Environmental Policy and Misallocation: The Productivity Effect of Intensity Standards

Trevor Tombe
University of Calgary

Jennifer Winter
University of Calgary

This Version: April 2015

Abstract

Firm-level idiosyncratic policy distortions misallocate resources between firms, lowering aggregate productivity. Many environmental policies create such distortions; in particular, output-based intensity standards (which limit firms' energy use or emissions per unit of output) are easier for high-productivity firms to achieve. We investigate the productivity effect of intensity standards using a tractable general-equilibrium model featuring multiple sectors and firm-level heterogeneity. Qualitatively, we demonstrate standards are always inferior to uniform taxes, as they misallocate both dirty and clean inputs across firms and sectors, which lowers productivity. Quantitatively, we calibrate the model to US data and show these productivity losses can be large.

JEL Classification: Q4, Q5, H2, E6

Keywords: Energy intensity, energy tax, productivity

Tombe (Corresponding Author): Assistant Professor, Department of Economics, University of Calgary, 2500 University Drive NW, Calgary, Alberta, T2N1N4. Email: `ttombe@ucalgary.ca`. Winter: Associate Director, Energy and Environmental Policy, The School of Public Policy, University of Calgary. Email: `jwinter@ucalgary.ca`. We would like to thank Younes Ahmadi and Akio Yamazaki for valuable research assistance. We also thank Jevan Cherniwchan, Andrew Leach, Derek Lemoine, Carol McAusland, Alaz Munzur, Scott Taylor, Jiangyang Zhang, seminar participants at Laval and Calgary and various conference participants at CREE 2012 and CEA 2013 for valuable comments and suggestions. We gratefully acknowledge the financial support of The School of Public Policy and the Social Sciences and Humanities Research Council of Canada (IDG 430-2012-421).

1 Introduction

Firm-level policy distortions – such as differences in tax rates or regulatory treatment across firms – misallocate resources and can substantially lower aggregate productivity (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Brandt et al., 2013; Bartelsman et al., 2013). Many environmental policies are firm-specific and therefore create such misallocations. In particular, intensity standards (or output-based regulations) place limits on energy use or emissions per unit of output, which are easier for high-productivity firms to achieve than for low-productivity firms. These policies are increasingly common and a growing literature compares standards to flat energy or emissions taxes. While taxes are often first-best, a number of factors favour standards: market power (Holland, 2009; Li and Shi, 2012); incomplete regulation or leakage (Holland, 2012); learning-by-doing in production (Gerlagh and van der Zwaan, 2006); certain pre-existing tax distortions (Parry and Williams, 2011); or unexpected business-cycle productivity shocks (Fischer and Springborn, 2011). To date, however, there is no quantitative exploration of the misallocation and productivity effects of environmental policy in general and intensity standards in particular. We fill this gap.

What do we mean by misallocation? And how do intensity standards misallocate resources and lower productivity? Inputs to production, such as labour or energy inputs, are misallocated if marginal revenue products differ across producers. With such differences, total output will grow if a worker moves from where she is valued less to where she is valued more. What matters for total output is therefore not only the underlying productivity of individual producers, but also the allocation of resources between them. Environmental policies in general will increase the costs of polluting or the costs of using of "dirty" inputs, but intensity standards in particular will increase those costs more for low-productivity firms than for high-productivity firms.

The existing literature, starting perhaps with Helfand (1991), highlights two distinct distortions that intensity standards create. First, standards impose an implicit tax on a producer's emissions or on its use of dirty inputs. Second, they provide an implicit subsidy to a producer's output. These taxes and subsidies are *implicit* as neither results in actual payments by or to producers. Instead, standards impose constraints on producer choices that can be equivalently achieved through a tax on dirty inputs with all proceeds rebated to the producer as a subsidy to its output. The output subsidy is important, as marginal costs of production under a standard are *lower* than under an explicit tax. Holland (2012) demonstrates this implies a social planner would choose a higher level of emissions with a standard than it would with an emissions tax. Instead of exploring the optimal level of emissions, we compare alternative policies that each achieve the same environmental objective. Our focus is therefore on sectoral and aggregate productivity (output per bundle of inputs), and *not* on welfare. We show the implicit tax on dirty inputs and the implicit output subsidy *both* misallocate resources across firms and sectors. Both distortions combine to misallocate dirty inputs – this is closely related to the widely understood requirement that optimal policy equalizes marginal abatement costs. In addition, the output subsidy misallocates clean inputs – this is a distinct (and previously unexplored) general-equilibrium effect of standards.

We investigate two broad types of intensity standards: (1) improvement targets and (2) level targets. Improvement targets require that firms achieve a given percentage reduction in their energy intensity (relative to their own baseline). Level targets, on the other hand, require that firms meet a common energy intensity, which is normally a percentage reduction in the *sector's* average energy intensity. Alternatively, improvement targets are *firm-specific* and level targets are *sector-specific*. A third type of intensity target is a variation of firm-specific targets that applies only to large emitters; we call these *threshold-based* targets. For example, the U.S. EPA regulates facilities emitting more than 100,000 tons per year of CO_2-equivalent. The Canadian province of Alberta does the same. Each of these policies will have different distortionary effects on energy input costs. Improvement targets increase energy input costs in energy-intensive industries by more than in other industries, while level targets increase energy input costs by more for low-productivity firms. Our model allows us to evaluate the effect of these distortions on allocations, and therefore productivity, across firms and sectors.

To perform this analysis, we build a tractable quantitative model that cleanly maps into readily available sector-level data on production and energy use. The model builds heavily on insights from macroeconomics (in particular, Hsieh and Klenow, 2009, and Jones, 2013) and provides a useful tool to analyze environmental policies. The broad features of the model can be quickly summarized. Aggregate GDP is a composite of N goods produced by multiple sectors. To produce these goods, each sector aggregates the output from a continuum of firms that produce horizontally differentiated varieties. At the individual firm level, production is similar to Copeland and Taylor (1994) where firms use labour for production and abatement activities. Firms also require a certain level of energy, that is increasing in production but decreasing in how many workers are used in abatement.[1] This structure results in a simple Cobb-Douglas production technology, where all firms are "dirty" but vary in their "dirtiness".

The structure allows for powerful qualitative results and tractable quantitative analysis. We provide expressions that transparently link labour and energy allocations to sectoral and aggregate productivity. We further provide analytic propositions that establish the inferiority of intensity standards due to resource misallocation. This holds for *any* admissible parameter values in the model. Building on these propositions, we undertake a variety of illustrative quantitative simulations to show the magnitude of these effects are often substantial. Sector-specific standards are particularly damaging, resulting in the most misallocation and (in the expanded model) firm exit. For example, a sector-specific target that lowers total energy use by 10 per cent will lower aggregate productivity by nearly 0.15 per cent while a uniform tax on energy has *zero* effect on productivity. If energy use and emissions are proportional, consistent with evidence in the next section, this is equivalent to the standard costing \$35 *more* per tonne of CO_2 abated than a uniform energy tax. Our quantitative analysis is only illustrative, but highlights that misallocation from

[1] In Copeland and Taylor (1994), emissions are a by-product of production. The two approaches are largely equivalent, except emissions are *free* while energy is *purchased*. This introduces some difficulties for our model calibration and quantitative exercises. In Appendix B we demonstrate our key qualitative results hold when emissions are a by-product. We explore the empirical relationship between energy use and emissions in Section 2.

intensity standards can be large. We also extend the model to allow (1) entry and exit of firms and (2) more general substitution possibilities between industries. Firm exit will amplify the negative effects of standards while more substitutability between sectors dampen them.

Of course, there are ways to improve the performance of intensity standards. If intensity standards are implemented as tradable permits, their economic costs in our framework are identical to a tax.[2] Similarly, fines for violating standards, the value of which would not differ across firms, can significantly lower the productivity costs of non-tradable intensity standards. With a fine equivalent to a 50% tax on energy use – chosen for illustrative purposes only – the productivity consequences of a standard falls to less than half our baseline estimates. Finally, as we show the costs of standards are even larger when firm exit is considered, policy-makers may consider transfers or other supports to prevent firm exit (though this is not our focus).

Our work fits within a substantial literature in environmental economics, especially those cited earlier that investigate the efficacy of intensity standards. While we are agnostic about the optimal level of energy reduction, we evaluate the efficiency of various policies in achieving a given objective. Our key contribution is to quantitatively evaluate the cost of environmental policies in the presence of firm-level heterogeneity in productivity. While firm heterogeneity has been introduced via limited cost heterogeneity, rich productivity differences across a continuum of firms has, until recently, not been investigated. Most similar to our paper are Li and Sun (2011) and Li and Shi (2012). They show firm heterogeneity is important to evaluate environmental policies, though we differ in important ways. First, our model is a *quantitative* tool that matches key features of production and energy use across multiple industries. Second, the standards we consider differ, and we explore many types. Third, our substantive focus differs, as we explore the misallocation and productivity costs of intensity standards. Finally, we analytically prove intensity standards *cannot* increase productivity, a possibility within their more stylized framework.

Finally, we must emphasize that our approach is conceptually distinct from the large literature on abatement-cost heterogeneity.[3] In that literature, marginal abatement cost curves are typically exogenously assumed to take a certain functional form. When the marginal costs of inputs differ across firms, there exists a reallocation of abatement responsibilities that lowers the total cost of achieving a given aggregate target. The data requirements to estimate such curves are often prohibitive. As mentioned, some of our analysis is related to abatement cost differences; we discuss this in Section 4.3. However, misallocation goes beyond this phenomenon. Output distortions in particular will have a distinct, general equilibrium effect on the allocations between firms and sectors.

The next section begins our analysis with a broad overview of modeling emissions, energy use, and intensity standards. We demonstrate clearly how intensity standards misallocate inputs,

[2]Firms exceeding a mandated intensity can purchase the "spare" capacity of other firms. In our framework, an efficient trading market would equalize the marginal cost of energy and therefore misallocation costs would be zero. For further reference, Goulder and Schein (2013) show permit trading is equivalent to an energy tax. For trading schemes that implement intensity standards, see McKitrick (2005). Recently, de Vries et al. (2014) find cap-and-trade can potentially dominate tradable intensity standards in a model with multiple heterogeneous sectors.

[3]See, for example, Newell and Stavins (2003) for a complete discussion.

how allocations can be equivalently captured by a system of implicit taxes and subsidies, and how inefficient allocations affect productivity. With this broad intuition, we develop our full model in Section 3, with qualitative and quantitative analysis in Sections 4 and 5. To explore firm exit and more general substitution possibilities across sectors, we expand the model and analysis in Section 6. Finally, Section 7 concludes.

2 Intensity Standards and Factor Misallocation

How do we model emissions and intensity standards? How does the allocation of resources affect aggregate productivity? This section explores the answers to these basic questions before proceeding to the heart of our analysis. We begin with a well-established modeling approach to incorporate emissions as a free input in production, following Copeland and Taylor (1994). We also argue an alternative approach, where emissions are proportional to costly energy inputs, proves convenient in many applications, especially the quantitative analysis to come.

With the production technologies broadly defined, we proceed to describe how and why resource allocations matter for aggregate productivity. Following the recent macroeconomics literature on the aggregate effect of firm- and sector-level distortions, we show how intensity standards (and, more broadly, a wide variety of environmental policies) can create implicit taxes that may vary across producers and how this variation lowers productivity. This will prove a powerful modeling approach, both for sharp qualitative results and for fully general equilibrium quantitative simulations.

2.1 Emissions as an Input: A Primer

Copeland and Taylor (1994) show, under certain conditions, emissions that are a by-product of production can be equivalently modeled as an input. Consider, a firm i that allocates labour L_i to production L_i^y or emissions abatement L_i^a. Output is produced using

$$Y_i = A_i \left(\frac{1-\alpha_i}{\alpha_i} \right)^{1-\alpha_i} L_i^y, \tag{1}$$

where A_i is productivity and $\alpha_i \in (0,1)$. One can (and should) interpret labour inputs broadly to include any primary input, such as physical or human capital. Emissions are a by-product of production, assumed to follow

$$E_i = \frac{1-\alpha_i}{\alpha_i} \left[\frac{L_i^y}{\left(L_i^y + L_i^a\right)^{\alpha_i}} \right]^{\frac{1}{1-\alpha_i}}. \tag{2}$$

Without abatement (where $L_i^a = 0$), emissions per worker are $(1-\alpha_i)/\alpha_i$. As labour is allocated to abatement, L_i^a increases and emissions decrease. The parameter α_i may be interpreted as the

effectiveness of the abatement technology. As α_i grows, emissions fall.[4]

In general, total employment must be allocated to either production or abatement, so $L_i = L_i^y + L_i^a$. With this, equations 1 and 2 combine to yield

$$Y_i = A_i L_i^{\alpha_i} E_i^{1-\alpha_i}. \tag{3}$$

Emissions can therefore be thought of as an input (i.e., "environmental services") in the production process along with labour L_i. Intuitively, as more labour L_i is allocated to abatement, output declines. This effect is captured by a lower E_i, which lowers output Y_i.

A powerful alternative to emissions as a by-product is to consider emissions as resulting from purchased energy inputs. This would literally interpret E_i as an input purchased at some positive price from an energy supplier. The Copeland and Taylor (1994) structure above is valid, with abatement interpreted as activity of workers that do not contribute directly to production but instead lowers total energy use.[5] The energy-inputs interpretation dramatically sharpens our qualitative results and eases our calibration and simulation exercises – we maintain this interpretation throughout. In Appendix B, we illustrate how our analysis would differ in a model with emissions as a by-product, and confirm all our key results hold. In any case, the basic Cobb-Douglas production technology of equation 3 will be used throughout this paper, though it is instructive to keep equations 1 and 2 in mind. As E_i/L_i falls, for example, labour used in abatement L_i^a grows: falling E_i/L_i is therefore synonymous with increased abatement activity.

This interpretation is also reasonable, as certain emissions at the sector level are roughly proportional to energy use. To see this, consider the World Input-Output Database's environmental accounts, which provides two measures of energy inputs: total and emissions-relevant. Importantly, emissions-relevant energy use does not include raw material inputs (such as crude oil used by refineries) but only energy consumed in the production process, not energy transformed from one form to another. For the United States, the correlation between sectoral emissions-relevant energy use and CO_2 emissions is 0.998. Figure 1 illustrates this strong relationship. In panel (a) the roughly proportional relationship is clear. This suggests a 10 per cent change in energy use will lead to a 10 per cent change in emissions, for example. In panel (b), we compare total energy inputs with emissions-relevant energy inputs. The two are nearly perfectly correlated, with important deviations for only three sectors. We will be careful to calibrate the model using only data for emissions-relevant energy use.

[4]The Copeland and Taylor (1994) abatement technology is related to that found in Li and Shi (2012). They are identical if the Li and Shi (2012) parameters (b, γ) are $b = \gamma = (1 - \alpha_i)/\alpha_i$. The difference is Li and Shi allow emissions per worker without abatement (b) to differ from the inverse effectiveness of the abatement technology (γ).

[5]More broadly, one might loosely interpret abatement as any costly activity that lowers the use of emissions-relevant energy, such as substitution between different fuel types. We remain focused on general macroeconomic relationships, though such firm-specific considerations may be an interesting avenue for future research.

Figure 1: Sectoral Energy Use and Carbon Emissions (WIOD 2009)

(a) Emissions-Relevant Energy Use vs Emissions

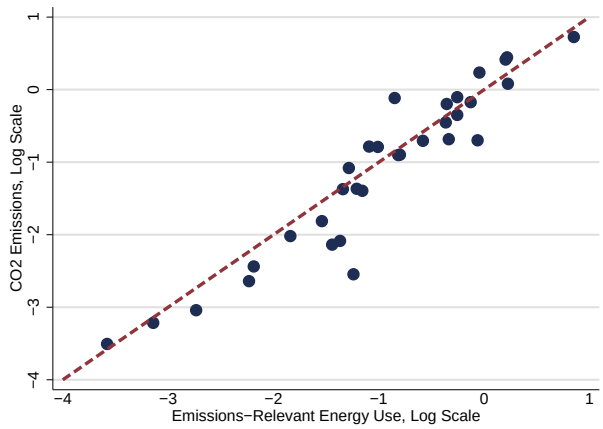

(b) Energy Use, Total vs Emissions-Relevant

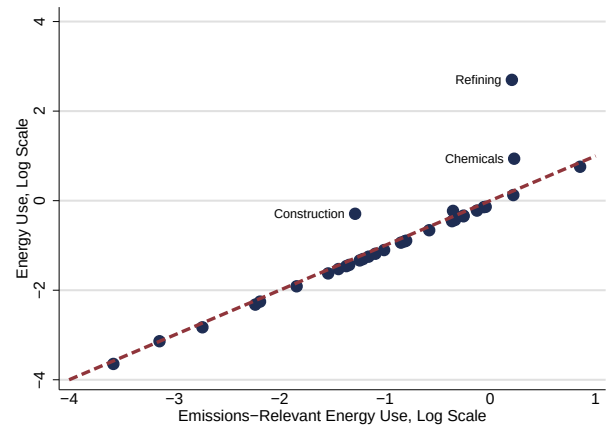

Displays a strong relationship between sectoral emissions-relevant energy use and emissions. The correlation coefficient between the two is 0.998 (0.975 when in logs). Panel (b) illustrates emissions-relevant energy use and total energy use are nearly identical for most sectors, with construction, refining, and chemicals as the exceptions. The dotted line is the 45-degree line. Data is for the United States from the 2009 environmental accounts of the World Input-Output Database.

2.2 Misallocation and Productivity, The Basics

How does input allocations affect productivity? An economy's aggregate productivity depends not only on the productivity of its individual firms but also on the allocation of resources across those firms. With multiple producers, there will be some optimal allocation between producers. When allocations deviate from that optimum, aggregate productivity falls.

To illustrate this, consider maximizing some aggregate $Y = F(y_1, \ldots, y_N)$ of output y_i from N producers. Further suppose all producers have identical α_i, though they may differ in productivity. How should we allocate inputs (L_i, E_i), given some fixed aggregate supply of both, such that Y is maximized and input markets clear? From the first order conditions of the constrained maximization problem, it is straightforward to show the marginal rates of technical substitution must equalize across firms,

$$\frac{dy_i}{dL_i} / \frac{dy_i}{dE_i} = \frac{dy_j}{dL_j} / \frac{dy_j}{dE_j} \quad \forall\,(i,j)\,.$$

For the production technology of equation 3, or indeed for any technology with constant (and common) elasticities of substitution between inputs, the above implies

$$\frac{E_i}{L_i} = \frac{E_j}{L_j} \quad \forall\,(i,j)\,.$$

Deviations from equality in input ratios therefore result in lower aggregate output.

Of course, if input elasticities α_i differ across firms then input ratios will optimally differ. But, the general insight is sound: policies that distort firm input decisions away from some optimal

allocation will lower aggregate productivity. When the full model is developed in Section 3 we will derive a clear expression mapping allocations to productivity, both in aggregate and by sector. Before getting there, we must explore how intensity standards create differences in input ratios between producers. We do this in the next section.

2.3 The Effect of Intensity Standards

Baseline energy intensity, from equations 2 and 3, is

$$\frac{E_i}{Y_i} = \frac{1}{A_i}\left(\frac{E_i}{L_i}\right)^{\alpha_i} = \frac{1}{A_i}\left(\frac{1-\alpha_i}{\alpha_i}\right)^{\alpha_i}, \tag{4}$$

which is clearly decreasing in productivity A_i. What is the effect of a binding limit on energy intensity? Imagine the government does not allow E_i/L_i to exceed some intensity standard s. There is no change to firms that are already below this threshold, but firms above it must substitute labour for energy to bring E_i/L_i below s. Precisely,

$$\frac{E_i}{Y_i} = \min\left\{s, A_i^{-1}\left(\frac{1-\alpha_i}{\alpha_i}\right)^{\alpha_i}\right\}.$$

If the standard is binding, then $E_i = sY_i$ and $Y_i = A_i L_i^{\alpha_i} E_i^{1-\alpha_i} = A_i^{\frac{1}{\alpha_i}} s^{\frac{1-\alpha_i}{\alpha_i}} L_i$. This implies the energy per worker for firms bound by the standard is $(A_i s)^{1/\alpha_i}$ and energy per worker for firms not bound by the standard is $(1-\alpha_i)/\alpha_i$. Across all firms,

$$\frac{E_i}{L_i} = \min\left\{(A_i s)^{1/\alpha_i}, \frac{1-\alpha_i}{\alpha_i}\right\}. \tag{5}$$

We plot these results in Figure 2. Low-productivity firms have high energy-intensity absent policy (the dotted line labeled "baseline intensity"). Intensity standards that set limits on energy per unit of output will bind on these low-productivity firms. In response, these firms will shift their inputs towards labour, with lower-productivity firms disproportionately affected.

There are other types of standards that set firm-specific limits s_i. We refer to these as firm-specific standards to distinguish them from the sector-specific standards where s is common to all firms within a sector. Suppose the government requires firms lower their energy intensity to x per cent of the firm's own baseline. From equation 4, the standard would be $s_i = x\left((1-\alpha_i)/\alpha_i\right)^{\alpha_i}/A_i$, where $x < 1$ reflects the required improvement. In this case, all firms would find the standard binding and energy per worker would be $x^{1/\alpha_i}(1-\alpha_i)/\alpha_i$ (identical for firms with similar α_i). Keep this in mind. We will see later that sector-specific standards misallocate resources both across firms and across sectors while firm-specific standards misallocate resources only across sectors.

The next – and final – section provides an equivalent way in which we can capture the effect of intensity standards. This approach closely mirrors the misallocation literature in macroeconomics. It provides a powerful and tractable way in which full general equilibrium analysis of standards

Figure 2: Stylized Effect of Intensity Standards

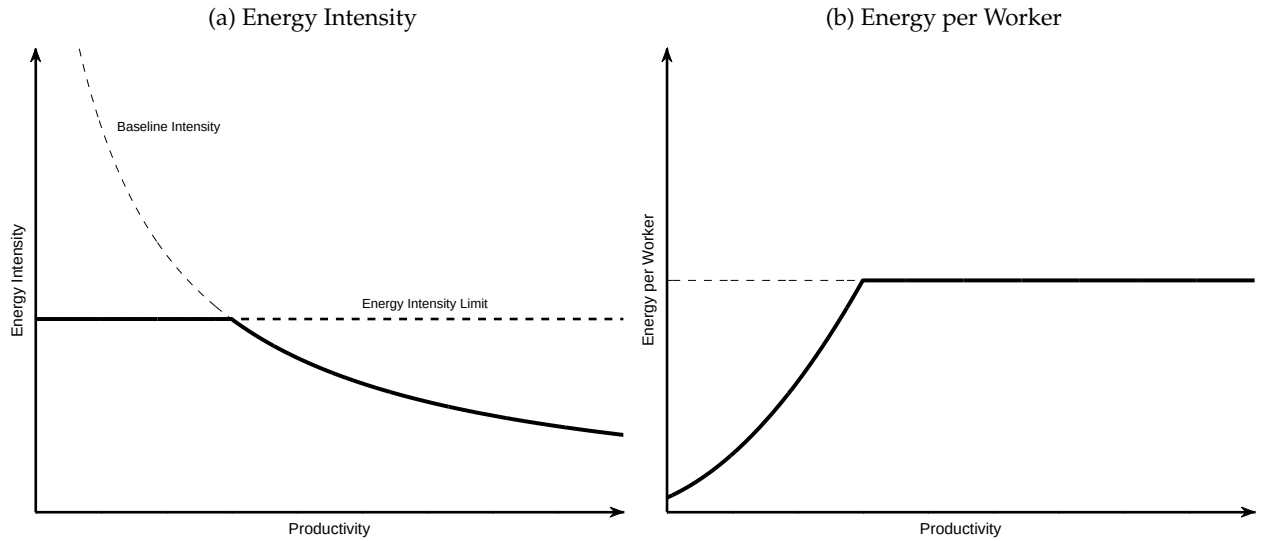

Displays the stylized effect of a binding energy intensity standard on (1) energy intensity and (2) energy per worker. The parameter values used are not relevant. A strict limit on energy intensity causes firms for which the limit is binding to shift input use towards labour and away from energy. The extent of the shift depends on the firm's productivity.

and taxes may be done, both qualitatively and quantitatively.

2.4 Intensity Standards as Market Distortions

A firm maximizing profits will make input choices such that the marginal revenue product of each input equals the marginal cost. For labour, producers will hire until $\alpha_i P_i Y_i / L_i = w$, where w is the wage. In all that follows, we treat labour as the numeraire and therefore normalize $w = 1$. For energy inputs, the decision rule is similar: purchase energy until the marginal revenue product equals energy's price. If aggregate supply is perfectly elastic, we can define units such that energy's price is one, so $(1-\alpha_i)P_i Y_i / E_i = 1$.

Input market distortions, however, change these expressions. If producers face a distortion to the price of energy, denoted $\tau_i^e \geq 1$, then $(1-\alpha_i)P_i Y_i / E_i = \tau_i^e$. Combined with the expression for optimal labour inputs, energy per worker is

$$\frac{E_i}{L_i} = \frac{1-\alpha_i}{\alpha_i}\frac{1}{\tau_i^e}. \tag{6}$$

Combined with equation 5, we have

$$\tau_i^e = \begin{cases} (A_i s)^{-1/\alpha_i}\left(\frac{1-\alpha_i}{\alpha_i}\right) & \text{if } A_i < \left(\frac{1-\alpha_i}{\alpha_i}\right)^{\alpha_i}\frac{1}{s} \\ 1 & \text{otherwise} \end{cases}$$

This demonstrates that one can use energy-input distortions to mirror the effect of intensity stan-

Figure 3: Stylized Representation of the Two Key Distortions

(a) The Energy-Input Distortion

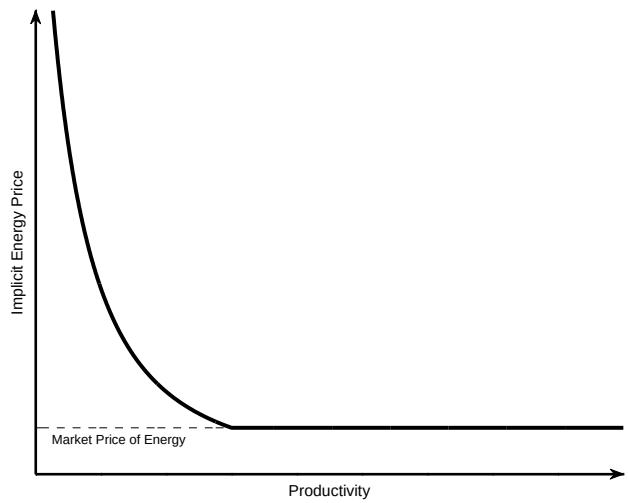

(b) The Output Distortion

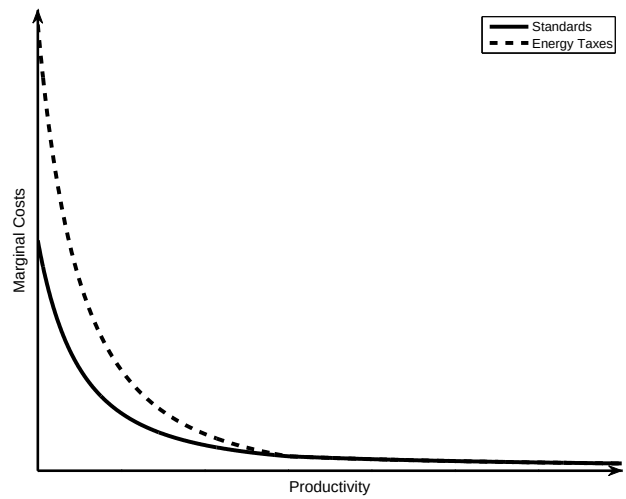

Displays the two key distortions in our analysis. First, the energy distortion raises the implicit price of energy (that is, firms make input decisions *as if* energy prices are τ_i^e). Consider this an implicit energy tax that induces identical energy per worker compared to an intensity standard. Second, the output distortion creates differences in marginal costs between standards (*implicit* increases in energy's price) and *explicit* energy taxes (which actually increase energy's price). Standards do not actually create payments to a government, which is why the resulting marginal costs are lower than taxes.

dards on input decisions. That is, the energy per worker induced by intensity standards is equivalent to the optimal energy per worker chosen by a firm facing an energy price of τ_i^e that is decreasing in firm productivity. We plot this in panel (a) of Figure 3. For firms unaffected by the standard, $\tau_i^e = 1$, while firms for which the standard is binding face $\tau_i^e > 1$.[6]

There is one complication, though, as intensity standards do not actually raise the price of energy, they do so only *implicitly*. Such input market distortions in the misallocation literature are often referred to as *implicit* taxes as there no actual payments to a government. Firms make input decisions *as if* energy's price is $\tau_i^e \geq 1$, though the actual price is still one. Payments for energy are therefore E_i under standards compared to $\tau_i^e E_i$ if τ_i^e was an explicit energy tax.

To see the effect of this, it is useful to derive marginal costs of production. From equation 6 and labour costs of L_i, total costs with explicit energy taxes are $TC_i = L_i + \tau_i^e E_i = L_i/\alpha_i$. Equations 3 and 6 give conditional labour demand of $L_i = \frac{Y_i}{A_i}\left(\frac{\tau_i^e \alpha_i}{1-\alpha_i}\right)^{1-\alpha_i}$. All together, marginal costs are

$$\frac{\partial TC_i}{\partial Y_i} \equiv \frac{1}{A_i}\left(\frac{1}{\alpha_i}\right)^{\alpha_i}\left(\frac{\tau_i^e}{1-\alpha_i}\right)^{1-\alpha_i},$$

which is a familiar expression for firms with Cobb-Douglas production functions operating in competitive markets. But, if the distortion τ_i^e is not an explicit tax, total costs are $TC_i = L_i + E_i$

[6]One could equivalently capture the energy input distortion as an implicit subsidy to labour rather than an implicit tax on energy. After all, what matters is the *relative* price of inputs. All results hold with either $\tau_i^e \geq 1$ as above or with $\tau_i^L = 1/\tau_i^e$, where $\tau_i^L \leq 1$ is an implicit subsidy to labour.

and marginal costs are

$$c_i = \frac{1}{A_i}\left(\frac{1}{\alpha_i}\right)^{\alpha_i}\left(\frac{\tau_i^e}{1-\alpha_i}\right)^{1-\alpha_i}\left(\alpha_i + \frac{1-\alpha_i}{\tau_i^e}\right). \tag{7}$$

The term $\alpha_i + (1-\alpha_i)/\tau_i^e \leq 1$, since $\tau_i^e \geq 1$ and $\alpha_i \in (0,1)$, and therefore marginal costs are *lower* than when τ_i^e is an explicit tax. We plot the two marginal costs in panel (b) of Figure 3. It is *as if* an intensity standard is an explicit tax combined with an output subsidy that lowers marginal costs by a factor $\alpha_i + (1-\alpha_i)/\tau_i^e$. It will be convenient to define this factor as an *output distortion* $\tau_i^y = \alpha_i + (1-\alpha_i)/\tau_i^e$. Intensity standards result in $\tau_i^y < 1$, while explicit taxes with $\tau_i^e = \tau$ have no output distortion, so $\tau_i^y = 1$. This output distortion will not affect energy per worker (equation 6), as it affects the marginal revenue products of energy and labour in the same way.

This is closely related to the "output subsidy distortion" identified in the environmental literature. It is a principle reason why taxes dominate standards – standards create artificially low marginal costs. Typically, the literature models a firm as choosing L_i and E_i to minimize costs subject to a constraint on intensity. Increasing output will relax this constraint, which effectively lowers the marginal cost of production. The output subsidy is precisely the extent to which marginal costs fall due to the intensity standard.[7] These differences in marginal costs matter, as they lead to different optimal emissions levels – that is, a social planner constrained to use standards will choose a higher level of emissions than would be the case under taxes (and only the latter achieves the first-best). This result will hold in our framework as well, but it is not our focus.

Importantly, either policy *could* implement any particular level of aggregate energy use and our paper explores the productivity consequences of alternative policies that achieve the same aggregate conservation objective. It turns out, the output distortion is very relevant for general equilibrium allocations and therefore productivity *independent* of its effect on the optimal level of total energy use. This is our focus and will be a key theme throughout the rest of the paper.

3 A Multi-Sector Model with Heterogeneous Firms

In this section, we outline a tractable, multi-sector, and multi-input model that can incorporate a wide variety of policies. Before proceeding to the specifics, we outline the model's key features. There are multiple sectors within which a continuum of firms produce differentiated goods and may also differ in productivity. Each sector's output is an aggregate of these individual varieties. Across sectors, labour and energy are used with different intensities. A government sector can impose environmental policies that govern the use of energy. Finally, energy has a perfectly elastic supply and we define units such that its price is one.

Our focus is not on evaluating optimal environmental objectives but on the distortions that

[7] Holland (2012), for example, defines γ as the shadow price of emissions (which is positive when the intensity constraint binds) and the output subsidy as $s\gamma$. In our framework, it is straightforward to show $\gamma_i = (\tau_i^e - 1)$ and therefore the output subsidy is $E_i \cdot (\tau_i^e - 1)/Y_i$.

policy may introduce. We express equilibrium sectoral and aggregate productivity as a function of labour and energy *allocations*, independent of the level of aggregate energy use. We do not consider any externalities associated with energy-related emissions and merely take energy reduction goals as given. We therefore remain focused on productivity rather than welfare.

3.1 Households and Production Technologies

The household is simple. Given income I, a representative household consumes a final good Y with price P to generate utility $U(Y)$. The precise form of utility is unimportant, but must be such that $U'(Y) > 0$. The household is endowed with labour L, which it supplies inelastically. We continue to treat labour as the numeraire and therefore $w = 1$. Households also receive lump-sum transfers from total firm profits, government revenue, and all payments to energy. With total income I, the household's decisions are subject to a budget constraint $PY \leq I$.

The final good is produced by a perfectly competitive firm with output Y using as inputs the output Y_i from multiple sectors indexed $i = \{1, ..., N\}$ with the following technology,

$$Y = \prod_{i=1}^{N} Y_i^{\beta_i}, \tag{8}$$

where $\sum_{i=1}^{N} \beta_i = 1$. As its profits are zero, total expenditures $\sum_{i=1}^{N} P_i Y_i$ equals total revenues I. The allocation of spending across sectors to maximize PY subject to $\sum_{i=1}^{N} P_i Y_i \leq I$ is trivially

$$\frac{P_i Y_i}{PY} = \beta_i. \tag{9}$$

Presuming a Cobb-Douglas technology for aggregate output is not overly restrictive, admits explicit expressions for aggregate productivity, and simplifies expressions for the allocation of resources across sectors. In Section 6.2, we demonstrate our quantitative results are similar using a more general CES aggregator for Y.

3.2 Sectoral Output and Intermediate Varieties

For sectoral output, a perfectly competitive firm produces Y_i by combining a continuum of horizontally differentiated intermediate varieties using a CES technology

$$Y_i = \left(\int y_i(\nu)^{(\sigma-1)/\sigma} \mathrm{d}\nu \right)^{\sigma/(\sigma-1)},$$

where $\sigma \geq 0$ is the elasticity of substitution across varieties ν. Each variety-ν is produced by a different firm and they (potentially) differ in their productivity. In the next section, we will see the above production function is identical to equation 3, if sectoral productivity A_i is expressed in a particular way.

Individual firms use the production technologies detailed in Section 2.1; specifically, a firm

with productivity φ produces

$$y_i(\varphi) \quad = \quad \varphi L_i(\varphi)^{\alpha_i} E_i(\varphi)^{1-\alpha_i}, \tag{10}$$

using labour ($L_i(\varphi)$) and energy ($E_i(\varphi)$). As different varieties are symmetric in Y_i, and firms with identical productivity make identical output and input choices, we conveniently re-express sectoral output as

$$Y_i \quad = \quad \left(\int y_i(\varphi)^{(\sigma-1)/\sigma} g_i(\varphi) \mathrm{d}\varphi \right)^{\sigma/(\sigma-1)}, \tag{11}$$

where $y_i(\varphi)$ is the production of a firm with productivity φ and $g_i(\varphi)$ is the distribution of productivity across firms in sector i. Note that multiple firms producing distinct varieties may have identical productivity; the function $g_i(\varphi)$ is the mass of firms (varieties) with productivity φ. It is convenient to presume the distribution of firm productivity within sectors is Pareto, with CDF $F_i(\varphi) = 1 - \varphi^{-\theta}$. The variance depends inversely on θ and the density of this distribution is $g_i(\varphi) = \theta\varphi^{-\theta-1}$.[8] This particular distribution is useful for both compact and closed-form solutions to many model expressions, and also interacts with equation 11 to generate empirically plausible firm size distributions. We will elaborate on this point in Section 5.1.

Given a price $p_i(\varphi)$ for output from firms with productivity φ, the cost-minimizing demand for individual intermediate varieties to produce a given amount of Y_i takes a well-known form,

$$y_i(\varphi) \quad = \quad Y_i \left(\frac{p_i(\varphi)}{P_i} \right)^{-\sigma}. \tag{12}$$

This and equation 11 give the sector price index

$$P_i \quad = \quad \left[\int p_i(\varphi)^{1-\sigma} g_i(\varphi) \mathrm{d}\varphi \right]^{1/(1-\sigma)}. \tag{13}$$

As firms produce differentiated goods and sell to an aggregator with an elasticity of substitution σ, they will charge a markup over costs and earn positive profit. To see this, maximize profit $\pi_i(\varphi) = [p_i(\varphi) - c_i(\varphi)]\, y_i(\varphi)$ using demand from equation 12. The result is simple: a firm with productivity φ charges a price $p_i(\varphi) = mc_i(\varphi)$, where the markup is $m = \sigma/(\sigma-1) > 1$.[9] Finally, all of the results from Section 2.4 hold; in particular, marginal costs are

$$c_i(\varphi) \quad = \quad \varphi^{-1} \left(\frac{1}{\alpha_i} \right)^{\alpha_i} \left(\frac{\tau_i^e(\varphi)}{1-\alpha_i} \right)^{1-\alpha_i} \tau_i^y(\varphi), \tag{14}$$

where $\tau_i^y(\varphi) = 1$ with taxes and $\tau_i^y(\varphi) = \alpha_i + (1-\alpha_i)/\tau_i^e(\varphi) \leq 1$ with standards.

[8] The lower bound of φ is one in all sectors; it plays no role in our analysis until we allow firm entry and exit.

[9] Importantly, common elasticities of substitution within sectors implies markups are the same in all sectors. While this is at odds with existing estimates, it simplifies much of our analysis. Our results change little if we use sector-specific elasticity estimates from Broda and Weinstein (2006).

This completes the model, as we have enough detail to fully solve for general equilibrium allocations of labour and energy. The next section shows how these allocations map directly into sectoral and aggregate productivity.

3.3 Sectoral and Aggregate Productivity

It is straightforward to derive the overall productivity of each sector. Define sector i's productivity as $A_i = Y_i / L_i^{\alpha_i} E_i^{1-\alpha_i}$, where $L_i = \int L_i(\varphi) g_i(\varphi) \mathrm{d}\varphi$ and $E_i = \int E_i(\varphi) g_i(\varphi) \mathrm{d}\varphi$ are the total labour and energy inputs used by firms within sector i. We can solve for sectoral productivity A_i as a function of individual firm productivity and the allocation of labour and energy across firms. Let $l_i(\varphi) = L_i(\varphi)/L_i$ and $e_i(\varphi) = E_i(\varphi)/E_i$ be the shares of a sector's labour and energy going to firms with productivity φ. With these shares, combine sectoral output from equation 11 with firm production technologies from equation 10 and the definition of A_i to yield

$$A_i = \left(\int \left[\varphi l_i(\varphi)^{\alpha_i} e_i(\varphi)^{1-\alpha_i} \right]^{\frac{\sigma-1}{\sigma}} g_i(\varphi) \mathrm{d}\varphi \right)^{\frac{\sigma}{\sigma-1}}. \tag{15}$$

Aggregate (economy-wide) productivity can be derived similarly. Define sectoral labour and energy allocations $l_i = L_i/L$ and $e_i = E_i/E$ and define aggregate productivity (TFP) as $A = Y/L^{\bar{\alpha}} E^{1-\bar{\alpha}}$, where $\bar{\alpha} = \sum_{i=1}^{N} \alpha_i \beta_i$ is the average labour share across sectors. As sectoral output is $Y_i = A_i L_i^{\alpha_i} E_i^{1-\alpha_i}$ and aggregate output is from equation 8,

$$A = \prod_{i=1}^{N} \left(A_i l_i^{\alpha_i} e_i^{1-\alpha_i} \right)^{\beta_i}. \tag{16}$$

What matters for aggregate productivity is not only the underlying productivity of producers A_i but also the allocation of labour and energy across those producers.

The overall use of energy does not affect aggregate productivity. That is, it is not the stringency of environmental policies that affects productivity but the effect of those policies on the allocation of resources across firms and sectors. Of course, the overall use of energy in the economy still matters. A policy that affects neither sectoral productivity nor sectoral resource allocations will leave aggregate productivity unchanged. If the policy lowers energy use, then E will fall and aggregate output Y will also fall (by how much depends on $1 - \bar{\alpha}$). So, policies can lower GDP without lowering productivity. For policies that achieve similar levels of total energy use E, only differences in the effect of these policies on aggregate productivity matter. This is our focus.

4 Qualitative Analysis of Intensity Standards

To see how taxes and standards differ, we begin with a qualitative analysis. We will show that taxes do not lower sectoral or aggregate productivity while standards *may* lower sectoral productivity (depending on their type) but *will* lower aggregate productivity.

4.1 Within-Sector Allocations and Sectoral Productivity

As intermediate firms charge a markup m over costs, total revenue $p_i(\varphi)y_i(\varphi)$ implies total costs are $p_i(\varphi)y_i(\varphi)/m$. Labour's share of total costs are $\alpha_i/\tau_i^y(\varphi)$ and therefore firm demand for labour is such that $\alpha_i p_i(\varphi)y_i(\varphi)/m\tau_i^y(\varphi) = L_i(\varphi)$. As prices are proportional to marginal costs, equation 12 implies $p_i(\varphi)y_i(\varphi) \propto c_i(\varphi)^{1-\sigma}$ and therefore

$$p_i(\varphi)y_i(\varphi) \propto \left(\frac{\varphi}{\tau_i^e(\varphi)^{1-\alpha_i}\tau_i^y(\varphi)}\right)^{\sigma-1} \equiv r_i(\varphi),$$

where $r_i(\varphi)$ is defined for later convenience. Dividing $L_i(\varphi)$ by its integral across firms in sector i, and using the above expression for firm revenue, gives the equilibrium labour allocation

$$l_i(\varphi) \quad = \quad \frac{r_i(\varphi)/\tau_i^y(\varphi)}{\int r_i(\varphi)/\tau_i^y(\varphi)g_i(\varphi)\mathrm{d}\varphi}. \tag{17}$$

A similar derivation gives the equilibrium energy allocation

$$e_i(\varphi) \quad = \quad \frac{r_i(\varphi)/\tau_i^e(\varphi)\tau_i^y(\varphi)}{\int r_i(\varphi)/\tau_i^e(\varphi)\tau_i^y(\varphi)g_i(\varphi)\mathrm{d}\varphi}. \tag{18}$$

These shares, along with equation 15, determine sectoral productivity A_i.

To see how the equilibrium allocations affect sectoral productivity, it is easiest to derive the *optimal* allocations. That is, let $l_i^*(\varphi)$ and $e_i^*(\varphi)$ be the allocations that maximize sectoral productivity A_i. Comparing equilibrium allocations to the optimal allocations is informative for determining a policy's effect on productivity, as deviations from optimal allocations lower A_i. We derive the optimal allocations in the following proposition.

Proposition 1 *The unique labour and energy allocations $l_i(\varphi)$ and $e_i(\varphi)$ that maximize sectoral productivity A_i from equation 15, subject to $\int l_i(\varphi)g_i(\varphi)\mathrm{d}\varphi = 1$ and $\int e_i(\varphi)g_i(\varphi)\mathrm{d}\varphi = 1$, are*

$$l_i(\varphi)^* = e_i(\varphi)^* = \frac{\varphi^{\sigma-1}}{\int \varphi^{\sigma-1}g_i(\varphi)\mathrm{d}\varphi}.$$

To implement these these allocations in equilibrium, $\tau_i^e(\varphi)$ must not vary across firms within sectors.

Proof: See Appendix A.

Beyond providing the optimal (productivity-maximizing) allocations, this proposition shows that whether different policies result in equilibrium allocations that equal the optimal allocations depends on whether energy distortions vary with firm productivity. Simply put, policies that affect energy prices differently for different firms lead to sub-optimal allocations of labour and energy; consequently, sectoral productivity falls.

From this result, it is straightforward to compare taxes to standards. First, taxes do not result in an implicit output subsidy, so $\tau_i^y(\varphi) = 1$. Taxes are also common across firms and sectors, so $\tau_i^e(\varphi) = \tau$. From equations 17 and 18, and the definition of $r_i(\varphi)$, the equilibrium allocations achieve the optimal allocations. With uniform energy taxes, sectoral productivity is as high as it can be. With intensity standards, however, this may not be the case. Standards create an output distortion $\tau_i^y(\varphi) = \alpha_i + (1-\alpha_i)/\tau_i^e(\varphi)$. For this not to differ across firms, the energy distortion must not differ across firms either. As we saw in Section 2.4, standards that set a limit on intensity that is common across firms will result in $\tau_i^e(\varphi)$ that is decreasing in firm productivity. Therefore, equilibrium labour and energy allocations will differ from the optimal allocation. In this case, intensity standards will result in sectoral productivity A_i that is less than what taxes achieve.

Not all standards will lower sectoral productivity. For example, intensity standards that specify a certain improvement over a firm's own baseline energy intensity will *not* cause $\tau_i^e(\varphi)$ to differ across firms within a sector. As we saw earlier, the energy intensity of a firm with productivity φ is $\left[(1-\alpha_i)/\alpha_i\tau_i^e(\varphi)\right]^{\alpha_i}/\varphi$. A standard that constrains this firm to an intensity 90 per cent of its own initial level, for example, creates an implicit energy tax

$$\tau_i^e(\varphi) = 0.9^{-1/\alpha_i},$$

which varies only across firms with different α_i (i.e., those in different sectors). As energy distortions $\tau_i^e(\varphi)$ do not vary across firms within a sector, neither do the output distortions $\tau_i^y(\varphi)$. These standards therefore result in equilibrium labour and energy allocations that are optimal within sectors so sectoral productivity will achieve it's maximum.

4.2 Between-Sector Allocations and Aggregate Productivity

As before, solve for the labour and energy allocations. Define $L_i/L = l_i$ and $E_i/E = e_i$ as the share of total labour and energy used in sector i. Further define $\bar{\tau}_i^y$ and $\bar{\tau}_i^e$ as the average output and energy distortions across all firms in sector i.

Proposition 2 *The equilibrium labour and energy allocations are*

$$l_i = \frac{\alpha_i\beta_i/\bar{\tau}_i^y}{\sum_{j=1}^N \alpha_j\beta_j/\bar{\tau}_j^y} \qquad \textit{and} \qquad e_i = \frac{(1-\alpha_i)\beta_i/\bar{\tau}_i^y\bar{\tau}_i^e}{\sum_{j=1}^N (1-\alpha_j)\beta_j/\bar{\tau}_j^y\bar{\tau}_j^e},$$

where the average distortions in each sector are weighted harmonic means,

$$\bar{\tau}_i^y = \frac{\int r_i(\varphi)g_i(\varphi)\mathrm{d}\varphi}{\int r_i(\varphi)\frac{1}{\tau_i^y(\varphi)}g_i(\varphi)\mathrm{d}\varphi} \qquad \textit{and} \qquad \bar{\tau}_i^e = \frac{\int r_i(\varphi)\frac{1}{\tau_i^y(\varphi)}g_i(\varphi)\mathrm{d}\varphi}{\int r_i(\varphi)\frac{1}{\tau_i^y(\varphi)}\frac{1}{\tau_i^e(\varphi)}g_i(\varphi)\mathrm{d}\varphi}.$$

Proof: See Appendix A.

The intuition behind these expressions is simple. First, we should not be surprised that average distortions are harmonic means across firms, as a harmonic mean is (generally) the proper way to

average *rates*. They are very useful objects. Imagine a representative firm produces each sector's output. Further imagine this firm faces distortions $\bar{\tau}_i^y$ and $\bar{\tau}_i^e$. How would the economy's labour and energy be allocated across sectors? With competitive markets, the representative firm would hire labour and energy such that $L_i = \alpha_i P_i Y_i / \bar{\tau}_i^y$ and $E_i = (1-\alpha_i) P_i Y_i / \bar{\tau}_i^y \bar{\tau}_i^e$. Summing across producers, taking ratios, and using $\beta_i = P_i Y_i$, we have the above labour and energy allocations. So, sectors are like a representative firm facing distortions $\bar{\tau}_i^y$ and $\bar{\tau}_i^e$.

To see how policies affect allocations across sectors and consequently aggregate productivity, we proceed as before by solving for the *optimal* allocations to which we compare the above equilibrium allocations. The following proposition provides the optimal shares of labour and energy for each sector and the distortions $\bar{\tau}_i^y$ and $\bar{\tau}_i^e$ that implement optimal allocations in equilibrium.

Proposition 3 *The labour and energy allocations l_i and e_i that maximize aggregate productivity $A = \prod_{i=1}^N A_i^{\beta_i} l_i^{\alpha_i \beta_i} e_i^{(1-\alpha_i)\beta_i}$ subject to $\sum_{i=1}^N l_i = 1$ and $\sum_{i=1}^N e_i = 1$ are*

$$l_i^* = \frac{\alpha_i \beta_i}{\sum_{j=1}^N \alpha_j \beta_j} \qquad \text{and} \qquad e_i^* = \frac{(1-\alpha_i)\beta_i}{\sum_{j=1}^N (1-\alpha_j)\beta_j}.$$

The output and energy distortions that implement these allocations in equilibrium are such that $\bar{\tau}_i^y$ does not vary across sectors and $\bar{\tau}_i^e = E/E'$, where $E' \leq E$ is the energy use implemented by a policy. Finally, if $\alpha_i \neq \alpha_j$ for at least two sectors $i \neq j$, then standards cannot implement the optimal allocations.

Proof: See Appendix A.

This proposition combines nicely with Proposition 1. To achieve maximum productivity, not only must $\tau_i^e(\varphi)$ not vary across firms in each sector (to ensure sectoral productivity A_i is maximized) but so also must average distortions $\bar{\tau}_i^y$ and $\bar{\tau}_i^e$ not vary. Furthermore, Proposition 3 establishes a particular energy distortion $\bar{\tau}_i^e = E/E'$ that implements some desired level of total energy use $E' \leq E$. For example, if $E' = 0.9E$ then $\bar{\tau}_i^e = 1/0.9$. An energy tax of $100 \times (0.9^{-1} - 1)$ will do just that. As taxes do not have output distortions, aggregate productivity will be maximized.

Standards, on the other hand, must result in lower aggregate productivity, even if $\bar{\tau}_i^e$ (somehow) equaled E/E', unless α_i are identical across sectors. This is straightforward to see, though we leave the details to Appendix A. Comparing equilibrium allocations from Proposition 2 with optimal allocations from Proposition 3, we must have neither $\bar{\tau}_i^y$ nor $\bar{\tau}_i^y \bar{\tau}_i^e$ vary across sectors. It is straightforward to show standards result in average output distortions of $\bar{\tau}_i^y = \alpha_i + (1-\alpha_i)/\bar{\tau}_i^e$, so optimal allocations requires $1 - \alpha_i + \alpha_i \bar{\tau}_i^e = 1 - \alpha_j + \alpha_j \bar{\tau}_j^e$ must hold for all sectors $i \neq j$. As energy distortions must be $\bar{\tau}_i^e = E/E'$, any policy that lowers total energy use will violate this condition unless α_i are identical across sectors. Thus, standards cannot achieve the optimal allocations if α_i differ across sectors. How much allocations differ and how much productivity falls are quantitative questions to which we turn shortly, in Section 5.

4.3 Differences in Marginal Abatement Costs

The previous two sections establish that resource allocations affect productivity. *Mis*allocations lower productivity, and standards generate those misallocations. Is this different than the existing (and well-known) result in environmental economics that optimal policy requires equalized marginal abatement costs across producers? Yes.

To see this in more detail, consider each sector as a representative firm facing energy and output distortions $\bar{\tau}_i^e$ and $\bar{\tau}_i^y$. Marginal abatement costs are $\partial Y/\partial E_i \propto \beta_i(1-\alpha_i)/E_i$ with total energy use $E_i = (1-\alpha_i)L_i/\alpha_i\bar{\tau}_i^e$. If labour allocations are fixed at their optimal allocation $L_i \propto \alpha_i\beta_i$, marginal abatement costs will be $\partial Y/\partial E_i \propto \bar{\tau}_i^e$. So, in this sense, marginal abatement costs are closely related to average energy distortions. But, as we saw, labour allocations depend on $\bar{\tau}_i^e$ in the case of standards; specifically, $L_i \propto \alpha_i\beta_i/\bar{\tau}_i^y = \alpha_i\beta_i/(\alpha_i + (1-\alpha_i)/\bar{\tau}_i^e)$. Thus, marginal abatement costs are

$$\frac{\partial Y}{\partial E_i} \propto \frac{\beta_i(1-\alpha_i)}{E_i} = 1 - \alpha_i + \alpha_i\bar{\tau}_i^e.$$

For marginal abatement costs to equalize across sectors, so must $1-\alpha_i+\alpha_i\bar{\tau}_i^e$ which implies labour allocations are $L_i \propto \beta_i\alpha_i\bar{\tau}_i^e$. This is not optimal. Indeed, this says that to equalize abatement costs, labour must re-allocate towards sectors with tight standards (where $\bar{\tau}_i^e$ is higher). Intuitively, standards increase labour's share of total costs, especially in energy-intensive sectors, and therefore changes the relative demand for labour. Sectors that see an above-average increase in labour's share of total costs will increase their spending on labour relative to other sectors, thereby increasing their share of aggregate employment. This leads to sub-optimal labour allocations, lowering aggregate productivity.

We can see this another way. The above expression is equivalent to $\partial Y/\partial E_i \propto \bar{\tau}_i^y\bar{\tau}_i^e$. This, along with Propositions 2 and 3, means that if marginal abatement costs equalize then energy is optimally allocated. Labour, however, is not optimally allocated as output distortions will not equalize unless both $\bar{\tau}_i^e$ and α_i are the same across all sectors. Policies that merely equalize marginal abatement costs are insufficient to maximize productivity. Energy may be optimally allocated but labour will not be. Taxes, on the other hand, do not have output distortions so both marginal abatement costs are equalized and labour is optimally allocated. This is a novel general-equilibrium effect that our framework tractably captures.

5 Quantitative Analysis of Intensity Standards

With the qualitative results now well established we proceed to illustrate the quantitative importance of the negative productivity effects of standards. In all that follows, we approximate the integrals with summations across 100,000 firms in each sector drawn such each firm has equal weight $g_i(\varphi)$ under the Pareto distribution. This is a sufficiently large number of firms, as increases beyond this level have little effect on our results but increase the computational burden.

We begin in the next section with a calibration of our model before proceeding to the full

quantitative analysis. For all of the simulations, we examine the effect of standards and taxes that both reduce total energy use by 10 and 20 per cent. This allows a comparison of the economic costs of alternative policies without taking a stand on the optimal amount of energy conservation.

5.1 Model Calibration

Here we calibrate the model to match key moments in sector-level data on energy spending and output for forty-four sectors in the United States. We use the U.S. Bureau of Economic Analysis' KLEMS data.

The parameters to calibrate include the sectoral productivity parameter (θ), the within-sector elasticity of substitution (σ), each sector's share of total spending (β_i), and labour's share of output (α_i). The latter two parameters have straightforward and observable counterparts in data. We set $1-\alpha_i$ to the ratio of spending on energy inputs to spending on energy, labour, and capital (we interpret L_i broadly to include both physical and human capital inputs). Importantly, this data only looks at energy inputs used up in the process of production, not raw material inputs that are themselves transformed (such as oil used by refineries to make gasoline). From this same data, we set β_i to equal the share of total economy-wide spending on energy, labour, and capital accounted for by each sector. We report both values in in Table 1.

The parameters θ and σ are slightly more difficult to set. From equation 12, firm sales are $\beta_i P_i^{\sigma-1}\left(mc_i/\varphi\right)^{1-\sigma}$ and the CDF of sales is therefore $G[\beta_i P_i^{\sigma-1}\left(mc_i/\varphi\right)^{1-\sigma} \leq z]$. With a slight rearrangement and using Pareto-distributed productivity, we have

$$G\left(\varphi \leq (z/\beta_i)^{\frac{1}{\sigma-1}} \frac{mc_i}{P_i}\right) = 1-(\beta_i/z)^{\frac{\theta}{\sigma-1}} A_i^{-\theta}.$$

The probability that sales *exceed* some level z is therefore proportional to $z^{-\theta/(\sigma-1)}$, which means firm sales follows a power law with slope $\theta/(\sigma-1)$, which is consistent with data (Axtell, 2001). In recent work, di Giovanni et al. (2011) estimate this slope precisely for a variety of sectors. The mid-point of their estimates – and the preferred estimate of di Giovanni and Levchenko (2013) – is $\theta/(\sigma-1) = 1.06$. Given estimates of the within-sector elasticity of substitution σ we can set θ. The within-sector elasticity of substitution σ is set to 4, which is the mean value of the elasticities in Broda and Weinstein (2006).[10]

5.2 Firm-Specific and Threshold-Based Standards

Firm-specific standards require intensity reductions of some percentage below each firm's own baseline intensity. In Section 2.3, we saw energy per worker under these standards is $x^{1/\alpha_i}(1-$

[10]Though their estimate applies only to tradable goods, there are no estimates of this elasticity for non-tradables. We also explore the effect of using sector-specific σ, and find little depends on this. We opt for a common σ as identical markups across sectors simplifies our expressions for labour and energy allocations.

Table 1: Sector-Level Summary Statistics and Calibrated Parameters (2012)

Sector	NAICS Code	From BEA KLEMS Data: Output Share β_i	Energy's Share $1-\alpha_i$
Agriculture, Forestry, and Fishery	11	0.020	0.092
Oil and Gas Extraction	211	0.019	0.009
Mining, except Oil and Gas	212	0.006	0.064
Support Activities for Mining	213	0.007	0.023
Utilities	22	0.025	0.196
Construction	23	0.053	0.052
Food, Beverage, Tobacco	311-312	0.020	0.052
Textiles	313-314	0.002	0.054
Apparel and Leather and Products	315-316	0.001	0.009
Wood Products	321	0.002	0.066
Paper Products	322	0.005	0.170
Printing and Related Support Activities	323	0.003	0.025
Petroleum and Coal Products	324	0.017	0.048
Chemical Products	325	0.034	0.055
Plastics and Rubber Products	326	0.006	0.050
Nonmetallic Mineral Products	327	0.004	0.141
Primary Metals	331	0.006	0.190
Fabricated Metal Products	332	0.012	0.034
Machinery	333	0.013	0.019
Computer and Electronic Products	334	0.021	0.005
Electrical Equipment	335	0.005	0.016
Transportation Equipment	336	0.021	0.016
Furniture and Related Products	337	0.002	0.022
Miscellaneous Manufacturing	339	0.007	0.010
Wholesale Trade	42	0.069	0.013
Retail Trade	44-45	0.067	0.023
Transportation and Warehousing	48-49	0.052	0.267
Information	51	0.064	0.006
Finance, Insurance, and Real Estate	52-53	0.158	0.040
Legal Services	5411	0.018	0.002
Computer Systems and Services	5415	0.020	0.005
Misc. Professional, Scientific, Tech. Services	5412-5414, 5416-5419	0.060	0.011
Management of Companies and Enterprises	55	0.028	0.024
Administrative and Support Services	561	0.039	0.015
Waste Management and Remediation Services	562	0.004	0.034
Educational Services	61	0.006	0.136
Ambulatory Health Care Services	621	0.044	0.011
Hospitals, Nursing, and Residential Care	622-623	0.016	0.102
Social Assistance	624	0.004	0.044
Cultural Activities	711-712	0.005	0.018
Amusements, Gambling, and Recreation	713	0.005	0.071
Accommodation	721	0.008	0.036
Food Services and Drinking Places	722	0.022	0.037
Other Services, except Government	81	0.022	0.036

Note: Data extracted from the BEA KLEMS data for 2012. Output share is a sector's value-added plus energy spending relative to the whole economy. Energy's share of output is the ratio of energy spending to spending on energy, labour, and capital.

$\alpha_i)/\alpha_i$. The implied energy distortion to implement this is therefore

$$\tau_i^e(\varphi) = x^{-1/\alpha_i},$$

where $x < 1$ reflects the required improvement. For example, if firms are required to set energy intensity to no more than 90% of their initial level, then $x = 0.9$. Importantly, this varies only across firms with different α_i (i.e., those in different sectors).

In our qualitative analysis, we found these standards do not misallocate labour or energy across firms within sectors, though they do between sectors. The quantitative effect of this, how-

Figure 4: Threshold-Based Standards, More Costly as Fewer Firms Covered

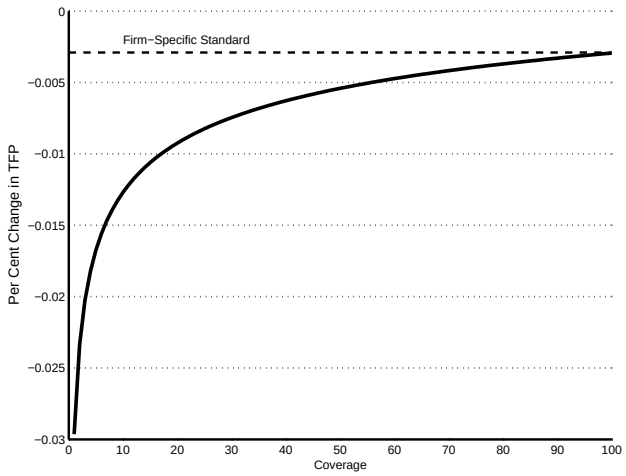

Displays the change in aggregate productivity incurred to lower total energy use by 10 per cent using a threshold-based standard, for various thresholds. The horizontal axis displays the share of initial total energy use by firms covered by the standard. As more firms are covered, the costs of the standard falls. The range is selected such that sectors are not entirely covered or uncovered.

ever, is minor. To lower total energy use by 10 per cent results in an aggregate TFP loss of 0.003 per cent, which is very small. We calculate the implied cost per tonne of carbon, presuming that lowering energy use by 10 per cent results in 10 per cent lower emissions, and find a 0.003 per cent TFP loss for the United States is equivalent to \$0.73/tonne in excess abatement costs relative to a uniform tax on energy. To lower energy use by 20 per cent, the costs are larger: a 0.012 per cent loss in TFP or \$1.48/tonne. By comparison, a tax has no TFP consequence and GDP falls by just over 0.5 per cent to achieve a 10 per cent reduction in energy and by just over one per cent to achieve a 20 per cent reduction.

Of course, this standard is highly infeasible, as it requires knowledge of each firm's baseline level of energy intensity. With close to 7.5 million establishments in the U.S. (according to the U.S. Business Census), it is hard to imagine firm-specific standards in practice. Instead, consider firm-specific standards that apply only to "sufficiently large" firms. While one may suspect that as fewer firms are covered, productivity losses will be smaller, the opposite is the case. As firms within a sector are treated differently, there will be misallocations within sectors and consequently reductions in sectoral productivity. This will amplify the negative effect of the standard on aggregate TFP.

We explore a range of possible thresholds and display the resulting costs in Figure 4. As the threshold falls and more firms are covered, the productivity costs are smaller. In the limit where all firms are covered, the productivity costs are the smallest and equal the costs associated with firm-specific standards. To lower energy use by 10 per cent under a firm-specific intensity standard that binds only firms in the top decile of firm size, for example, results in an aggregate TFP loss of 0.015 per cent – almost five times more costly than if all firms were covered.

5.3 Sector-Specific Targets

We turn now to intensity standards that apply to all firms. Each firm must meet or exceed a specified level of energy use per unit of output, based on (say) some improvement over the sector's average intensity. The rationale for this type of policy is simple: ease of administration and enforcement. While enforcement is still firm-by-firm, the common baseline makes administration and enforcement of this target easier. Unlike firm-specific targets, which require knowledge of each firm's prior *and* current intensity, sector-specific targets require only knowledge of the current intensity. Random audits with high penalties may be sufficient to enforce this standard.

We simulate a requirement that each firm not exceed some limit to energy intensity s_i that is common across firms within each sector but may differ across sectors. The standard is set so each sector's limit is the same improvement over the sector's initial average intensity.[11] Importantly, this type of standard imposes differential implicit energy taxes on firms – with low-productivity firms paying higher implicit energy taxes $\tau_i^e(\varphi)$. Firms with sufficiently high productivity have energy intensities less than their sector's target and therefore will not be subject to the standard. Qualitatively, we saw this results in misallocation of labour and energy within and between sectors. Quantitatively, these effects are meaningful.

5.3.1 Sector-Specific Standards with a Single Sector

We begin with a simple exercise to highlight the effect of an intensity standard on sectoral productivity, abstracting from between-sector considerations. Consider the case where there is only one sector, or that we are looking at one of many sectors but the total allocation of labour to sectors is fixed.

From Section 2.4, an intensity standard s generates an implicit energy tax of

$$\tau_i^e(\varphi) = \begin{cases} (\varphi s)^{-\frac{1}{\alpha_i}} \left(\frac{1-\alpha_i}{\alpha_i}\right) & \text{if } \varphi < \left(\frac{1-\alpha_i}{\alpha_i}\right)^{\alpha_i} \frac{1}{s} \\ 1 & \text{otherwise} \end{cases}$$

Total energy use in the sector is

$$E_i = \frac{1-\alpha_i}{\alpha_i} L_i \int \frac{1}{\tau_i^e(\varphi)} l_i(\varphi) g_i(\varphi) \mathrm{d}\varphi$$

and therefore to lower total energy use by 10 per cent, set s such that $0.9 = \int \frac{1}{\tau_i^e(\varphi)} l_i(\varphi) g_i(\varphi) \mathrm{d}\varphi$, where labour allocations are determined as in equation 17.

In Figure 5 we illustrate the change in a sector's TFP from lowering energy use by 10 per cent using a sector-specific standard. Without differences in productivity across firms, standards will not create misallocations. In panel (a), we plot productivity losses for various α_i both when

[11] In our framework, the simple average energy intensity is $\int (e_i(\varphi)/y_i(\varphi)) g_i(\varphi) \mathrm{d}\varphi = ((1-\alpha_i)/\alpha_i)^{\alpha_i} \int \varphi^{-1} g_i(\varphi) \mathrm{d}\varphi$. Using the Pareto-distributed firm productivity within sectors, this becomes $[(1-\alpha_i)/\alpha_i]^{\alpha_i} \cdot \theta/(1+\theta)$.

Figure 5: Results for Single Sector Lowering Energy Use

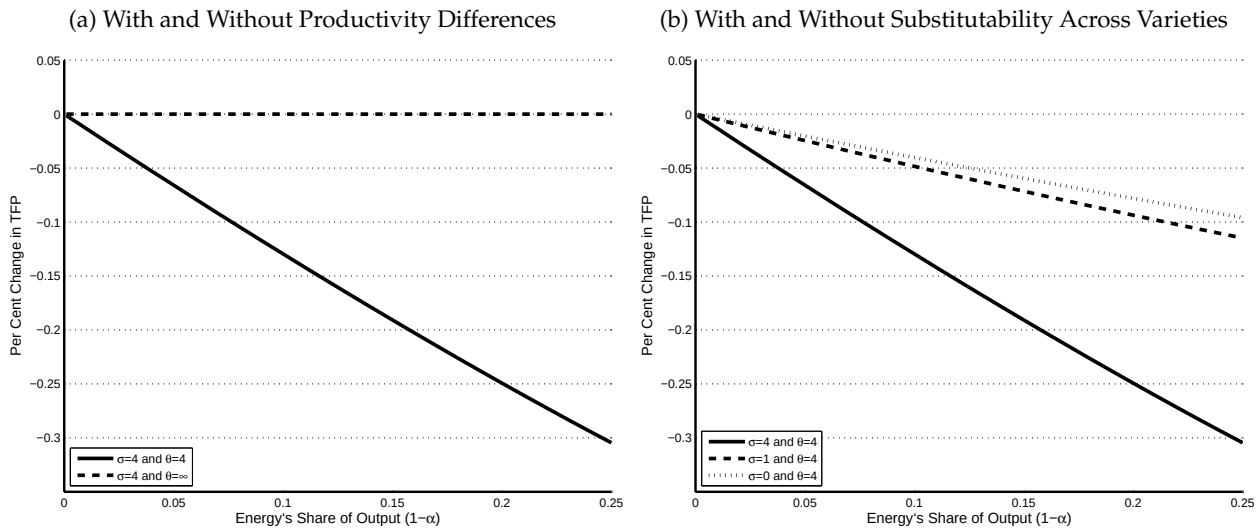

Displays the reduction of a single sector's TFP resulting from an intensity standard that lowers energy use by 10 per cent. The x-axis is the importance of energy for production $1-\alpha$; sectors with higher energy use will have larger productivity consequences from intensity standards. The productivity effect depends on the existence of productivity heterogeneity across firms ($1/\theta$) and on the degree of substitutability between varieties within the sector (σ).

$\theta = 4$ compared to when $\theta = \infty$ (homogenous firms). For a sector with $1-\alpha_i = 0.25$, productivity declines by over 0.3 per cent if $\theta = 4$; productivity does not change at all if firms are identical. This follows immediately from the above expression for the implicit energy tax. Without differences in productivity, $\tau_i^e(\varphi)$ and $\tau_i^y(\varphi)$ are the same for all firms and allocations are optimal within the sector. Of course, productivity *does* differ substantially between firms; this exercise highlights the importance of considering heterogeneity across firms.

One may wonder how much of the within-sector misallocation is due to sector-specific standards shifting expenditures towards relatively cleaner (i.e., more productive) firms. Standards raise the costs of firms that are covered, and therefore also raises the price of their output. The price of output from firms that are not covered becomes lower in comparison. From equation 12, sales of uncovered firms will increase when $\sigma > 1$.[12] The quantitative effect of this shifting of expenditures across firms is important, as illustrated in panel (b) of Figure 5. With fixed expenditure shares across firms (which is the case when $\sigma = 1$), there are lower productivity effects than when $\sigma = 4$. Depending on the value of α, the overall reduction in sectoral TFP is approximately two to three times as large when $\sigma = 4$ as when $\sigma = 1$. If there is no substitutability across firm output at all (when $\sigma = 0$), the change in TFP is the smallest although very close to when $\sigma = 1$.

Finally, we investigate how standards affect labour allocations. We illustrate in Figure 6 the effect of $\tau_i^e(\varphi)$ and $\tau_i^y(\varphi)$ on firm costs and on the equilibrium labour allocations. We do this for

[12]This is a distinct mechanism from Holland et al. (2009). In that paper the substitution was a choice made a firm between different fuel types within the boundaries of the firm. In this case, the implicit subsidy is due to the effect intensity standards have on the price index.

Figure 6: Effect of Distortions on Costs and Allocations Across Firms

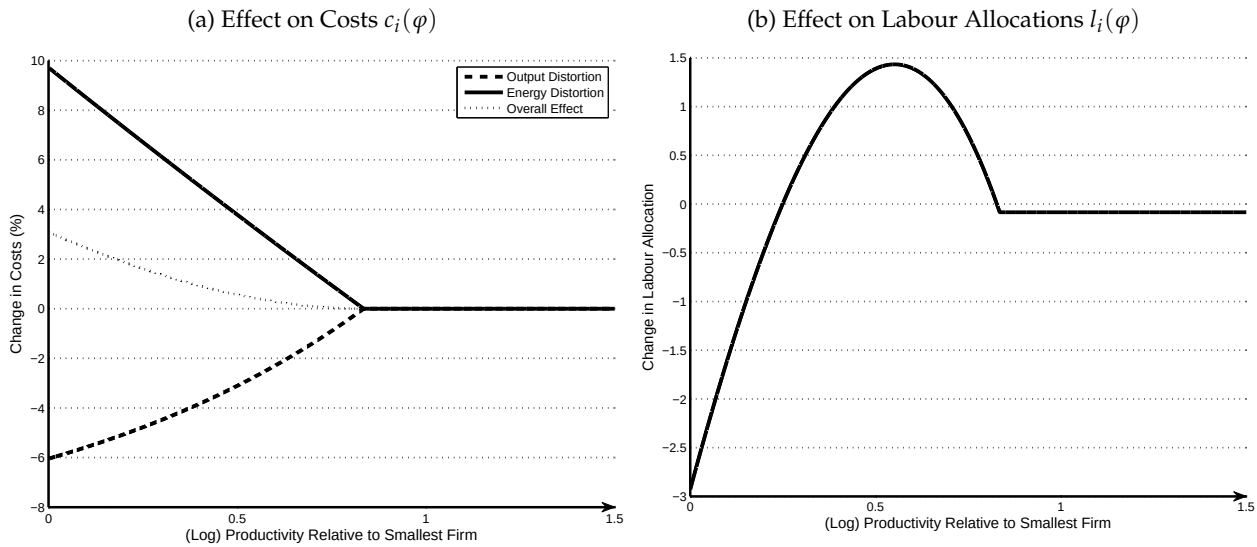

Displays the effect of distortions on firm costs and labour allocations. Examines stylized industry with $\alpha = 0.9$ and imposes a common standard s such that energy use declines to 90 per cent of initial levels. The output distortion $\tau_i^y(\varphi) \leq 1$ lowers costs and increases labour's share of total costs. The energy distortion $\tau_i^e(\varphi)^{1-\alpha_i} \geq 1$ increases costs. There are two offsetting effects. First, costs always rise for distorted firms and therefore their sales decline and they demand fewer workers. Second, labour's share of costs rise. For low-productivity firms, the former effect dominates; for medium-productivity firms, the latter effect dominates. For high-productivity firms, neither distortion matters, though their labour allocation shrinks somewhat as labour markets must clear.

a sector with $\alpha_i = 0.9$, though the qualitative results are identical for other values. Two notable patterns emerge. First, standards raise costs for low-productivity firms, though the output distortion offsets some of the direct effect of the energy distortions (recall, energy distortions are implicit taxes on energy while output distortions are implicit subsidies to output). The second pattern is less intuitive: medium-productivity firms expand their share of a sector's employment, while low- and high-productivity firms reduce their share. Why? It is instructive to consider two alternative forces. First, energy and output distortions both affect firm sales through $\left[\tau_i^e(\varphi)^{1-\alpha_i}\tau_i^y(\varphi)\right]^{\sigma-1}$. But, the output distortion further affects labour's share of total costs $\alpha_i/\tau_i^y(\varphi)$. The former effect lowers demand for labour while the latter effect raises it. For low-productivity firms, the first effect dominates; for medium-productivity firms, the second effect dominates. High-productivity firms are not affected except indirectly through the labour market; in this case, their allocations fall.

5.3.2 Sector-Specific Standards for Multiple Sectors

We now simulate the effect of sector-specific standards for all sectors. We report the results of this policy in Table 2. The differences between a tax and a standard are dramatic: aggregate GDP declines by 0.65 per cent under the standard but by only 0.51 per cent under the uniform tax that achieves an identical reduction in energy use. To lower energy use by 20 per cent, aggregate GDP falls by 1.52 per cent with a standard and just over 1 per cent with a tax. In terms of aggregate

Table 2: Sector-Specific Standards vs Energy Taxes

	Lower Energy Use by 10%		Lower Energy Use by 20%	
	Intensity Standards	Energy Taxes	Intensity Standards	Energy Taxes
Main Results				
Change in Aggregate GDP	-0.649%	-0.506%	-1.516%	-1.068%
Change in Aggregate TFP	-0.144%	0.000%	-0.453%	0.000%
... Within-Sector Productivity Change	-0.143%	0.000%	-0.453%	0.000%
... Between-Sector Allocations	-0.001%	0.000%	0.000%	0.000%
With Option to Pay a Fine (Equivalent to 50% Energy Tax)				
Change in Aggregate GDP	-0.569%	-0.506%	-1.143%	-1.068%
Change in Aggregate TFP	-0.063%	0.000%	-0.076%	0.000%
... Within-Sector Productivity Change	-0.065%	0.000%	-0.078%	0.000%
... Between-Sector Allocations	0.002%	0.000%	0.002%	0.000%

Displays results of lower total energy use with sector-specific standards, where all firms must improve their energy intensity at or beyond a target that is based on the sector's prior average. We compare these changes to uniform energy taxes. The top panel are our main results. The bottom panel allow firms the option to pay a fine.

productivity, there is no effect from taxes but large negative effects from standards. To lower energy use by 10 per cent, aggregate productivity declines by nearly 0.15 per cent with intensity standards. To lower energy use by 20 per cent, this cost rises to 0.45 per cent. These effects are large, equivalent to \$35.64 and \$56.04 per tonne of carbon abated (again, presuming energy use and carbon emissions are proportional).

Decomposing the aggregate GDP reduction yields a starkly different result from facility-specific targets: within-industry effects dominate. Specifically, as one can see in the last two rows of Table 2, between-sector allocations contribute only a tiny amount to the overall productivity consequences of standards. Given the importance of within-sector TFP effects, we list each industry in Table 3. So far, only the "no exit" columns of the table are relevant, as we have yet to expand the model to allow firm entry and exit decisions. Energy-intensive industries see a much larger decline in productivity. In Paper Products, for example, TFP declines by over 0.5 per cent. These energy-intensive industries, however, expand their allocations of labour, offsetting some of the contractionary effect of their lower energy use. Why? The output distortion τ_i^y increases the share of total costs allocated to labour. Sectors where energy inputs are important (low α_i) therefore see an increase in their allocations relative to other sectors. The quantitative effect of these between-sector allocations, however, pales in comparison to the within-sector effect and the resulting large reduction in sectoral TFP.

Table 3: Response to Sector-Specific Intensity Targets, by Industry

Industry	Energy Share	No Exit			With Exit			
		Output	TFP	Allocation	Output	TFP	Allocation	Active Firms
Agriculture, Forestry, and Fishery	0.092	0.18%	-0.26%	0.44%	-0.11%	-0.55%	0.44%	-7.20%
Oil and Gas Extraction	0.009	-0.41%	-0.02%	-0.39%	-0.44%	-0.05%	-0.39%	-0.70%
Mining, except Oil and Gas	0.064	-0.02%	-0.17%	0.16%	-0.22%	-0.37%	0.16%	-4.94%
Support Activities for Mining	0.023	-0.31%	-0.06%	-0.25%	-0.38%	-0.13%	-0.25%	-1.73%
Utilities	0.196	0.86%	-0.62%	1.49%	0.24%	-1.24%	1.50%	-15.98%
Construction	0.052	-0.10%	-0.14%	0.04%	-0.26%	-0.30%	0.04%	-4.02%
Food, Beverage, Tobacco	0.052	-0.10%	-0.14%	0.04%	-0.26%	-0.30%	0.04%	-4.00%
Textiles	0.054	-0.08%	-0.15%	0.06%	-0.25%	-0.32%	0.06%	-4.20%
Apparel and Leather and Products	0.009	-0.41%	-0.02%	-0.39%	-0.44%	-0.05%	-0.39%	-0.69%
Wood Products	0.066	-0.01%	-0.18%	0.17%	-0.21%	-0.38%	0.18%	-5.08%
Paper Products	0.170	0.69%	-0.52%	1.22%	0.16%	-1.06%	1.23%	-13.68%
Printing and Related Support Activities	0.025	-0.30%	-0.06%	-0.23%	-0.37%	-0.14%	-0.23%	-1.88%
Petroleum and Coal Products	0.048	-0.13%	-0.13%	0.00%	-0.28%	-0.28%	0.00%	-3.71%
Chemical Products	0.055	-0.08%	-0.15%	0.07%	-0.25%	-0.32%	0.07%	-4.25%
Plastics and Rubber Products	0.050	-0.11%	-0.14%	0.02%	-0.27%	-0.29%	0.02%	-3.88%
Nonmetallic Mineral Products	0.141	0.51%	-0.42%	0.93%	0.06%	-0.86%	0.94%	-11.25%
Primary Metals	0.190	0.82%	-0.60%	1.42%	0.22%	-1.20%	1.43%	-15.43%
Fabricated Metal Products	0.034	-0.23%	-0.09%	-0.14%	-0.34%	-0.19%	-0.14%	-2.58%
Machinery	0.019	-0.34%	-0.05%	-0.29%	-0.40%	-0.11%	-0.29%	-1.42%
Computer and Electronic Products	0.005	-0.44%	-0.01%	-0.43%	-0.46%	-0.03%	-0.43%	-0.39%
Electrical Equipment	0.016	-0.36%	-0.04%	-0.32%	-0.41%	-0.09%	-0.32%	-1.21%
Transportation Equipment	0.016	-0.36%	-0.04%	-0.32%	-0.41%	-0.09%	-0.31%	-1.25%
Furniture and Related Products	0.022	-0.32%	-0.06%	-0.26%	-0.38%	-0.13%	-0.26%	-1.68%
Miscellaneous Manufacturing	0.010	-0.40%	-0.03%	-0.38%	-0.43%	-0.06%	-0.37%	-0.78%
Wholesale Trade	0.013	-0.39%	-0.03%	-0.35%	-0.42%	-0.07%	-0.35%	-0.96%
Retail Trade	0.023	-0.31%	-0.06%	-0.25%	-0.38%	-0.13%	-0.25%	-1.72%
Transportation and Warehousing	0.267	1.26%	-0.92%	2.20%	0.43%	-1.76%	2.23%	-22.30%
Information	0.006	-0.43%	-0.02%	-0.42%	-0.45%	-0.04%	-0.41%	-0.48%
Finance, Insurance, and Real Estate	0.040	-0.18%	-0.11%	-0.08%	-0.31%	-0.23%	-0.08%	-3.10%
Legal Services	0.002	-0.47%	0.00%	-0.46%	-0.47%	-0.01%	-0.46%	-0.14%
Computer Systems and Services	0.005	-0.45%	-0.01%	-0.43%	-0.46%	-0.03%	-0.43%	-0.35%
Misc. Professional, Scientific, Tech. Services	0.011	-0.40%	-0.03%	-0.37%	-0.43%	-0.06%	-0.37%	-0.81%
Management of Companies and Enterprises	0.024	-0.31%	-0.06%	-0.24%	-0.38%	-0.14%	-0.24%	-1.80%
Administrative and Support Services	0.015	-0.37%	-0.04%	-0.33%	-0.41%	-0.09%	-0.33%	-1.16%
Waste Management and Remediation Services	0.034	-0.23%	-0.09%	-0.15%	-0.34%	-0.19%	-0.14%	-2.57%
Educational Services	0.136	0.48%	-0.40%	0.88%	0.05%	-0.83%	0.89%	-10.84%
Ambulatory Health Care Services	0.011	-0.40%	-0.03%	-0.37%	-0.43%	-0.06%	-0.37%	-0.85%
Hospitals, Nursing, and Residential Care	0.102	0.25%	-0.29%	0.54%	-0.07%	-0.61%	0.54%	-7.98%
Social Assistance	0.044	-0.16%	-0.12%	-0.04%	-0.29%	-0.26%	-0.04%	-3.41%
Cultural Activities	0.018	-0.34%	-0.05%	-0.30%	-0.40%	-0.11%	-0.29%	-1.40%
Amusements, Gambling, and Recreation	0.071	0.03%	-0.20%	0.23%	-0.19%	-0.42%	0.23%	-5.53%
Accommodation	0.036	-0.22%	-0.10%	-0.12%	-0.33%	-0.21%	-0.12%	-2.75%
Food Services and Drinking Places	0.037	-0.21%	-0.10%	-0.11%	-0.32%	-0.21%	-0.11%	-2.84%
Other Services, except Government	0.036	-0.21%	-0.10%	-0.12%	-0.33%	-0.21%	-0.12%	-2.77%

Displays results for each industry in response to sector-specific energy intensity targets that achieve an aggregate reduction in energy use of 10 per cent. The intensity target is relative to each sector's average energy intensity. The ratio of the standard to the initial average intensity is the same across sectors. Energy's share of output is taken directly from the industry data table. The change in resource allocations represents the change in labour and energy allocations $\hat{l}_i^{\alpha_i}\hat{e}_i^{(1-\alpha_i)}$.

5.3.3 Sector-Specific Standards with the Option to Pay a Fine

The implicit increase in energy prices for low-productivity firms facing sector-specific targets is substantial. A small tweak to the policy can limit some of the large productivity losses associated with this type of standard. Consider providing firms an option: meet the standard or pay a fine. For simplicity, the fine will be an explicit tax on energy inputs. Revenue earned through this fine is rebated back to the firm in proportion to output – this allows for a clean comparison between the standard with and without the option to pay a fine. It essentially imposes a limit on the implicit energy tax that standards create. As this exercise is meant to be illustrative only, we arbitrarily set a large fine equivalent to a 50 per cent energy tax. Only firms for which $\tau_i^e > 1.5$ will opt to pay the fine – these are the lowest productivity firms.

The option to pay a fine dramatically reduces the costs of imposing energy-intensity standards. Among firms that opt to pay it, there is no misallocation within the sector – after all, $\tau_i^e = 1.5$ and $\tau_i^y = \alpha_i + (1-\alpha)/1.5$ for those firms. This limits the extent of misallocation and dampens the negative productivity effect of a standard. We display the results of this exercise in the bottom panel of Table 2. To lower energy use by 10 per cent, the option to pay the fine mitigates over half of the negative effect of the standard on aggregate productivity. This drop is equivalent to the excess productivity costs falling from \$35.64 per tonne of carbon abated to \$15.65 per tonne. While only illustrative, these results suggest allowing firms to pay fines can significantly lower the economic costs of standards relative to taxes.

6 Expanding the Basic Model

Our final set of exercises expands the model to incorporate entry and exit decisions of firms and a more general CES aggregate production function.

6.1 The Effect of Firm Entry and Exit

The analysis so far has abstracted from the potential for firms to shut down in the face of increased costs. A quick glance back at Figure 3 shows that intensity standards that apply to all firms and impose the same target are particularly costly for the least productive firms. Consequently, under sector-specific standards, some firms may shut down. To quantify this margin requires additional model structure; namely, we must add fixed costs.

Define a fixed cost of operating $f_i > 0$ that potentially differs across sectors. Firm profits are $\pi_i(\varphi) = p_i(\varphi)y_i(\varphi)/\sigma - f_i$ and therefore total profit of sector i is $\Pi_i = \beta_i/\sigma - M_i f_i$, where M_i is the number of operating firms. Firms operate only if profits are non-negative; the threshold level of productivity to operate is $\varphi_i^* = \pi_i^{-1}(0)$. A firm with this productivity will have revenues $p_i(\varphi_i^*)y_i(\varphi_i^*) = \sigma f_i$. Together with $p_i(\varphi)y_i(\varphi) = \beta_i P_i^{\sigma-1}\left(mc_i(\varphi)\right)^{1-\sigma}$ and $c_i(\varphi) \propto$

$\varphi^{-1}\tau_i^e(\varphi)^{1-\alpha_i}\tau_i^y(\varphi)$, we have counterfactual price changes

$$\hat{P}_i = \frac{\tau_i^e(\varphi_i^*)^{1-\alpha_i}\tau_i^y(\varphi_i^*)}{\hat{\varphi}_i^*}, \tag{19}$$

where $\hat{x} = x'/x$ denotes a counterfactual equilibrium value of a variable relative to its initial equilibrium value.

We know how policies affect distortions, but we must now consider how they affect a sector's price index. An alternative expression for sectoral productivity A_i is useful.

Proposition 4 *Sectoral productivity A_i based on firm allocations (equation 15) is equivalently*

$$\left(\int \left[\varphi l_i(\varphi)^{\alpha_i} e_i(\varphi)^{1-\alpha_i}\right]^{\frac{\sigma-1}{\sigma}} g_i(\varphi)\mathrm{d}\varphi\right)^{\frac{\sigma}{\sigma-1}} = \bar{\tau}_i^{e\,1-\alpha_i}\bar{\tau}_i^y\tilde{\varphi}_i,$$

where $\bar{\tau}_i^y$ and $\bar{\tau}_i^e$ are as defined in Proposition 2 and

$$\tilde{\varphi}_i = \left[\int \left(\frac{\varphi}{\tau_i^e(\varphi)^{1-\alpha_i}\tau_i^y(\varphi)}\right)^{\sigma-1} g_i(\varphi)\mathrm{d}\varphi\right]^{1/(\sigma-1)}.$$

Proof: See Appendix A.

This proposition, along with equation 13 and optimal firm pricing $p_i(\varphi) = mc_i(\varphi)$, implies

$$P_i = \frac{m\bar{c}_i\bar{\tau}_i^{e\,1-\alpha_i}\bar{\tau}_i^y}{A_i}, \tag{20}$$

where $\bar{c}_i = \alpha_i^{-\alpha_i}(1-\alpha_i)^{-(1-\alpha_i)}$. With this and equation 19,

$$\hat{A}_i = \left(\frac{\bar{\tau}_i^e}{\tau_i^e(\varphi_i^*)}\right)^{1-\alpha_i}\frac{\bar{\tau}_i^y}{\tau_i^y(\varphi_i^*)}\hat{\varphi}_i^*, \tag{21}$$

which allows us to solve for how policies affect exit through $\hat{\varphi}_i^*$.[13]

Before proceeding to our full quantitative simulations with exit, we should discuss how exit can affect sectoral productivity. There are two potential effects. First, as firms exit there are fewer goods available to produce sectoral output. This loss of variety directly lowers the productivity of the sectoral aggregator. On the other hand, firms that exit have lower productivity than firms that do not. As firms exit, sectoral productivity will tend to increase through this selection effect.

We can get a sense of the quantitative magnitude of this effect if energy and output distortions do not vary across firms. In this case, it is straightforward to show that Pareto distributed

[13]The algorithm to compute $\hat{\varphi}_i^*$ is simple. Given an initial and a counterfactual threshold φ_i^* and $\varphi_i^{*\prime}$ the change in productivity $\hat{A}_i$ is from equation 15, with $l_i(\varphi) = e_i(\varphi) = 0$ for firms below the threshold, and the average distortions $\bar{\tau}_i^c$ and $\bar{\tau}_i$ are from Proposition 2. We then back out a new $\hat{\varphi}_i^*$ implied by $\hat{A}_i$ and the distortions from equation 21, and iterate until convergence. The value of the initial threshold does not matter, so long as it exceeds $\underline{\varphi}_i$ sufficiently such that the new counterfactual productivity does as well.

productivity $g_i(\varphi)$ implies

$$A_i = \left(\frac{\theta}{\theta - (\sigma - 1)}\right)^{1/(\sigma-1)} M_i^{\frac{1}{\sigma-1} - \frac{1}{\theta}}, \tag{22}$$

where M_i is the mass of operating firms (a familiar result in trade models based on Melitz, 2003). If $\theta > \sigma - 1$, then the first effect – lower productivity through loss of variety – dominates. Of course, as we discussed, the firm size distribution follows a power law distribution with a parameter close to (but larger than) one. So θ is close to (but larger than) $\sigma - 1$. In our calibration, $\theta = 1.06 \times (\sigma - 1)$ and $\sigma = 4$ so the strength of changes in M_i on sectoral TFP A_i is minor – a 10 per cent decrease in the number of operating firms lowers A_i by less than 0.2 per cent. Of course, this only holds when there are no distortions. If there are, then no closed-form solution to A_i as above exists; so, we must simulate the full model.

We display alternative results in Table 4. Our quantitative results are amplified significantly as firms exit. To lower total energy use by 10 per cent, aggregate TFP declines by nearly 0.3 per cent – more than double the productivity effect without exit and roughly equivalent to \$72 per tonne of abated emissions. For a 20 per cent reduction in energy use, aggregate productivity declines by over 0.7 per cent, or over \$91 per tonne. The number of firms that shut down in this experiment is also large. We report the change in the number of operating firms, $100 \times (\hat{M}_i - 1)$, for each industry in the last column of Table 3. We also report the change in sectoral TFP and allocations when firm exit is possible; almost all of the effects are amplified. These results strongly suggest that to the extent that sector-specific standards disproportionally burden low-productivity firms, firm exit can dramatically increase the productivity costs of standards. The option to pay a fine, as before, dramatically limits the negative effect of standards and shrinks the number of firms wishing to exit.

6.2 Upper-Level CES Aggregator

So far, the share of aggregate expenditures allocated to each sector was fixed at β_i. This limits the extent to which policies can affect allocations between sectors. We now generalize aggregate production to a CES aggregator

$$Y = \left(\sum_{i=1}^{N} \beta_i Y_i^{(\rho-1)/\rho}\right)^{\rho/(\rho-1)}, \tag{23}$$

where ρ is the elasticity of substitution between goods from different sectors. Now, sectoral revenue is proportional to prices and final goods weights β_i according to $P_i Y_i \propto \beta_i^{\rho} P_i^{1-\rho}$. From equation 20,

$$P_i Y_i \propto \beta_i^{\rho} \left(\frac{\bar{c}_i \bar{\tau}_i^{e\,1-\alpha_i} \bar{\tau}_i^{y}}{A_i}\right)^{1-\rho} \equiv r_i,$$

Table 4: Sector-Specific Standards vs Energy Taxes, with Entry and Exit

	Lower Energy Use by 10%		Lower Energy Use by 20%	
	Intensity Standards	Energy Taxes	Intensity Standards	Energy Taxes
Main Results				
Change in Aggregate GDP	-0.797%	-0.506%	-1.796%	-1.068%
Change in Aggregate TFP	-0.292%	0.000%	-0.736%	0.000%
... Within-Sector Productivity Change	-0.295%	0.000%	-0.742%	0.000%
... Between-Sector Allocations	0.002%	0.000%	0.006%	0.000%
With Option to Pay a Fine (Equivalent to 50% Energy Tax)				
Change in Aggregate GDP	-0.592%	-0.506%	-1.156%	-1.068%
Change in Aggregate TFP	-0.087%	0.000%	-0.089%	0.000%
... Within-Sector Productivity Change	-0.089%	0.000%	-0.091%	0.000%
... Between-Sector Allocations	0.002%	0.000%	0.002%	0.000%

Displays results of lower total energy use with sector-specific standards when there are fixed costs of operating. Contrast these results with those in Table 2, where fixed costs are zero.

where r_i is proportional to the equilibrium share of total expenditures allocated to sector i. Similarly, define $r_i^* = \beta_i^\rho \left(\bar{c}_i / A_i^*\right)^{1-\rho}$ as the *optimal* expenditure share for sector i in the case of no distortions. Notice that if $\rho > 1$, distortions will affect expenditures and this adds another channel through which policies can affect the allocation of labour and energy.

Following the derivations for equilibrium allocations in Section 4.2, equilibrium allocations are

$$l_i = \frac{\alpha_i r_i / \bar{\tau}_i^y}{\sum_{j=1}^N \alpha_j r_j / \bar{\tau}_j^y} \qquad \text{and} \qquad e_i = \frac{(1-\alpha_i) r_i / \bar{\tau}_i^e \bar{\tau}_i^y}{\sum_{i=1}^N (1-\alpha_j) r_j / \bar{\tau}_j^e \bar{\tau}_j^y},$$

and optimal allocations are

$$l_i^* = \frac{\alpha_i r_i^*}{\sum_{j=1}^N \alpha_j r_j^*} \qquad \text{and} \qquad e_i^* = \frac{(1-\alpha_i) r_i^*}{\sum_{j=1}^N (1-\alpha_j) r_j^*}.$$

Taxes no longer implement optimal allocations, as variation in $\tau^{1-\alpha_i}$ causes variation in r_i. Given sectoral revenue, allocations within sectors are as in the baseline model. We also return to our baseline case of no entry or exit. So, equations 17 and 18, along with average distortions from Proposition 2, still hold.

Finally, we must define a new notion of aggregate productivity. From equation 23, changes in aggregate output equal a kind of average change in sectoral output. As a sector's output change depends on the change in its productivity and allocations of labour and energy,

Table 5: Sector-Specific Standards vs Energy Taxes

	Lower Energy Use by 10%		Lower Energy Use by 20%	
	Intensity Standards	Uniform Taxes	Intensity Standards	Uniform Taxes
Change in Aggregate GDP	-0.540%	-0.425%	-1.251%	-0.862%
Change in Aggregate TFP	-0.118%	-0.003%	-0.363%	-0.005%
... Within-Sector Productivity Change	-0.143%	0.000%	-0.444%	0.000%
... Between-Sector Allocations	0.075%	-0.003%	0.148%	-0.008%

Displays effect of using sector-specific standards and taxes to lower aggregate (economy-wide) energy use when $\rho = 1.685$. Under standards, all firms must improve their energy intensity at or beyond a target that is based on the sector's prior average. GDP changes are due to lower energy use and from lower TFP (aggregate labour supply is unchanged). We further decompose TFP changes into (1) sectoral productivity changes and (2) resource allocations between sectors. These components are not additive. See text for details.

$$\hat{Y} = \left(\sum_{i=1}^{N} \omega_i \left(\hat{A}_i \hat{L}_i^{\alpha_i} \hat{E}_i^{1-\alpha_i} \right)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)},$$

where $\omega_i = \beta_i Y_i^{(\rho-1)/\rho} / \sum_{j=1}^{N} \beta_j Y_j^{(\rho-1)/\rho}$. Now, define aggregate TFP as

$$\hat{A} = \hat{Y} / \left(\sum_{i=1}^{N} \omega_i \left(\hat{E}^{1-\alpha_i} \right)^{(\rho-1)/\rho} \right)^{\rho/(\rho-1)}, \tag{24}$$

which is the equilibrium change in output $\hat{Y}$ relative to what the change would have been if labour and energy were optimally allocated, sectoral productivity did not change, and the aggregate energy change was $\hat{E}$. When $\rho = 1$, equation 24 is equivalent to the baseline model's equation 16.

If $\rho > 1$, price increases in one sector will lead demand for the other goods to increase. If $\rho < 1$, the reverse is true. We follow Acemoglu and Guerrieri (2008) in calibrating this parameter to be consistent with time-series patterns in real and nominal output by sector. As $P_i Y_i \propto \beta_i^{\rho} P_i^{1-\rho}$, take logs and rearrange to yield

$$ln\left(P_{it} Y_{it}\right) \quad = \quad \delta_i + \delta_t + \beta ln\left(Y_{it}\right) + \epsilon_{it},$$

where δ_i and δ_t are industry and year fixed-effects to capture β and common (P_t, Y_t) terms. We estimate this using U.S. sector-level annual data between 1970 and 2005 from the EU-KLEMS dataset for the United States (which provides a longer panel than our primary data source). The data provides a direct measure of a sector's nominal output ($P_{it} Y_{it}$) but not its real output (Y_{it}). Instead, we have a measure of each sector's real quantity *index* of output. It is sufficient to use this quantity index for Y_{it} as the sectoral fixed effects will capture differences across sectors in their base year prices. The coefficient estimate that results is precisely estimated: $\hat{\beta} = 0.406$ with a standard error for this estimate of 0.0169. Since $\beta = (\rho - 1)/\rho$, we have that $\hat{\rho} = 1/(1 - \hat{\beta}) = 1.685$, with a two-standard-error confidence interval of [1.599, 1.780]. This estimate is larger

than $\hat{\rho} = 0.76$ found by Acemoglu and Guerrieri (2008). We find larger substitutability since Acemoglu and Guerrieri (2008) aggregate all sectors into two while we use 48 (it is well known that substitutability increases with the level of disaggregation).

We display the results of this expanded model in Table 5 for sector-specific standards. Having higher substitutability between sectors ($\rho > 1$) dampens the negative effect of standards relative to our baseline results in Table 2. The productivity effect within sectors is unchanged, as expected, but now between-sector reallocation of labour offsets some of the negative effect on aggregate TFP. This should not be surprising. If sectoral output was perfectly substitutable for producing aggregate output (the case of $\rho \to \infty$), then small price changes in a sector's output would induce a large reallocation and therefore only a minor change in the cost of producing aggregate output. If $\rho < 1$, the reverse is true and standards have an even larger negative effect than our baseline results. In any case, incorporating between-sector reallocation does not substantially change our initial results, qualitatively or quantitatively.

7 Conclusion

Productivity differences between firms has important implications for the economic costs of intensity standards. If policies disproportionately increase energy prices for some producers relative to others, as intensity standards tend to do, then a misallocation of resources results and aggregate and sectoral productivity is consequently lowered. To investigate this, both qualitatively and quantitatively, we develop a new multi-sector model of energy use with heterogeneous firms. The model admits highly tractable and intuitive expressions that map allocations (of labour and energy) to productivity. It also cleanly maps to available sector-level data on production and input use and allows for tractable and transparent policy analysis. Despite its richness, quantitative simulations are straightforward.

We specifically consider three types of standards: (1) targets that are firm-specific; (2) sector-specific targets that are common across firms within each sector; and (3) targets that apply only to sufficiently large emitters. We find intensity standards significantly lower productivity. This is especially true with sector-specific standards, as this disproportionately increases costs for low-productivity firms. Uniform taxes that achieve the same environmental objective have no effect on productivity. We further explore means of lowering the negative effects of standards. Allowing firms to pay fines if intensity standards are not met is particularly effective, so long as fines do not vary across firms. While our quantitative results are only illustrative, the general result is clear: strict enforcement of intensity standards is far from the optimal means of lowering energy consumption or reducing emissions.

References

Acemoglu, Daron and Veronica Guerrieri, "Capital Deepening and Nonbalanced Economic Growth," *Journal of Political Economy*, 2008, *116* (3), 467–498.

Aitchison, J. and J. A. C. Brown, *The log-normal distribution with special reference to its uses in economics*, New York: Cambridge University Press, 1957.

Axtell, Robert, "Zipf Distribution of U.S. Firm Sizes," *Science*, 2001, *293* (5536), 1818–1820.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, 2013, *103* (1), 305–334.

Brandt, Loren, Trevor Tombe, and Xiadong Zhu, "Factor Market Distortions Across Time, Space, and Sectors in China," *Review of Economic Dynamics*, 2013, *16* (1), 39–58.

Broda, Christian and David E. Weinstein, "Globalization and the Gains from Variety," *The Quarterly Journal of Economics*, 2006, *121* (2), 541–585.

Copeland, Brian and M. Scott Taylor, "North-South Trade and the Environment," *The Quarterly Journal of Economics*, 1994, *109*, 755–787.

de Vries, Frans, Bouwe Dijkstra, and Matthew McGinty, "On Emissions Trading and Market Structure: Cap-and-Trade versus Intensity Standards," *Environmental & Resource Economics*, 2014, *58* (4), 665–682.

di Giovanni, Julian and Andrei Levchenko, "Firm entry, trade, and welfare in Zipf's world," *Journal of International Economics*, 2013, *89* (2), 283–296.

_ , _ , and Romain Ranciere, "Power laws in firm size and openness to trade: Measurement and implications," *Journal of International Economics*, 2011, *85* (1), 42–52.

Fischer, Carolyn and Michael Springborn, "Emissions targets and the real business cycle: Intensity targets versus caps or taxes," *Journal of Environmental Economics and Management*, 2011, *62* (3), 352–366.

Gerlagh, Reyer and Bob van der Zwaan, "Options and Instruments for a Deep Cut in CO2 Emissions: Carbon Dioxide Capture or Renewables, Taxes or Subsidies?," *The Energy Journal*, 2006, *27* (3), 25–48.

Goulder, Lawrence and Andrew Schein, "Carbon Taxes vs. Cap and Trade: A Critical Review," *NBER Working Paper*, 2013, (19338).

Helfand, Gloria E., "Standards versus Standards: The Effects of Different Pollution Restrictions," *The American Economic Review*, 1991, *81* (3), 622–634.

Holland, Stephen, Jonathan Hughes, and Christopher Knittel, "Greenhouse Gas Reductions under Low Carbon Fuel Standards?," *American Economic Journal: Economic Policy*, 2009, *1* (1), 106–46.

Holland, Stephen P., "Taxes and Trading versus Intensity Standards: Second-Best Environmental Policies with Incomplete Regulation (Leakage) or Market Power," NBER Working Paper 15262, National Bureau of Economic Research 2009.

_ , "Emissions taxes versus intensity standards: Second-best environmental policies with incomplete regulation," *Journal of Environmental Economics and Management*, 2012, *63* (3), 375–387.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

Jones, Charles, "Misallocation, Economic Growth, and Input-Output Economics," in "Advances in Economics and Econometrics," Vol. 2, Cambridge University Press, 2013.

Li, Zhe and Jianfei Sun, "Environmental Policy with Heterogeneous Plants," 2011.

_ **and Shouyoung Shi**, "Emission Tax or Standard? The Role of Productivity Dispersion," 2012.

McKitrick, Ross, "Decentralizing a Regulatory Standard Expressed in Ratio or Intensity Form," *The Energy Journal*, 2005, *26* (4), 43–51.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Newell, Richard G. and Robert N. Stavins, "Cost Heterogeneity and the Potential Savings from Market-Based Policies," *Journal of Regulatory Economics*, 2003, *23* (1), 43–59.

Parry, Ian W.H. and Roberton C. Williams, "Moving U.S. Climate Policy Forward: Are Carbon Taxes the Only Good Alternative?," *Resources For the Future, Discussion Papers*, 2011, (dp-11-02).

Restuccia, Diego and Richard Rogerson, "Policy Distortions and Aggregate Productivity with Heterogeneous Plants," *Review of Economic Dynamics*, 2008, *11* (4), 707–720.

Appendix A: Proofs of Propositions

Proof of Proposition 1: To show standards lower sectoral productivity A_i, derive the optimal labour and energy allocations and then show standards result in different shares. Maximize A_i subject to $\int l_i(\varphi)g_i(\varphi)d\varphi = 1$ and $\int e_i(\varphi)g_i(\varphi)d\varphi = 1$. The first-order condition for labour is

$$\frac{\sigma}{\sigma-1}A_i^{\frac{1}{\sigma}}\frac{\sigma-1}{\sigma}y_i(\varphi)^{-\frac{1}{\sigma}}\varphi\alpha_i l_i^{\alpha_i-1}e_i^{1-\alpha_i}g_i(\varphi) = \lambda,$$

and for energy is

$$\frac{\sigma}{\sigma-1}A_i^{\frac{1}{\sigma}}\frac{\sigma-1}{\sigma}y_i(\varphi)^{-\frac{1}{\sigma}}\varphi(1-\alpha_i)l_i^{\alpha_i}e_i^{-\alpha_i}g_i(\varphi) = \mu.$$

Rearrange the condition for labour to yield,

$$\alpha_i A_i \frac{y_i(\varphi)^{\frac{\sigma-1}{\sigma}}}{\int y_i(\varphi)^{\frac{\sigma-1}{\sigma}}g_i(\varphi)\mathrm{d}\varphi} = \lambda l_i.$$

Integrate across firms to get $A_i\alpha_i = \lambda$, which then implies

$$l_i(\varphi) = \frac{y_i(\varphi)^{\frac{\sigma-1}{\sigma}}}{\int y_i(\varphi)^{\frac{\sigma-1}{\sigma}}g_i(\varphi)\mathrm{d}\varphi}.$$

Similar derivations for energy gives $A_i(1-\alpha_i) = \mu$ and therefore

$$e_i(\varphi) = \frac{y_i(\varphi)^{\frac{\sigma-1}{\sigma}}}{\int y_i(\varphi)^{\frac{\sigma-1}{\sigma}} g_i(\varphi)\mathrm{d}\varphi}.$$

So, the optimal allocations for labour and energy are the same and output is then $y_i(\varphi) \propto \varphi l_i(\varphi)^{\alpha_i} e_i(\varphi)^{1-\alpha_i} = \varphi l_i(\varphi) = \varphi e_i(\varphi)$. The above becomes,

$$l_i(\varphi) \quad = \quad \frac{\varphi^{\frac{\sigma-1}{\sigma}} l_i(\varphi)^{\frac{\sigma-1}{\sigma}}}{\int \varphi^{\frac{\sigma-1}{\sigma}} l_i(\varphi)^{\frac{\sigma-1}{\sigma}} g_i(\varphi)\mathrm{d}\varphi} \Rightarrow l_i(\varphi) \propto \varphi^{\sigma-1},$$

where the constant of proportionality is common across all firms in the sector. The same holds for energy. Integrating across firms and dividing to eliminate the constant, we have our result

$$l_i(\varphi)^* = e_i(\varphi)^* = \frac{\varphi^{\sigma-1}}{\int \varphi^{\sigma-1} g_i(\varphi)\mathrm{d}\varphi}.$$

To implement these optimal shares in the decentralized economy, note that equilibrium firm output is $y_i(\varphi) \propto p_i(\varphi)^{-\sigma}$ and therefore

$$l_i = \frac{p_i(\varphi)^{-(\sigma-1)}}{\int p_i(\varphi)^{-(\sigma-1)} g_i(\varphi)\mathrm{d}\varphi}.$$

So, we require $p_i(\varphi) \propto \varphi^{-1}$ to implement the first best. As prices are a fixed markup over marginal costs, $c_i(\varphi) = \frac{1}{\varphi_i}\left(\frac{1}{\alpha_i}\right)^{\alpha_i}\left(\frac{\tau_i^e(\varphi)}{1-\alpha_i}\right)^{1-\alpha_i}\tau_i^y(\varphi)$. Further, since $\tau_i^y(\varphi) = \alpha_i + (1-\alpha_i)/\tau_i^e(\varphi)$ with standards we require $\tau_i^e(\varphi)$ not depend on firm productivity to implement the optimal allocations. Any standards for which this is not the case will result in allocations that differ from $l_i(\varphi)^*$ and $e_i(\varphi)^*$ and therefore sectoral productivity A_i will fall. ■

Proof of Proposition 2: Total spending on labour by all firms in sector i is $L_i = \int L_i(\varphi) g_i(\varphi)\mathrm{d}\varphi$. Optimal firm labour decisions require $L_i(\varphi) = \alpha_i p_i(\varphi) y_i(\varphi)/m\tau_i^y(\varphi)$. Integrate across all firms within the sector and then sum across all sectors to get

$$\begin{aligned}
\frac{L_i}{L} \equiv l_i \quad &= \quad \frac{\alpha_i \int \left(p_i(\varphi)y_i(\varphi)/\tau_i^y(\varphi)\right) g_i(\varphi)\mathrm{d}\varphi}{\sum_{j=1}^N \alpha_j \int \left(p_j(\varphi)y_j(\varphi)/\tau_j^y(\varphi)\right) g_j(\varphi)\mathrm{d}\varphi}, \\
&= \quad \frac{\alpha_i \left[\int \left(p_i(\varphi)y_i(\varphi)\right) g_i(\varphi)\mathrm{d}\varphi\right]/\bar{\tau}_i^y}{\sum_{j=1}^N \alpha_j \left[\int \left(p_j(\varphi)y_j(\varphi)\right) g_j(\varphi)\mathrm{d}\varphi\right]/\bar{\tau}_i^y}, \\
&= \quad \frac{\alpha_i \beta_i/\bar{\tau}_i^y}{\sum_{j=1}^N \alpha_j \beta_j/\bar{\tau}_j^y}.
\end{aligned}$$

The above relies on $\bar{\tau}_i^y = \left[\int p_i(\varphi)y_i(\varphi)g_i(\varphi)\mathrm{d}\varphi\right] / \left[\int \left(p_i(\varphi)y_i(\varphi)/\tau_i^y(\varphi)\right) g_i(\varphi)\mathrm{d}\varphi\right]$. We will justify

this shortly. First, the allocation of energy is similarly derived. As $E_i = \int E_i(\varphi) g_i(\varphi) \mathrm{d}\varphi$ and $E_i(\varphi) = (1-\alpha_i) p_i(\varphi) y_i(\varphi) / m \tau_i^y(\varphi) \tau_i^e(\varphi)$, we have

$$\begin{aligned} \frac{E_i}{E} \equiv e_i &= \frac{(1-\alpha_i) \int \left(p_i(\varphi) y_i(\varphi) / \tau_i^y(\varphi) \tau_i^e(\varphi) \right) g_i(\varphi) \mathrm{d}\varphi}{\sum_{j=1}^N (1-\alpha_j) \int \left(p_j(\varphi) y_j(\varphi) / \tau_j^y(\varphi) \tau_j^e(\varphi) \right) g_j(\varphi) \mathrm{d}\varphi}, \\ &= \frac{(1-\alpha_i) \beta_i / \bar{\tau}_i^y \bar{\tau}_i^e}{\sum_{j=1}^N (1-\alpha_j) \beta_j / \bar{\tau}_j^y \bar{\tau}_j^e}, \end{aligned}$$

where the second line uses

$$\bar{\tau}_i^y \bar{\tau}_i^e = \frac{\int p_i(\varphi) y_i(\varphi) g_i(\varphi) \mathrm{d}\varphi}{\int p_i(\varphi) y_i(\varphi) \frac{1}{\tau_i^y(\varphi)} \frac{1}{\tau_i^e(\varphi)} g_i(\varphi) \mathrm{d}\varphi}.$$

Now, we justify the use of harmonic means for the average output and energy distortions of each sector. Begin with the average output distortion $\bar{\tau}_i^y$. From above, we see that for each firm the output distortion $\tau_i^y(\varphi)$ is proportional to the revenue per worker; specifically, rearrange the above expression for $L_i(\varphi)$ to find

$$\tau_i^y(\varphi) = \frac{\alpha_i}{L_i(\varphi) / \left[p_i(\varphi) y_i(\varphi) / m \right]}.$$

Intuitively, it is a measure of by how much labour's share of total costs deviate from the optimal share α_i. For an entire industry, spending on labour is $\alpha_i \int \left(p_i(\varphi) y_i(\varphi) / m \tau_i^y(\varphi) \right) g_i(\varphi) \mathrm{d}\varphi$ and total costs are $\int p_i(\varphi) y_i(\varphi) g_i(\varphi) \mathrm{d}\varphi / m$. So, labour's share of a sector's total costs is $\frac{\alpha_i \int \left[p_i(\varphi) y_i(\varphi) / \tau_i^y(\varphi) \right] g_i(\varphi) \mathrm{d}\varphi}{\int p_i(\varphi) y_i(\varphi) g_i(\varphi) \mathrm{d}\varphi}$ and therefore

$$\bar{\tau}_i^y = \frac{\int p_i(\varphi) y_i(\varphi) g_i(\varphi) \mathrm{d}\varphi}{\int p_i(\varphi) y_i(\varphi) \frac{1}{\tau_i^y(\varphi)} g_i(\varphi) \mathrm{d}\varphi}.$$

For the average energy distortion, the most intuitive derivation is slightly different. Think of $\tau_i^e(\varphi)$ as an explicit tax. We are interested in the average tax rate within a sector. A firm's pre-tax energy spending is $E_i(\varphi) = (1-\alpha_i) p_i(\varphi) y_i(\varphi) / m \tau_i^y(\varphi) \tau_i^e(\varphi)$. For the overall average energy tax rate across all active firms within an industry, take the ratio of taxes paid to pre-tax energy spending

$$\begin{aligned} \bar{\tau}_i^e &\equiv \frac{\int E_i(\varphi) (\tau_i^e(\varphi) - 1) g_i(\varphi) \mathrm{d}\varphi}{\int E_i(\varphi) g_i(\varphi) \mathrm{d}\varphi} + 1, \\ &= \frac{\int E_i(\varphi) \tau_i^e(\varphi) g_i(\varphi) \mathrm{d}\varphi}{\int E_i(\varphi) g_i(\varphi) \mathrm{d}\varphi}, \\ &= \frac{\int p_i(\varphi) y_i(\varphi) \frac{1}{\tau_i^y(\varphi)} g_i(\varphi) \mathrm{d}\varphi}{\int p_i(\varphi) y_i(\varphi) \frac{1}{\tau_i^y(\varphi) \tau_i^e(\varphi)} g_i(\varphi) \mathrm{d}\varphi}, \end{aligned}$$

which is a weighted harmonic mean of $\tau_i^e(\varphi)$. ■

Proof of Proposition 3: This is a simple constrained maximization problem. Maximize aggregate productivity $A = \prod_{i=1}^{N} A_i^{\beta_i} l_i^{\alpha_i \beta_i} e_i^{(1-\alpha_i)\beta_i}$ subject to $\sum_{i=1}^{N} l_i = 1$ and $\sum_{i=1}^{N} e_i = 1$. The first order conditions for l_i are

$$\alpha_i \beta_i \frac{1}{l_i} = \lambda,$$

where λ is the shadow price of labour. Summing across producers means $\lambda = \sum_{i=1}^{N} \alpha_i \beta_i$ and therefore the optimal labour allocation is $l_i = \alpha_i \beta_i / \sum_{j=1}^{N} \alpha_j \beta_j$. Similarly, the first order conditions for e_i are

$$(1-\alpha_i)\beta_i \frac{1}{e_i} = \mu,$$

and therefore we also have $\mu = \sum_{i=1}^{N} (1-\alpha_i)\beta_i$. As with labour, this implies $e_i = (1-\alpha_i)\beta_i / \sum_{j=1}^{N} (1-\alpha_j)\beta_j$.

There are a number of ways to find the distortions that implement these shares. We have previously seen optimal labour decisions require $\alpha_i P_i Y_i / \bar{\tau}_i^y = L_i$ and therefore

$$l_i = \frac{\alpha_i \beta_i / \bar{\tau}_i^y}{\sum_{j=1}^{N} \alpha_j \beta_j / \bar{\tau}_j^y} = \frac{\alpha_i \beta_i}{\sum_{j=1}^{N} \alpha_j \beta_j}$$

only if $\bar{\tau}_i^y$ is the same for all i.

For energy, we have $E_i = \frac{1-\alpha_i}{\alpha_i} \frac{1}{\bar{\tau}_i^e} l_i L$. At the optimal allocation of labour, $E_i = \frac{1}{\bar{\tau}_i^e} \frac{(1-\alpha_i)\beta_i}{\sum_{j=1}^{N} \alpha_j \beta_j} L$ and therefore

$$e_i = \frac{(1-\alpha_i)\beta_i / \bar{\tau}_i^e}{\sum_{j=1}^{N} (1-\alpha_j)\beta_j / \bar{\tau}_j^e} = \frac{(1-\alpha_i)\beta_i}{\sum_{j=1}^{N} (1-\alpha_j)\beta_j}$$

only if $\bar{\tau}_i^e$ is the same for all i.

There is an alternative (completely equivalent) way to establish this proposition that may be instructive. Maximize output per worker

$$\frac{Y}{L} = \prod_{i=1}^{N} \left(\varphi_i \left(\frac{1-\alpha_i}{\alpha_i} \frac{1}{\bar{\tau}_i^e} \right)^{1-\alpha_i} l_i \right)^{\beta_i},$$

by choosing labour allocations and energy distortions $\bar{\tau}_i^e$, subject to $E = \sum_{i=1}^{N} \frac{1-\alpha_i}{\alpha_i} \frac{1}{\bar{\tau}_i^e} l_i L$ and $\sum_{i=1}^{N} l_i = 1$. The first-order conditions for labour shares l_i are

$$\beta_i = \lambda \frac{1-\alpha_i}{\alpha_i} \frac{1}{\bar{\tau}_i^e} l_i L + l_i \mu,$$

and therefore $(1-\mu)E^{-1} = \lambda$. The first-order conditions for the energy distortions are

$$\beta_i (1-\alpha_i) = \lambda \frac{1-\alpha_i}{\alpha_i} \frac{1}{\bar{\tau}_i^e} l_i L.$$

Together, these imply

$$\beta_i = \beta_i(1-\alpha_i) + l_i\mu,$$

or $\mu = \alpha_i\beta_i/l_i$ and $\mu = \sum_{i=1}^N \alpha_i\beta_i$. This gives the optimal allocation of labour as above. Use this value for μ in the condition $(1-\mu)E^{-1} = \lambda$ derived earlier to get $E^{-1}\sum_{i=1}^N(1-\alpha_i)\beta_i = \lambda$. The first-order condition for energy distortions then implies

$$\bar{\tau}_i^e = \left(\frac{\sum_{j=1}^N(1-\alpha_j)\beta_j}{\sum_{j=1}^N \alpha_j\beta_j}\right)\frac{L}{E},$$

which is the same for all sectors. In the initial equilibrium, $\bar{\tau}_i^e = 1$. So, to lower energy use to $E' < E$ the energy distortion must be $\bar{\tau}_i^e = E/E'$. Notice with these distortions, we have

$$e_i = \frac{(1-\alpha_i)\beta_i}{\sum_{j=1}^N(1-\alpha_j)\beta_j},$$

which is the first-best.

It remains to show that standards cannot achieve the optimal allocations if α_i differ across sectors. We already established that $\bar{\tau}_i^e$ and $\bar{\tau}_i^y$ cannot vary across sectors, so neither does $\bar{\tau}_i^y\bar{\tau}_i^e$. It is straightforward to show $\bar{\tau}_i^y = \alpha_i + (1-\alpha_i)/\bar{\tau}_i^e$.[14] Given this, we have $\bar{\tau}_i^y\bar{\tau}_i^e = 1-\alpha_i+\alpha_i\bar{\tau}_i^e$ and therefore optimal allocations require $1-\alpha_i+\alpha_i\bar{\tau}^e = 1-\alpha_j+\alpha_j\bar{\tau}^e$ for all $i \neq j$ and some common output distortion $\bar{\tau}^e$. Equivalently, they imply $\alpha_j - \alpha_i = (\alpha_j-\alpha_i)\bar{\tau}^e$, which is only true if $\bar{\tau}^e = 1$ or if α_i do not vary across sectors. As $\bar{\tau}_i^e = E/E'$, any policy that lowers total energy use we will have $\bar{\tau}^e > 1$. Thus, standards cannot achieve the optimal allocations if α_i differ across sectors. With taxes, there is no output distortion (so $\bar{\tau}_i^y = 1$) and energy distortions are common $\bar{\tau}_i^e = \tau$; so, taxes implement the optimal allocations. ■

Proof of Proposition 4: To show the equivalence of $A_i = \bar{\tau}_i^{e\,1-\alpha_i}\bar{\tau}_i^y\tilde{\varphi}$ with the shares definition in equation 15, simply use the allocations from equations 17 and 18 with the average distortion measures from Proposition 2. Specifically, re-write them to show labour allocations are

$$l_i(\varphi) = \bar{\tau}_i^c\frac{r_i(\varphi)/\tau_i^y(\varphi)}{\int r_i(\varphi)g_i(\varphi)\mathrm{d}\varphi} = \bar{\tau}_i^c\frac{r_i(\varphi)/\tau_i^y(\varphi)}{\tilde{\varphi}^{\sigma-1}},$$

and, since

$$\bar{\tau}_i^e\bar{\tau}_i^y = \frac{\int r_i(\varphi)g_i(\varphi)\mathrm{d}\varphi}{\int r_i(\varphi)/\tau_i^e(\varphi)\tau_i^y(\varphi)g_i(\varphi)\mathrm{d}\varphi},$$

[14] To see this, note $E_i = (1-\alpha_i)P_iY_i/m\bar{\tau}_i^y\bar{\tau}_i^e$ and $L_i = \alpha_iP_iY_i/m\bar{\tau}_i^y$. Labour's share of a sector's total cost is $\frac{\alpha_iP_iY_i/m\bar{\tau}_i^y}{P_iY_i/m} = \alpha_i/\bar{\tau}_i^y$. Labour costs are $L_i = E_i\frac{\alpha_i}{1-\alpha_i}\bar{\tau}_i^e$. So, with standards, $\frac{L_i}{L_i+E_i} = \frac{\frac{\alpha_i}{1-\alpha_i}\bar{\tau}_i^e}{1+\frac{\alpha_i}{1-\alpha_i}\bar{\tau}_i^e} = \frac{\alpha_i}{\bar{\tau}_i^y}$ which yields the result.

energy allocations are

$$e_i(\varphi) = \bar{\tau}_i^e \bar{\tau}_i^y \frac{r_i(\varphi)/\tau_i^e(\varphi)\tau_i^y(\varphi)}{\int r_i(\varphi) g_i(\varphi) \mathrm{d}\varphi} = \bar{\tau}_i \bar{\tau}_i^c \frac{r_i(\varphi)/\tau_i^e(\varphi)\tau_i^y(\varphi)}{\tilde{\varphi}^{\sigma-1}}.$$

Insert these into equation 15 and simplify,

$$\begin{aligned}
A_i &= \left(\int \left[\varphi_i l_i(\varphi)^{\alpha_i} e_i(\varphi)^{1-\alpha_i} \right]^{(\sigma-1)/\sigma} g_i(\varphi) \mathrm{d}\varphi \right)^{\sigma/(\sigma-1)} \\
&= \bar{\tau}_i^{e\,1-\alpha_i} \bar{\tau}_i^y \left(\int \left[\frac{1}{\tilde{\varphi}^{\sigma-1}} \left(\frac{\varphi_i}{\tau_i^y(\varphi)\tau_i^e(\varphi)^{1-\alpha_i}} \right) \left(\frac{\varphi_i}{\tau_i^y(\varphi)\tau_i^e(\varphi)^{1-\alpha_i}} \right)^{\sigma-1} \right]^{(\sigma-1)/\sigma} g_i(\varphi) \mathrm{d}\varphi \right)^{\sigma/(\sigma-1)}, \\
&= \bar{\tau}_i^{e\,1-\alpha_i} \bar{\tau}_i^y \frac{1}{\tilde{\varphi}^{\sigma-1}} \left(\int \left(\frac{\varphi_i}{\tau_i^y(\varphi)\tau_i^e(\varphi)^{1-\alpha_i}} \right)^{\sigma-1} g_i(\varphi) \mathrm{d}\varphi \right)^{\sigma/(\sigma-1)}, \\
&= \bar{\tau}_i^{e\,1-\alpha_i} \bar{\tau}_i^y \tilde{\varphi}. \blacksquare
\end{aligned}$$

Appendix B: Emissions are a Free Input

If emissions are a by-product of production, the Copeland and Taylor (1994) result $Y_i = A_i L_i^{\alpha_i} E_i^{1-\alpha_i}$ holds with emissions as a free input. Total costs are therefore only L_i. Labour's share of total costs in the initial equilibrium is one. Standards, as they are *implicit* rather than *explicit* taxes, do not change this. This introduces some complications to our base model, but our key results still hold.

For simplicity, consider the case where the production technology of equation 3 applies to all firms and the exponents $\alpha_i = \alpha$ for all i. The only differences are therefore differences in productivity φ_i. Conditional labour demand is still $L_i = \frac{Y_i}{\varphi_i} \left(\frac{\tau_i^e \alpha}{1-\alpha} \right)^{1-\alpha}$ when standards impose an energy distortion τ_i^e. As emissions are free, marginal costs are

$$c_i = \frac{1}{\varphi_i} \left(\frac{1}{\alpha} \right)^{\alpha} \left(\frac{\tau_i}{1-\alpha} \right)^{1-\alpha} \alpha.$$

This is equivalent to a firm actually paying a tax on emissions but with all tax payments refunded in proportion to output. The output distortion here is $\tau_i^y = \alpha$. In the case of an explicit tax on emissions $\tau_i^e = \tau$ for all i, and marginal costs are $\partial(L_i + (\tau - 1)E_i)/\partial Y_i$, where $(\tau - 1)E_i$ are explicit tax payments. Marginal costs are

$$c_i = \frac{1}{\varphi_i} \left(\frac{1}{\alpha} \right)^{\alpha} \left(\frac{\tau}{1-\alpha} \right)^{1-\alpha} \left(\alpha + (1-\alpha) \frac{\tau - 1}{\tau} \right),$$

which reveal an output distortion $\tau_i^y = \alpha + (1-\alpha)\frac{\tau-1}{\tau}$. So, unlike the baseline model where emissions result from purchased energy inputs, explicit emissions taxes result in an output distortion. This distortion, however, declines as taxes rise. In the limit as $\tau \to \infty$, the output distortion

$\tau_i^y \to 1$. For standards, regardless of the strength of the environmental policy, the output distortion remains at $\tau_i^y = \alpha$.

With this new expression for the output distortion, we can proceed largely as in the main text. Labour's share of total costs is still α/τ_i^y as in the earlier analysis. For standards, this means labour's share of total costs is always one. For taxes, labour's share depends on the magnitude of the tax. In the limit, as $\tau \to \infty$ labour's share of total costs approaches α. The equilibrium allocation of labour and emissions are as before:

$$l_i = \frac{\alpha\beta_i/\tau_i^y}{\sum_{j=1}^N \alpha\beta_j/\tau_j^y} \qquad \text{and} \qquad e_i = \frac{(1-\alpha)\beta_i/\tau_i^y\tau_i^e}{\sum_{j=1}^N (1-\alpha)\beta_j/\tau_i^y\tau_i^e}.$$

As α are identical, it is straightforward to show aggregate productivity is

$$A = \prod_{i=1}^N \left(\varphi_i \beta_i^\alpha \left(\frac{\beta_i/\tau_i^e}{\sum_{j=1}^N \beta_j/\tau_j^e} \right)^{1-\alpha} \right)^{\beta_i}.$$

Back to the main question: How do standards affect aggregate productivity? The following proposition will prove useful.

Proposition 5 *Any policy that leaves labour allocations unchanged and creates implicit emissions taxes τ_i will change aggregate TFP according to*

$$\hat{A} = \prod_{i=1}^N \left(\frac{1/\tau_i^e}{\sum_{j=1}^N \beta_j/\tau_i^e} \right)^{\beta_i(1-\alpha)} \leq 1,$$

where hats ($\hat{x}$) will always refer to the ratio of a counterfactual equilibrium value of variable to the initial equilibrium value. The inequality is strict if and only if τ_i^e is identical across all i.

Proof: The change in TFP $\hat{A}$ can be re-expressed as $\hat{A} = \prod_{i=1}^N (\bar{\tau}^e/\tau_i^e)^{\beta_i(1-\alpha)} = (\bar{\tau}^e/\prod_{i=1}^N \tau_i^{e\,\beta_i})^{(1-\alpha)}$, where $\bar{\tau}^e$ is the weighted (by β_i) harmonic mean of τ_i^e; specifically, $\bar{\tau}^e = 1/\sum_{i=1}^N \beta_i/\tau_i^e$. As harmonic means are always less than geometric means, we have $\hat{A} \leq 1$. This holds strictly so long as at least two τ_i^e's differ. ■

We can go further if firm productivity is distributed log-normal.

Proposition 6 *If $\beta_i = 1/N$ for all i, productivity φ_i is distributed log-normal, $ln(\varphi) \sim N(\mu, \sigma^2)$, with mean μ and standard deviation σ, and if an emissions intensity standard binds on all firms then as $N \to \infty$,*

$$\hat{A} = exp\left\{ -\frac{1-\alpha}{2} \left(\frac{\sigma}{\alpha}\right)^2 \right\} \leq 1,$$

where the equality holds if and only if $\sigma = 0$ (homogeneous firms).

Proof: Input distortions are given by $\tau_i = (\varphi_i s)^{-\frac{1}{\alpha}} \frac{1-\alpha}{\alpha}$. Taking logs, $log(\tau_i) = C - \frac{1}{\alpha} log(\varphi_i)$, where C is a constant common across firms. For distortions to be distributed log-normal is therefore

Figure 7: Quantitative Example of TFP Effect of Standards

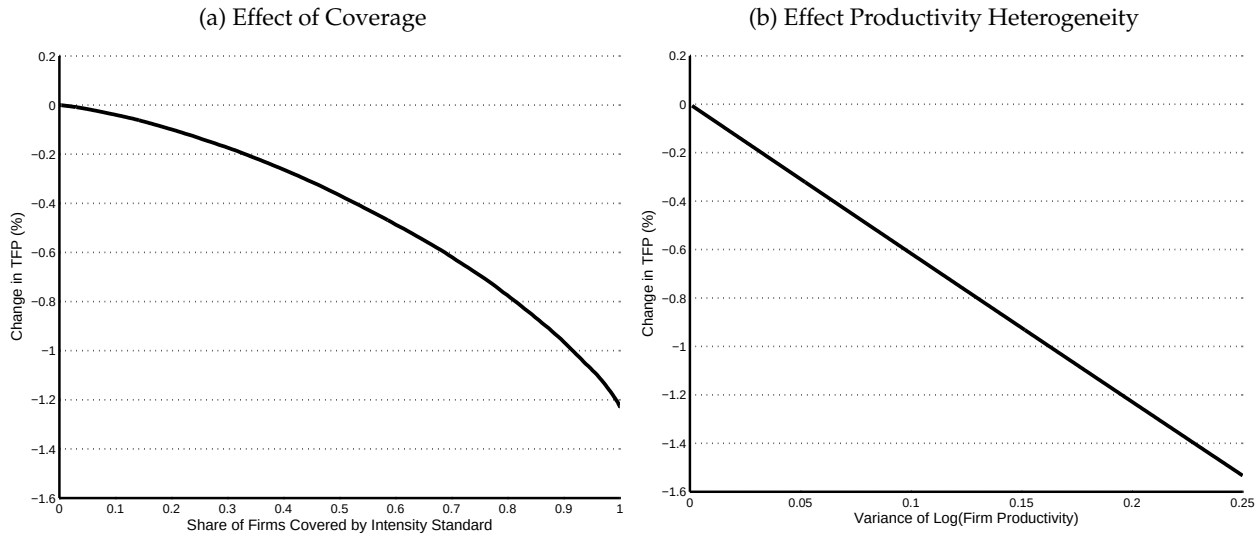

Displays the reduction of TFP for simulated set of firms with log-Normal distributed productivity, with $\alpha = 0.9$. In panel (a), the standard is gradually lowered until all firms are covered. In panel (b), the standard covers all firms, but variance in firm productivity differs along the x-axis.

sufficient that productivity is. If $ln(\varphi) \sim N(\mu, \sigma^2)$ then $ln(\tau) \sim N(\theta, \nu^2)$ where $\theta = -\mu/\alpha + C$ and $\nu = \sigma/\alpha$. As $N \to \infty$, aggregate productivity relative to the optimal is

$$\begin{aligned} ln\left(\frac{A}{A*}\right) &= \frac{1-\alpha}{N}\sum_{i=1}^{N} ln\left(\bar{\tau}/\tau_i\right), \\ &= \frac{1-\alpha}{N}\sum_{i=1}^{N} ln\left(\bar{\tau}\right) - \frac{1-\alpha}{N}\sum_{i=1}^{N} ln\left(\tau_i\right), \\ &= (1-\alpha)ln\left(\bar{\tau}\right) - (1-\alpha)\theta, \\ &= -\frac{1-\alpha}{2}\left(\frac{\sigma}{\alpha}\right)^2 \end{aligned}$$

where the third line follows from the Law of Large Numbers, and the fourth line from the harmonic mean of a log-normal distribution being $e^{\theta - \frac{1}{2}\nu^2}$ (Aitchison and Brown, 1957). ■

This is very useful.[15] If $\sigma^2 = 0.2$ (roughly what we see in the data) and $\alpha = 0.9$ then TFP losses are about 1.2 per cent if the standard binds on all firms. In Figure 7, we plot the TFP consequences of emissions intensity standards on this example industry. In panel (a), the standard is gradually lowered and this binds more and more firms. The tighter the standard, the greater the TFP loss. In panel (b), we simulate a standard that binds on all firms but for varying degree of productivity heterogeneity. When all firms have identical productivity, the standard has no TFP effect. As productivity differences increase, standards have increasingly negative effects on productivity. Models with homogeneous firms are ill-equipped to analyze the productivity effects of standards.

[15] Inspiration for this proposition is from Proposition 4 of Jones (2013).

What about explicit taxes on emissions? If a government imposes an actual tax on emissions, then all of the above analysis holds with $\tau_i^e = \tau > 1$ and only one small complication. Taxes mean that total labour costs no longer equal total costs. As we saw, taxes create a cost distortion $\tau_i^y < 1$, though as $\alpha_i = \alpha$ for all i the output distortion does not vary across producers, so $l_i = \beta_i$ from above. Moreover, as $\tau_i^e = \tau$ emissions allocations are $e_i = \beta_i$. From Proposition 5, we have $\hat{A} = 1$. So, emission intensity standards lower productivity while emissions taxes do not.

The above results are so sharp because producers all share the same abatement technology. To look at variation in abatement technologies, consider a case where α_i varies. Equilibrium allocations above still hold, although the α_i's no longer cancel out. To simplify this analysis consider two sectors: one is clean (no emissions at all; or, $\alpha_{clean} = 1$) while the other is dirty. This exercise is instructive. If there is only one dirty sector, then it is responsible for all emissions and in no sense does equalizing marginal abatement costs matter.

Let β be the share of aggregate expenditure going to the dirty sector and α as the dirty sector's abatement effectiveness parameter. Let the share of labour going to the dirty sector be l_d and the share in the clean sector be l_c. We previously saw the change in aggregate productivity is $\hat{A} = \prod_{i=1}^{N} \hat{l}_i^{\alpha_i \beta_i} \hat{e}_i^{(1-\alpha_i)\beta_i}$ and therefore in the two-sector case,

$$\hat{A} = \hat{l}_c^{(1-\beta)} \hat{l}_d^{\alpha\beta} \hat{e}^{(1-\alpha)\beta}.$$

Under a standard, labour allocations do not change, so $\hat{A}^{std} = \hat{e}^{(1-\alpha)\beta}$. Since all emissions are from the dirty sector, $\hat{e} = 1$ and therefore $\hat{A}^{std} = 1$. So, in the two-sector case, standards do not affect aggregate productivity.

In the case of a tax, output distortions will lead labour allocations to change. Specifically, since $\hat{l}_i = \frac{1/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}{\sum_{i=1}^{N}\beta_i/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}$ we have $\hat{l}_d = \frac{1/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}{(1-\beta)+\beta/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}$ and $\hat{l}_c = \frac{1}{(1-\beta)+\beta/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}$. The change in aggregate productivity is then

$$\begin{aligned}
\hat{A}^{tax} &= \left(\frac{1}{(1-\beta)+\beta/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha_i}{\alpha_i}\right)}\right)^{(1-\beta)} \left(\frac{1/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}{(1-\beta)+\beta/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}\right)^{\alpha\beta}, \qquad (25)\\
&> \left(\frac{1/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}{(1-\beta)+\beta/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)}\right)^{\alpha\beta+(1-\beta)},\\
&> 1,
\end{aligned}$$

where the second inequality follows from $\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right) > 1$; the third from $1/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right) > (1-\beta)+\beta/\left(1+\frac{\tau-1}{\tau}\frac{1-\alpha}{\alpha}\right)$, given $0 < \beta < 1$. This may be a surprising result: an explicit emissions tax *increases aggregate productivity*. This results from a more efficient allocation of labour. In the case where energy inputs are purchased, the baseline allocation of labour was already optimal so an explicit tax did not affect productivity.

This simple two-sector example demonstrates that equilibrium allocations between producers matter independently of abatement costs. Sub-optimal allocations lower productivity and differ-

ent policies have different effects on allocations. We graphically illustrate the productivity effect of taxes relative to standards in a two sector in Figure 8 for all values of β and α.

Figure 8: Productivity Gain from Tax in Simple Two-Sector Model

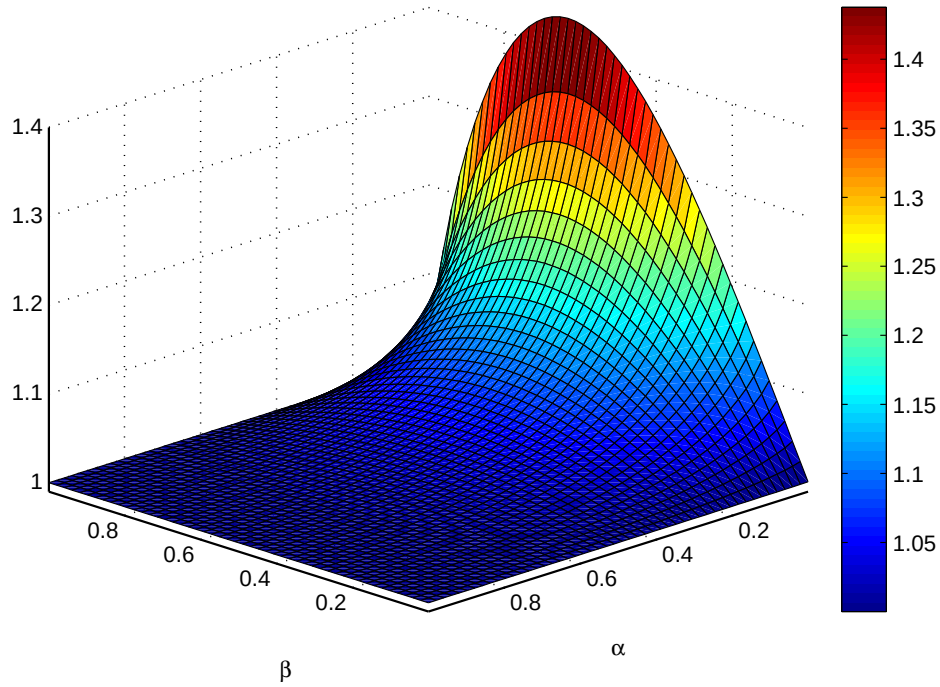

Displays the productivity effect $\hat{A}$ from equation 25 of a tax on emissions in a simple model with two sector, one clean (no emissions) and one dirty. The tax is $\tau = 1.5$, though the results hold for any τ.

Can we make similar statements about the inferiority of standards outside of this two-sector case? The optimal allocations of Proposition 3 still hold, as do the distortions that implement those allocations: $\tau_i^e = E'/E$ (as before) and τ_i^y is common across i. Sector-specific standards will not be able to implement optimal allocations, as emissions distortions will depend on firm productivity. Moreover, the share of labour to producer i is always β_i with standards. Thus, standards will never achieve the optimal labour allocation.

Can taxes? They always have $\tau_i^e = \tau$, which satisfies one of the optimal allocation conditions. Sufficiently high taxes will also satisfy the restriction on output distortions. As τ grows large, τ_i^y approaches 1. For small taxes, we do not have a general proof, though extensive computational simulations reveal taxes are always superior to standards. This reveals an important reason why we prefer the emissions from energy inputs case of the main paper: not just is it more realistic, but we can analytically establish the superiority of taxes over intensity standards. When emissions are a free input, we can only do this for "sufficiently high" levels of emissions reductions (high τ).