

Asset Prices and Portfolios with Externalities*

Steven D. Baker[†] Burton Hollifield[‡] Emilio Osambela[§]

February 18, 2019

Abstract

Elementary portfolio theory implies that environmentalists optimally hold more shares of polluting firms than non-environmentalists, and that polluting firms attract more investment than otherwise identical non-polluting firms. These results reflect the demand to hedge against high pollution states. Pigouvian taxation can reverse the aggregate investment results, but environmentalists still overweight polluters. We introduce countervailing motives for environmentalists to underweight polluters, comparing the implications when environmentalists coordinate to internalize pollution, or have nonpecuniary disutility from holding polluter stock. With nonpecuniary disutility, introducing a green derivative product may dramatically alter who invests most in polluters, but has no impact on aggregate pollution.

*We thank Aytek Malkhozov, Edward Van Wesep, and participants at various seminars and conferences for comments.

[†]McIntire School of Commerce, University of Virginia, sdb7e@virginia.edu

[‡]Tepper School of Business, Carnegie Mellon University, burtonh@cmu.edu

[§]Board of Governors of the Federal Reserve System, emilio.osambela@frb.gov

1 Introduction

Humanity's choice as regards the threat of global warming is often cast in stark terms: either dramatically reduce emissions of greenhouse gasses, or expect a future with severe environmental damage, leading to extreme weather events, crop failures, and potentially catastrophic loss of biodiversity (Pachauri et al., 2014). Achieving large emissions reductions would likely require the abandonment of fossil fuel power generation from sources such as coal, crude oil, or their derivatives, with obvious negative consequences for the financial performance of related industries. Should such industries remain large and economically essential, then we should expect to suffer the environmental consequences.

A key obstacle to addressing global warming is a source of market failure recognized since at least Pigou (1920): greenhouse gas emissions create a negative externality. Producers and consumers of carbon intensive products typically make small individual contributions to the global stock of greenhouse gasses, and carbon dioxide (CO^2) emissions would rarely have immediate or localized negative impacts. Hence such emissions are a particularly difficult externality to address, e.g., through Pigouvian taxation, because of their global scope and cumulative, long term, negative effects.

In the absence of an effective, global carbon taxation regime that would attenuate the returns to CO^2 intensive industries, investors may be internalizing the societal costs of environmental damage. Chava (2014), for example, notes that \$1 in every \$8 invested is under socially responsible investment (SRI) restrictions, and finds that firms engaged in environmentally damaging activities have a higher cost of capital in both equity and debt markets. He further argues that such increased costs are not plausibly attributable to increased risk of future regulation alone.

A natural alternative explanation, as modeled by Heinkel et al. (2001), is that sufficient divestment of polluter's stock by green investors can raise the polluter's cost of capital, although Davies and Van Wesep (2017) argue that this may, perversely, reward the managers of polluting firms. The decision to advocate divestment, and forgo higher average

returns, likely reflects above average aversion to environmental damage. In the United States, Democrats are more typically associated with environmentalism than Republicans, and Hong and Kostovetsky (2012) find that fund managers who donate more heavily to Democrats than Republicans are more likely to manage SRI funds, or to pursue greener investment unofficially.

So far as we are aware, no existing papers rationalize green portfolios or SRI based on conventional portfolio choice objectives, e.g., maximizing expected returns subject to risk tolerance. In fact the motives for green investing are often left a bit of a mystery. Bénabou and Tirole (2010) emphasize image concerns as motivation for socially responsible behavior, although this does not apply when portfolio allocations are private. In Davies and Van Wesep (2017), some investors are simply more moral than others, and internalize the externality even though their individual investment decisions would make a negligible contribution to it. The implicit baseline is that, absent such moral compulsions, investors would simply ignore the externality in their investment decisions.

We show that this is not the case. In short, serious negative externalities of uncertain magnitude, such as global warming, represent undiversifiable risk. Assets that hedge such risk will trade at a premium. In the stark terms of our opening paragraph, stock in greenhouse gas emitting industries offers such a hedge. Environmentalists who suffer the greatest utility loss from, e.g., reduced biodiversity, have the strongest hedging motive, and should load up on polluters. If society in general has heightened sensitivity to environmental damage, then the cost of capital to polluters falls, leading to higher investment in polluting industries, which leads to greater environmental damage. The private hedging benefit creates the public cost of increased global warming.

We do not suggest that the hedging motive reflects how investors behave in practise, or even how they ought to behave. Rather our message is that elementary asset pricing and portfolio theory mechanisms should not be neglected in the analysis of firms that generate negative environmental externalities, even if moral incentives or behavioral biases are ulti-

mately found to play a dominant role. For example, consider the question of an optimal carbon tax. Bénabou and Tirole (2010) note that, if externalities are partially internalized without taxation, the theoretically optimal Pigouvian tax rate is too high. There has been relatively sophisticated analysis of optimal environmental taxation, e.g., Bovenberg and Goulder (1996), which takes account of industry input-output linkages and existing distortionary taxes. But households are assumed to have perfect foresight, so our simple hedging motive is absent, and asset pricing implications are not emphasized in the analysis. If green investors underweight polluters relative to the market, then the strength of their moral incentives should be estimated relative to an overweight rather than a neutral position, since they must first overcome the hedging motive.

To understand potential explanations for the popularity of SRI and the reluctance of environmentalists to hold polluter stocks, we nest two countervailing motives within our framework. In the first case, environmentalists internalize their collective contribution to pollution, and coordinate their investment strategy, optimally reducing their investment in polluters. This strategy is effective in reducing aggregate investment in dirty production, and may produce a positive spillover: if non-environmentalists in the population still care somewhat about pollution, it is optimal for them to also reduce investment in polluters, because divestment by environmentalists reduces everyone's hedging demand. Because environmentalists who internalize are motivated by the aggregate impact of their actions, they take stronger action when they are a larger share of the population. Therefore internalizing, as we define it, cannot explain unilateral divestment by environmentalists when they are a small share of the population.

Our second candidate explanation is that environmentalists suffer nonpecuniary disutility from investment in polluters. We show that this explanation is isomorphic to the case where environmentalists disagree with non-environmentalists about the expected returns on polluter stock. In this case environmentalists substantially reduce their individual investment in polluters even if their action is futile, as for example when they are a small share of

the population. Disutility from investment is similar to taxing environmentalists and non-environmentalists at different rates. As a consequence, aggregate investment in polluters decreases approximately linearly with the environmentalists' share of the population.

Nonpecuniary disutility from investment presumably stems from the identity of the issuers (i.e., polluters), not from the statistical characteristics of the investment. This presents an opportunity for financial innovation: through marketing, investments that are close statistical substitutes could offer different nonpecuniary utility. For example green bonds might offer environmentalists positive nonpecuniary utility, even if they are statistically similar to conventional bonds issued by non-polluting companies. Or it may be possible to “greenwash” investment in fundamentally dirty industries, repackaging such investment to avoid objections by environmentalists. Such subterfuge may not even be necessary: purchasing carbon emissions allowance futures could avoid disutility or even been seen as virtuous, although Aatola et al. (2013) suggests the resale price of such futures rises when coal is cheap and heavily utilized in power generation. Virtuous emissions allowance prices and sinful dirty investment are statistically intertwined.

One might expect the introduction of green bonds to increase aggregate investment in non-polluting companies, whereas greenwashing or substituting carbon emissions futures as an environmental hedge could indirectly increase aggregate investment in polluters. We model a green alternative to polluter stock, styled as a carbon emissions forward contract, as a pure financial innovation, leaving the underlying productive technologies unchanged. Environmentalists shift away from direct investment in polluters, preferring the forward contract instead, whereas non-environmentalists short the forward and increase their direct investment in polluters: essentially, they perform statistical arbitrage. Curiously, we prove that aggregate investment is unchanged by introducing the green alternative, despite dramatically altering direct investment in polluters by individual market participants.

Our work is connected to a large literature on SRI and corporate social responsibility (CSR), but adopts a different focus. Most work on SRI studies investment and asset pricing

effects via the corporate governance channel, taking demand for SRI funds as given. Dimson et al. (2015) study engagements with public companies by SRI activists to address negative environmental and social externalities, and find that successful engagements are followed by positive abnormal returns. By contrast Krüger (2015) finds that both negative and positive news connecting firms to externalities is followed by negative stock returns. However he isolates an important exception: when CSR is undertaken to correct previous bad behavior, the return is positive. This highlights the key theoretical question addressed by extant work on CSR: whether managers act to reduce negative externalities because it increases a firm's value, or because of agency problems, i.e., a manager's personal desire to reduce pollution at shareholder cost.

Davies and Van Wesep (2017) study how divestment campaigns interact with managerial compensation. Perversely, they show that managers may benefit financially from divestment campaigns, even though such campaigns depress short term stock prices. While they motivate divestment by moral investors as the result of internalizing negative externalities, their theoretical aim is the modeling of management incentives, and the reasons for divestment are not explicitly modeled: moral investors are simply constrained not to hold stock in immoral firms. In a model without agency frictions, Kelsey and Milne (2006) show that it is not optimal for firms to pursue a profit-maximizing strategy when there are externalities, and that shareholders will demand a reduction in the externality relative to the profit-maximizing case. To tractably analyze the firm's optimal decision, they abstract from risk or trade in financial assets, which are central features of our analysis. Albuquerque et al. (forthcoming) study the investment and asset pricing implications when firms may choose to pursue CSR for purposes of product market differentiation. They show that CSR leads to higher profit margins and higher firm valuation in equilibrium, but it is consumers, not investors, who have a taste for CSR products in their model.

We instead focus on investor attitudes towards externalities and the firms that produce them. We study the incentives – and disincentives – for moral or environmentally conscious

investors to favor SRI, taking the externalities produced by some firms as intrinsic to their production technologies. In contrast to Roth Tran (2018), who considers the portfolio choice problem of a philanthropy in partial equilibrium, we consider a heterogeneous general investor population, and endogenize prices, investment, and the amount of the externality in equilibrium.

2 Model

To convey the theoretical intuition in the simplest possible setting, we present a static model of portfolio choice with fixed capital. Analogous results from a model with flexible capital are summarized in Section 2.1. While greenhouse gas emissions provide a motivating example, the model applies to externalities generally.

There are two periods, $t \in \{0, 1\}$. Each of two firms, $i \in \{C, D\}$, issue one unit of stock at price P_i in period zero, and produce a period 1 dividend \tilde{Y}_i . The dividends are normally distributed, with identical mean μ and variance σ^2 , and correlation ρ . The distinguishing feature is that firm C uses a clean technology that produces no externality, whereas firm D uses a dirty technology that produces a negative externality \tilde{X} . We have in mind, for example, two electrical power generating companies, one of which uses wind, the other coal. We interpret the dividend of each firm as reflecting the quantity of electricity supplied, such that the negative externality is proportional to dividends of firm D : $\tilde{X} = \tilde{Y}_D$.¹ The government may levy a tax τ on dividends of dirty firm D , such that after tax dividends are $(1 - \tau)\tilde{Y}_D$. Tax revenues $\tilde{T} = \tau\tilde{Y}_D$ are redistributed equally amongst households.

There are two types of atomistic household, $j \in \{L, M\}$, with population mass $\eta \in (0, 1)$ of type M and mass $1 - \eta$ of type L . Households are distinguished by their sensitivity to environmental damage, or pollution. Type L suffers relatively less from pollution \tilde{X} , whereas type M suffers more from \tilde{X} . The difference could reflect propensity to enjoy

¹We assume one-to-one proportions without loss of generality. It is also possible to add idiosyncratic noise to \tilde{X} without altering results: the key assumption is positive correlation between polluter dividends and pollution.

outdoor activities, enjoyment of natural species diversity, or susceptibility to illness related to \tilde{X} , in addition to an abstract moral concern for the environment. Households derive utility from composite consumption

$$\tilde{W}_j - \lambda_j \tilde{X} + \tilde{T}, \quad (1)$$

where \tilde{W}_j is terminal wealth, $-\lambda_j \tilde{X}$ captures disutility from pollution, with $0 \leq \lambda_L < \lambda_M$, and \tilde{T} is government transfers. Here we assume that material and environmental consumption are perfect substitutes, i.e., enjoyment of a natural world relatively free from pollution does not require material wealth.

Households solve a portfolio optimization problem over terminal wealth, trading shares in each firm $\{\theta_{C,j}, \theta_{D,j}\}$, and a riskless asset. Without loss of generality, we assume that $r_f = 0$ and that households have zero initial wealth, so terminal wealth is

$$\tilde{W}_j = \theta_{C,j}(\tilde{Y}_C - P_C) + \theta_{D,j}((1 - \tau)\tilde{Y}_D - P_D). \quad (2)$$

Households of each type $j \in \{L, M\}$ maximize CARA utility over terminal consumption,

$$\begin{aligned} & \max_{\theta_{C,j}, \theta_{D,j}} \mathbb{E}[-\exp\{-\alpha(\tilde{W}_j - \lambda_j \tilde{X} + \tilde{T})\}] \\ & \equiv \max_{\theta_{C,j}, \theta_{D,j}} \mathbb{E}[-\exp\{-\alpha(\tilde{W}_j - \lambda_j \tilde{Y}_D + \tau \tilde{Y}_D)\}]. \end{aligned} \quad (3)$$

All households have identical coefficient of absolute risk aversion $\alpha > 0$, and differ only in their sensitivity to pollution λ_j . The second line in Equation (3) makes the following assumption explicit: households are aware that pollution is worse when the dirty technology produces more electricity, which is also when the dirty firm is able to pay a higher dividend \tilde{Y}_D . Households are also aware that government revenues and transfers are dependent upon taxation of the dirty firm. Atomistic households are powerless to alter the dynamics of either the externality or transfers, and simply take them as given. Individual households

can, however, choose portfolio allocations that hedge risk of environmental damage.

After solving Equation (3) from its first order conditions, the portfolio rule for each household type $j \in \{L, M\}$ is

$$\theta_{C,j} = \frac{\mu - P_C - \rho(\mu - P_D/(1 - \tau))}{\alpha\sigma^2(1 - \rho^2)}, \quad (4)$$

$$\theta_{D,j} = \frac{1}{1 - \tau} \left(\frac{\mu - P_D/(1 - \tau) - \rho(\mu - P_C)}{\alpha\sigma^2(1 - \rho^2)} - \tau + \lambda_j \right). \quad (5)$$

The allocation to the clean firm $\theta_{C,j}$ is the same across types, because \tilde{Y}_C is unrelated to the externality. Examining the difference between the two types' allocations to the dirty firm provides our first result.

Proposition 1 *Households of type M, who are hurt most by pollution, choose to hold more shares of polluting firm D than do households of type L, who are hurt less by pollution. The difference in shareholdings, $\theta_{D,M} - \theta_{D,L}$, increases as firm D is more heavily taxed.*

Proof. Recall our assumption $\lambda_L < \lambda_M$ distinguishing less and more environmentally sensitive households. From Equation (5),

$$\theta_{D,M} - \theta_{D,L} = \frac{\lambda_M - \lambda_L}{1 - \tau} > 0. \quad (6)$$

■

The intuition behind Proposition 1 is simple: households wish to hedge environmental risk, by substituting material consumption for lost environmental consumption. To do so, they buy stock in firm D , which pays off when pollution is high. Because type M suffers most from pollution, he has the strongest hedging motive, and so buys more of stock D than does type L . A tax increase scales down dividends, so households scale up their trade to maintain the same after-tax hedge.²

²With fixed capital, the tax does not reduce the externality.

Imposing market clearing allows us to solve for stock prices. Each firm issues one unit of stock per capita, and recalling that all households choose the same $\theta_{C,j}$,

$$(1 - \eta)\theta_{C,L} + \eta\theta_{C,M} = 1 \Rightarrow \theta_{C,L} = \theta_{C,M} = 1. \quad (7)$$

Substituting optimal portfolios into the market clearing conditions and solving for equilibrium stock prices yields

$$P_C = \mu - \alpha\sigma^2(1 + \rho - \rho(\eta\lambda_M + (1 - \eta)\lambda_L)), \quad (8)$$

$$P_D = (1 - \tau) (\mu - \alpha\sigma^2(1 + \rho - \eta\lambda_M - (1 - \eta)\lambda_L)). \quad (9)$$

The fact that stock in polluting firm D hedges environmental risk leads to the following counterintuitive result.

Proposition 2 *Suppose $\tau = 0$ and $\rho < 1$. Then $P_D > P_C$: stock in the polluting firm is worth more than stock in an otherwise identical firm that does not pollute.*

Proof.

$$P_D - P_C = \alpha\sigma^2(1 - \rho)(\eta\lambda_M + (1 - \eta)\lambda_L) > 0.$$

■

Naturally the above proposition continues to hold when τ and ρ are sufficiently small.

2.1 Frictionless capital adjustment

The previous section illustrated the main theoretical intuition in a simple and familiar portfolio choice framework with exogenous dividends. Analogous results hold when capital may be frictionlessly allocated to dirty or clean production technologies. Capital adjustment frictions are important for the effects of taxes in particular. With fixed capital, a tax on the polluting firm's dividends reduces the value of its stock but does nothing to reduce pollution.

With frictionless capital adjustment, Pigouvian taxation works to reduce equilibrium pollution, and the value of each firm is always identical to its capital stock, i.e., Tobin's $q = 1$. In reality capital is neither fixed nor completely liquid; the fact that our results hold at both extremes suggest the mechanism at work is general.

As before there are two competitive firms, C and D . Each firm has access to a risky technology that produces the same final consumption good. In period 0, households of type $j \in \{L, M\}$ choose investment $I_{C,j}$ and $I_{D,j}$, which they finance by borrowing at $r_f = 0$. Production is complete in period 1. Output from production is

$$\tilde{\mu}_i I_i, \quad i \in \{C, D\}, \quad (10)$$

where per capita investment in each firm is a weighted average reflecting the share of less (L) and more (M) environmentally sensitive households in the population,

$$I_i = (1 - \eta)I_{i,L} + \eta I_{i,M}, \quad (11)$$

and $\tilde{\mu}_i$ is stochastic productivity. Analogous to the previous section, we assume $\tilde{\mu}_D$ and $\tilde{\mu}_C$ are exogenous normal random variables with mean μ , standard deviation σ , and correlation ρ .

In the production setting it is more convenient to take investment $I_{i,j}$ as the choice variable for each household. But in order to compare results from the production setting to the endowment setting, think of each firm as paying a liquidation dividend

$$\tilde{Y}_i = \tilde{\mu}_i I_i. \quad (12)$$

Normalize the number of shares issued by firm i to 1. Because we assume constant returns to scale and no capital adjustment costs, the price of firm i stock is equal to investment in

firm i : $P_i = I_i$. The shareholdings of agent j in firm i are

$$\theta_{i,j} = \frac{I_{i,j}}{I_i}. \quad (13)$$

As before, the dirty firm's technology produces a negative externality,

$$\tilde{X} = \tilde{\mu}_D I_D. \quad (14)$$

One interpretation, if the consumption good is electricity, is that I_D determines how many coal power plants are built, whereas $\tilde{\mu}_D$ captures intensity of utilization, and hence the rate at which each plant consumes coal. The externality is proportional to the amount of dirty electricity generated.³

Dividends of the dirty firm are subject to a Pigouvian tax, so the after tax dividend is $(1 - \tau)\tilde{Y}_D$. In our setup it is not important whether the proportional tax is levied on gross sales of the final good, on investment, or on the dividend: the result is the same. Tax receipts are distributed equally among households.

Otherwise we assume the same setup as before, such that each agent solves

$$\max_{I_{C,j}, I_{D,j}} \mathbb{E}[-\exp\{-\alpha(I_{C,j}(\tilde{\mu}_C - 1) + I_{D,j}((1 - \tau)\tilde{\mu}_D - 1) - \lambda_j \tilde{X} + \tilde{T})\}], \quad (15)$$

to maximize expected utility. Solving from first order conditions for each agent $j \in \{L, M\}$,

$$I_{C,j} = \frac{\mu - 1 - \rho(\mu - 1/(1 - \tau))}{\alpha\sigma^2(1 - \rho^2)}, \quad (16)$$

$$I_{D,j} = \frac{1}{1 - \tau} \left(\frac{\mu - 1/(1 - \tau) - \rho(\mu - 1)}{\alpha\sigma^2(1 - \rho^2)} - (\tau - \lambda_j)I_D \right). \quad (17)$$

Each type of household allocates the same capital to the clean firm. Notice that $I_{D,j}$ depends on per-capita investment in the dirty firm I_D , which individual households take as

³As in the previous section, the constant of proportionality is not important, since this is captured later by the choices of λ_j .

given when solving their optimization problems, as well as on sensitivity to pollution λ_j . Employing equilibrium condition Equation (11),

$$I_{D,L} = \frac{(1 - \tau + \eta(\lambda_L - \lambda_M))(\mu - 1/(1 - \tau) - \rho(\mu - 1))}{(1 - (1 - \eta)\lambda_L - \eta\lambda_M)\alpha\sigma^2(1 - \rho^2)(1 - \tau)}, \quad (18)$$

$$I_{D,M} = \frac{(1 - \tau + (1 - \eta)(\lambda_M - \lambda_L))(\mu - 1/(1 - \tau) - \rho(\mu - 1))}{(1 - (1 - \eta)\lambda_L - \eta\lambda_M)\alpha\sigma^2(1 - \rho^2)(1 - \tau)}. \quad (19)$$

We can now state propositions analogous to those in the previous section. We assume parameter values such that equilibrium investment in both firms is strictly positive.

Proposition 3 *Households of type M, who are hurt most by pollution, invest more in polluting firm D than do households of type L, who are hurt less by pollution.*

Proof. From Equation (18) and Equation (19),

$$\frac{I_{D,L}}{I_{D,M}} = \frac{1 - \tau + \eta(\lambda_L - \lambda_M)}{1 - \tau - (1 - \eta)(\lambda_L - \lambda_M)} < 1,$$

since $\lambda_L - \lambda_M < 0$. Given parameters such that investment in firm D is strictly positive, the result follows. ■

Unlike the case with fixed capital, increasing the Pigouvian tax τ will decrease investment in the dirty firm. This will also cause the absolute difference $I_{D,M} - I_{D,L}$ to decrease, although the relative difference $I_{D,M}/I_{D,L}$ is increasing in τ . Equivalently, raising Pigouvian taxes will cause the environmentalist to hold an increasing fraction of the shares in the polluting firm, even though the value of those shares will decrease.

Combining Equations (11, 18, 19), we can write

$$I_C = \frac{\mu - 1 - \rho(\mu - 1/(1 - \tau))}{\alpha\sigma^2(1 - \rho^2)}, \quad (20)$$

$$I_D = \left(\frac{1}{1 - (1 - \eta)\lambda_L - \eta\lambda_M} \right) \left(\frac{\mu - 1/(1 - \tau) - \rho(\mu - 1)}{\alpha\sigma^2(1 - \rho^2)} \right). \quad (21)$$

This leads to a proposition for investment analogous to Proposition 2 for prices in the

previous section.

Proposition 4 *Suppose $\tau = 0$ and $\rho < 1$. Then $I_D > I_C$: investment in the polluting firm is greater than investment in an otherwise identical firm that does not pollute.*

Proof. For $\tau = 0$,

$$\frac{I_D}{I_C} = \frac{1}{1 - (1 - \eta)\lambda_L - \eta\lambda_M} > 1,$$

since $\eta\lambda_M > 0$, $(1 - \eta)\lambda_L \geq 0$, and we assume parameters such that investment is greater than zero in both firms. ■

In summary, our main results hold for either fixed or flexible capital allocation across firms.

3 Discussion and extensions

While the hedging motive in our theoretical models reflects basic portfolio theory, the implications are counterintuitive. They are also at odds with available empirical work. Chava (2014) shows that firms excluded by environmental screens, for example due to high emissions of pollutants or greenhouse gasses, have a higher cost of capital than comparable firms that meet the screens – the opposite of Proposition 2. Hong and Kacperczyk (2009) show similar results for “sin stocks,” such as tobacco and firearms firms, which might also be seen as producers of negative externalities in the form of second hand smoke or gun crime.

These papers also highlight the role of socially responsible investment (SRI) constraints in determining stock ownership: investment funds committed to SRI, analogous to investing exclusively in our clean firm, are marketed to those most sensitive to the externality – the opposite of Proposition 1. In the context of US political affiliation, Hong and Kostovetsky (2012) show that, relative to Republican fund managers, Democratic fund managers tilt their portfolios towards SRI, whether or not SRI is a declared objective of the fund. On the assumption that SRI stocks perform relatively better on average under Democratic

administrations – e.g., due to more stringent enforcement of environmental regulations – conventional hedging motives dictate the opposite approach. Republicans should hedge an undesirable election outcome by pursuing SRI; yet they do not.

In this section we model possible explanations: countervailing motives that could overwhelm the desire to hedge bad realizations of the externality. We consider two possibilities: that the more sensitive (M) agents internalize the public impact of their collective investment decisions, or that the more sensitive agents get nonpecuniary disutility from investing in the dirty firm. The key distinction between these explanations concerns coordinated versus competitive behavior. Do more sensitive agents coordinate with each other, e.g., reducing their investment in polluting firms to achieve some reduction in aggregate pollution? Or do they avoid such investment because it would cause them direct personal discomfort, regardless of whether their action is effective in reducing aggregate pollution or not? If the latter, then we can maintain our assumption of competitive, atomistic agents. In reality such explanations are often blurred together: the social norms that motivate Hong and Kacperczyk (2009) are not optimally coordinated decisions oriented towards a specific objective, but clearly they have aspects of group coordination.

Therefore we take this opportunity to cleanly differentiate the two explanations within our theoretical framework, and to highlight their different implications. We also wish to compare the impact of these countervailing motives with what government intervention would achieve via Pigouvian taxation. Since such taxation is only effective in the production setting, we use that setting as our baseline for the extensions. While the interaction of Pigouvian taxation with our new disinvestment motives could be interesting, we assume no Pigouvian taxation in the extensions ($\tau = 0$, hence $T = 0$), to simplify expressions. Table 2 summarizes the model characteristics considered in each section or in related appendices.

3.1 Internalizing the externality via coordinated investment

Agents of type M internalize their contribution to the externality by thinking as follows: “if everyone of my type behaves as I do, how will that affect the externality?” That is, agent M coordinates with others of his type. Since the externality is

$$\tilde{X} = \tilde{\mu}_D I_D = \tilde{\mu}_D [(1 - \eta)I_{D,L} + \eta I_{D,M}], \quad (22)$$

M can affect the externality through the choice of $I_{D,M}$. We assume he cannot coordinate with the type L , and takes the investment decision of type L agents as given. The type M investment problem is

$$\max_{I_{C,M}, I_{D,M}} \mathbb{E}[-\exp\{-\alpha(I_{D,M}((1 - \eta)\lambda_M)\tilde{\mu}_D - 1) + I_{C,M}(\tilde{\mu}_C - 1) - I_{D,L}(1 - \eta)\lambda_M\tilde{\mu}_D)\}]. \quad (23)$$

Note that the above problem incorporates knowledge of M 's population share, η . Some shorthand notation simplifies expressions for equilibrium investment, and also has a natural interpretation. Agent M effectively taxes the polluting firm's output at rate $\tau_M = \eta\lambda_M$, reflecting his sensitivity to the externality and his type's population share. The analogous expression for agent L is $\tau_L = (1 - \eta)\lambda_L$. We also define net-of-tax coefficients $T_L = 1 - \tau_L$ and $T_M = 1 - \tau_M$.

Optimal investment for type M conditional on the investment choice of type L is

$$I_{C,M} = \frac{\mu - 1 - \rho(\mu - \frac{1}{T_M})}{\alpha\sigma^2(1 - \rho^2)}, \quad (24)$$

$$I_{D,M} = \frac{1}{T_M} \left[\frac{(1 - \rho)\mu + \rho - \frac{1}{T_M}}{\alpha\sigma^2(1 - \rho^2)} + (1 - \eta)\lambda_M I_{D,L} \right]. \quad (25)$$

Since internalizing is like privately taxing investment in firm D at rate τ_M , it is not surprising that M 's investment strategy resembles that with Pigouvian taxation at rate $\tau = \tau_M$. In fact, M 's investment in the clean firm, $I_{C,M}$, is identical to the case with $\tau = \tau_M$: see

Equation (20).

The effect on $I_{D,M}$ is more interesting, because it depends recursively on agent L 's investment decision. Solving for equilibrium requires a further assumption: does L also internalize, coordinating with other type L agents, or not? Here we will focus on the case where L does not internalize, with the case where both agents internalize in Appendix A.

If L does not internalize, equilibrium investment in the dirty firm is

$$I_{D,L} = \frac{(1 + (\lambda_L - \lambda_M)\eta)(\mu - \rho(\mu - 1)) - T_M - \frac{\lambda_L\eta}{T_M}}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}, \quad (26)$$

$$I_{D,M} = \frac{(1 + (\lambda_M - \lambda_L)(1 - \eta))(\mu - \rho(\mu - 1)) - \frac{T_L}{T_M} - \lambda_M(1 - \eta)}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}, \quad (27)$$

$$I_D = \frac{\mu - \rho(\mu - 1) - \frac{1 - (1 - \eta)\tau_L}{T_M}}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}. \quad (28)$$

The above results differ from the case with Pigouvian taxation for several reasons. The private tax τ_M only applies to type M , since type L does not internalize.⁴ And while internalizing is conceptually similar to a tax, it generates no revenues to redistribute. Finally there is the equilibrium hedging effect: any reduction in I_D reduces the hedging motive, because it reduces the variance of \tilde{X} . Provided $\lambda_L > 0$, then $I_{D,L}$ will decrease somewhat in response to internalizing by type M , because the actions of type M are effective in reducing the mean and variance of \tilde{X} .

Comparison of expressions for equilibrium investment confirm the following intuitive results: when M internalizes, $I_{D,L}$, $I_{D,M}$, and I_D all decrease relative to the case without internalizing or taxation, but the decrease is smaller than what a Pigouvian tax of $\tau = \tau_M$ would achieve, provided $\eta < 1$. But is the wedge between each of these cases large or small, and how does it change with type M 's population share?

We present a numerical example to answer these questions, using parameter values in Table 1. In this example, $\lambda_L = 1\%$, whereas $\lambda_M = 3\%$. From an individual perspective,

⁴While τ_L and τ_M have a natural interpretation as agent-specific private tax rates, these terms arise in the expressions for equilibrium investment even when agents do not internalize.

investment in firm D offers the same expected return as firm C , and also hedges bad realizations of the externality. But aggregate investment in firm D has effectively 1% lower expected return from the standpoint of agent L , or 3% lower expected return from the standpoint of agent M , due to the negative externality.

In the base case where both agents behave competitively and without Pigouvian taxation, Figure 1 shows that type M invests more in firm D than type L , and aggregate investment I_D rises with M 's population share, in line with Proposition 3. This is the hedging motive at work. However the magnitude of the difference $I_{D,M} - I_{D,L}$ is small, and as a consequence I_D rises only slightly with η . Imposing a fixed Pigouvian tax of $\tau = \lambda_M = 3\%$ reduces $I_{D,L}$, $I_{D,M}$, and I_D substantially, by about 1/3.⁵ Otherwise the introduction of taxation changes little: M still invests relatively more than L , and the difference is still small.

When M internalizes, Figure 1 shows that $I_{D,M}$ and I_D become very sensitive to η . The two previously described reference cases bound investment when M internalizes. If $\eta = 0$, then $I_D = I_{D,L}$, and $I_{D,M}$ is the same as the case without Pigouvian taxation, because even as a group type M 's investment does not contribute to the externality. When $\eta = 1$, then $I_D = I_{D,M}$, and $I_{D,M}$ is the same as the case with Pigouvian taxation at rate λ_M , because the externality is entirely attributable to group M 's investment decision. Finally, internalizing produces a more rapid drop in I_D as η nears 1, due to the combined effects of M 's increasing population share, and the greater group incentive for type M to reduce investment per capita.

Figure 1 also hints at the practical importance of competing effects in explaining observed portfolio choices of socially conscious investors. First, demand for hedging the externality may be a second order consideration: in our example, $I_{D,M} - I_{D,L}$ is small unless M internalizes, despite three-times the sensitivity to the externality for agent M relative to agent L . Second, despite the small difference in hedging demand, internalizing cannot plausibly explain why socially conscious investors (type M) would invest substantially less than others

⁵Note that $\tau = \lambda_M$ is the largest tax that could be justified based on either agent's utility loss from the externality. So I_D in this case could be viewed as a lower bound for socially optimal investment in firm D .

(type L) unless socially conscious investors are a large portion of the investor population. When η is small, say 5% of the population, type M still invests about the same as type L . Only when collective action would be very effective – large η – does M overcome his greater hedging motive and disinvest from the dirty firm.

3.2 Nonpecuniary disutility from investment

It is possible that socially responsible investors coordinate to achieve a common objective, thus overcoming their individual hedging motives. An even simpler explanation is that such investors have a visceral distaste for investment in polluting companies that overcomes pecuniary motives for doing so. Similar behavior is documented by Morewedge et al. (2016), who investigate why sports fans bet on their own teams to win, rather than hedging the disutility from a loss by betting on the opposing team. The hypothesized source of disutility is a conflict with the bettor’s identity as a supporter of his team. Similarly, an investor who cares greatly for the environment would find investment in a polluting firm contrary to his identity as an environmentalist, even if such investment hedges dirty environmental outcomes that are beyond his control.

Suppose type M receives nonpecuniary disutility proportional to his personal investment in firm D . As in Section 3.1, our starting point is the investment economy without Pigouvian taxation, and we assume that type L ’s problem is unchanged from this baseline. Incorporating the new disutility, agents of type M solve

$$\max_{I_{C,M}, I_{D,M}} \mathbb{E}[-\exp\{-\alpha(I_{D,M}(\tilde{\mu}_D - 1) + I_{C,M}(\tilde{\mu}_C - 1) - \zeta_D I_{D,M} - \lambda_M \tilde{X})\}]. \quad (29)$$

We nest the new disutility within total utility, and make it linear in $I_{D,M}$, with slope $-\zeta_D < 0$. While many alternative approaches are justifiable, this parsimonious choice fits nicely within our framework, and allows for simple comparisons with results from previous sections. It highlights what we view as the key distinction relative to Sections 2.1 and 3.1: disutility is

now proportional to $I_{D,M}$, not I_D or $\eta I_{D,M}$. Hence nonpecuniary disutility from investment, while an unusual ingredient in a portfolio choice problem, is within the scope of a conventional competitive agent utility maximization problem: it is not proportional to aggregate choice variables, and it does not require coordination.

While we have in mind that parameter ζ_D scales a nonpecuniary penalty, in our static framework it is equivalent to say type M perceives expected net return

$$\mu - \zeta_D - 1 \tag{30}$$

on investment in firm D , whereas type L perceives expected net return

$$\mu - 1. \tag{31}$$

So it is as if the types have heterogeneous beliefs about expected stock returns.⁶ Hong and Kostovetsky (2012) discuss non-pecuniary incentives for fund managers to pursue SRI – the perk of increased approval from Democratic peers – as distinct from pecuniary explanations – such as different models of stock market behavior. They acknowledge that the two motives may have similar implications. Our simple example shows that the motives may be isomorphic.

Incorporating the new disutility, agent M chooses investment

$$I_{C,M} = \frac{\mu - 1 - \rho(\mu - \zeta_D - 1)}{\alpha\sigma^2(1 - \rho^2)}, \tag{32}$$

$$I_{D,M} = \frac{\mu - \zeta_D - 1 - \rho(\mu - 1)}{\alpha\sigma^2(1 - \rho^2)} + \lambda_M I_D. \tag{33}$$

The key difference, relative to Section 3.1, is that ζ_D is not scaled by η : whether or not his type contributes much to the externality, M feels badly about investing in firm D . Investment

⁶Note that this equivalence would not carry over to a dynamic setting, because the agents would agree about wealth dynamics with nonpecuniary disutility, whereas they would disagree about wealth dynamics with heterogeneous beliefs.

by agent L is unchanged from Equation (17) and Equation (16), conditional on a given level of aggregate investment I_D in firm D .

Investment in the dirty firm is

$$I_{D,L} = \frac{(1 + (\lambda_L - \lambda_M)\eta)(\mu - 1 - \rho(\mu - 1)) - \lambda_L\eta\zeta_D}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}, \quad (34)$$

$$I_{D,M} = \frac{(1 + (\lambda_M - \lambda_L)(1 - \eta))(\mu - 1 - \rho(\mu - 1)) - T_L\zeta_D}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}, \quad (35)$$

$$I_D = \frac{\mu - 1 - \rho(\mu - 1) - \eta\zeta_D}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}, \quad (36)$$

for types L and M , and on an average per-capita basis, respectively. Similar to the results when M internalizes, agent L decreases his investment $I_{D,L}$ somewhat, because the decrease in I_D reduces the variance of the externality \tilde{X} , which reduces hedging demand.

Continuing our numerical example introduced in the previous section, we use common parameters in Table 1, and choose disutility parameter $\zeta_D = \frac{\lambda_M}{1 - \lambda_M} = 3.09\%$.⁷ Figure 2 shows $I_{D,L}$, $I_{D,M}$, and I_D incorporating disutility. We again include the frictionless investment case without Pigouvian taxation ($\tau = 0$), and with Pigouvian taxation at rate $\tau = \lambda_M = 3\%$, to provide points of reference.

Similar to Figure 1, aggregate investment I_D decreases substantially as M 's population share η increases. But in Figure 2 the rate of decrease is approximately linear in η , because $I_{D,M}$ and $I_{D,L}$ are nearly constant over the domain of η . Once again, the hedging effect is second order: even though I_D is nearly 1/3 lower at $\eta = 1$ than at $\eta = 0$, such that the standard deviation of the externality \tilde{X} is also about 1/3 lower, $I_{D,M}$ and $I_{D,L}$ change by only a percent or two.

The example in Figure 2 suggests that disutility from investment could easily overcome the hedging motive, and that it could explain, e.g., environmentalists' preference for socially responsible investment, provided that nonpecuniary disutility from investment is of approx-

⁷ ζ_D is chosen such that $I_{D,M}$ has the same value when $\eta = 1$ as the case where M internalizes, or there is Pigouvian taxation at rate λ_M . The slight adjustment reflects that disutility is proportional to investment, whereas the tax is levied on gross output.

imately the same magnitude as disutility from pollution. The example also suggests that, contrary to the case where environmentalists internalize, the population share of environmentalists should have little impact on how they invest.

4 Green financial innovation: derivatives

If investors have a visceral distaste for holding shares of firm D , such feelings probably stem from the identity of the company and the fact that it causes the externality, rather than from the return characteristics of firm D 's equity. Furthermore, in our example the incentive to hold shares in firm D is not only to hedge the externality, but also for optimal portfolio diversification.

Suppose there is a derivative contract offering a similar payoff to investment in firm D , but without the disutility: a greener alternative. Specifically, consider a forward contract with payoff

$$\tilde{S} - F, \tag{37}$$

where F is the equilibrium forward price that clears the market with zero net supply of forward contracts, and \tilde{S} is the underlying spot price, which we take as exogenous with distribution $\tilde{S} \sim \mathcal{M}(\mu, \sigma)$. For now we leave \tilde{S} abstract: it is statistically related to $\tilde{\mu}_D$, but type M can trade the forward without bad feelings. The spot price \tilde{S} has correlation $\rho_S \neq 0$ with the externality $\tilde{X} = \tilde{\mu}_D I_D$ or, equivalently, with returns to investment in firm D , but it is uncorrelated with $\tilde{\mu}_C$.

In this section we also assume for simplicity that $\rho = 0$, i.e., firms D and C have uncorrelated stock returns, noting that this matches the running numerical example with parameters in Table 1. With $\rho = 0$, optimal investment in firm C is unrelated to the externality, and is unaffected by the introduction of forwards at either the individual or aggregate level. A similar thought experiment, relevant to investment in the clean firm, might be the introduction of a green bond. In Appendix B we consider a green bond that

gives agent M positive nonpecuniary utility, and that has a pecuniary payoff positively correlated with firm C , capturing the idea that a subset of the clean firm's activities might be marketed as helping the environment.

Here we focus on the introduction of a forward contract correlated with the dirty firm payoff, and how it affects investment in the dirty firm. Define f_j , the forward position of agent j . In the economy with nonpecuniary disutility and forwards, agent L solves

$$\max_{I_{C,L}, I_{D,L}, f_L} \mathbb{E}[-\exp\{-\alpha(I_{D,L}(\tilde{\mu}_D - 1) + I_{C,L}(\tilde{\mu}_C - 1) + f_L(\tilde{S} - F) - \lambda_L \tilde{X})\}], \quad (38)$$

and agent M solves

$$\max_{I_{C,M}, I_{D,M}, f_M} \mathbb{E}[-\exp\{-\alpha(I_{D,M}(\tilde{\mu}_D - 1) + I_{C,M}(\tilde{\mu}_C - 1) + f_M(\tilde{S} - F) - \lambda_M \tilde{X} - \zeta_D I_{D,M})\}]. \quad (39)$$

Agents optimally take forward positions

$$f_L = \frac{\mu - F}{\alpha\sigma^2} - \rho_S(I_{D,L} - I_D\lambda_L), \quad (40)$$

$$f_M = \frac{\mu - F}{\alpha\sigma^2} - \rho_S(I_{D,M} - I_D\lambda_M). \quad (41)$$

In equilibrium forwards are in zero net supply, such that $\eta f_M + (1 - \eta)f_L = 0$. From this we pin down the market-clearing forward price

$$F = \mu - \rho_S(\mu - 1 - \eta\zeta_D). \quad (42)$$

Investment by each agent and average per-capita investment are

$$I_{D,L} = \frac{(1 + (\lambda_L - \lambda_M)\eta)(\mu - 1 - \rho_S(\mu - F)) - \lambda_L\eta\zeta_D}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho_S^2)}, \quad (43)$$

$$I_{D,M} = \frac{(1 + (\lambda_M - \lambda_L)(1 - \eta))(\mu - 1 - \rho_S(\mu - F)) - T_L\zeta_D}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho_S^2)}, \quad (44)$$

$$I_D = \frac{\mu - 1 - \eta\zeta_D}{(1 - \tau_L - \tau_M)\alpha\sigma^2}. \quad (45)$$

The following proposition summarizes the effect on investment of introducing the forward contract.

Proposition 5 *For a given $\eta \in (0, 1)$, when agent M has nonpecuniary disutility from investment in firm D , introducing the forward contract decreases $I_{D,M}$ relative to the case without the forward, and increases $I_{D,L}$ relative to the case without the forward. Both differences are increasing in η . Introducing the forward does not change average per-capita investment I_D .*

Proof. The difference between $I_{D,M}$ with versus without forwards is Equation (44) subtract Equation (35), which reduces to

$$\frac{-\rho_S^2\zeta_D(1 - \eta)}{\alpha\sigma^2(1 - \rho_S^2)} < 0,$$

given $\rho_S \neq 0$, $\eta \in (0, 1)$, and all other parameters positive, per prior assumptions. The expression is linearly increasing in η .

The difference between $I_{D,L}$ with versus without forwards is Equation (43) subtract Equation (34), which reduces to

$$\frac{\rho_S^2\zeta_D\eta}{\alpha\sigma^2(1 - \rho_S^2)} > 0,$$

with is also linearly increasing in η .

The difference between I_D with versus without forwards is Equation (45) subtract Equation (36), which is zero when $\rho = 0$, as assumed for this section.

■

These results have implications for the equilibrium effect of financial innovation on pollution, e.g., through the introduction of green alternative investments. We stress that the example concerns the introduction of a derivative in zero net supply, rather than the introduction of new productive green technology: in this sense it is a good example of a pure financial innovation. Perhaps surprisingly, the externality is unaffected by the innovation, since I_D is unchanged. Introduction of a green bond has analogous implications for investment in the clean firm; see Appendix B. However the forward does alter *who* invests in the polluting firm: environmentalists (type M) invest relatively less in polluters, whereas the remaining investors (type L) invest relatively more. But the greater the environmentalist’s share of the population, the “less effective” the green alternative is, in the sense that $I_{D,M}$ falls only slightly with the introduction of the forward when η is large. This parallels the result that the forward is neutral with respect to aggregate investment I_D : the forward only enables environmentalists to divest in equilibrium to the extent that non-environmentalists are available to take the other side of the forward contract, which they in turn hedge by increasing investment in the polluting firm.

Once again our running numerical example may aid intuition, and give some idea as to magnitudes. Our forward contract is an abstraction, but we have in mind carbon emissions futures contracts as an example. The long-term relationship between carbon emissions allowance (EA) prices and the stock returns of heavy CO^2 emitters is potentially complicated and has not, to our knowledge, been studied extensively. However Aatola et al. (2013) documents a significant and negative relationship between the price of coal and the price of EU carbon EAs, whereas the (comparatively clean alternative) natural gas price shows a significant and positive relationship with EA prices. Therefore it seems plausible that $\rho_S > 0$, and we choose $\rho_S = 0.8$. Remaining parameter values are in Table 1. Note that Proposition 5 requires only $\rho_S \neq 0$.⁸

⁸The sign of ρ_S affects which agents take long or short positions in the forward, but does not affect productive investment. The extent to which $I_{D,L}$ and $I_{D,M}$ are impacted by introduction of the forward is

Figure 3 illustrates the magnitude of the effects described in Proposition 5 with our example parameters. When his type is a small share of the population ($\eta \approx 0$), agent M invests almost nothing in firm D , with his investment $I_{D,M}$ increasing linearly as his share of the population increases. Relative to the case with disutility only, introducing the forward flips the sign of the relationship between η and $I_{D,M}$. When his population share $1 - \eta$ is large, agent L invests far more than M in firm D , and $I_{D,L}$ increases even further as $1 - \eta$ declines. However the pace of L 's decline in population share is more rapid than the increase in his investment, so I_D is still decreasing in η in equilibrium. Indeed this decline in I_D exactly matches Figure 2, the case without forwards. The difference between the figures is entirely in who invests: individual agent investment was approximately invariant to η in Figure 2, whereas individual investment is very sensitive to η in Figure 3. The implications for individual investment are also very different from Figure 1, where M correctly views divestment by his type as more impactful as η increases, and so $I_{D,M}$ is decreasing in η , the opposite of Figure 3.

Price changes and trade in the forward market actualize changes to individual investment, as illustrated in Figure 4. For $\eta \in (0, 1)$, agent L sells forwards, whereas agent M buys them. As η increases, the forward price F rises, and L sells more forwards to M . However L hedges his increased short position by investing more in firm D , which has returns positively correlated with \tilde{S} , and hence negatively correlated with a short forward position. While agent L sells more forwards on a per-capita basis, his share of the population is declining, so from the perspective of agent M the forward market is drying up as $\eta \rightarrow 1$. Eventually the forward price rises to the point where M views it as too costly an alternative to investment in firm D , so $f_M = 0$, and $I_{D,M}$ peaks, as previously illustrated in Figure 3.⁹

In our example there is always a risk premium to the long side of the contract, $\mu - F > 0$. This result is parameter dependent. Per Equation (42), the risk premium may be negative

also sensitive to ρ_S^2 , and would be small for $\rho_S \approx 0$.

⁹Note that forward positions f_L and f_M are defined on a per-capita basis. Adjusted for population shares, type L always sells as many forwards as type M buys.

if instead $\rho_S < 0$, obviously, but it may also turn negative if η is sufficiently large and type M has sufficiently high disutility ζ_D from investment in firm D . Counterintuitively, the risk premium is independent of λ_M and λ_L , which would ordinarily determine whether high realizations of $\tilde{\mu}_D$ are good, because output is high, or bad, because the externality is high. These parameters are absent because the motive for trade in forwards is not that L is relatively more tolerant of the externality than M , but rather that M has nonpecuniary disutility from investment in firm D . Without such disutility, forwards would not trade, because the agents can already adjust their exposure to $\tilde{\mu}_D$ by adjusting their investment in firm D . Thus the risk premium on forwards has, in a sense, nothing to do with risk.

5 Welfare

Up to this point we have emphasized determinants of individual and aggregate investment. Because investment in firm D increases the externality in expectation, the tendency is to consider increases in I_D as bad, whereas decreases are good. But of course the externality is only one component of expected utility, and even when different mechanisms have identical implications for I_D , welfare implications may differ greatly.

It is possible to characterize some interventions or structural economic changes as Pareto improving, but for interesting cases the welfare implications are parameter-or-state-dependent. So instead we give a brief but representative illustration of the effects based on our running numerical example, with baseline parameters in Table 1.

Figure 5 illustrates the welfare effects when type M internalizes pollution, as in Section 3.1, conditional on M 's population share η . Points of comparison are an economy without internalizing or Pigouvian taxes ($\tau = 0$), and one without internalizing but with Pigouvian taxation at M 's marginal cost ($\tau = \lambda_M$). We report expected utility for types L and M , and also a population weighted average expected utility. There are many alternatives to the latter representation of aggregate social welfare; this is merely one example of

interest. Appendix C discusses the investment allocation that a social planner would choose to maximize this welfare criterion.

Relative to the case without taxes, internalizing is Pareto improving for any η : it raises expected utility for all agents. In our simple production setting this result generalizes to other parameter values, provided agents agree that \tilde{X} is a bad. Relative to the case with taxation the effects are ambiguous: with tax $\tau = \lambda_M$ and $\eta < 1$, agent M has higher expected utility than when he coordinates with others of his type, because agent L also feels the effect of the tax. But since the tax is above agent L 's marginal cost from the externality, his expected utility is the lowest in this scenario. On a population weighted average basis, whether internalizing improves upon Pigouvian taxation depends upon η , and of course it would also depend upon the level of the Pigouvian tax.

Perhaps the outcome of this comparison is unsurprising, but it still provides a useful lesson. Given that the socially optimal Pigouvian tax is ambiguous or difficult to assess, it may be desirable to instead facilitate coordination among like-minded investors, to internalize their group contribution to pollution. This should be Pareto improving if the status quo is no taxation. For discussion of the optimal Pigouvian tax under the population-weighted welfare criterion, see Appendix C.

By contrast, the welfare implications of nonpecuniary disutility are negative for agent M , regardless of his population share η , as illustrated in the middle panel of Figure 6. M 's expected utility is always lower than in the baseline model, with or without taxation. Recall that the amount of disutility, ζ_D , is chosen such that, at the boundary $\eta = 1$, investment implications are identical to the case with Pigouvian taxation $\tau = \lambda_M$. In this sense, disutility achieves the optimal reduction in the externality when $\eta = 1$. But because disutility acts like a dissipative tax on investment in firm D , M 's expected utility still falls on net. Scenarios in which disutility could be welfare improving are ones where both the disutility ζ_D and the sensitivity to the externality λ_M are so large that $I_{D,M}$ would optimally be close to zero, i.e., scenarios in which complete divestment from the dirty firm is warranted.

Figure 6 also shows that introduction of the forward alleviates M 's utility loss somewhat, especially when η is low and forwards trade on terms more favorable to M . But M 's expected utility still remains well below that in the baseline model. Ironically, the beneficiary of agent M 's disutility from dirty investment is agent L , who sees an increase in expected utility due to the reduction in the externality (top panel, Figure 6). His expected utility increases even more with the introduction of the forward contract, as he shares the surplus generated by its introduction. The advantages to L are nevertheless small relative to the loss in utility experienced by M , so weighted average expected utility, shown in the bottom panel, is typically below that of the baseline model with $\tau = 0$.

It follows that disutility from investment is a double-edged sword. On the one hand, when investors feel bad about funding dirty industries, they will fund them less, and this reduces pollution. On the other hand, the cost in bad feelings may outweigh the benefit. From a welfare standpoint, campaigns to reduce externalities by making people feel bad about dirty investment can be counterproductive, particularly if some level of investment in the technology is optimal or essential.

6 Conclusion

We present an equilibrium model of portfolios and asset prices when one firm produces a negative externality — pollution — and the other does not. Households are heterogeneous, with some more sensitive to the externality than others. More sensitive agents tilt their portfolios towards the polluting firm, and the polluting firm is more highly valued than the non-polluting firm. Observed preferences for green investments among environmentally conscious investors, and observed higher returns to polluting firms, are more at odds with conventional theory than has previously been recognized. We discuss the implications of various modifications to our elementary model, which might reconcile theory with extant empirical work. We also show that financial innovation in the form of a green derivative has

an effect on welfare, but has no effect on investment in the polluting firm.

References

- Aatola, P., Ollikainen, M., Toppinen, A., 2013. Price determination in the eu ets market: Theory and econometric analysis with market fundamentals. *Energy Economics* 36, 380–395.
- Albuquerque, R. A., Durnev, A., Koskinen, Y., forthcoming. Corporate social responsibility and firm risk: Theory and empirical evidence. *Management Science*.
- Allen, K., February 20, 2018. Strict us market rules limit corporate sellers of green bonds. *Financial Times*.
URL <https://www.ft.com/content/baa217c4-157c-11e8-9376-4a6390addb44>
- Bénabou, R., Tirole, J., 2010. Individual and corporate social responsibility. *Economica* 77 (305), 1–19.
- Bovenberg, A. L., Goulder, L. H., 1996. Optimal environmental taxation in the presence of other taxes: general-equilibrium analyses. *The American Economic Review* 86 (4), 985–1000.
- Chava, S., 2014. Environmental externalities and cost of capital. *Management Science* 60 (9), 2223–2247.
- Davies, S. W., Van Wesep, E. D., 2017. The unintended consequences of divestment. *Journal of Financial Economics*.
- Dimson, E., Karakaş, O., Li, X., 2015. Active ownership. *The Review of Financial Studies* 28 (12), 3225–3268.
- Heinkel, R., Kraus, A., Zechner, J., 2001. The effect of green investment on corporate behavior. *Journal of financial and quantitative analysis* 36 (4), 431–449.
- Hong, H., Kacperczyk, M., 2009. The price of sin: The effects of social norms on markets. *Journal of Financial Economics* 93 (1), 15–36.
- Hong, H., Kostovetsky, L., 2012. Red and blue investing: Values and finance. *Journal of Financial Economics* 103 (1), 1–19.
- Kelsey, D., Milne, F., 2006. Externalities, monopoly and the objective function of the firm. *Economic Theory* 29 (3), 565–589.
- Krüger, P., 2015. Corporate goodness and shareholder wealth. *Journal of financial economics* 115 (2), 304–329.
- Morewedge, C. K., Tang, S., Larrick, R. P., 2016. Betting your favorite to win: Costly reluctance to hedge desired outcomes. *Management Science*.

Pachauri, R. K., Allen, M. R., Barros, V. R., Broome, J., Cramer, W., Christ, R., Church, J. A., Clarke, L., Dahe, Q., Dasgupta, P., et al., 2014. Climate change 2014: synthesis report. Contribution of Working Groups I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change. IPCC.

Pigou, A. C., 1920. The economics of welfare, 1920. McMillan&Co., London.

Roth Tran, B., 2018. Divest, disregard, or double down? philanthropic endowment investments in objectionable firms. Tech. rep., Federal Reserve Board.

A Model solution when both types of agent internalize pollution

Section 3.1 solves a model in which agents of type M internalizes their collective contribution to pollution, but agents of type L still treat pollution as an externality. This section summarizes the model solution under the alternative assumption that both L and M internalize.

Investment in the polluter is

$$I_{D,L} = \frac{(1 + \eta(\lambda_L - \lambda_M))(\mu - \rho(\mu - 1)) - \frac{T_M}{T_L} - \frac{\eta\lambda_L}{T_M}}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)} \quad (46)$$

$$I_{D,M} = \frac{(1 + (1 - \eta)(\lambda_M - \lambda_L))(\mu - \rho(\mu - 1)) - \frac{T_L}{T_M} - \frac{(1-\eta)\lambda_M}{T_L}}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}. \quad (47)$$

Relative to the case where only agent M internalizes, both agents invest less in firm D if agent L also internalizes. $I_{D,L}$ is reduced by

$$\frac{T_M \left(\frac{1}{T_L} - 1 \right)}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)},$$

while $I_{D,M}$ is reduced by

$$\frac{\lambda_M(1 - \eta) \left(\frac{1}{T_L} - 1 \right)}{(1 - \tau_L - \tau_M)\alpha\sigma^2(1 - \rho^2)}.$$

Agent M 's investment decision reflects less demand to hedge the externality, because the externality is less severe once L internalizes. Provided $\eta < 1$ and $\lambda_L > 0$, the reductions are strictly positive.

Investment in the clean firm is

$$I_{C,L} = \frac{\mu - 1 - \rho\left(\mu - \frac{1}{T_L}\right)}{\alpha\sigma^2(1 - \rho^2)}, \quad (48)$$

$$I_{C,M} = \frac{\mu - 1 - \rho\left(\mu - \frac{1}{T_M}\right)}{\alpha\sigma^2(1 - \rho^2)}, \quad (49)$$

for agents L and M , respectively. In general, internalizing pollution only affects investment in the clean firm when $\rho \neq 0$. For example, if $\rho > 0$, then internalizing increases investment in firm C , because the agents invest less in positively correlated firm D .

B Green financial innovation: bonds

Green bonds are intended to finance or refinance projects with small or positive environmental impact, including energy efficient buildings, renewable energy, etc. Legal requirements for green bonds are still evolving, and vary by jurisdiction (Allen, 2018). A subset of investors prefer environmentally friendly projects, and a green certification allows easy identification of related investments. We model this preference in agent M as positive nonpecuniary utility from green bond holdings.

Define risky green bond payoff \tilde{V} , which has correlation ρ_V with $\tilde{\mu}_C$, the clean firm's productivity. We have in mind $\rho_V > 0$, since the clean firm also uses technology without environmental impact. However we show that ρ_V is irrelevant for aggregate productive investment, counter to intuition.

P_V is the price of the bond, which is set to clear the market with

$$(1 - \eta)v_L + \eta v_M = 0, \tag{50}$$

where v_L and v_M are positions of agents L and M , respectively. Green bonds are in zero net supply: they are a pure financial innovation, not a new productive technology.

Equivalent to the section introducing the forward contract, we assume in this section that $\rho = 0$ (the clean and dirty firms' productivities are uncorrelated with each other), \tilde{V} is independent of $\tilde{\mu}_D$, and $\tilde{V} \sim \mathcal{M}(\mu, \sigma)$.

Agent M solves:

$$\max_{I_{C,M}, I_{D,M}, v_M} \mathbb{E}[-\exp\{-\alpha(I_{D,M}(\tilde{\mu}_D - 1) + I_{C,M}(\tilde{\mu}_C - 1) + v_M(\tilde{V} + \zeta_C - P_V) - \lambda_M \tilde{X} - \zeta_D I_{D,M})\}]. \quad (51)$$

In Section 4, investing in firm D caused M nonpecuniary disutility $-\zeta_D < 0$, whereas trading forwards didn't, even though forwards were a statistical substitute for such investment. In this section, investing in firm C produces no nonpecuniary utility for agent M , whereas investing in green bonds produces positive nonpecuniary utility $\zeta_C > 0$.

Agent L solves the analogous problem without nonpecuniary utility of any sort from investment holdings,

$$\max_{I_{C,L}, I_{D,L}, v_L} \mathbb{E}[-\exp\{-\alpha(I_{D,L}(\tilde{\mu}_D - 1) + I_{C,L}(\tilde{\mu}_C - 1) + v_L(\tilde{V} - P_V) - \lambda_L \tilde{X})\}]. \quad (52)$$

Because we assume $\rho = 0$, the solutions for $I_{D,M}$ and $I_{D,L}$ are unaffected by introduction of the green bond. The interesting question is how firm C investment will be affected. Investment in the clean firm, and positions in the green bond, are

$$I_{C,L} = \frac{\mu - 1 - \rho_V(\mu - P_V)}{\alpha\sigma^2(1 - \rho_V^2)}, \quad (53)$$

$$I_{C,M} = \frac{\mu - 1 - \rho_V(\mu + \zeta_C - P_V)}{\alpha\sigma^2(1 - \rho_V^2)}, \quad (54)$$

$$v_L = \frac{\mu - P_V - \rho_V(\mu - 1)}{\alpha\sigma^2(1 - \rho_V^2)}, \quad (55)$$

$$v_M = \frac{\mu + \zeta_C - P_V - \rho_V(\mu - 1)}{\alpha\sigma^2(1 - \rho_V^2)}. \quad (56)$$

Imposing the market clearing condition pins down the bond's price,

$$P_V = \mu + \eta\zeta_C - \rho_V(\mu - 1). \quad (57)$$

Substituting the solution for the equilibrium price into the investment and bond positions

gives reduced expressions in terms of the model parameters,

$$I_{C,L} = \frac{(1 - \rho_V^2)(\mu - 1) + \rho_V \eta \zeta_C}{\alpha \sigma^2 (1 - \rho_V^2)}, \quad (58)$$

$$I_{C,M} = \frac{(1 - \rho_V^2)(\mu - 1) - \rho_V (1 - \eta) \zeta_C}{\alpha \sigma^2 (1 - \rho_V^2)}, \quad (59)$$

$$v_L = \frac{-\eta \zeta_C}{\alpha \sigma^2 (1 - \rho_V^2)}, \quad (60)$$

$$v_M = \frac{(1 - \eta) \zeta_C}{\alpha \sigma^2 (1 - \rho_V^2)}. \quad (61)$$

The above results highlight several features of the green bond. First, trade in the green bond is motivated purely by agent M 's nonpecuniary utility ζ_C , as $v_L = v_M = 0$ for $\zeta_C = 0$. Provided $\zeta_C > 0$, i.e., the environmentalist likes to hold green bonds, then M takes a long position in the bond, and L takes a short position. When the bonds are highly correlated with firm C productivity ($|\rho_V| \approx 1$), more bonds trade, because the agents are able to hedge their bond exposure by trading firm C 's equity. In the most relevant case where $\rho_V > 0$, agent L 's trade can be thought of as refinancing investment in firm C by issuing green bonds: he invests ρ_V in firm C per unit of green bond sold. To compensate agent L for residual unhedged risk, a portfolio that shorts the green bond and goes long firm C pays a positive spread relative to investment in firm C ,

$$E[(\tilde{\mu}_C - 1) - (\tilde{V} - P_V)] = (1 - \rho_V)(\mu - 1) + \eta \zeta_C, \quad (62)$$

that is decreasing in the amount of unhedged risk. In the limiting case where $\rho_V \uparrow 1$, the spread is $\eta \zeta_C$: the joint surplus created when agent L intermediates investment in firm C , by issuing green bonds that generate nonpecuniary utility for agent M , is split according to the population shares of each agent.

It is also evident in the limiting case that agent L increases firm C investment one-for-one with his green bond issues, whereas agent M decreases firm C investment one-for-one with green bond purchases. In fact green bonds have no effect on aggregate green investment for

any ρ_V , since

$$I_C = (1 - \eta)I_{C,L} + \eta I_{C,M} = \frac{\mu - 1}{\alpha\sigma^2}, \quad (63)$$

which is unchanged from the case without the green bond in Equation (20) with $\rho = 0$. Therefore introduction of green bonds may improve welfare, by labeling investment in a way that makes certain investors feel good, but equilibrium investment is unchanged from the case without green bonds.

C A social planner's problem

As an example of a socially optimal benchmark, consider a planner's Pareto problem with population shares as weights:

$$\max_{I_{C,L}, I_{C,M}, I_{D,L}, I_{D,M}} \mathbb{E}[-(1 - \eta)e^{-\alpha\xi_L} - \eta e^{-\alpha\xi_M}], \text{ where} \quad (64)$$

$$\xi_L = I_{D,L}((1 - (1 - \eta)\lambda_L)\tilde{\mu}_D - 1) + I_{C,L}(\tilde{\mu}_C - 1) - I_{D,M}\eta\lambda_L\tilde{\mu}_D, \quad (65)$$

$$\xi_M = I_{D,M}((1 - \eta\lambda_M)\tilde{\mu}_D - 1) + I_{C,M}(\tilde{\mu}_C - 1) - I_{D,L}(1 - \eta)\lambda_M\tilde{\mu}_D. \quad (66)$$

ξ_L and ξ_M are composite consumption for each agent. The above problem corresponds to our basic setting with frictionless capital adjustment, as described in Section 2.1.

Conditional on dirty investment, the planner's solution for clean investment takes the same form as the competitive solution, reflecting the fact that clean investment produces no externality:

$$I_{C,L} = \frac{\mu - 1}{\alpha\sigma^2} - \rho(I_{D,L} - ((1 - \eta)I_{D,L} + \eta I_{D,M})\lambda_L),$$

$$I_{C,M} = \frac{\mu - 1}{\alpha\sigma^2} - \rho(I_{D,M} - ((1 - \eta)I_{D,L} + \eta I_{D,M})\lambda_M).$$

However the planner's solution for dirty investment is more involved. Therefore we solve for $I_{D,L}$ and $I_{D,M}$ numerically, for the example parameters in Table 1.

Since the planner's investment policy is not decentralizable, we also consider the Pigouvian tax rate τ^* that the planner would choose to maximize his objective in the competitive equilibrium. τ^* is the solution to

$$\max_{\tau} \mathbb{E}[-(1 - \eta)e^{-\alpha\xi_L^*(\tau)} - \eta e^{-\alpha\xi_M^*(\tau)}], \quad (67)$$

where $\xi_L^*(\tau)$ and $\xi_M^*(\tau)$ are the consumption plans chosen by each agent in the competitive equilibrium subject to Pigouvian tax τ .¹⁰

Figure C1 illustrates socially optimal investment as a function of agent M population share η . As in previous figures, we include competitive solutions without Pigouvian taxation, $\tau = 0$, and with taxation at rate $\tau = \lambda_M$, for comparison.

Three main points emerge from the figure. First, for any given η , it is socially optimal, as well as individually optimal, for M to invest more in the dirty firm than L . Therefore the competitive outcome that $I_{D,M} > I_{D,L}$, summarized in Proposition 3, is consistent with constrained-efficiency, since hedging against high realizations of the externality is actually more valuable to agent M than to agent L .¹¹ Second, in the competitive equilibrium both agents invest too much in firm D relative to the social optimum, unless that investment is taxed. Relative to the social optimum, the amount of overinvestment is increasing in η . Third, neither $\tau = 0$ nor $\tau = \lambda_M$ closely approximate the socially optimal outcome. τ^* , which is contingent on η , does approximate the socially optimal outcome, although it does not precisely achieve the optimum except when $\eta = 0$ or $\eta = 1$. Any tax proportional to dirty investment will fall more heavily on agent M than agent L , such that $I_{D,M}$ is too low relative to $I_{D,L}$, in comparison with the social optimum. Hence for $\eta \in (0, 1)$ and Pigouvian tax τ^* , $I_{D,L}$ lies slightly above the planner's solution, whereas $I_{D,M}$ lies slightly below the planner's solution.

Figure C2 reports expected utility (EU) for the planner's solution, again with competitive

¹⁰That is, the agents choose investment as per Equations (16), (18) and (19).

¹¹By contrast the outcome in Figure 1, where $I_{D,M} < I_{D,L}$ except for very small values of η , is constrained-inefficient.

solutions $\tau = 0$ and $\tau = \lambda_M$ for comparison. Since our main analysis focuses on the motives of the environmentalist agent M , the investment policy most beneficial to agent L has been relatively neglected. For small η , the planner's policy approximates what is optimal for L , achieving higher EU than for $\tau = 0$, which is too low a tax, or $\tau = \lambda_N$, which is too high a tax. However for higher η the planner depresses $I_{D,L}$ relative to what agent L would choose, as the marginal cost of the externality to agent M exceeds the marginal benefit of investment returns to agent L . By contrast agent M always enjoys higher EU under the planner's solution than for the competitive no-tax regime, and for some values of η he finds the planner's solution preferable to taxation at his marginal cost λ_M . It is socially optimal, in these cases, to make the non-environmentalist shoulder more of the burden of reducing aggregate dirty investment, because the environmentalist derives greatest hedging benefit from dirty investment. By contrast Pigouvian taxation cannot achieve this shifting of the divestment burden, which would require the effective tax rate on agent M to be lower than that on agent L . Finally, maximizing population average EU, in the bottom panel, is the planner's objective, hence the planner's solution reflects an upper bound for this measure of welfare. Note that population average EU for τ^* is slightly below what the planner achieves for $\eta \in (0, 1)$, although the difference is visually indistinguishable.

	Parameter	Value
Mean gross productivity of investment	μ	1.1
Standard deviation of productivity	σ	0.15
Correlation, productivity of firms B and G	ρ	0
Coefficient of absolute risk aversion	α	3
Sensitivity to externality, agent L	λ_L	0.01
Sensitivity to externality, agent M	λ_M	0.03

Table 1: Parameter values. The table reports the baseline parameter values used in our numerical examples.

Section	Frictionless capital adjustment	Pigouvian taxation	Coordination	Nonpecuniary utility	Derivatives
2		✓			
2.1	✓	✓			
3.1	✓		✓		
3.2	✓			✓	
4	✓			✓	✓
A	✓		✓		
B	✓			✓	✓
C	✓		✓		

Table 2: Summary of model variants. High level characteristics of the model variants considered in the paper are summarized by section number in the table above.

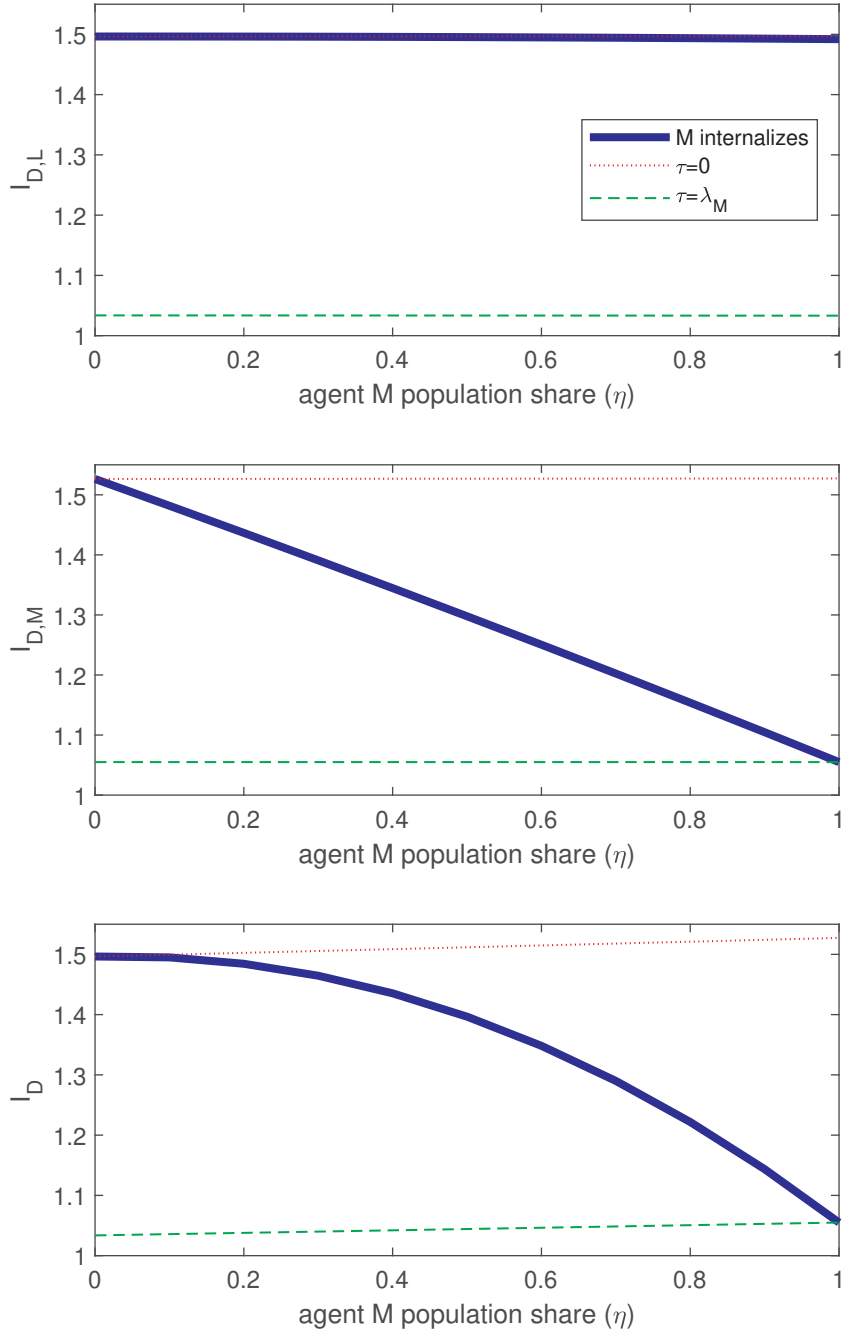


Figure 1: Investment: effect when M internalizes. The figure shows investment in dirty firm D by type L (top), type M (middle), and as a per-capita average (bottom), as a function of type M population share η . We compare three cases: without Pigouvian taxation, with Pigouvian taxation at fixed rate $\tau = \lambda_M$, and when type M internalizes, taking account of the externality generated by the collective investment of type M .

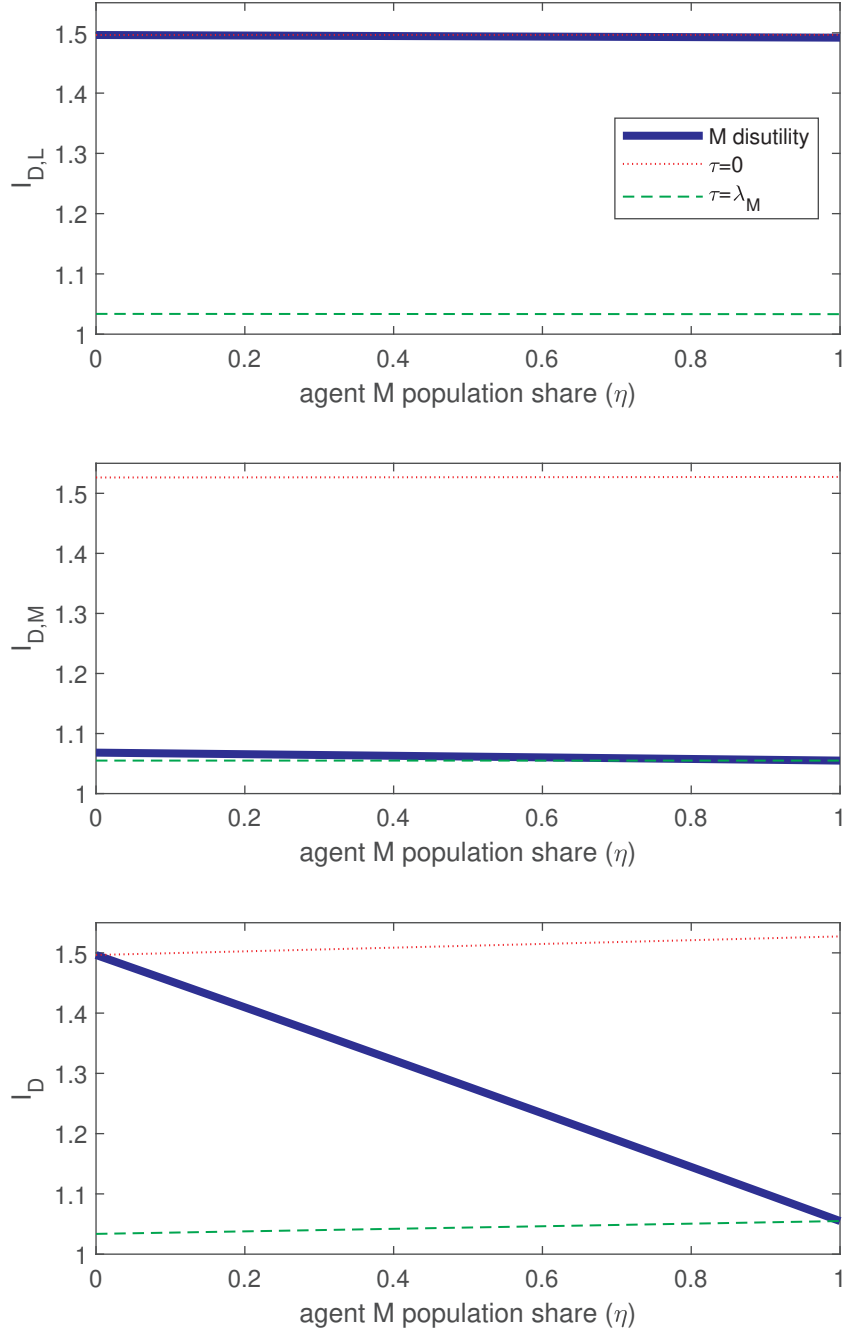


Figure 2: Investment: effect when M has disutility. The figure shows investment in dirty firm D by type L (top), type M (middle), and as a per-capita average (bottom), as a function of type M population share η . We compare three cases: without Pigouvian taxation, with Pigouvian taxation at fixed rate $\tau = \lambda_M$, and when type M has nonpecuniary disutility over investment in firm D with parameter $\zeta_D = \frac{\lambda_M}{1-\lambda_M}$.




Figure 3: Investment: effect with forward trading. The figure shows investment in dirty firm D by type L (top), type M (middle), and as a per-capita average (bottom), as a function of type M population share η . We compare three cases: without Pigouvian taxation, with Pigouvian taxation at fixed rate $\tau = \lambda_M$, and when type M has nonpecuniary disutility over investment in firm D with parameter $\zeta_D = \frac{\lambda_M}{1-\lambda_M}$, but may trade an alternative forward contract without suffering disutility. Correlation of the forward's payoff with the externality is $\rho_S = 0.8$.




Figure 4: Forwards. The figure shows forward positions of type L (top) and type M (middle), and the forward price (bottom), as a function of type M population share η . Type M has nonpecuniary disutility over investment in firm D with parameter $\zeta_D = \frac{\lambda_M}{1-\lambda_M}$, but may trade an alternative forward contract without suffering disutility. Correlation of the forward's payoff with the externality is $\rho_S = 0.8$.

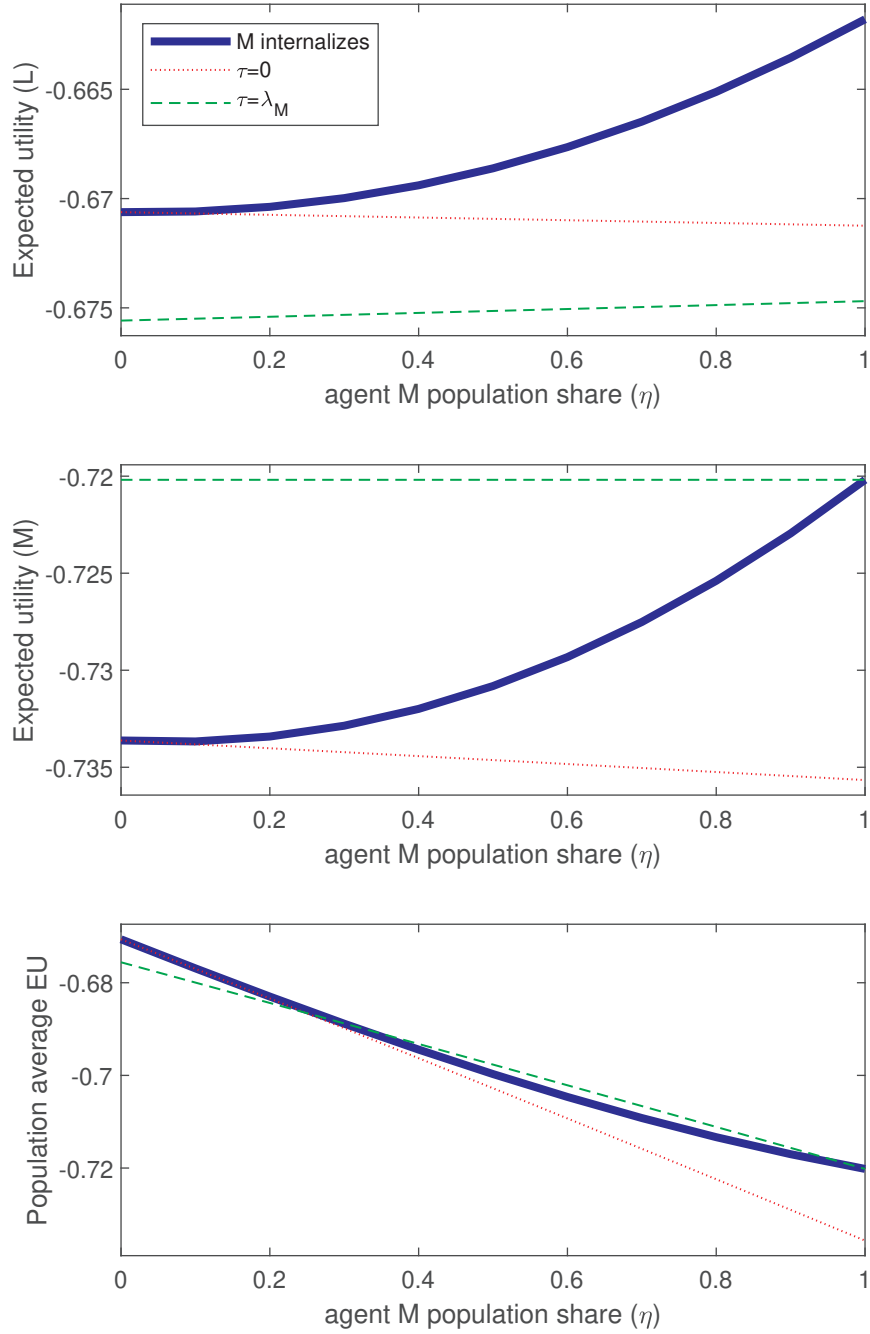


Figure 5: Expected utility: effect when M internalizes. The figure shows expected utility for type L (top), type M (middle), and as a per-capita average (bottom), as a function of type M population share η . We compare three cases: without Pigouvian taxation, with Pigouvian taxation at fixed rate $\tau = \lambda_M$, and when type M internalizes, taking account of the externality generated by the collective investment of type M .

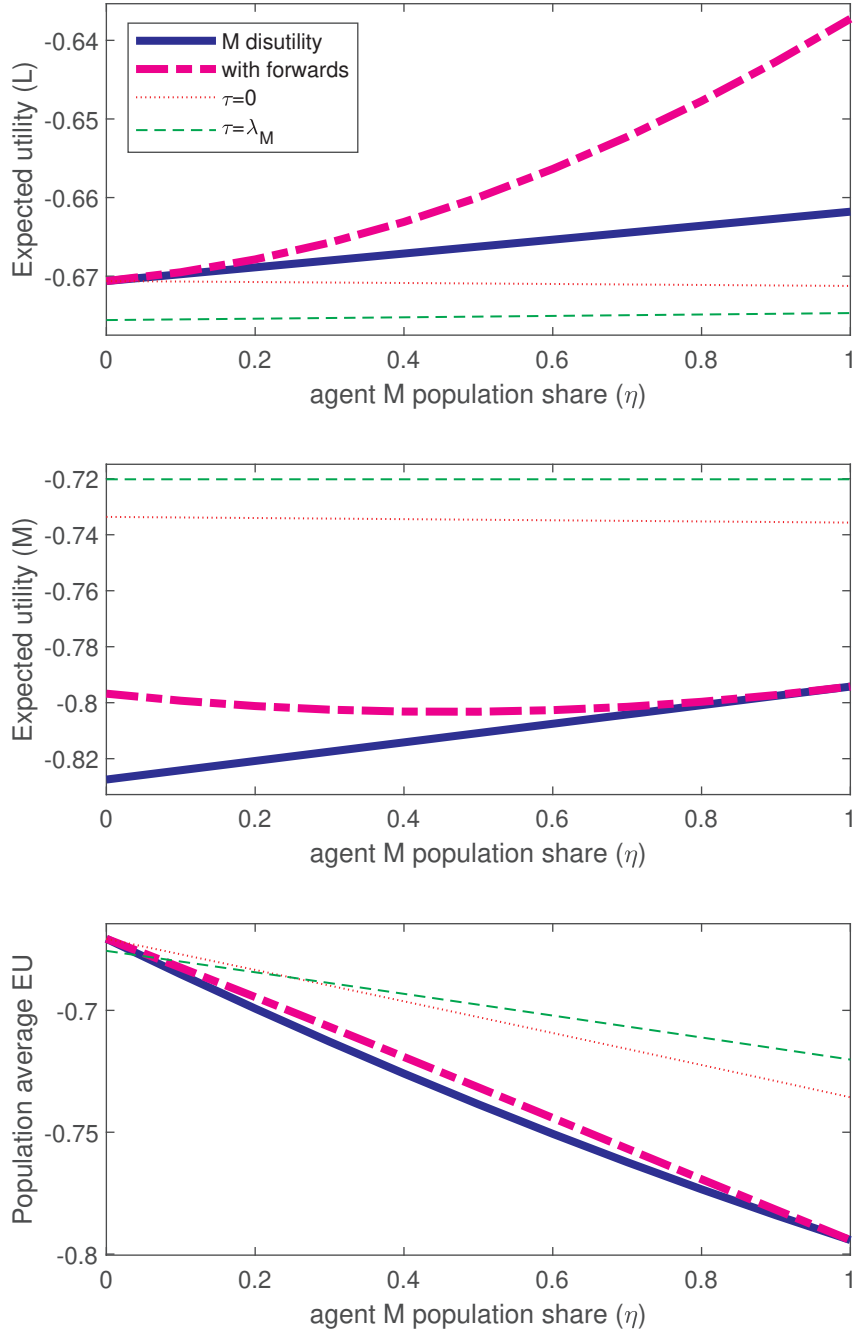


Figure 6: Expected utility: effect when M has disutility. The figure shows expected utility for type L (top), type M (middle), and as a per-capita average (bottom), as a function of type M population share η . The figure shows the case where type M has nonpecuniary disutility over investment in firm D with parameter $\zeta_D = \frac{\lambda_M}{1-\lambda_M}$, and the effect of introducing an alternative forward contract. For comparison, the figure also includes expected utility for the baseline model without disutility, without ($\tau = 0$) or with ($\tau = \lambda_M$) Pigouvian taxation. Correlation of the forward's payoff with the externality is $\rho_S = 0.8$.

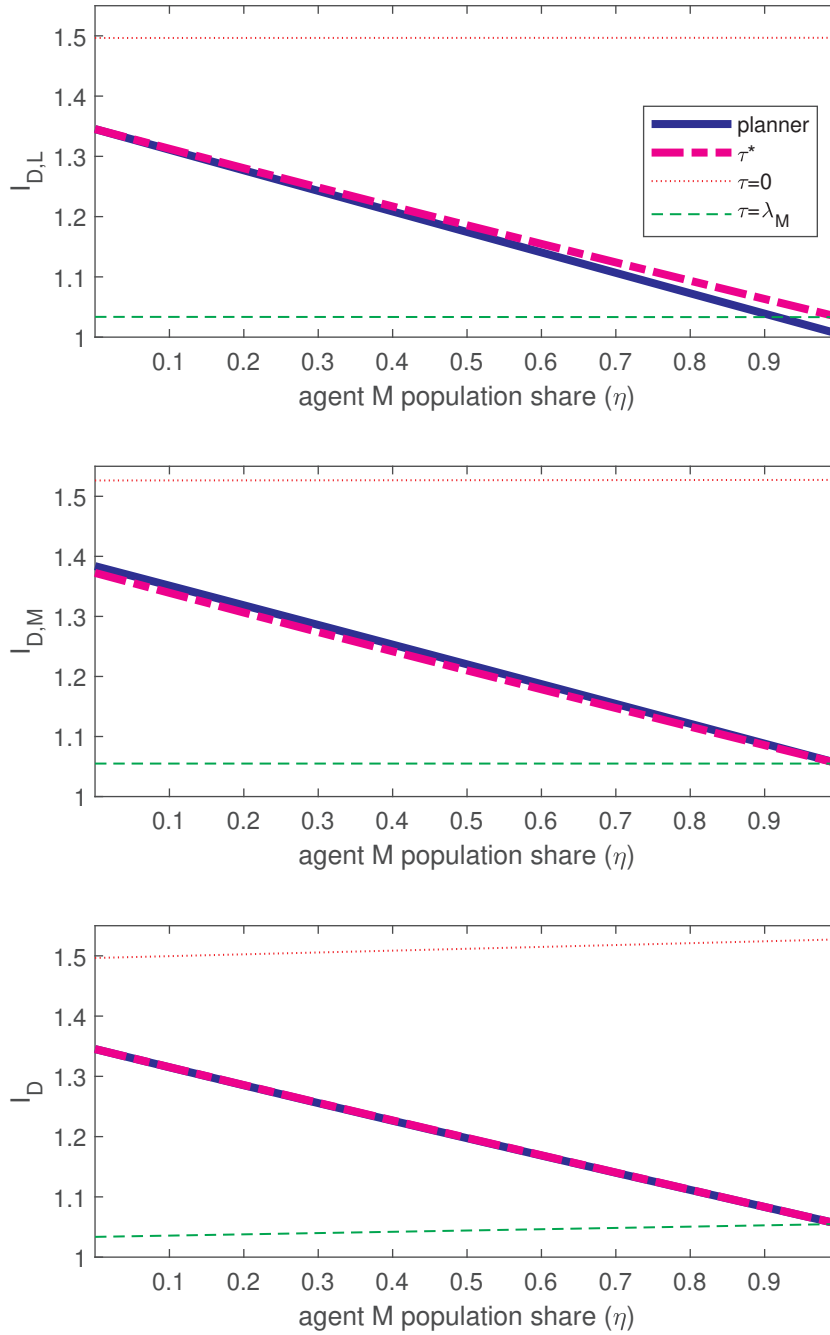


Figure C1: Investment: social planner’s solution. The figure shows investment in dirty firm D by type L (top), type M (middle), and as a per-capita average (bottom), as a function of type M population share η . We compare four cases: competitive equilibrium without Pigouvian taxation ($\tau = 0$), competitive equilibrium with Pigouvian taxation at fixed rate $\tau = \lambda_M$, a social planner’s solution, and the Pigouvian tax rate τ^* that maximizes the planner’s objective when agents behave competitively. The planner uses population shares as Pareto weights, and takes account of the externality in his solution.




Figure C2: Expected utility: social planner’s solution. The figure shows expected utility for type L (top), type M (middle), and as a per-capita average (bottom), as a function of type M population share η . We compare four cases: competitive equilibrium without Pigouvian taxation ($\tau = 0$), competitive equilibrium with Pigouvian taxation at fixed rate $\tau = \lambda_M$, a social planner’s solution, and the Pigouvian tax rate τ^* that maximizes the planner’s objective when agents behave competitively. The planner uses population shares as Pareto weights, and takes account of the externality in his solution.