

A Framework for Valuing the Employment Consequences of Environmental Regulation

Robert Shimer*

February 10, 2013

Abstract

I develop a two sector model in which one sector produces a good that generates pollution, a negative externality. I show that even if it takes time for workers to switch sectors, the optimal tax on the dirty good depends only on the marginal rate of substitution between private consumption of the dirty good and pollution. The time it takes workers to switch sectors and the number of workers near the margin for switching affects the employment response to the optimal tax but not the tax itself.

*University of Chicago, Department of Economics, 1126 East 59th Street, Chicago, IL, 60637 (e-mail: robert.shimer@gmail.com). This research was supported by a grant from the Environmental Protection Agency. The paper was prepared for a conference on the “Employment Effects of Environmental Regulation” on October 26, 2012. I am grateful to Ivan Werning and conference participants for comments. The views expressed in this paper reflect that of the author and not necessarily of the Environmental Protection Agency.

1 Introduction

The goal of this paper is to develop a framework for evaluating the welfare consequences of environmental regulation, with an explicit focus on the possibility that these regulations may temporarily boost the unemployment rate. This approach draws on the modern theory of unemployment pioneered by Lucas and Prescott (1974). I consider an economy with two sectors. One produces a clean good, while the other produces a dirty good that generates a negative externality, “pollution.” Everyone in the economy would like to consume both goods but suffers from the pollution caused by other people’s consumption of the dirty good. I explore how a tax on the consumption of the dirty good and subsidy to the consumption of the clean good shifts individuals’ consumption behavior and hence the production of the two goods.

I am particularly interested in situations in which a worker’s human capital is specific to the production of one of the goods. It is possible for the worker to produce the other good, but doing so entails undergoing an unemployment spell with a consequent loss of income. In this environment, a tax on a subset of the goods in the economy hurts the workers with a comparative advantage in producing those goods because it reduces the pre-tax price of the good and hence the value of those workers’ human capital. Conversely, a subsidy to a good improves the welfare of the producers of that good. Therefore any effort to tax goods that create negative externalities will have distributional consequences, which makes a welfare analysis of such policies tricky.

There are also some interesting dynamic aspects to the tax policy. Over time, a pollution tax will cause some workers who produce the dirty good to

leave that sector and move to the clean sector of the economy, enduring a spell of unemployment. This gradually shifts the number of workers able to produce each of the goods, again with potential consequences for the optimal tax.

I show that despite these considerations, the optimal tax on the dirty good depends only on individuals' preferences, their marginal rate of substitution between the private consumption of the dirty good and pollution. In particular, it is independent of how costly is an unemployment spell and how specific is human capital.

I first consider a static model in which the dirty sector is already in decline so that some workers would be leaving it for the clean sector even without any tax. In this case, an increase in the tax on the dirty good does not affect workers' relative income, although it does induce more workers to exit the dirty sector. Under the assumption that the willingness to reduce private consumption in return for a reduction in pollution does not depend on an individual's wealth, I find that everyone agrees on the optimal tax rate. Moreover, that tax rate can be expressed as a function of the marginal rate of substitution between private consumption of the dirty good and pollution. In particular, it does not depend directly on how much unemployment workers experience as a result of the tax policy change.

I then develop a dynamic model in which workers are continuously moving back and forth between the two sectors of the economy because of idiosyncratic shocks to their human capital, or more precisely to their ability to produce each of the goods. In this case a tax on one of the goods must change workers' relative income and so has real distributional consequences. I abstract from

these distributional issues by looking at an economy with complete financial markets. Alternatively, one could allow taxes on the winners (workers with a strong comparative advantage at producing the clean good) and subsidies to the losers (workers with a strong comparative advantage at producing the dirty good). I therefore focus on the optimal policy from the perspective of an individual with the mean level of income. Such an individual's preferred tax on the dirty good is again a simple function of the marginal rate of substitution between his private consumption of the dirty good and pollution. Once again, this does not depend on how time-consuming it is for workers to switch sectors of the economy, nor does it depend on how strong is the comparative advantage that some workers have for working in one sector of the economy.

These results may seem surprising, so it is worth emphasizing that they do not imply that the optimal size of the two sectors is independent of the strength of comparative advantage or the duration of unemployment. To be concrete, suppose that there are many workers with a strong comparative advantage at producing the dirty good or who would face a long unemployment spell before finding a job producing the clean good. Although this fact would not affect the optimal tax, it would imply that workers are unresponsive to the optimal tax. As a result, if workers are strongly attached to the dirty sector, then it is optimal for more workers to continue producing the dirty good. Nevertheless, a policy maker who can tax the production of the dirty good does not need to understand the strength of comparative advantage or the unemployment consequences of his policy in order to compute the optimal tax. In contrast, a policy maker choosing an optimal quantity restriction would need to know

this information.

I focus throughout this paper on a scenario in which there are two goods, one clean and one dirty, and no technology for abating the pollution generated by the dirty technology. Other scenarios are certainly empirically relevant. For example, it may be possible to reduce pollution by expanding the labor devoted to producing a third good, say an abatement technology. One could model unemployment of workers moving into this sector as well using similar tools to the ones I develop here. I would expect that optimal policy does not directly depend on how hard it is to train a worker to use the abatement technology, but instead could be expressed in terms of simple formulae that do not explicitly acknowledge the existence of unemployment.

The paper proceeds by first developing a static model in which workers are initially allocated to one sector of the economy and must decide whether to move to the other sector at the cost of foregoing some of their income. In Section 3 then develop a dynamic model in which workers continually move across sectors in response to idiosyncratic productivity shocks, experiencing unemployment when they move. In both models, the decentralized equilibrium would be Pareto optimal in the absence of any pollution. The externality creates a role for taxes and in both models I focus on developing simple expressions for the optimal tax rate. Section 4 discusses the robustness of my main findings and concludes.

2 Static Model

2.1 Environment

This section develops a simple static general equilibrium model with unemployment. I consider an economy which uses labor to produce two goods, a clean good and a dirty good. The economy is inhabited by a large number of individuals $i \in [0, 1]$. The assumption that there is a continuum of individuals formalizes the notion that each individual acts as if his own actions affect neither the level of pollution nor the prices in the economy.

Each individual supplies a unit of labor inelastically, consumes the two goods, and also cares about how much of the dirty good is produced. In particular, for a particular individual i , let c_i denote his consumption of the clean good, d_i denote his consumption of the dirty good, and D denote the total production (and hence consumption) of the dirty good, $D \equiv \int_0^1 d_i di$. I assume the individual's preferences are ordered by the utility function $V(u(c_i, d_i), D)$ where V is increasing in its first argument and weakly decreasing in its second. In addition, I assume that the subutility function u is positive-valued, increasing, concave, and has constant returns to scale.¹ This implies that the marginal rate of substitution between the two private goods, $u_c(c_i, d_i)/u_d(c_i, d_i)$, is a decreasing function of the ratio c_i/d_i . I also assume u satisfies Inada conditions, so this ratio approaches infinity when $c_i/d_i = 0$ and it approaches 0 at the opposite extreme. This ensures that both goods are always consumed in equilibrium.

¹The assumption of constant returns to scale is equivalent to assuming u is homothetic, since V is an arbitrary increasing function.

Each individual maximizes utility subject to a budget constraint $p_c c_i + p_d d_i = y_i$, where y_i is his income, discussed further below, and p_c and p_d are the price of the two goods in terms of an arbitrary numeraire. When choosing his consumption, the individual takes as given the price of the two goods as well as the total production of the dirty good D . It follows that he sets the marginal rate of substitution u_c/u_d equal to the ratio of prices $p_c/p_d \equiv q$. Equivalently, homotheticity of the utility function u implies that each individual chooses $d_i = c_i f(q)$, where f is an increasing function. The budget constraint then implies

$$p_c c_i = \frac{q y_i}{q + f(q)} \text{ and } p_d d_i = \frac{f(q) y_i}{q + f(q)}$$

Now let $C \equiv \int_0^1 c_i di$ denote total consumption and production of the clean good. Integrating the previous expressions across individuals implies that the ratio of total consumption of the dirty good to total consumption of the clean good is also an increasing function of the relative price q :

$$\frac{D}{C} = f(q). \tag{1}$$

For example, if $u(c, d) = c^a d^{1-a}$ for some $a \in (0, 1)$, $f(q) = (1 - a)q/a$ and so the expenditure share on good C , $p_c C / (p_c C + p_d D)$, is a constant a . This is the case where the marginal rate of substitution between the two goods is 1. If the goods are poorer substitutes, then an increase in the relative price of the clean goods raises the expenditure share on clean goods.

I next turn to the worker's income. This is the product of his wage, which depends on the sector he works in, and the amount of time he works, which

depends on whether he switches sectors. More precisely, the wage per unit of labor input for a worker producing the clean good is w_c and the wage per unit of labor input for a worker producing the dirty good is w_d . Now assume that prior to the single time period, a fraction n_c^0 of the workers were engaged in producing the clean good and the remaining $n_d^0 = 1 - n_c^0$ were engaged in producing the dirty good. If a worker continues to produce the same good, he can supply one unit of labor. If he switches industry, he spends a fraction $1 - \phi$ of the period unemployed before he finds a job, and so can only supply ϕ units of labor.² Each worker will stay in his original sector if moving reduces his income and move if it raises his income. Let n_c denote the number of workers who actually work in the clean sector and $n_d = 1 - n_c$ denote the number of workers who actually work in the dirty sector. Then assuming some workers produce in each sector, workers' mobility decisions imply

$$n_c \begin{matrix} \geq \\ < \end{matrix} n_c^0 \Rightarrow \begin{cases} w_d = \phi w_c \\ \phi^{-1} w_c \geq w_d \geq \phi w_c \\ w_d = \phi^{-1} w_c. \end{cases}$$

That is, if workers move between sectors, $n_c \neq n_c^0$, then it must be the case that movers are indifferent about doing so, earning the same labor income in either sector, while workers in the growing sector strictly prefer to stay in that sector.³ If wage gaps are not large enough to cover the cost of unemployment,

²I assume that the individual does not enjoy any additional leisure while he is unemployed, but it is straightforward to modify this assumption.

³This indifference condition holds in equilibrium and uses the fact that both goods are always consumed. Otherwise it would be possible to drive all workers into one sector while

then workers remain in their old sector.

A worker in sector $s \in \{c, d\}$ produces one unit of good s if he works full time and ϕ units of that good if he is moving into the sector and so spends some time unemployed. It follows that production of the two goods is

$$C = (1 - \phi) \min\{n_c, n_c^0\} + \phi n_c \text{ and } D = 1 - (1 - \phi) \max\{n_c, n_c^0\} - \phi n_c \quad (2)$$

For example, if workers do not switch sectors, $n_c = n_c^0$ and so $C = n_c^0$ and $D = 1 - n_c^0$. If workers move into the clean sector, $n_c > n_c^0$ and output of the clean good is boosted for a fraction ϕ of the period by the $n_c - n_c^0$ who move into the sector, while output of the dirty good is reduced by $n_c - n_c^0$ throughout the period.

I assume that the market for the two goods is competitive, but the government levies a tax at rate τ_c on the sale of the clean good and τ_d on the sale of the dirty good. As a result, the price of the two goods is equal to the after-tax cost of producing them, $p_c = (1 + \tau_c)w_c$ and $p_d = (1 + \tau_d)w_d$. The government runs a balanced budget, which requires $\tau_c w_c C + \tau_d w_d D = 0$, so one of the taxes is negative. These two tax instruments are complex enough to allow the government to obtain the first best allocation, but I consider other equivalent tax systems below.

To complete the characterization of equilibrium, I now look at two cases. In the first, there is no reallocation of workers across sectors, $n_c = n_c^0$, and all workers prefer to stay in the original sector rather than enduring an unemployment spell but possibly higher wages in the other sector. In the second, still keeping a sufficiently large wage gap to encourage workers to move.

some workers switch sectors and all workers in the shrinking sector are indifferent about moving. The equilibrium always takes one of these two forms. I show that the impact of tax changes depends on which configuration the equilibrium has. In the first case, a marginal change in taxes has distributional consequences but no impact on pollution. In the second, a marginal change in taxes has no distributional consequences and instead gives rise to a simple formula for the optimal pollution tax.

2.2 Equilibrium when there is no reallocation

I start my analysis with the case in which there is no reallocation in equilibrium, so $n_c = n_c^0$. In that case, consumption of the two goods is simply given by the initial allocation of labor, $C = n_c^0$ and $D = 1 - n_c^0$. Then equation (1) pins down the relative price of the two goods q as a function of the initial labor share n_c^0 :

$$f(q) = \frac{1 - n_c^0}{n_c^0}.$$

For this to be an equilibrium, it must be the case that workers do not want to move,

$$\frac{q}{\phi} \geq \frac{1 + \tau_c}{1 + \tau_d} \geq \phi q,$$

which may or may not hold. In particular, a large enough tax on the dirty good and subsidy to the clean good will always lead to a violation of the second inequality and induce some workers to move to the clean sector, the alternative configuration that I turn to next.

In the case with no equilibrium reallocation, a change in the taxes τ_c and

τ_d does not change the relative price of the two goods because it does not change the production of the two goods and relative prices must clear the goods market. In particular, such a tax also does nothing to abate pollution. It does, however, have distributional consequences. Combine the government budget constraint $\tau_c w_c C + \tau_d w_d D = 0$ with the expression for the relative wage $w_c/w_d = q(1 + \tau_d)/(1 + \tau_c)$ to get

$$\tau_c = -\frac{D/C\tau_d}{q + \tau_d(D/C + q)},$$

decreasing in the tax on dirty goods τ_d since q , C , and D are all independent of the tax rate. It follows then that the relative wage satisfies

$$\frac{w_c}{w_d} = q + \tau_d(q + D/C).$$

Since aggregate output is unchanged and workers producing the clean good are relatively wealthier, they are made better off by an increase in the tax on dirty goods. Conversely, workers producing the dirty good are made worse off. The bottom line is that when there is no reallocation in equilibrium, labor is supplied inelastically and so pollution taxes have distributional effects but do not affect pollution. An increase in the tax on the dirty good helps the workers producing the clean good at the expense of the workers producing the dirty good.

2.3 Equilibrium when there is reallocation

I turn next to the configuration in which some workers are moving between sectors. To be concrete, suppose that they are moving from the dirty sector to the clean sector, $n_c > n_c^0$. For these workers to be willing to move and others to be willing to stay in the dirty sector, it follows that $w_d = \phi w_c$. Since prices satisfy $p_c = (1 + \tau_c)w_c$ and $p_d = (1 + \tau_d)w_d$, the relative price $q = p_c/p_d = \frac{1+\tau_c}{(1+\tau_d)\phi}$. Then equation (1) pins down the ratio of the production of dirty and clean goods, $D = Cf(q)$, while the production function (2) pins down D and C as functions of n_c . Solving for n_c gives

$$n_c = \frac{1 - (1 - \phi)n_c^0 f\left(\frac{1+\tau_c}{(1+\tau_d)\phi}\right)}{1 + \phi f\left(\frac{1+\tau_c}{(1+\tau_d)\phi}\right)}$$

For this to be an equilibrium, it is necessary and sufficient that $n_c > n_c^0$ or equivalently $\frac{C}{C+D} > n_c^0$. It is straightforward to verify that this is true if and only if the tax on dirty goods is too large for the first type of equilibrium to obtain.

In an equilibrium with mobility, an increase in the tax on dirty goods and commensurate decrease in the tax on clean goods lowers the relative price of the clean good, $q = \frac{1+\tau_c}{(1+\tau_d)\phi}$, thereby reducing the demand for the dirty good and inducing more workers to migrate out of the industry. To understand the welfare consequences of this, it is useful to first think about a case in which there is no pollution externality, $V(u, D) = u$. Take a typical individual i with income y_i , either w_c or $w_d = \phi w_c$. Since average income is $Y = w_c n_c^0 + w_d(1 - n_c^0)$ and preferences are homothetic, it is easy to verify that she consumes

$c_i = (y_i/Y)((1 - \phi)n_c^0 + \phi n_c)$ and $d_i = (y_i/Y)(1 - n_c)$, i.e. her share of the production of the two goods.

Now by varying the taxes τ_c and τ_d , the government can change n_c and hence the equilibrium level of consumption of the two goods; however, as long as there is some mobility, it cannot change anyone's relative income y_i/Y . This is $\frac{1}{n_c^0 + \phi(1 - n_c^0)}$ for the n_c^0 individuals who start in the clean sector and $\frac{\phi}{n_c^0 + \phi(1 - n_c^0)}$ for the $1 - n_c^0$ individuals who start in the dirty sector. Therefore, individual i would like the government to set taxes so that n_c maximizes $u(c_i, d_i) = (y_i/Y)u((1 - \phi)n_c^0 + \phi n_c, 1 - n_c)$. The solution to this problem sets the marginal rate of substitution equal to the relative productivity of the marginal worker in the two sectors, $u_c(c, d)/u_d(c, d) = 1/\phi$, or equivalently $d/c = f(1/\phi)$. To achieve this objective, each individual prefers that the government not levy a distortionary tax, that is it should set $\tau_c = \tau_d = 0$. This result is not particularly surprising. Absent any externality, there is no role for distortionary taxes.

I now reintroduce the assumption that the dirty good causes a negative externality, $V_D(u, D) < 0$. This creates an obvious role for a Pigouvian tax to reduce the production of the dirty good. Yet whether all workers will agree that a such a tax is beneficial is unclear since differences in wealth may induce individuals to value the negative externality differently. That is, the marginal rate of substitution between their private consumption of dirty goods and their external consumption of pollution,

$$\sigma^{d,D} \equiv \frac{\partial V(u(c, d), D)/\partial d}{\partial V(u(c, d), D)/\partial D},$$

may differ across individuals. While this is potentially a real issue, it is not central to this paper. To circumvent this problem, I make a particular assumption on preferences which ensures individuals have a common interest about taxes. I assume that preferences over private consumption and pollution take the form

$$V(u, D) = \Psi(u/v(D)) \tag{3}$$

for some increasing functions Ψ and v .⁴ In this case, the marginal rate of substitution between dirty goods and pollution for an individual consuming the average amount (C, D) is

$$\sigma^{d,D} = \frac{v'(D)/v(D)}{u_d(C, D)/u(C, D)}.$$

Under this restriction, every individual prefers the same tax rate.

To prove this, first note that each individual recognizes that his relative income is unaffected by small changes in the tax rate: an individual who is initially working in the dirty sector earns ϕ times as much as an individual who is initially working in the clean sector as long as there is some mobility from the dirty sector to the clean sector. As a result, an individual's income relative to average income, y_i/Y , is still either $\frac{\phi}{n_c^0 + \phi(1 - n_c^0)}$ or $\frac{1}{n_c^0 + \phi(1 - n_c^0)}$, depending on his initial sector. Now homothetic preferences imply that all individuals allocate the same share of their income to dirty consumption, $d_i/c_i = f(q)$. It follows that the ratio of dirty consumption d_i to pollution D for any individual i is

⁴I also assume that v satisfies appropriate conditions which ensure that the social optimum has an interior level of production of the dirty good. Convexity is sufficient but not necessary.

simply equal to their income relative to average income y_i/Y . That is, a pollution tax changes supply of clean and dirty goods without any distributional impact.

Now consider the optimal level of the tax. A marginal increase in the tax that induces one worker to move out of the dirty sector reduces production of the dirty good and of pollution by 1 and raises production of the clean good by ϕ . Thus an individual whose relative income is y_i/Y wants employment in the dirty sector to solve

$$\begin{aligned} \max_{n_c, c, d, D} \quad & \Psi(u(c, d)/v(D)) \\ \text{s.t.} \quad & c = \frac{y_i}{Y}((1 - \phi)n_c^0 + \phi n_c), \\ & d = \frac{y_i}{Y}(1 - n_c), \\ & \text{and } D = 1 - n_c. \end{aligned}$$

The first constraint recognizes that his consumption of the clean good is a fraction y_i/Y of aggregate production of that good, the second constraint recognizes the same property for the dirty good, and the third constraint equates production of the dirty good to pollution. Eliminating c , d , and D using the constraints and the homogeneity of u , I get that an individual with relative income y_i/Y chooses n_c to solve

$$\max_{n_c} \Psi \left(\frac{y_i}{Y} u((1 - \phi)n_c^0 + \phi n_c, 1 - n_c) / v(1 - n_c) \right).$$

The choice of n_c is obviously independent of y_i/Y , as I asserted earlier. Opti-

mal production of the two goods then satisfies the first order condition

$$\frac{v'(D)}{v(D)} = \frac{u_d(C, D) - \phi u_c(C, D)}{u(C, D)}.$$

Equivalently, since the marginal rate of substitution between clean and dirty goods, u_c/u_d , is equal to the relative price q , which in turn equals $\frac{1+\tau_c}{(1+\tau_d)\phi}$, I get

$$\frac{\tau_d - \tau_c}{1 + \tau_d} = \sigma^{d,D}. \quad (4)$$

That is, the optimal pollution tax is simply a function of the marginal rate of substitution between dirty goods and pollution for a hypothetical individual with the average level of income. If such an individual would be unwilling to give up any of his dirty goods in return for a reduction in pollution, then the optimal tax is zero. As he becomes more willing to make this substitution, the optimal tax wedge is larger. This is the key result from the static model.

A curious aspect of the optimal tax formula (4) is that unemployment does not directly appear in it. That is, the optimal tax wedge $\frac{\tau_d - \tau_c}{1 + \tau_d}$ is independent of the amount of time it takes workers to switch from the dirty industry to the clean one, $1 - \phi$, and how many workers need to switch industries, $n_c - n_c^0$. Indeed, the tax wedge would be unchanged if there were no unemployment in the model, $\phi = 1$. Intuitively, it is simply necessary to tax the dirty good by enough to equate the private cost of purchasing the good to the social cost of consuming it.

This analysis is a bit misleading for two reasons. First, for general preferences the marginal rate of substitution between dirty goods and pollution,

$\sigma^{d,D}$, is not a constant but depends on the production of both the clean and dirty goods, and hence on the ease of mobility ϕ . Put differently, to know the optimal level of pollution, it is necessary to understand the tradeoff between the dirty good and pollution not only at the current level of production but at the purported optimal level. While in practice it may be difficult to learn this key parameter, this issue is not made any more difficult by the presence of unemployment.

The second reason the analysis is misleading is that the government budget constraint links the two tax rates. To see this, combine the government budget constraint $\tau_c w_c C = \tau_d w_d D$ with the wage ratio $w_d = \phi w_c$ to get $\tau_c = -\frac{\tau_d D}{\phi C}$. Then using $q = \frac{1+\tau_c}{(1+\tau_d)\phi}$ and the optimal tax formula (4), I get $q = \frac{1-\sigma^{d,D}}{\phi}$ at the optimum. Finally, the consumer problem implies $D/C = f(q)$ and so

$$\frac{D}{C} = f\left(\frac{1-\sigma^{d,D}}{\phi}\right). \quad (5)$$

Assuming that $\sigma^{d,D}$ is constant, this is decreasing in the fraction of time that a worker who switches sectors is employed ϕ . It follows that if ϕ is small, it is optimal to allow more production of the dirty good. Despite this, the initial condition n_c^0 still does not enter into the optimal tax calculation. This is because, once workers are moving across sectors, the marginal cost of reallocation is constant.⁵

The last few paragraphs may appear to be contradictory, so it is worth emphasizing why they are not. Equation (4) states that the optimal tax wedge

⁵The entire analysis in this subsection is of course predicated on the assumption that $n_c > n_c^0$. If equation (5) implies $D/C > (1 - n_c^0)/n_c^0$, then there is no reallocation.

depends only the marginal of substitution between dirty goods and pollution. To compute it, it is not necessary to know the severity of the consequent unemployment problem. Equation (5) states that the desired ratio of dirty to clean consumption depends on how severe the unemployment problem is. The reconciliation is simple. A given tax schedule (τ_c, τ_d) will induce more workers to reallocate, and hence a smaller ratio D/C , if ϕ is larger, so mobility is less costly.

There is a formal equivalence between tax and quantity regulation in this environment. Nevertheless, if the government does not understand the unemployment consequences of sectoral reallocation, taxes offer a clear advantage. The optimal tax formula depends only on preferences, while the optimal quantity restriction requires understanding both preferences and technology.

Equation (4) pins down the optimal tax wedge between consumption of the dirty and clean goods. I then pinned down the level of the two taxes with the government budget constraint. It is worth noting that the government can accomplish the same objective with other tax instruments. For example, suppose the government taxes the consumption of the dirty good at $\tau_d = \sigma^{d,D}/(1-\sigma^{d,D})$ and rebates the proceeds lump-sum to households. It is easy to verify that the equilibrium allocation is unchanged. Alternatively, the government can use the proceeds to compensate the workers who were initially employed in the dirty sector, lessening the redistributive consequences of the optimal policy.

The bottom line is that when workers are moving from the dirty sector to the clean sector, an increase in the tax on the dirty good induces more workers to move and raises their consumption. Under particular assumptions,

it is possible to abstract from the distributional consequences of this policy and focus on the tax rate that all workers find optimal. Curiously, the formula for the optimal tax can be expressed in terms of preferences, without reference to how much unemployment the optimal pollution tax causes. Nevertheless, if it is harder for workers to reallocate across sectors, a given tax induces fewer workers to reallocate and so it optimal to allow for more pollution.

2.4 Discussion

The simple model is useful for illustrating some principles and organizing thoughts but is too stylized to be taken seriously. One assumption that seems particularly problematic is that all workers in a given sector either strictly prefer to stay in their sector or are indifferent about moving out of the sector. This gave rise to two distinct cases, one with no reallocation where taxes were purely redistributive (Section 2.2) and one with reallocation where taxes abated pollution but did not affect the wealth distribution (Section 2.3). Moreover, the model predicts that if some workers are moving from the dirty sector to the clean sector, there are no workers moving in the opposite direction.

In reality, there are always workers moving in both directions between any two sectors of the economy. Moreover, it seems likely that even if some workers find it optimal to exit a sector, there are other workers who would find leaving to be very painful. An important next step is therefore to develop a model that has the features of both cases. One way to do this is to introduce idiosyncratic shocks that affect the costs and benefits of switching sectors for each worker. Depending on how many workers are near the margin of

indifference, such a model will give rise to results that look more like one or the other of the two cases I have analyzed so far. In particular, if the distribution of idiosyncratic shocks is not too disperse and most workers are initially quite happy to stay in their sector, a small increase in pollution taxes will primarily redistribute wealth from workers in the dirty sector to those in the clean sector, while a larger increase will reduce pollution with little additional distributional consequences. With a more disperse shock distribution, any change in taxes will have both effects, hurting workers who are far from the indifference margin while also inducing workers who are at the margin to pay the cost of moving to the clean sector.

A model with worker heterogeneity will also predict that the initial distribution of employment will matter for the optimal tax because the marginal cost of reallocating workers across sectors will naturally be increasing in the amount of reallocation. That is, the first few workers to exit a dirty industry might have been on the margin of exiting in any case and so will find the cost of exiting to be relatively small. But a larger tax will induce a larger contraction in demand for the dirty good and hence in employment, which will cause more workers to exit. The cost for these inframarginal workers will naturally be larger. A correct tax formula will therefore have to account for the initial distribution of employment as well.

A second weakness of this simple model is that it lacks any real dynamics. It takes time for a displaced worker to find a new job. As some workers leave a sector, wages increase and the remaining workers find it more attractive to stay. It is conceptually straightforward to extend the model to allow for

switching sectors to take a real amount of time and thus to explore the dynamic employment consequences of environmental regulation. In particular, as some workers exit an industry, the remaining workers will be selected to be those who are more attached to the industry, making further reductions in employment more costly. To understand the importance of these forces, I turn to a dynamic model with idiosyncratic shocks.

3 Dynamic Model

3.1 Environment

I consider a discrete time environment. Let $t = 0, 1, \dots$ denotes the time period. The economy is again inhabited by a unit measure of individuals $i \in [0, 1]$. Each individual is infinitely-lived, discounts the future with factor $\beta \in (0, 1)$, and has preferences over consumption of the clean good, $c_{i,t}$, consumption of the dirty good, $d_{i,t}$, and pollution D_t given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_{i,t}, d_{i,t}, D_t).$$

I assume u is increasing in its first two arguments, decreasing in its third argument, and strictly concave. Each individual is uncertain about his own future private consumption $(c_{i,t}, d_{i,t})$ because he does not know whether he will be employed and how much he will earn in the future. The expectations operator \mathbb{E}_0 captures this by assuming individuals seek to maximize their expected lifetime utility.

An individual's productivity depends on where he works and evolves stochastically over time. At any point in time t , assume that individual i can potentially produce either $\ell_{c,i,t}$ units of the clean good or $\ell_{d,i,t}$ of the dirty good. There are two reasons a worker cannot achieve this productivity. First, individual i can only work in one sector. Second, if an individual attempts to switch sectors, there is a chance he will fail and instead be unemployed.

More precisely, if a worker was last employed in sector $s \in \{c, d\}$, he is free to work in that sector in the current period. In that case, his labor income is $w_{s,t}\ell_{s,i,t}$, the product of his wage and his productivity. If he instead attempts to switch to sector s' , he obtains a job with probability ϕ , in which case he earns $w_{s',t}\ell_{s',i,t}$. Otherwise he is unemployed during period t .

Thus at the start of a period t , worker i observes $(\ell_{c,i,t}, \ell_{d,i,t})$ and the sector where he was last employed, $s_{i,t-1} \in \{c, d\}$. He then decides whether to work in sector $s_{i,t-1}$. If he does, $s_{i,t} = s_{i,t-1}$ and he earns $w_{s_{i,t},t}\ell_{s_{i,t},i,t}$. If he attempts to switch to sector s' , he succeeds with probability ϕ , in which case $s_{i,t} = s'$ and he earns $w_{s_{i,t},t}\ell_{s_{i,t},i,t}$. Otherwise he fails, is unemployed, earns nothing, and remains attached to his old sector, $s_{i,t} = s$. Finally, let $n_{s,i,t} = 1$ if a worker i succeeds in working in sector s in period t .

Potential productivity follows a first order Markov process conditional on current employment. More precisely, denote current potential productivity by $\ell \equiv (\ell_c, \ell_d)$ and current employment status by $s \in \{c, d, \emptyset\}$, where $s = c$ represents a worker employed in the clean sector, $s = d$ represents a worker employed in the dirty sector, and $s = \emptyset$ denotes an unemployed worker. Then potential productivity next period takes value $\ell' = (\ell'_c, \ell'_d)$ with probability

$\pi(\ell'|\ell, s)$. This notation is quite general. It recognizes that productivity may be persistent, that employment may enhance productivity, and that the enhancement may be sector specific, so a worker employed in sector s becomes more productive only in sector s . I assume the realization of the idiosyncratic productivity shock is independent across individuals and over time. This means that there are no aggregate shocks in the model.

Each worker employed in sector s produces one unit of good s per unit of potential productivity. Thus the aggregate output of the clean and dirty goods are

$$C_t = \int_0^1 \ell_{c,i,t} n_{c,i,t} di \text{ and } D_t = \int_0^1 \ell_{d,i,t} n_{d,i,t} di. \quad (6)$$

As in the static model, the two goods are sold in competitive markets and so the wage per unit of productivity in the two sectors is related to the output prices via

$$p_{c,t} = (1 + \tau_{c,t})w_{c,t} \text{ and } p_{d,t} = (1 + \tau_{d,t})w_{d,t}. \quad (7)$$

The government rebates the proceeds from any tax receipts lump-sum, with T_t denoting the lump-sum transfer.

I turn next to a description of financial markets. Individuals are risk-averse and face an uncertain income stream due both to unemployment risk and human capital risk. For reasons that I discuss below, I abstract from this idiosyncratic uncertainty by assuming financial markets are complete. Formally, each individual belongs to a household that seeks to maximize the average member's utility. The household observes each individual's potential productivity and tells him where to seek work in each period. It then pools the income

and uses the proceeds to provide each member with a common level of consumption of the clean good, c_t , and dirty good d_t . Critically, I assume that although the household is large enough to pool risk, it is still small relative to the size of the economy and so treats the aggregate production of the dirty good, D_t , as fixed.

It follows that the household seeks to maximize

$$\sum_{t=0}^{\infty} \beta^t u(c_t, d_t, D_t),$$

subject to the budget constraint

$$\sum_{t=0}^{\infty} (p_{c,t}c_t + p_{d,t}d_t) = \sum_{t=0}^{\infty} \left(w_{c,t} \int_0^1 \ell_{c,i,t} n_{c,i,t} di + w_{d,t} \int_0^1 \ell_{d,i,t} n_{d,i,t} di + T_t \right), \quad (8)$$

where the integrals reflect the fact that the household has many members, with individual i producing $\ell_{j,i,t}$ units of good j if he works in that sector. In addition, the household faces the laws of motion for the potential productivity of each member i , $\pi(\ell'|\ell, s)$, conditional on her search strategy s . Note that I have dropped the expectations operator, since the large family does not face any uncertainty.

The assumption that large households insure individuals against all idiosyncratic risk is extreme. I make it for two reasons. First, without complete markets, tax policy would have a redistributive effect, as in the static model without sectoral reallocation in Section 2.2. Workers with a comparative advantage in producing the clean good would like to tax the dirty good both because of the pollution externality and because it moves relative prices in

favor of the good they produce. Workers with a comparative advantage in producing the dirty good conversely may prefer to subsidize production of that good despite the pollution externality. In an environment with complete markets, I can abstract from this tension and focus on the average worker's preferences.

Second, without the large household assumption or some other assumption that ensures markets are complete, characterizing individual behavior is much more complicated.⁶ Individuals will wish to save when their productivity is temporarily high and borrow when it is low or when they are unemployed. While this consumption-savings problem is interesting, I view it as detracting from the main message of the paper. For example, incomplete markets would potentially be relevant even in an environment without unemployment ($\phi = 1$), if individuals cannot insure themselves against idiosyncratic fluctuations in their productivity.

Finally, I assume that the government must run a balanced budget in each period,

$$\tau_{c,t}w_{c,t}C_t + \tau_{d,t}w_{d,t}D_t = T_t.$$

Together with the goods market clearing conditions in equation (6) and the relationship between prices and labor costs in equation (7), this implies

$$p_{c,t}C_t + p_{d,t}D_t = w_{c,t} \int_0^1 \ell_{c,i,t} n_{c,i,t} di + w_{d,t} \int_0^1 \ell_{d,i,t} n_{d,i,t} di + T_t,$$

⁶An exception is the case when $u(\lambda c, \lambda d, D) = \lambda u(c, d, D)$. In this case, individuals have an infinite intertemporal elasticity of substitution in consumption and the consumption-savings problem is trivial to solve. My general formulation allows for a finite intertemporal elasticity of substitution but at the cost of having to assume markets are complete.

so the household budget constraint (8) must hold in every period.

In equilibrium, households optimally choose where each individual should attempt to work and how much of each good to consume, taking as given taxes, prices, and wages. In addition, the government and private budget constraints hold at each date. The equilibrium depends on the initial conditions, including the joint distribution of potential productivity and initial sector, $(\ell_{c,i,0}, \ell_{d,i,0}, s_{i,-1})$. It also depends on the tax and transfer policy $(\tau_{c,t}, \tau_{d,t}, T_t)$.

I proceed in three steps. First, I describe an alternative formulation of the household's problem. Second, I use that problem to describe individuals' decisions about where to work. Finally, I compare these outcomes to that of a hypothetical social planner who wishes to maximize the utility of the representative household and use this to compute the optimal tax policy.

3.2 Alternative Formulation

A typical household produces $\int_0^1 \ell_{c,i,t} n_{c,i,t} di$ units of the clean good and $\int_0^1 \ell_{d,i,t} n_{d,i,t} di$ units of the dirty good in period. Because of taxes, it can only afford to consume $\int_0^1 \ell_{c,i,t} n_{c,i,t} di / (1 + \tau_{c,t})$ units of the clean good and $\int_0^1 \ell_{d,i,t} n_{d,i,t} di / (1 + \tau_{d,t})$ units of the dirty good. The government collects the remaining output and rebates it lump-sum to households.

Of course, in equilibrium all households behave the same and so their consumption is equal to their production. Nevertheless, the tax distortion affects their behavior because each household imagines that it can consume a different amount than it produces. That is, in equilibrium each household

attempts to maximize

$$\sum_{t=0}^{\infty} \beta^t u(c_t, d_t, D_t),$$

subject to the production constraints

$$c_t = \frac{\int_0^1 \ell_{c,i,t} n_{c,i,t} di + C_t \tau_{c,t}}{1 + \tau_{c,t}}$$

and

$$d_t = \frac{\int_0^1 \ell_{d,i,t} n_{d,i,t} di + D_t \tau_{d,t}}{1 + \tau_{d,t}}.$$

The households treat C_t and D_t as fixed. They understand that the evolution of potential productivity ℓ depends on how a worker is allocated between the two sectors. They also recognize that mobility frictions limit the possibility of reallocation. They then choose how to allocate their workers in order to maximize utility. In equilibrium, those choices imply $c_t = C_t$ and $d_t = D_t$, so all households behave identically.

This formulation simplifies the household's problem by avoiding the need to discuss the determination of wages and prices. Of course, it is possible to find the wages and prices that decentralize the equilibrium; however, since I am mainly concerned with equilibrium allocations, it is not necessary to compute these prices.

3.3 Mobility

I focus instead on the key mobility decision. A typical individual starts period t with levels of potential productivity $\ell = (\ell_c, \ell_d)$ and a connection to sector s .

Denote his expected lifetime contribution to the household's utility by $J_t(\ell, s)$. To analyze mobility, it is easiest to express this recursively. I start with the value to the household of a worker who previously worked in the clean sector:

$$J_t(\ell, c) = \max \left\{ \begin{aligned} & \frac{\ell_c u_c(C_t, D_t, D_t)}{1 + \tau_{c,t}} + \beta \sum_{\ell'} \pi(\ell' | \ell, c) J_{t+1}(\ell', c), \\ & \phi \left(\frac{\ell_d u_d(C_t, D_t, D_t)}{1 + \tau_{d,t}} + \beta \sum_{\ell'} \pi(\ell' | \ell, d) J_{t+1}(\ell', d) \right) \\ & + (1 - \phi) \beta \sum_{\ell'} \pi(\ell' | \ell, \emptyset) J_{t+1}(\ell', c) \end{aligned} \right\}. \quad (9)$$

The first term in the maximization shows her contribution to household utility if she remains in the clean sector. She produces ℓ_c units of the clean good. The government takes a fraction $\tau_{c,t}/(1 + \tau_{c,t})$ of that, while the remainder increments household utility by the marginal utility of consumption of the clean good, $u_c(C_t, D_t, D_t)$. Her continuation value depends on the new productivity draw, which in turn depends on the fact that she worked in the clean sector.

The second term in the maximization shows her expected contribution to utility if she attempts to switch sectors. She succeeds with probability ϕ , in which case she produce ℓ_d units of the dirty good. The household keeps a fraction $1/(1 + \tau_{d,t})$ of that, incrementing marginal utility by $u_d(C_t, D_t, D_t)$ per unit of consumption. Finally, the worker has a continuation value that reflects the fact that she has switched sectors. If the worker fails to find a job, she is unemployed, contributes nothing to current utility, and has a continuation value that reflects her status as an unemployed worker with attachment to the clean sector.

The value of a worker who previously worked in the dirty sector is symmetric:

$$J_t(\ell, d) = \max \left\{ \frac{\ell_d u_d(C_t, D_t, D_t)}{1 + \tau_{d,t}} + \beta \sum_{\ell'} \pi(\ell' | \ell, d) J_{t+1}(\ell', d), \right. \\ \left. \phi \left(\frac{\ell_c u_c(C_t, D_t, D_t)}{1 + \tau_{c,t}} + \beta \sum_{\ell'} \pi(\ell' | \ell, c) J_{t+1}(\ell', c) \right) \right. \\ \left. + (1 - \phi) \beta \sum_{\ell'} \pi(\ell' | \ell, \emptyset) J_{t+1}(\ell', d) \right\}. \quad (10)$$

The interpretation of each term is unchanged.

Given a path of taxes and transfers (or aggregate production), it is conceptually straightforward to solve for the individual's mobility decision by iterating the value function. Given mobility decisions, it is possible to compute the supply of the two goods using equation (6) and hence verify that this is consistent with the conjectured level of production. If so, we have an equilibrium. Otherwise it is necessary to update the guess about aggregate production.

3.4 Optimal Allocation

Rather than solve the household's problem directly, I seek to find the tax rate that maximizes the utility of a hypothetical individual who consumes the average amount of output produced,

$$\sum_{t=0}^{\infty} \beta^t u(C_t, D_t, D_t).$$

I answer this by examining a social planner's problem, where the social planner recognizes that production of the two goods satisfies

$$C_t = \int_0^1 \ell_{c,i,t} n_{c,i,t} di \text{ and } D_t = \int_0^1 \ell_{d,i,t} n_{d,i,t} di$$

The social planner internalizes the impact of production of the dirty good. He also understands that the evolution of potential productivity ℓ depends on how workers are allocated between the two sectors in the current period, while mobility frictions limit the possibility of reallocation.

At the start of each period, every worker is described by her potential productivity $\ell = (\ell_c, \ell_d)$ and by the sector she last worked in $s \in \{c, d\}$. The worker can then either produce ℓ_s units of good s or attempt to produce $\ell_{s'}$ units of the other good $s' \neq s$, succeeding with probability ϕ . Let $H_t(\ell, s)$ denote the marginal value to the average individual's utility of a worker (ℓ, s) at the start of period t . For a worker with past employment in the clean sector, this solves

$$\begin{aligned} H_t(\ell, c) = \max & \left\{ u_c(C_t, D_t, D_t) \ell_c + \beta \sum_{\ell'} \pi(\ell' | \ell, c) H_{t+1}(\ell', c), \right. \\ & + \phi \left((u_d(C_t, D_t, D_t) + u_D(C_t, D_t, D_t)) \ell_d + \beta \sum_{\ell'} \pi(\ell' | \ell, d) H_{t+1}(\ell', d) \right) \\ & \left. + (1 - \phi) \beta \sum_{\ell'} \pi(\ell' | \ell, \emptyset) H_{t+1}(\ell', c) \right\}. \quad (11) \end{aligned}$$

If the worker is assigned to the clean sector, she works for sure and produces ℓ_c units of the clean good, valued at the marginal utility of the clean good.

If the worker is assigned to the dirty sector, she only works with probability ϕ , producing ℓ_d units of the dirty good, valued at its marginal utility which incorporates both the consumption benefit u_d and the pollution cost u_D . Otherwise, the worker does not produce any output and remains attached to the clean sector. Likewise, for a worker with past employment in the dirty sector, the social value of the worker solves

$$H_t(\ell, d) = \max \left\{ \begin{aligned} & (u_d(C_t, D_t, D_t) + u_D(C_t, D_t, D_t))\ell_d + \beta \sum_{\ell'} \pi(\ell'|\ell, d)H_{t+1}(\ell', d), \\ & + \phi \left(u_c(C_t, D_t, D_t)\ell_c + \beta \sum_{\ell'} \pi(\ell'|\ell, c)H_{t+1}(\ell', c) \right) \\ & + (1 - \phi)\beta \sum_{\ell'} \pi(\ell'|\ell, \emptyset)H_{t+1}(\ell', d) \end{aligned} \right\}. \quad (12)$$

Optimal mobility decisions solve these two equations.

Rather than solve the planner's problem directly, I find the tax and transfer system that decentralizes the allocation chosen by the planner. Once again, let

$$\sigma_t^{d,D} \equiv \frac{-u_D(C_t, D_t, D_t)}{u_d(C_t, D_t, D_t)}$$

denote the marginal rate of substitution between the dirty good and pollution at time t . Set

$$\tau_{c,t} = 0 \text{ and } \tau_{d,t} = \frac{\sigma_t^{d,D}}{1 - \sigma_t^{d,D}}. \quad (13)$$

That is, a household consumes its production of the clean good but only consumes a fraction $1 - \sigma_t^{d,D}$ of its production of the dirty good. The rest is taxed and rebated lump-sum back to all households. Then it is easy to verify that

equations (9) and (10) from the decentralized equilibrium are equivalent to equations (11) and (12) from the social optimum. That is, if the time path of production of the two goods is socially optimal and taxes satisfy equation (13), then all individuals attempt to produce the good that the social planner would like them to produce.

Equation (13) implies

$$\frac{\tau_{d,t} - \tau_{c,t}}{1 + \tau_{d,t}} = \sigma_t^{d,D}, \quad (14)$$

equivalent to the static optimal tax formula. In the static economy, it was unimportant how the pollution tax revenue was rebated to households. I showed it could be done either by subsidizing the consumption of the clean good or by providing a lump-sum transfer. This was because there were only two goods and hence one relative price that needed to be corrected through taxes. In the dynamic model, there are two goods at each date. It is necessary to get the right relative price of the two goods at each date and the right relative price of the clean good at two different dates. The former is accomplished by any tax satisfying the static equation (14). The latter requires that the tax on the clean good is constant at different dates, which I accomplish by setting it equal to zero. Any time-varying tax on the consumption of the clean good distorts intertemporal decisions about when to produce and when to switch sectors. It will therefore be inefficient. Since in general the revenue from taxing the dirty good is time-varying, it is necessary to have time-varying lump-sum transfer in order to ensure efficiency.

3.5 Wages and Prices

If so desired, we could also back out the equilibrium prices and wage rates. Since all households are identical, there is no room for trade, so this is simply a question of finding the prices and wages that ensure there is no desire to trade. One can verify that for households to be satisfied with their static consumption of the two goods, we require that

$$\frac{u_c(C_t, D_t, D_t)}{p_{c,t}} = \frac{u_d(C_t, D_t, D_t)}{p_{d,t}}.$$

This states that the marginal rate of substitution between the two goods is equal to the price ratio.

In addition, for the households to be satisfied with the timing of their consumption, we require that

$$\frac{u_c(C_t, D_t, D_t)}{p_{c,t}} = \beta \frac{u_c(C_{t+1}, D_{t+1}, D_{t+1})}{p_{c,t+1}}.$$

This states that the intertemporal marginal rate of substitution between the clean good in consecutive periods is again equal to the price ratio.

Finally, competition among firms ensures $p_{c,t} = (1 + \tau_{c,t})w_{c,t}$ and $p_{d,t} = (1 + \tau_{d,t})w_{d,t}$. Using these equations and fixing a numeraire, e.g. consumption of the clean good in period 0, it is straightforward to solve for the entire time path of wages and prices in any equilibrium.

4 Discussion and Conclusion

The models I have developed in this paper are deliberately stylized but capture an essential feature of many recent theories of unemployment: unemployment is a necessary consequence of the reallocation of workers across sectors of the economy. If the economy did not reallocate workers, it would be unable to take advantage of new technologies. Thus, while unemployment is a costly outcome for an individual worker, enduring some unemployment is still optimal for the economy as a whole. Indeed, the models in this paper share a common feature that, in the absence of pollution, the decentralized equilibrium without taxes would be socially optimal. This is a useful benchmark because it allows me to explore how policy optimally deals with a pollution problem alone.

Nevertheless, it is worth stressing that there are other recent theories of unemployment in which the equilibrium without pollution or taxes is still inefficient. Many of the papers build on the Mortensen and Pissarides (1994) model of unemployment and examine the role of wage rigidities (Shimer, 2005; Hall, 2005; Shimer, 2010). These papers focus on the behavior of wages over the business cycle, but wage rigidities can arise in the cross-section as well. In particular, an increase in the tax on dirty goods should reduce the wage for workers producing those goods and so should induce some workers to move out of the sector. Suppose that for some reason the wage does not fall and instead workers are rationed in their ability to supply labor to the dirty sector. Workers who fail to find a job are unemployed but be induced to stay attached to the dirty sector in the hope of eventually finding a job.

Although working out that model goes beyond the scope of this paper,

I would not expect the that key formula in this paper is stil applicable in this rigid wage environment. Instead, taxes would optimally address the labor market friction and perhaps rise less than would otherwise be expected. While this is potentially a real issue, it does not seem that environment policy is well suited to dealing with problems that originate in the labor market. That is, rigid wages would cause in an inefficient allocation of labor across sectors even absent any pollution problem. While there may be a role for the government to address this inefficiency, its connection with environmental regulation is tenuous.

References

- Hall, Robert E.**, “Employment fluctuations with equilibrium wage stickiness,” *American Economic Review*, 2005, *95* (1), 50–65.
- Lucas, Robert E. and Edward C. Prescott**, “Equilibrium search and unemployment,” *Journal of Economic Theory*, 1974, *7* (2), 188–209.
- Mortensen, Dale T. and Christopher A. Pissarides**, “Job creation and job destruction in the theory of unemployment,” *The Review of Economic Studies*, 1994, *61* (3), 397–415.
- Shimer, Robert**, “The cyclical behavior of equilibrium unemployment and vacancies,” *American Economic Review*, 2005, *95* (1), 25–49.
- , *Labor Markets and Business Cycles*, Princeton University Press, 2010.