path. For example, stronger technology spillovers imply that the long-run equilibrium levels of fossil and green technology must be closer together. However, data on energy innovation (and, thus, on technology) are not available on the balanced growth path of the 1960s.

However, the data do provide suggestive evidence that the strength of the spillovers, ϕ , is sufficiently large such that the interior balanced growth path in which agents innovate in both energy sectors is stable. If this was not the case, then the only stable balanced growth paths are corner solutions in which agents innovate in only the fossil or the green energy sector, which would imply that green innovation was zero along the balanced growth path of the 1960s. However, in the early 1970s, green innovation expenditures were over half of all energy innovation expenditures. It is highly unlikely that green innovation would go from nonexistent to over half of all energy innovation in such a short time frame. Thus, the spillovers must be sufficiently strong so that positive innovation occurred in both fossil and green energy along the 1960s' balanced growth path. I set $\phi = 0.5$ in the main specification. In Appendix F.1, I report the main results for a range of values of $\phi > 0.2$, where 0.2 is the cutoff value for which the interior balanced growth path in which agents innovate in both energy sectors is stable.

Labor share in green energy is 0.09, implying that green energy is a very capital-intensive sector. Consistent with this calibration, green energy technologies, such as nuclear, solar, and, particularly energy efficiency, are all very capital-intensive. The elasticity of substitution between fossil energy and oil imports, ε_f , is considerably higher than that between green energy and the composite comprised of fossil energy and oil imports, ε_e , (6.24 compared to 1.5), suggesting that fossil energy is a better substitute for oil imports than green energy. This is intuitive since one component of fossil energy is domestically produced oil, which is a perfect substitute for imported oil. The diversity of the energy sectors, ρ_f and ρ_g , are both small compared to the non-energy sector. However, ρ_g is slightly larger than ρ_f (0.011 compared to 0.0095), indicating that the green energy sector is more diverse than the fossil sector.

The calibrated value of the distribution parameter, δ_y , is very small. This is because the

elasticity of substitution between energy and non-energy inputs in the production of output, ε_y , is also very small, ($\varepsilon_y = 0.05$). The value of the quantity that is economically relevant for the optimal allocation between energy and non-energy inputs, $\left(\frac{\delta_y}{1-\delta_y}\right)^{\varepsilon_y}$, (see equation (16) in Appendix A), is substantially larger $\left(\left(\frac{\delta_y}{1-\delta_y}\right)^{\varepsilon_y} = 0.01\right)$.

Additionally, the calibrated size of the productivity shock, ν , is 0.64, suggesting that the combined effects of the series of environmental, health, and safety regulations, the distortions created by price controls and import policies, and the declining capacity of US oil fields reduced productivity in fossil energy. This result relates to the literature linking the productivity slowdown to increased environmental regulation in the 1970s (e.g., Gray (1987)). To comply with the regulations, firms must divert resources away from output production.

4.5 Comparison to empirical studies

As an additional check on both the calibration and the model specification, it is useful to compare the implications of the calibrated quantitative model with the empirical literature on energy prices and innovation. Both Popp (2002) and Aghion et al. (2015) calculate the elasticity of green energy patents with respect to a change in energy prices. Popp (2002) estimates this elasticity from aggregate US time series data from 1970-1994 on fossil energy prices and green energy patents in 11 energy technologies. Six of these technologies relate to energy supply (such as solar) and five to energy demand (such as the reuse of industrial waste heat). Aghion et al. (2015) focus on the automobile industry. They use a cross-country, firm-level panel on green car patents (e.g., hybrid vehicle technologies) and tax-inclusive gas prices to estimate the price elasticity of R&D in green car technologies. The five-year price elasticity of green patents is 0.21 in Popp (2002) and is 3.7 in Aghion et al. (2015).²⁷

To compare these empirical results with the present paper, I rewrite the technology accumulation equation for green technology (equation (4)) as the sum of the existing green technology stock and new green ideas, \mathcal{I}_{gt}

$$A_{gt} = A_{gt-1} + \mathcal{I}_{gt} \quad \text{where} \quad \mathcal{I}_{gt} = \gamma \left(\frac{S_{gt}}{\rho_g}\right)^\eta A_{t-1}^\phi A_{gt-1}^{1-\phi}. \quad (13)$$

²⁷Both Popp (2002) and Aghion et al. (2015) estimate a dynamic specification which make it possible to compute the five year elasticity. See Table 4 in Popp (2002) and Table 10 Appendix C in Aghion et al. (2015).

New green ideas are the flow input into technology and, thus, correspond to green patents in the data. Let $P_{\tilde{F}}$ be the tax-inclusive price of the composite comprised of fossil energy and foreign oil, (see equation (17) in Appendix A). The (one-period) elasticity of new green ideas with respect to a change in $P_{\tilde{F}}$ is given by,

$$\epsilon = \left(\frac{\mathcal{I}_{gt^*} - \mathcal{I}_{gt^*-1}}{\mathcal{I}_{gt^*-1}} \right) \left(\frac{P_{\tilde{F}t^*-1}}{P_{\tilde{F}t^*} - P_{\tilde{F}t^*-1}} \right), \quad (14)$$

where t^* is the period in which the tax is introduced.²⁸

The value of this elasticity in the model is 1.7. This estimate is between the estimates in Popp (2002) and Aghion et al. (2015). One explanation for why the model value of the elasticity is larger than the estimate in Popp (2002) is that the green innovation in the sectors covered in Popp's study is less responsive to changes in $P_{\tilde{F}}$ than is average green innovation. However, a second explanation for the different elasticity is the source of the change in $P_{\tilde{F}}$. Popp's calculation uses aggregate variation in fossil energy prices due to oil shocks (or similar macroeconomic events) instead of from a carbon tax. While both oil shocks and carbon taxes increase incentives for green innovation, oil shocks also increase incentives for fossil innovation. If there is crowd-out between fossil and green innovation, then the price elasticity of green innovation will be smaller when the price change is caused by an oil shock than when it is caused by a carbon tax. Consistent with this hypothesis, the model elasticity of green ideas from an increase in $P_{\tilde{F}}$ from an oil shock is 1.3, approximately 25 percent smaller than the elasticity from a carbon tax. In a related empirical patent study, Popp and Newell (2012) find suggestive evidence of this crowd-out within energy supply technologies (such as oil refining and solar).

One explanation for why the model value of the elasticity is smaller than the estimate in Aghion et al. (2015), is that innovation in green car technologies is more responsive than average green innovation to changes in the fossil energy price. Some of the variation in gas prices comes from differences in the gas tax and some comes from oil shocks. However, since the automobile industry does not supply fossil energy, price changes from oil shocks and carbon taxes should create similar incentives for innovation in green car technologies. Thus, differences in the elasticity estimates due to crowd-out are not as likely in this case.

²⁸Prior to period t^* , the economy is on a balanced growth path.

5 Results

I perform two exercises to fully explore the interactions between endogenous innovation and climate policy. In both exercises, the economy begins on the same balanced growth path, but innovation is endogenous in the first and exogenous in the second. I introduce constant carbon taxes in the 2016-2020 period which are chosen to achieve the Power Plan target of a 30-percent reduction in emissions from the balanced growth value by 2030 for each model. The size of the tax necessary to achieve the 2030 reduction is different for economies with endogenous versus exogenous innovation.

The endogenous-innovation model is my baseline model. Machines, workers, and scientists all adjust in response to the tax. The exogenous-innovation model has the endogenous innovation channel shut down. Unlike in the endogenous-innovation model, only machines and workers adjust in response to the tax; the scientists (and hence the levels and growth rates of technology) are fixed at their balanced growth values.

5.1 Carbon tax: the role of endogenous innovation

The carbon taxes required to achieve the Power Plan target are 30.3 and 24.5 in 2013 dollars per ton of CO₂ in the exogenous- and endogenous-innovation models, respectively.²⁹ The required carbon tax is 19.2 percent lower when innovation is endogenous. Regardless of whether innovation is endogenous, the carbon tax operates through prices to shift demand from fossil to green energy, reducing emissions. However, when innovation is endogenous, this shift in demand spurs green innovation and reduces fossil innovation. Over time, this change in innovation reduces the marginal cost of producing green energy relative to fossil. This lowers the relative price of green to fossil energy, creating stronger incentives for the final-good producer to switch from fossil to green. Thus endogenous innovation amplifies the price incentives created by the carbon tax, implying that the same reduction in emissions can be achieved with a smaller tax.

An analogous interpretation of this result is that endogenous innovation increases the emissions reduction from a given-sized carbon tax. In particular, if the carbon tax is 30.3

²⁹The values of these taxes correspond to an increase in the price of oil imports of 20.8 percent in the exogenous innovation and 16.8 percent in the endogenous innovation case. Such changes are slightly smaller than the 28 percent increase in the price of oil imports used for the calibration.

dollars per ton, then endogenous innovation increases the percent reduction in emissions by close to five percentage points (from 30 percent to 34.6 percent). A policy implication of these results is that if the government designs a cap and trade system to achieve a target permit price (perhaps because a carbon tax is politically infeasible), then endogenous innovation implies that the government should issue fewer permits in order to achieve its price target. This implication is particularly relevant for the EU Emissions Trading System (ETS) where, for several reasons, governments over-allocated permits and the price fell below the desired level.

The finding that the carbon tax necessary to achieve the Power Plan target is 19.2 percent lower when innovation is endogenous is sensitive to both the size of the targeted reduction in emissions and the time frame in which the reduction must be achieved. Appendix E analyses the effects of innovation for different-sized emissions targets and different time frames. In particular, the effect of endogenous innovation on the size of the carbon tax is smaller if the target is more stringent or the time frame is shorter. More stringent targets and shorter time frames force agents to rely less on technological advances and more on shifts in production factors (i.e., workers and machines) to achieve the emissions goal. This switch reduces the role of endogenous innovation and its accompanying effects on the size of the required carbon tax.

Table 3 provides more details on the mechanisms driving the effects of endogenous innovation. Column 2 of Table 3 reports the values on the baseline balanced growth path. Each row in columns 3-5 reports a measure of the treatment effect; they show the percentage difference from the baseline in each of the variables in 2030 and on the long-run balanced growth path under the carbon tax. For example, the first row of column four implies that when innovation is endogenous, fossil energy scientists are 60.5 percent lower in 2030 under the tax than in the baseline. Table 3 does not include a column for endogenous innovation on the long-run balanced growth path because there are no transitional dynamics when innovation is exogenous, so the values in 2030 equal the values in the long-run balanced growth path under the carbon tax.

Table 3: Effects of a Constant Carbon Tax Which a Achieves 30 Percent Reduction in Emissions by 2030

	Levels on Baseline BGP	Percent Difference From the Baseline		
		Exogenous Innovation (Year 2030)	Endogenous Innovation (Year 2030)	Endogenous Innovation (Long-run BGP)
Innovation				
Fossil Scientists: S_f	1.5e-04	0	-60.5	-29.9
Green Scientists: S_g	1.0e-04	0	53.3	23.8
Non-Energy Scientists: S_n	0.010	0	0.4	0.2
Relative Technology				
Green to Fossil: $\frac{A_g}{A_f}$	0.4	0	44.5	144.6
Green to Non-energy: $\frac{A_g}{A_n}$	0.9	0	16.9	39.3
Relative Prices				
Green to Fossil: $\frac{P_g}{P_f}$	1.1	0.2	-7.0	-17.1
Green to Non-Energy: $\frac{P_g}{P_n}$	1.4	0.6	-1.0	-2.6
Energy to Non-Energy: $\frac{P_e}{P_n}$	3.9	14.9	13.0	14.6
Green to Fossil With Tax: $\frac{P_g}{P_f + \tau_f}$	1.1	-30.0	-30.0	-36.1
Relative Production				
Green to Fossil: $\frac{G}{F}$	1.4	78.0	79.2	112.9
Energy to Non-Energy: $\frac{E}{N}$	0.01	-0.7	-0.6	-0.7
Relative Wages				
Scientist to Worker: $\frac{w_s}{w_l}$	2.7	-	-0.1	-0.1
Climate				
Emissions	-	-30.0	-30.0	-36.9
Carbon Stock	-	-2.2	-2.1	-36.8

Note: the baseline balanced growth path is the balanced growth path with no carbon tax. This balanced growth path is the same in the endogenous- and exogenous-innovation models. The values on the long-run balanced growth path under exogenous innovation are the same as the 2030 values because there are no transitional dynamics when innovation is exogenous. Variable P_e (defined in equation (17) in Appendix A) is the tax-inclusive price of the CES composite of fossil energy, green energy, and oil imports, E .

The carbon tax leads to large shifts in fossil and green innovation and relatively small movements in non-energy innovation (innovation segment of Table 3). In 2030, the tax reduces fossil innovation by 60.5 percent, increases green innovation by 53.3 percent, while non-energy innovation only increases by 0.4 percent.³⁰ These results suggest that the increased green innovation comes at the expense (i.e. crowds out) fossil innovation and not

³⁰While, the overall number of scientists is fixed, the percentage change in the number of scientists in each sector do not sum to zero because the baseline levels are very different.

non-energy innovation. Since fossil and green energy are gross substitutes, the tax shifts demand from fossil to green energy, increasing the green innovation incentives. In contrast, because the energy and non-energy inputs are almost perfect complements, the effects of the tax on the value of non-energy production and the corresponding innovation incentives are small. These movements in innovation affect relative technology. By 2030, the ratio of green to fossil technology is 44.5 percent higher than in the baseline. On the long-run balanced growth path, this ratio is more than double its baseline value.

The prices segment of Table 3 shows the effects of the tax on relative prices. When innovation is endogenous, the relative price of green compared to fossil energy falls by 7.0 percent in 2030 and by 17.1 percent on the long-run balanced growth path. The fall in the relative price of green to fossil energy results from both increases in green innovation and decreases in fossil innovation. Increases in green innovation reduce the marginal cost of green energy production, reducing the relative price of green to fossil energy. Decreases in fossil innovation raise the marginal cost of fossil energy production (relative to the baseline), raising its price, and, thus, further reducing the relative price of green to fossil energy. In contrast, when innovation is exogenous, there is almost no change in the relative marginal costs of the different inputs, and relative prices are almost the same as on the baseline balanced growth path.³¹ Therefore, almost all of the change in the energy prices in the endogenous-innovation model results from changes in technology.

The production segment of Table 3 reports the effects of the carbon tax on the production of the different intermediate inputs. In 2030, the changes in the relative quantities of green compared to fossil energy production are similar between the endogenous- and exogenous-innovation models because the carbon tax achieves the same reduction in emissions. However, the long-run effects are very different; the tax increases the ratio of green to fossil production by 112.9 percent on the new balanced growth path when innovation is endogenous and by only 78.0 percent when innovation is exogenous. This difference arises because green technology relative to fossil keeps growing after 2030, further decreasing the relative price of green energy and, thus, increasing the final-good producer's green energy demand. Unlike fossil and green energy production, the changes in the ratios of non-energy to energy production

³¹Energy prices under the tax in the exogenous model are not identical to their baseline values because the general equilibrium effects lead to small changes in the wage rate. These changes have different effects on the marginal cost of production in the different sectors, which, in turn, affect relative prices.

are almost zero. Since the elasticity of substitution between energy and non-energy inputs is close to zero, the final-good producer must substitute green for fossil energy and oil imports to reduce emissions, instead of substituting non-energy inputs for energy inputs.

The relative wage segment of Table 3 reports the effects of the carbon tax on the return to supplying labor as a scientist (w_s) relative to the return to supplying labor as a worker (w_l). The carbon tax has almost no effect on these relative returns; in the new long-run balanced growth path, the return to scientists relative to the return to workers is only 0.1 percent smaller than its value in the initial equilibrium. This near constancy suggests that the carbon tax would not lead to substantial changes in the relative quantities of scientists and workers, supporting the assumption of fixed supplies of scientists and workers.

Finally, I calculate the consumption equivalent variation (CEV) to quantify the gross welfare costs of the policy. The CEV is the uniform percentage increase in an agent's consumption in the baseline that is necessary to make him indifferent between the baseline and the carbon tax scenarios. The CEVs are -0.5 percent and -0.6 percent with endogenous and exogenous innovation, respectively.³² Total consumption for all individuals in the United States from 2008-2012 was approximately \$53,671 billion (2012 dollars), so the CEVs with and without endogenous innovation equal approximately -\$270 and -\$320 billion, respectively.³³

Endogenous innovation reduces the gross welfare cost of the policy by 0.1 percentage points. Endogenous innovation affects the gross welfare costs through three partially offsetting channels. First, the carbon tax is smaller when innovation is endogenous; hence, the accompanying gross distortionary cost is smaller. Second, green energy is technologically behind fossil energy when the government implements the tax. Thus, the tax shifts energy production to a less productive sector. Endogenous innovation reduces these productivity losses as green technology catches up to fossil. Third, the shift in innovation from fossil to green energy reduces the aggregate growth rate along the transition path to a new long-run equilibrium. This temporary reduction in growth raises consumption costs and mutes the gross welfare gain from endogenous innovation.

AABH find that climate policy and endogenous innovation tip the economy to a new

³²To calculate the CEV, I set the annual rate of time preference to 1.5 percent and the intertemporal elasticity of substitution to the standard value of one half, $\theta = \frac{1}{2}$.

³³See BEA personal consumption expenditures.

long-run equilibrium where green technology grows and fossil technology is constant. The results in the present paper indicate somewhat smaller effects of endogenous innovation on climate policy outcomes than in AABH. These different findings are primarily due to two key parameters: the diminishing returns to innovation, η , and the strength of the across-sector technology spillovers, ϕ . Stronger diminishing returns to innovation (lower η) create incentives to spread scientists across both the fossil and green energy sectors. This spreading reduces the effect of a carbon tax on the direction of technical change. Stronger cross-sector spillovers (higher ϕ) reduce the path dependence in innovation. Green technology accumulates faster than fossil technology in response to the carbon tax. If some of the new green discoveries are applicable to fossil energy, then these spillovers indirectly encourage innovation in fossil energy. The calibration in the present paper uses middle values for both η and ϕ : $\eta = 0.79$, $\phi = 0.5$, while AABH use $\eta = 1$ and $\phi = 0$. This implicit parameter choice increases the role of endogenous innovation in AABH, relative to the present paper.

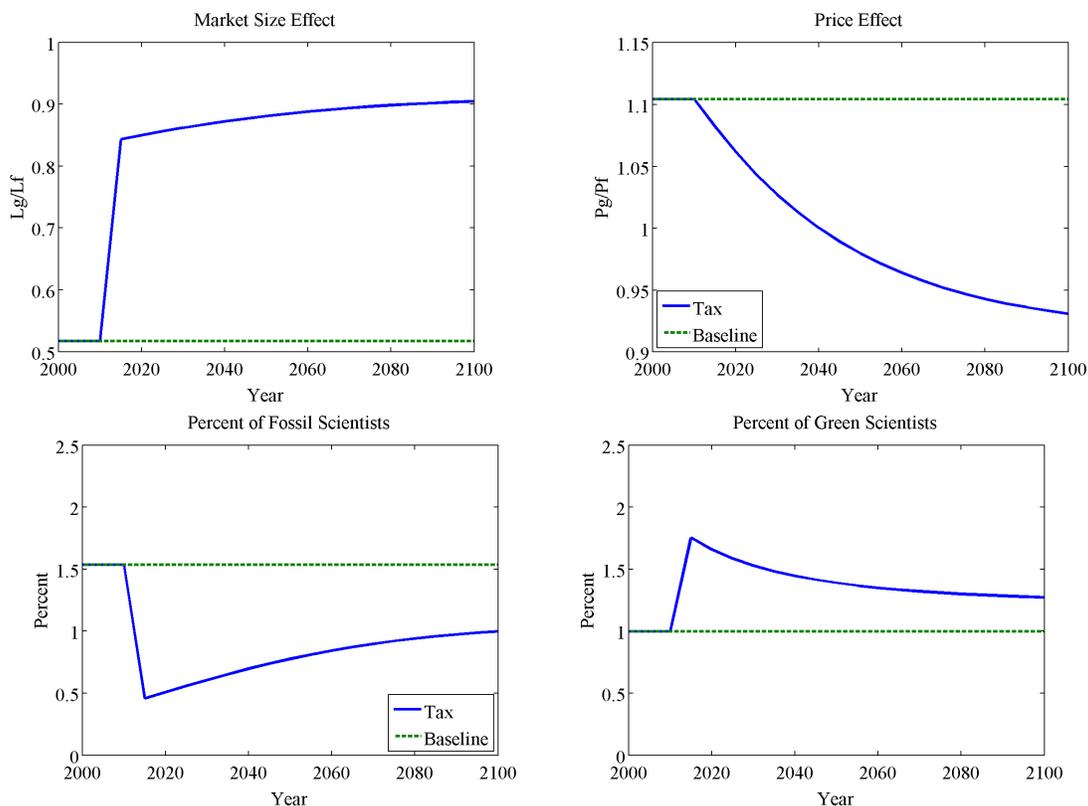
This paper is focused on the key mechanisms driving the interaction between innovation and climate policy. However, the results have interesting and opposing implications for the effects of endogenous innovation on the size of the optimal carbon tax. Endogenous innovation reduces the size of the tax necessary to achieve a given abatement target, implying a smaller optimal carbon tax. Working in the other direction, endogenous innovation also reduces the marginal abatement costs, raising the optimal abatement target, implying a larger carbon tax. Determining which of these two effects dominate is beyond the scope of this paper and is an interesting avenue for future research.

5.2 Dynamics

I discuss the dynamics along the transition to the new balanced growth path, focusing explicitly on the general equilibrium forces driving innovation. Figure 1 plots the dynamics with respect to four key variables in response to the carbon tax: (1) the market size of green energy relative to fossil energy (measured by the relative levels of employment, $\frac{L_g}{L_f}$), (2) the price of green energy relative to fossil, $\frac{P_g}{P_f}$, (3) the percent of fossil energy scientists, and (4) the percent of green energy scientists. The tax shifts demand from fossil to green energy, leading to an immediate jump in the green energy market size (top left panel of Figure 1) and in the percents of fossil and green energy scientists (bottom two panels of Figure 1).

The surge in green innovation relative to fossil leads to gradual improvements in the relative level of green technology, which reduce the relative price of green energy over time (top right panel of Figure 1). The fall in the relative price slowly reduces green innovation incentives, causing green energy scientists asymptote downward to a new equilibrium level, below their initial jump but above their value on the original balanced growth path.

Figure 1: Price and Market Size Effects From a Carbon Tax



6 Robustness analysis

Table 4 reports sensitivity analysis with respect to all the model parameters. As a summary statistic for the analysis, I analyze the percent that endogenous innovation reduces the size of the carbon tax required to achieve the Power Plan target.³⁴ The central column in Table 4 reports this summary statistic when the parameters equal their calibrated in values in the

³⁴Tables 8 and 9 in Appendix F.2 report the levels of the tax for each parameter perturbation. For most of the parameter values, the levels are not substantially different from the main specification.

main specification (given in Table 1). The high (low) column reports this summary statistic when the parameter is 25 percent bigger (smaller) than its value in Table 1. For example, reading from the first line of Table 4, if ε_y is 25 percent larger than its value in Table 1, then endogenous innovation reduces the carbon tax required to achieve the Power Plan target by 19.0 percent.

Given the prominence of ε_e in the literature, I conduct a sensitivity analysis over a wider range of values than for the other parameters in the model. Specifically, ε_e ranges from 1.1 to 3. There is reasonable consensus in the literature that fossil and green energy are substitutes, suggesting that this elasticity should exceed unity and empirical estimates range from 1.6-3 (Lanzi and Sue Wing (2010); Papageorgiou et al. (2013)). The model is not consistent with the targeted moments for values of this elasticity greater than 3. Additionally, I designed the model so that ε_y is near zero; the main specification uses $\varepsilon_y = 0.05$. Values of $\varepsilon_y < 0.05$, introduce kinks into the equilibrium conditions, making it very difficult to solve the model. Therefore, I only consider the high case for ε_y .

Table 4: Percent Change in the Carbon Tax from Endog. Innovation

Parameter	High	Central	Low
Imposed parameters			
Output elasticity of substitution: ε_y	19.0	19.2	-
Across sector spillovers: ϕ	18.4	19.2	20.2
Direct from data series			
Machine share in fossil energy: α_f	13.4	19.2	8.7
Machine share in non-energy: α_n	26.2	19.2	4.9
Energy elasticity of substitution: [†] ε_e	20.3	19.2	17.4
Sector diversity: ρ_f	19.5	19.2	18.9
Sector diversity: ρ_g	19.1	19.2	19.3
Population of scientists: S	19.2	19.2	19.2
Method of moments			
Machine share in green energy: α_g	20.2	19.2	26.6
Diminishing returns: η	35.9	19.2	10.9
Scientist efficiency: γ	21.7	19.2	16.0
Fossil elasticity of substitution: ε_f	18.3	19.2	20.1
Distribution parameter: δ_y	19.2	19.2	19.2
Distribution parameter: $\delta_{\bar{F}}$	17.0	19.2	32.9
1971-1975 productivity shock: ν	19.2	19.2	19.2

[†]The high case for ε_e is $\varepsilon_e = 3$ and the low case for ε_e is $\varepsilon_e = 1.1$.

The first two blocks of Table 4 report the robustness analysis for the eight parameters that I did not pin down using the 1970s oil shocks, $\{\varepsilon_y, \phi, \varepsilon_e, \alpha_f, \alpha_n, \rho_f, \rho_g, S\}$. For each perturbation of these eight parameters, I recalibrate the remaining six parameters $\{\alpha_g, \eta, \gamma, \varepsilon_f, \delta_y, \delta_{\tilde{F}}\}$ and the productivity shock, ν , to ensure that the model matches the moments described in Section 4. The third block of Table 4 reports robustness analysis for the six parameters and the productivity shock that I calibrated from the 1970s oil shocks. Since the original values of these parameters were determined by the method-of-moments procedure, the model will not match the targeted moments for the perturbations of these parameter values.

The results are particularly sensitive to changes in the diminishing returns to innovation, η . The diminishing returns reduce the amount agents increase green innovation in response to the carbon tax. As η rises, the diminishing returns to innovation fall, implying that agents increase green innovation more in response to the tax, thus increasing the effect of endogenous innovation on the size of the carbon tax.

The importance of endogenous innovation is increasing in ε_e . All else constant, higher values of ε_e increase the shift in demand from fossil to green energy in response to the tax. A larger demand shift leads to a larger change in innovation, which increases the effects of endogenous innovation on the size of the carbon tax. However the magnitude of these effects is reasonably small; even if \tilde{F} and G are almost Cobb-Douglas ($\varepsilon_e = 1.1$), endogenous innovation still reduces the size of the carbon tax by 17.4 percent.³⁵

Changes in the machine share in the fossil and green energy sectors have non-monotonic implications for the size of the required carbon tax. This is because the machine share has two potentially offsetting effects on the returns to innovation in a sector. First, all else constant, a lower machine share decreases machine demand in that sector. Since technology is embodied in the machines, this decrease in machine demand reduces the return to innovation in that sector. Second, a lower machine share also reduces the price elasticity of machine demand in the sector. Less elastic demand increases the machine producer's optimally chosen machine price, raising the returns to innovation. See equation (24) in Appendix A for a full discussion

³⁵One reason that ε_e does not have large implications for the effects of innovation on the size of the carbon tax is that matching the targeted moments with lower (higher) values of this elasticity requires weaker (stronger) diminishing returns to innovation. As the strength of the diminishing returns to innovation falls (η approaches unity), the effects of endogenous innovation on the size of the carbon tax increase, partially offsetting the decrease from the smaller substitution elasticity. Thus, the effects of this substitution elasticity are smaller than one might expect when the model is required to match the historical record.

of this intuition. Which effect dominates depends on the region of the parameter space and on the model's general equilibrium channels and thus, the sensitivity results are not always monotonic with respect to the machine share.

The low case for the machine share in non-energy implies only a 4.9 percent change in the carbon tax from endogenous innovation, the smallest response in the table. While at first this result might seem surprising, matching the moments with the low value of α_n requires very strong diminishing returns ($\eta = 0.32$), which substantially reduces the importance of endogenous innovation.

Increases in ϕ , the magnitude of the across-sector spillovers, dampen response of innovation to the carbon, reducing the role of endogenous innovation. However, this effect is relatively small. Appendix F.1 extends this robustness analysis to include values of ϕ ranging from 0.3 to 0.9. Even with very large spillovers, ($\phi = 0.9$) endogenous innovation still reduces the carbon tax by over 15 percent.

Higher values of γ increase the effect of changes in innovation on the relative levels of technology, thereby increasing the role of endogenous innovation. Increases in the distribution parameter, $\delta_{\tilde{F}}$, increase the final good producer's demand for fossil energy relative to oil imports, (see equation (16) in Appendix A). Since the carbon content of fossil energy is smaller than that of oil imports, this shift towards fossil energy results in a smaller increase in the after-tax price of \tilde{F} , the composite comprised fossil energy and oil imports. The smaller after-tax price change of \tilde{F} reduces the green innovation incentives in response to the tax. Finally, changes in the number of scientists, sector diversity, output distribution parameter, and productivity shock all have relatively small effects on the role of endogenous innovation.

Appendix F.3 analyzes an alternative summary statistic for the different parameter perturbations; the price elasticity of green ideas, defined in equations (13) and (14). As discussed in Section 4.5, the price elasticity implied by the quantitative model is larger than the empirical estimates in Popp (2002) but smaller than the empirical estimates in Aghion et al. (2015). This elasticity is endogenous to the model. However, parameter perturbations that bring the elasticity closer to the estimates in Aghion et al. (2015) increase the effects of endogenous innovation on the size of the carbon tax while perturbations that bring the elasticity closer to the estimates in Popp (2002) decrease its effects.

7 Conclusion

This paper develops a general equilibrium model to quantify the response of technology, prices, and other macroeconomic aggregates to climate policy. Building on the directed technical change literature, I model an economy in which scarce innovation resources can be allocated towards fossil energy, green energy, or non-energy intermediate inputs. I calibrate the model parameters using data from the natural experiment on energy prices and innovation from the oil shocks in the first half of the 1970s. I then use this empirically grounded model as a quantitative laboratory in which to study climate policy.

A key result is that endogenous innovation amplifies the price incentives created by the carbon tax. The tax operates through prices to shift innovation from fossil to green energy. This shift in innovation raises green technology compared to the baseline path, reducing the green energy price. Similarly, fossil innovation falls compared to the baseline path, raising the fossil energy price. These additional price movements reduce the size of the carbon tax required to attain a given abatement target. Specifically, endogenous innovation lowers the size of the carbon tax necessary to achieve a 30-percent reduction in emissions by 2030 by 19.2 percent.

Overall, the results imply that endogenous innovation has considerable effects on climate policy outcomes. Shifts in innovation lower the relative price of green energy compared to fossil by approximately 7 percent in the short run and 17 percent in the long run. Moreover, the relative level of green technology compared to fossil stabilizes at two and a half times its value on the baseline path.

References

- Acemoglu, D. (2002). Directed technical change. *The Review of Economic Studies* 69 (4), 781–809.
- Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012). The environment and directed technical change. *American Economic Review* 102(1), 131–166.
- Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2014). Transition to clean technology. *Working Paper*.

- Aghion, P., A. Dechezlepretre, D. Hemous, R. Martin, and J. Van Reenen (2015). Carbon taxes, path dependency and directed technical change: Evidence from the auto industry. *Journal of Political Economy* 124(1), 1–51.
- Berman, E. and L. Bui (2001). Environmental regulation and productivity: Evidence from oil refineries. *The Review of Economics and Statistics* 83(3), 498–510.
- Bohi, D. and M. Russell (1978). *Limiting Oil Imports*. Johns Hopkins University Press.
- Crabb, J. and D. Johnson (2007). Fueling the innovative process: Oil prices and induced innovation in automotive energy-efficient technology. *Working Paper*.
- Gerlagh, R. (2008). A climate-change policy induced shift from innovation in carbon-energy production to carbon-energy savings. *Energy Economics* 30, 425–448.
- Golosov, M., J. Hasler, P. Krusell, and A. Tsyvinski (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica* 82(1), 41–88.
- Goulder, L. and S. Schneider (1999). Induced technological change and the attractiveness of CO₂ abatement policies. *Resource and Energy Economics* 21, 211–253.
- Gourinchas, P.-O. and J. Parker (2002). Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Gray, W. (1987). The cost of regulation: OSHA, EPA and the productivity slowdown. *American Economic Review* 77(5), 998–1006.
- Grubb, M., M. Duong, and T. Chapuis (1994). Optimizing climate change abatement responses: On inertia and induced technology development. *Integrative Assessment of Mitigation, Impacts, and Adaptation to Climate Change*.
- Hamilton, J. (2009). Understanding crude oil prices. *Energy Journal* 30, 179–206.
- Hart, R. (2012). Directed technological change: It’s all about knowledge. *Working paper*.
- Hassler, J. and P. Krusell (2012). Economics and climate change: Integrated assessment in a multi-region world. *Journal of the European Economic Association* 10(5), 974–1000.

- Hassler, J., P. Krusell, and C. Olovsson (2012). Energy-saving technical change. *NBER working paper 18456*.
- Hemous, D. (2014). Environmental policy and directed technical change in a global economy: Is there a case for carbon tariffs? *Working paper*.
- Heutel, G. and C. Fischer (2013). Environmental macroeconomics: Environmental policy, business cycles, and directed technical change. *Annual Review of Resource Economics* 5, 197–210.
- Jones, C. and D. Vollrath (2013). *Introduction to Economic Growth*. W.W. Norton & Company.
- Jones, C. and J. Williams (1998). Measuring the social return to R&D. *The Quarterly Journal of Economics* 113(4), 1119–1135.
- Krusell, P. and A. Smith (2009). Macroeconomics and global climate change: Transition for a many region economy. *Working paper*.
- Lanzi, E. and I. Sue Wing (2010). Directed technical change in the energy sector: an empirical test of induced directed innovation. *Working paper*.
- Lichtenberg, F. (1987). Changing market opportunities and the structure of R&D investment. *Energy Economics* July, 154–158.
- Lucas, R. (2003). Macroeconomic Priorities. *American Economic Review* 92(1), 1–14.
- Mork, K. (1989). Oil and the macroeconomy when prices go up and down: an extension of Hamilton's results. *Journal of Political* 97(3), 740–744.
- Nelder, J. and R. Mead (1965). A simplex method for function minimization. *The Computer Journal*.
- Newell, R., A. Jaffe, and R. Stavins (1999). The induced innovation hypothesis and energy saving technological change. *Quarterly Journal of Economics* 114(3), 941–975.
- Nordhaus, W. (2002). Modeling induced innovation in climate change policy. *Resources for the Future Press*.

- Nordhaus, W. (2008). *A Question of Balance: Economic Modeling of Global Warming*. Yale University Press.
- Nordhaus, W. and J. Boyer (2000). *Warming the World: Economic Models of Global Warming*. The MIT Press.
- Papageorgiou, C., M. Saam, and P. Schulte (2013). Elasticity of substitution between clean and dirty energy inputs- a macroeconomic perspective. *Working paper*.
- Perlin, J. (2000). *From Space to Earth The Story of Solar Electricity*. Harvard University Press.
- Pew Center on Global Climate Change (2010). Economic Insights from Modeling Analyses of H.R. 2454- the American Clean Energy and Security Act (Waxman-Markey). *Innovative Policy Solutions to Global Climate Change*. January, 1–24.
- Popp, D. (2002). Induced innovation and energy prices. *American Economic Review* 92(1), 160–180.
- Popp, D. (2004). ENTICE: Endogenous technological change in the DICE model of global warming. *Journal of Environmental Economics and Management* 48(1), 742–768.
- Popp, D. and R. Newell (2012). Where does energy R&D come from? Examining crowding out from energy R&D. *Energy economics* 34, 980–991.
- Popp, D., R. Newell, and A. Jaffe (2009). Energy, the environment, and technical change. *NBER Working Paper 14832*.
- Priest, T. (2007). *The Offshore Imperative*. Texas A&M.
- Sandor, R., N. Clark, M. Kanakasabai, and R. Marques (2014). *Environmental Markets: A New Asset Class*. CFA Institute Research Foundation.
- Schempf, J. (2007). *Pioneering Offshore: the early years*. PennWell Custom Publishing.
- Smulders, J. and M. de Nooij (2003). The impact of energy conservation on technology and economic growth. *Resource and Energy Economics* 25, 59–79.

van der Zwaan, B., R. Gerlagh, G. Klaassen, and L. Schrattenholzer (2002). Endogenous technological change in climate change modeling. *Energy Economics* 24 (1), 1–19.

A Derivations of the main equations

I derive the main equations in the text. For ease of presentation, some of the equations are repeated. The final goods producer chooses F , G , N , and O^* to maximize profits taking prices as given. His optimization problem is (equation (9) in the text)

$$\max_{F_t, G_t, N_t, O_t^*} \{Y_t - (P_{ft} + \tau_f)F_t - P_{gt}G_t - (P_{ot}^* + \tau_o)O_t^* - P_{nt}N_t\}, \quad (15)$$

subject to the production technology defined in equations (1) and (2). The first order conditions imply the relative demands for the intermediate inputs are inversely related to their prices,

$$\begin{aligned} \frac{G_t^d}{\tilde{F}_t^d} &= \left(\frac{P_{\tilde{F}t}}{P_{gt}} \right)^{\varepsilon_e}, & \frac{F_t^d}{(O_t^*)^d} &= \left(\frac{P_{ot}^* + \tau_o}{P_{ft} + \tau_f} \right)^{\varepsilon_f} \left(\frac{\delta_{\tilde{F}}}{1 - \delta_{\tilde{F}}} \right)^{\varepsilon_f}, \\ \frac{E^d}{N^d} &= \left(\frac{P_n}{P_e} \right)^{\varepsilon_y} \left(\frac{\delta_y}{1 - \delta_y} \right)^{\varepsilon_y}. \end{aligned} \quad (16)$$

Variables $P_{\tilde{F}}$ and P_e denote the tax-inclusive prices of optimally chosen composites \tilde{F} and E , respectively. The first order and zero profit conditions implies that these prices are

$$\begin{aligned} P_{\tilde{F}t} &= (\delta_{\tilde{F}}^{\varepsilon_f} (P_{ft} + \tau_f)^{1-\varepsilon_f} + (1 - \delta_{\tilde{F}})^{\varepsilon_f} (P_{ot}^* + \tau_o)^{1-\varepsilon_f})^{\frac{1}{1-\varepsilon_f}} \\ P_{et} &= (P_{\tilde{F}t}^{1-\varepsilon_e} + P_{gt}^{1-\varepsilon_e})^{\frac{1}{1-\varepsilon_e}}. \end{aligned} \quad (17)$$

The final good is the numeraire. I normalize its price to unity. This yields the ideal price index

$$\delta_y^{\varepsilon_y} P_{et}^{1-\varepsilon_y} + (1 - \delta_y)^{\varepsilon_y} P_{nt}^{1-\varepsilon_y} = P_{yt} \equiv 1. \quad (18)$$

The intermediate-goods producers make fossil, green, and non-energy inputs which they sell to the final-good producer. I discuss the equations with respect to a representative fossil-

energy producer; the other sectors are symmetric. The fossil-energy producer chooses labor and machines to maximize profits taking prices as given,

$$\max_{L_{ft}, X_{fit}} P_{ft} L_{ft}^{1-\alpha_f} \int_0^1 X_{fit}^{\alpha_f} A_{fit}^{1-\alpha_f} di - w_{lft} L_{ft} - \int_0^1 P_{fit}^x X_{fit} di. \quad (19)$$

Variable P_{fi}^x denotes the price of machine i in sector F . The first order condition for X_{fit} implies the demand for machines

$$X_{fit} = \left(\frac{\alpha_f P_{ft}}{P_{fit}^x} \right)^{\frac{1}{1-\alpha_f}} A_{fit} L_{ft}, \quad (20)$$

where value $\frac{1}{1-\alpha_f}$ is the price elasticity of demand for machines. The first order condition for L_{ft} implies the wages to workers in sector F ,

$$w_{lft} = (1 - \alpha_f) P_{ft} X_{ft}^{\alpha_f} L_{ft}^{-\alpha_f} A_{ft}^{1-\alpha_f}. \quad (21)$$

In equilibrium, labor market clearing requires that workers' wages are equated across all sectors, $w_{lft} = w_{lgt} = w_{lnt}$, and that total labor demand equal the fixed, exogenous supply, $L_{ft} + L_{gt} + L_{nt} = L$.

The machine producers make machines which they sell to the intermediate-goods producers. The machines embody technology. Each machine, regardless of the sector and the level of technology, costs one unit of final good to produce. Each machine producer chooses price, quantity of machines, and the number of scientists to maximize profits subject to the machine demand from the intermediate producer (equation (20)). The optimization problem for fossil-energy machine producer i in period t is

$$\max_{P_{fit}^x, X_{fit}, S_{fit}} P_{fit}^x X_{fit} - X_{fit} - w_{sft} S_{fit} \quad (22)$$

subject to the demand for machines, (equation (20)) and the evolution of technology (equation (4) in the text).

The first order condition for the number of machines implies that the optimal machine

price is a constant markup over marginal cost

$$P_{fit}^x = \frac{1}{\alpha_f}. \quad (23)$$

This constant markup arises because the price-elasticity of machine demand, $\frac{1}{1-\alpha_f}$, is constant. Increases in the machine share increase the demand elasticity and decrease the markup.

Finally, the first order conditions for the number of scientists imply that the wage to a scientist in sector F is

$$w_{sft} = \frac{\eta\gamma A_{ft-1} \left(\frac{A_{t-1}}{A_{ft-1}}\right)^\phi P_{fit}^x X_{fit}}{\rho_f^\eta \left(\frac{1}{1-\alpha_f}\right) S_{ft}^{1-\eta} A_{fit}}. \quad (24)$$

Since the market for scientists is perfectly competitive, the wage of a scientist in the fossil energy sector equals the marginal return to innovation in that sector. As discussed in Section 6, changes in the machine share, α_f have potentially offsetting effects on the returns to fossil innovation. First, from equation (20), a decrease in α_f reduces fossil energy machine demand, lowering X_{fit} and the marginal returns to innovation (equation (24)). Working in the other direction, a decrease in α_f reduces the price elasticity of demand for machines, raising the optimal markup (equation (23)) and the returns to innovation (equation (24)). Which of these effects dominates depends on the region of the parameter space and on the model's general equilibrium channels (which effect the other endogenously determined quantities in equation (24)).

The equilibrium is symmetric across all machine producers within a sector, and, hence, $P_{fit}^x X_{fit} = P_{ft}^x X_{ft}$. To derive equation (10) in the text, observe that equation (20) implies that $P_{ft}^x X_{ft} = \alpha_f P_{ft} F_t$. Substitute this relationship into equation (24) to get equation (10) in the text.

B Time period and within-sector spillovers

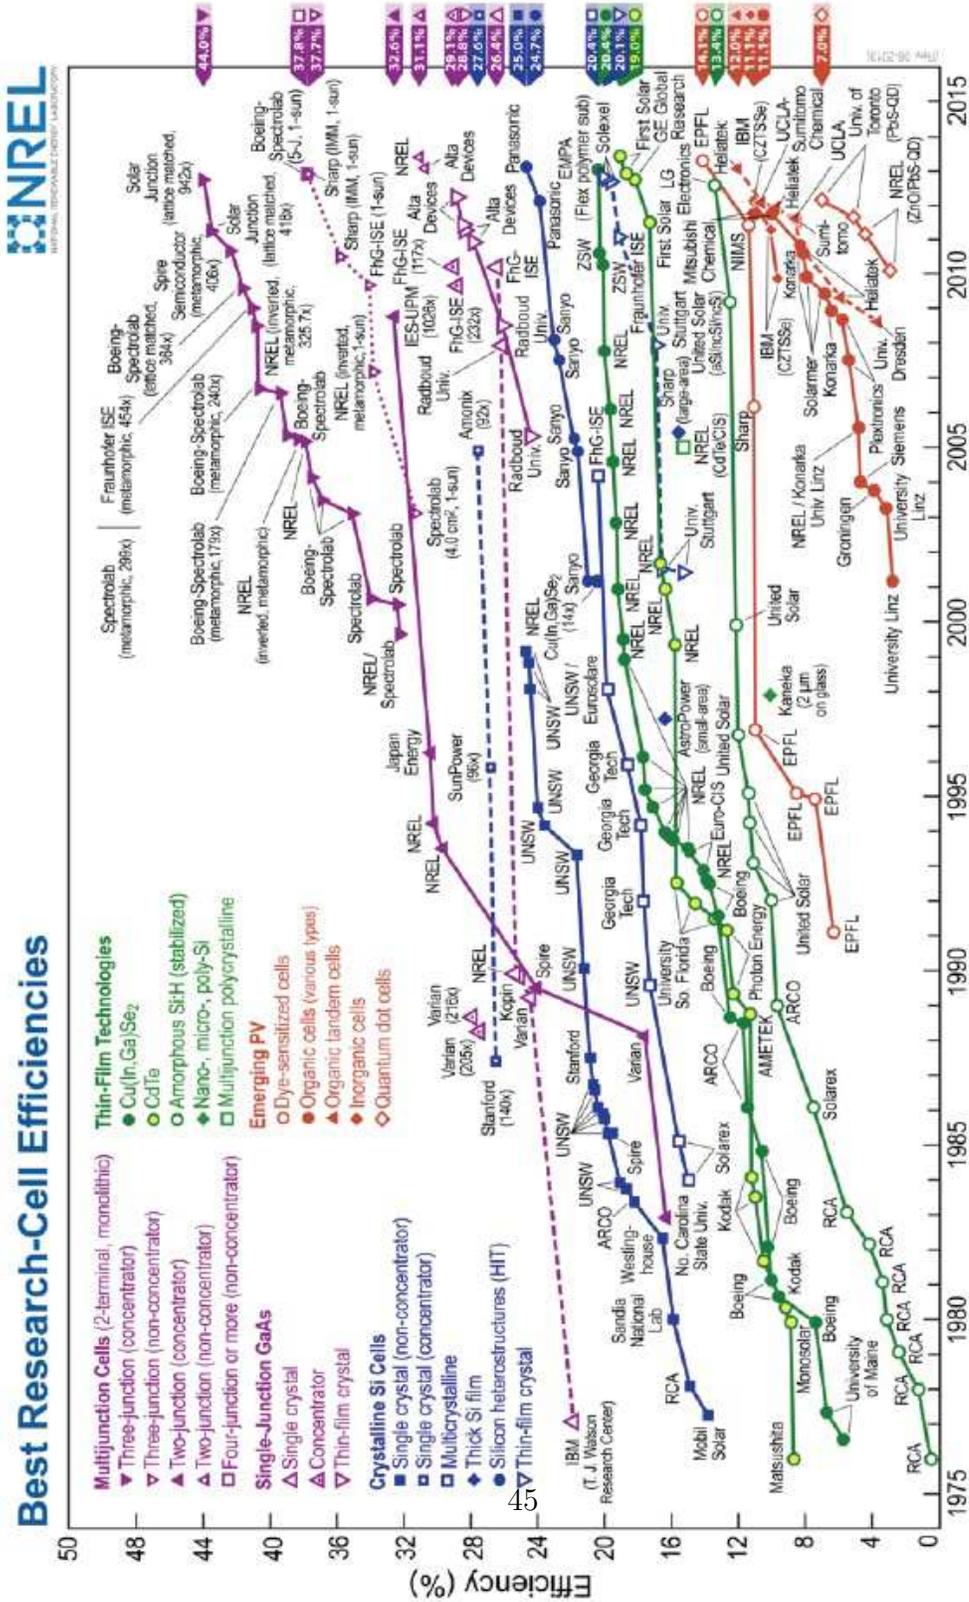
The time period in the model is five years. This choice implies that technology spillovers within a sector occur in five years. To determine this time period, I examine the rate of

technology spillovers experienced in a green and in a fossil industry. In particular, I focus on solar power and offshore drilling. In both cases, within-sector technology spillovers frequently occur in less than five years.

One form of technology embodied in a solar cell is the cell's efficiency. Cell efficiency measures the ratio of the cell's electrical output to incident energy from sunlight. Higher cell efficiency corresponds to higher technology. Figure 2, from the National Renewable Energy Laboratory, plots advances in cell efficiency from 1970-2010 and the company or research institution that achieved the advance. In most cases, the company with the leading cell efficiency is surpassed by a different company within five years. For example, in 1970, Mobile Solar had the leading efficiency in Single crystal non-concentrator Si cells. In 1978, Renewable Capital Assets (RCA) passed Mobile Solar; in 1980, Sandia National Laboratory passed RCA; and so on. The average length of time that a company or research institution maintains the leading efficiency is 3.84 years. This leapfrogging occurs in less than five years, on average, suggesting that within-sector spillovers over a five-year period are reasonable in the case of solar electricity.

As an example from the fossil energy sector, I consider the development of offshore drilling technology. An early technological advance in the offshore industry occurred in 1954, when the Offshore Drilling and Exploration Company (ODECO) developed the first submersible drilling barge, "Mr. Charlie." Mr. Charlie was designed to drill in what was considered deep water at the time (thirty feet). By 1957, just three years after Mr. Charlie's introduction, there were 23 such units in operation in the Gulf and 14 more under construction by numerous oil companies, including Zapata Offshore Company (founded by George H.W. Bush), ODECO, and others. Thus, in less than five years, the technology embodied in ODECO's Mr. Charlie spilled over to other major players in the industry (Schempf (2007), National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling).

Figure 2: Leading Solar Cell Efficiencies



A second major development in offshore drilling occurred in 1962, when Shell Oil launched Blue Water 1, a semi-submersible floating drilling platform that was equipped to operate in up to 600 feet of water (previous platforms could not exceed 150 feet). However, when Shell tried to lease the land for drilling, it was the only bidder on some of the deepwater tracts, and the government refused to honor bids without competition. Since no other companies could operate at those depths, “[Shell] realized that the only way [it] could ever have access to those frontier areas was to share [its] knowledge with the rest of the industry, to give them a base of technology from which they could expand” (Ron Geer, Shell mechanical engineer). In 1963, Shell hosted a “school for industries” in which it shared its frontier deep water technology with seven other companies. By 1968, these within-sector spillovers had led to the construction of 23 Blue-Water-like semi-submersibles and opened up deeper and rougher areas of the ocean to oil drilling and exploration (Priest (2007)).

Shell Oil continued its advance into deeper waters and, in 1976, constructed “Cognac,” a fixed platform connected to a well in 1000 feet of water in the Gulf. At the time, Cognac was the most costly and technologically advanced fixed platform installation ever completed. But within five years of its construction, other companies innovated on Cognac’s design and built similar platforms for much less money. To emphasize its cost savings compared to Cognac, Union Oil named its two 1000-foot platforms constructed from 1980-1981, “Cerveza” and “Cerveza-light.” But as energy historian Tyler Priest notes, “these beer-budget projects could not have happened without the deep water precedent established by Cognac” (Priest (2007)). Again, the development and diffusion of offshore drilling technology suggests that within-sector spillovers often occur in less than five years.

C Robustness with respect to the method of moments

I consider two robustness checks with respect to the method-of-moments calibration procedure. First, as is standard in much of the macroeconomic literature, I evaluate the fit of the model on five moments that were not directly targeted: the percent change in GDP per capita between the last 5 year period on the balanced growth path (1966-1970) and the shock period (1971-1975), the fraction of total labor in the fossil energy sector and the fraction of total capital (fixed assets) in the fossil energy sector in both the balanced growth path and the shock period. While matching these moments is not as crucial for capturing

the innovation incentives as matching the targeted moments, the model’s performance with respect to these not-targeted moments provides one measure of the model’s fit to the data. Table 5 reports the results from this robustness exercise. The values of these moments are relatively similar in the model and the data, suggesting that the model’s fit is reasonably strong.

Table 5: Not Targeted Moments

	Time Period	Model Value	Empirical Value
Fraction of labor in fossil	1961-1970	0.01	0.01
	1971-1975	0.01	0.01
Fraction of capital in fossil	1961-1970	0.05	0.03
	1971-1975	0.05	0.03
% Δ GDP in per capita	-	9.16	9.95

I consider a second robustness check of the method-of-moments procedure to address concerns that one limitation of the calibration strategy is that the early 1970s oil shocks occurred forty years ago. It’s possible that some of the parameter values could have changed over time. To mitigate this concern, I compare the responsiveness of innovation to the 2003 oil shock in the data and in the calibrated model. I do not change the calibration of the model in this exercise; the parameter values are the same as those listed in Table 1. The goal of this exercise is to determine if the model calibrated to the 1970s oil shock matches the innovation response to the 2003 oil shock.

As with the early 1970s oil shocks, I begin the simulation on a balanced growth path and then introduce an oil shock. I choose the size of the shock to match the average increase in the price of oil imports during the 2003-2007 time period compared to the previous five years (1998-2002). I calculate the percent change in fossil and green innovation relative to the percent change in the price of oil imports in both the simulation and in the data. I refer to this quantity as the elasticity of innovation with respect to the price of oil imports.³⁶ Table 6 reports these elasticities in the model and in the data.

³⁶A rigorous empirical estimate of the elasticity fossil and green energy innovation with respect to a change in the oil import price is beyond the scope of this paper. The goal of this robustness check is to show that the model generally matches a back-of-the-envelope calculation of the innovation responses to oil price changes in the more recent time periods.

Table 6: Elasticity of innovation with respect to the oil import price

	Model	Data
Fossil Innovation	0.8	0.7
Green Innovation	0.5	0.7
Total Energy Innovation	0.7	0.7

The empirical and model estimates are similar, suggesting that the parameters governing the responsiveness of energy innovation to price changes have not changed substantially over time. The largest discrepancy between the model and the data is that the price elasticity of green innovation is lower in the model. This difference could be partly explained by policies or expectations of policies which encouraged green innovation during this time. Over this period, an increasing number of states adopted renewable portfolio standards, Europe implemented the pilot phase of its carbon trading system (EU-ETS), and congress and the president began to lay the groundwork for the Energy Independence and Security Act of 2007. All these policy developments would encourage green innovation and lead to a higher empirical price elasticity.

This exercise has two caveats. First, energy prices and energy innovation are not stable for a sustained period preceding the 2003 oil shock, suggesting that the assumption that the economy was on a long-run balanced growth path prior to the shock is imperfect. Second, following the 1970s, agents learned that energy prices are uncertain and they formed expectations over future energy prices. I do not model expectations in this robustness analysis. The advantage of calibrating the model using the early 1970s' oil shock instead of the 2003 oil shock is that it avoids both of these caveats. Energy prices were relatively flat for the twenty years prior to the early 1970s oil shock, suggesting the economy could plausibly have been on a balanced growth path and that agents did not anticipate the shock.

D Standard error estimates

I jointly calibrate the parameters $\{\alpha_g, \varepsilon_f, \delta_{\bar{F}}, \delta_y, \eta, \gamma\}$ and the productivity shock, ν , to capture the relationships between energy prices, production, and innovation. To obtain empirical evidence of these relationships, I analyze the energy price increases triggered by historical

oil shocks in the early 1970s. Specifically, I calibrate the parameters to match moments generated by the following experiment in the model with moments generated by the oil and productivity shocks of the early 1970s in the US economy.

Initial balanced growth path (1961-1970): The economy is on a balanced growth path with respect to the price of oil imports and environmental and energy policies.

Shock period (1971-1975): Two unexpected shocks realize: (1) the price of oil imports increases from its value on the balanced growth path; and (2) a negative productivity shock affects domestic fossil energy production.

I target two moments from the initial balanced growth path: (1) average fossil energy production as a share of GDP and (2) the average value of oil imports as a share of GDP. I target four moments from the shock period: (1) average fossil energy production as a share of GDP, (2) the average value of oil imports as a share of GDP, (3) average fossil energy R&D expenditures as a fraction of total R&D expenditures, and (4) average green energy R&D expenditures as a fraction of total R&D expenditures. Table 2 reports the empirical values of these moments. Additionally, I target the long-run annual growth rate of GDP per capita of 2 percent. See Section 4 for more details on these moments and their selection process. I choose the parameter values that minimize the sum of the square of the residuals between the empirical and model values of the moments.

The provision of valid standard errors for my parameter estimates thus requires a model of the data generating process for these moments. This is difficult as the moments are realized only once and my model provides no guidance on their stochastic generating process. Furthermore, while I can reasonably assume that the values of the moments on the balanced growth path are constant between 1961-1970, this assumption is clearly invalid during the shock period, 1971-1975. Mindful of these concerns, I take the view that the standard errors should reflect uncertainty in my parameter estimation induced by idiosyncratic variation in the measured values of my moments. I propose a resampling procedure that uses annual variation over the balanced growth path to approximate this.

Specifically, I interpret any annual variation in the value of the balanced growth path moments as random noise. I resample each balanced growth path moment as its average

value during the balanced growth path plus a random error drawn from a normal distribution with mean zero and standard deviation equal to the standard deviation of the annual value of the moment along the balanced growth path.

I interpret any annual variation in the value of the shock period moments as both random error and the economic response to the shocks of the early 1970s. However, if each shock period moment was instead calculated over the balanced growth path period, its annual value should be constant. Thus, I resample each shock period moment as its average value during the shock period plus a random error drawn from a normal distribution with mean zero and standard deviation equal to the standard deviation of the annual value of the moment along the balanced growth path. Data on fossil and green R&D are not available during the balanced growth path period. For these two moments, I use the standard deviation of the fraction of total R&D in the petroleum refining and extraction sectors (SIC codes 13 and 29) on the balanced growth path instead of the standard deviation of the fractions of fossil and green R&D.³⁷

I resample the moments and reestimate the parameters 300 times. I then compute the bootstrapped standard error of each estimated parameter.³⁸ Table 7 reports the parameter values for the baseline specification with the bootstrapped standard errors in parentheses. With the exception of γ and δ_y , the parameter estimates demonstrate a reasonable degree of precision. However, these standard errors should be interpreted with caution. Their computation relies critically on the assumption that the idiosyncratic variation in moments in the balanced growth path is the same as in the shock period. Moreover, annual data over a ten year period is not necessarily representative of the true variance in the empirical distribution of these moments.

³⁷I rescale the estimate of the standard deviation of the fraction of petroleum R&D to account for small differences in the means between the fractions of fossil and green energy R&D during the shock period and the fraction of petroleum R&D on the balanced growth path.

³⁸For 76 of the 300 sets of moments, there did not exist parameter values such that the model was consistent with the set of moments. I removed these draws from the calculation of the standard error.

Table 7: Standard Errors for Parameter Estimates

Green energy machine share: α_g	0.9 (0.04)
Diminishing returns: η	0.8 (0.24)
Scientist efficiency: γ	4.0 (3.50)
Fossil elasticity of substitution: ε_f	6.2 (3.11)
Distribution parameter: δ_y	1.4e-38 (5.0e-38)
Distribution parameter: $\delta_{\bar{F}}$	0.5 (0.05)
Productivity shock: ν	0.6 (0.05)

E Target stringency and time frame

A key finding is that the carbon tax necessary to achieve the Power Plan target is 19.2 percent lower when innovation is endogenous. However, this result is sensitive to both the size of the targeted reduction in emissions and the time frame in which the reduction must be achieved. In this section, I analyze the effects of innovation for different-sized emissions targets and different time frames.

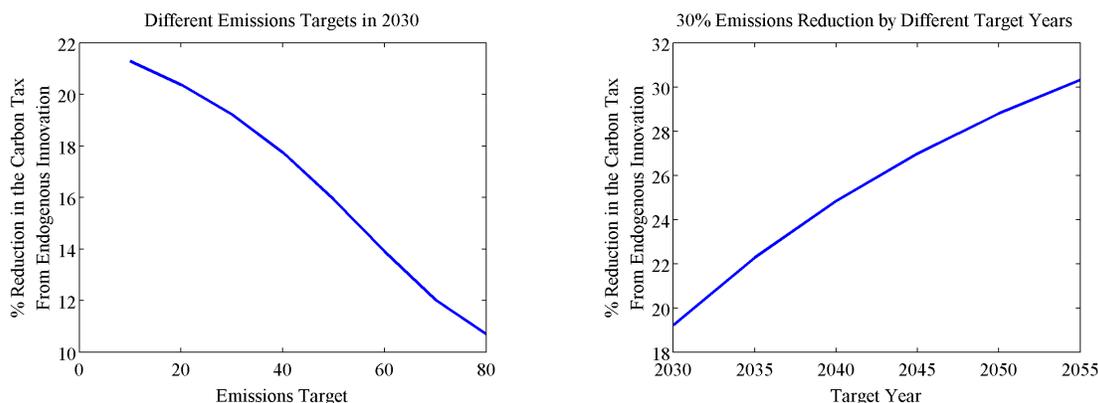
The left panel of Figure 3 plots the percent reduction in the carbon tax from endogenous innovation for different-sized emissions targets. For example, the carbon tax required to achieve a 10-percent reduction in emissions by 2030 is 21 percent lower when innovation is endogenous. The effect of endogenous innovation on the size of the carbon tax falls as the stringency of the emissions target increases (e.g., as the target goes from a 10-percent to a 20-percent reduction in emissions). Even with large changes in innovation, the relative technology stocks evolve slowly. A more stringent emissions target forces agents to rely less on technological advances and more on shifts in production factors (i.e., workers and

machines) to achieve the target. This switch reduces the role of endogenous innovation and its accompanying effects on the carbon tax.

The right panel of Figure 3 plots the percent decrease from endogenous innovation under a carbon tax designed to achieve a 30-percent reduction in emissions by different target years. For example, the carbon tax required to achieve a 30-percent reduction in emissions by 2035 is 22 percent lower when innovation is endogenous. The reduction in the carbon tax from endogenous innovation increases as the target year moves farther into the future. Again, even with large shifts in innovation, changes in the relative technology stocks occur gradually. More-distant target years provide time for technological change to occur and thereby allow agents to rely more heavily on innovation to reduce emissions.

Climate policy simulation models that exogenously assume large advances in green technological progress typically obtain lower carbon tax estimates for a given abatement target (Pew Center on Global Climate Change (2010)). As long as the time horizon for a given emissions target is sufficiently long, the results of this paper suggest that such technological advances are plausible and could lead to considerable reductions in the carbon tax. However, if policy makers strive to achieve large reductions in emissions quickly, then the potential for innovation to reduce the carbon tax is relatively small.

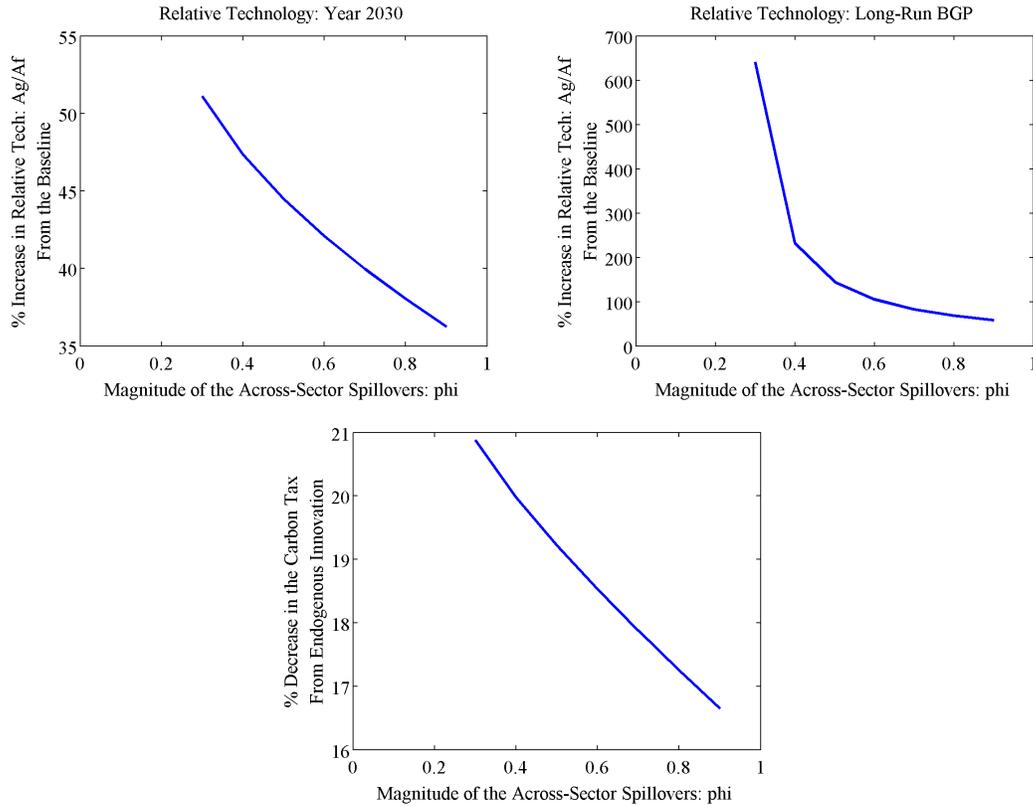
Figure 3: Effects of Innovation on the Size of the Carbon Tax



F Additional parameter robustness

F.1 Across sector spillovers: ϕ

Figure 4: Effects of ϕ on Changes in Relative Technologies and the Carbon Tax



I reexamine the main results when ϕ ranges from 0.3 to 0.9, the values such that the interior balanced growth path is stable. As with the robustness analysis over the parameters in Section 6, for each value of ϕ , I recalibrate the six parameters $\{\alpha_g, \eta, \gamma, \varepsilon_f, \delta_y, \delta_{\bar{F}}\}$ and the productivity shock, ν , to ensure that the model matches the moments described in Section 4. The top two panels of Figure 4 show the effects of ϕ on the response of relative technologies to a carbon tax that obtains the Power Plan target. Stronger spillovers dampen the change in the relative technology in response to the carbon tax. The long-run effects exceed the short-run effects because the TFP catchup ratio evolves slowly.

The bottom panel of Figure 4 plots the percent reduction in the carbon tax from endogenous innovation for different values of ϕ . Larger spillovers dampen the response of

innovation, decreasing the percent reduction in the carbon tax. However, even for very large spillovers, $\phi = 0.9$, endogenous innovation still reduces the carbon tax by over 15 percent.

F.2 Level of the carbon tax

For each parameter perturbation, Table 4 in Section 6 reports the percent change in the carbon tax required to achieve the Power Plan target from endogenous innovation. Tables 8 and 9 report the levels of the carbon tax required to achieve the target for each parameter perturbation in the endogenous- and exogenous-innovation models, respectively. In most of the perturbations, the level of the carbon tax is relatively similar to the tax in the main analysis. The largest differences occur for changes in ε_e , the elasticity of substitution between the green energy and the composite comprised of fossil energy and oil imports. Some of this effect is due to the larger range values for ε_e ; the low and high cases correspond to $\varepsilon_e = 1.1$ and 3, respectively as opposed to the 25 percent changes from the central value used for the other parameters. Even so, the result is intuitive; as the substitutability between these energy sources increases, the size of the carbon tax necessary to achieve the Power Plan target falls.

Table 8: Carbon Tax With Endog. Innovation (in 2013 \$ per ton CO_2)

Parameter	High	Central	Low
Imposed parameters			
Output elasticity of substitution: ε_y	24.5	24.5	-
Across sector spillovers: ϕ	24.5	24.5	24.5
Direct from data series			
Machine share in fossil energy: α_f	25.6	24.5	28.9
Machine share in non-energy: α_n	22.4	24.5	28.3
Energy elasticity of substitution: [†] ε_e	8.4	24.5	40.6
Sector diversity: ρ_f	24.9	24.5	23.9
Sector diversity: ρ_g	24.3	24.5	24.6
Population of scientists: S	24.5	24.5	24.5
Method of moments			
Machine share in green energy: α_g	24.1	24.5	20.4
Diminishing returns: η	19.6	24.5	26.7
Fossil elasticity of substitution: ε_f	25.4	24.5	23.5
Distribution parameter: δ_y	24.5	24.5	24.4
Distribution parameter: $\delta_{\tilde{F}}$	28.2	24.5	15.6
1971-1975 productivity shock: ν	24.5	24.5	24.5

[†]The high case for ε_e is $\varepsilon_e = 3$ and the low case for ε_e is $\varepsilon_e = 1.1$.

Table 9: Carbon Tax With Exog. Innovation (in 2013 \$ per ton CO_2)

Parameter	High	Central	Low
Imposed parameters			
Output elasticity of substitution: ε_y	30.2	30.3	-
Across sector spillovers: ϕ	30.0	24.5	30.7
Direct from data series			
Machine share in fossil energy: α_f	29.5	30.3	31.7
Machine share in non-energy: α_n	30.4	30.3	29.8
Energy elasticity of substitution: [†] ε_e	10.5	30.3	49.2
Sector diversity: ρ_f	30.9	30.3	29.5
Sector diversity: ρ_g	30.1	30.3	30.5
Population of scientists: S	30.3	30.3	30.3
Method of moments			
Machine share in green energy: α_g	24.1	30.3	27.8
Diminishing returns: η	30.6	30.3	29.9
Fossil elasticity of substitution: ε_f	31.2	30.3	29.4
Distribution parameter: δ_y	30.3	30.3	30.2
Distribution parameter: $\delta_{\tilde{F}}$	33.9	30.3	23.2
1971-1975 productivity shock: ν	30.3	30.3	30.3

[†]The high case for ε_e is $\varepsilon_e = 3$ and the low case for ε_e is $\varepsilon_e = 1.1$.

F.3 Price elasticity of green ideas

Table 10 analyzes an alternative summary statistic for the different parameter perturbations; the price elasticity of green ideas defined in equations (13) and (14). As discussed in Section 4.5, the green innovation elasticity implied by the quantitative model is larger than the empirical estimates in Popp (2002) but smaller than the empirical estimates in Aghion et al. (2015).

This elasticity is endogenous to the model. However, the parameter perturbations of the non-energy machine share can bring the elasticity close to the empirical values estimated in Popp (2002) and Aghion et al. (2015). In the high case, the elasticity is 3.2, smaller than but close to Aghion et al. (2015)'s value of 3.7, and in the low case, the elasticity is 0.2, approximately equal to Popp (2002)'s value of 0.21. The percent change in the carbon tax from endogenous innovation is 26.2 when the elasticity equals 3.2 and is 4.9 when the elasticity equals 0.2. Thus, an elasticity closer to Aghion et al. (2015)'s value would increase the importance of endogenous innovation while an elasticity closer to Popp (2002)'s value

would decrease its importance.³⁹

Table 10: Price Elasticity of Green Ideas

Parameter	High	Central	Low
Imposed parameters			
Output elasticity of substitution: ε_y	1.7	1.7	-
Across sector spillovers: ϕ	1.6	1.7	1.8
Direct from data series			
Machine share in fossil energy: α_f	5.2	1.7	0.3
Machine share in non-energy: α_n	3.2	1.7	0.2
Energy elasticity of substitution: [†] ε_e	0.5	1.7	1.6
Sector diversity: ρ_f	1.8	1.7	1.6
Sector diversity: ρ_g	1.7	1.7	1.7
Population of scientists: S	1.7	1.7	1.7
Method of moments			
Machine share in green energy: α_g	1.7	1.7	1.4
Diminishing returns: η	5.6	1.7	0.7
Scientist efficiency: γ	1.8	1.7	1.7
Fossil elasticity of substitution: ε_f	1.7	1.7	1.7
Distribution parameter: δ_y	1.7	1.7	1.7
Distribution parameter: $\delta_{\tilde{F}}$	1.9	1.7	1.5
1971-1975 productivity shock: ν	1.7	1.7	1.7

[†]The high case for ε_e is $\varepsilon_e = 3$ and the low case for ε_e is $\varepsilon_e = 1.1$.

³⁹The large changes in the elasticity in the high and low cases for α_f and α_n arise because recalibrating the model with alternative values for these parameters results in substantially different values for the diminishing returns to innovation, η . Similarly, the low elasticity when $\varepsilon_e = 3$ occurs because matching the targeted moments with the high value of ε_e requires much stronger diminishing returns to innovation than in the main specification. See footnote 35 for further discussion.