Dynamic Climate Policy with Both Strategic and Non-Strategic Agents

Taxes Versus Quantities

Larry Karp
Sauleh Siddiqui
Jon Strand

The World Bank
Development Research Group
Environment and Energy Team
October 2013

Abstract

This paper studies a dynamic game where each of two large blocs, of fossil fuel importers and exporters respectively, sets either taxes or quotas to exercise power in fossil-fuel markets. The main novel feature is the inclusion of a “fringe” of non-strategic (emerging and developing) countries which both consume and produce fossil fuels. Cumulated emissions over time from global fossil fuel consumption create climate damages which are considered by both the strategic importer and the non-strategic countries. Markov perfect equilibria are

examined under the four combinations of trade policies and compared with the corresponding static games where climate damages are given (not stock-related). The main results are that taxes always dominate quota policies for both the strategic importer and exporter and that “fringe” countries benefit from a tax policy as compared with a quota policy for the strategic importer, as the import fuel price then is lower, and the strategic importer’s fuel consumption is also lower, thus causing fewer climate damages.

This paper is a product of the Environment and Energy Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The authors may be contacted at jstrand1@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Dynamic climate policy with both strategic and non-strategic agents: Taxes versus quantities

Larry Karp

Sauleh Siddiqui

Jon Strand*

JEL classification numbers : C63; C73; Q41; Q54.

Sector boards: Environment; Energy

*Karp: Department of Agricultural and Resource Economics, University of California, Berkeley, and the Ragnar Frisch Center for Economic Research; email: karp@berkeley.edu. Siddiqui: Department of Civil Engineering and the Johns Hopkins Systems Institute, Johns Hopkins University, Baltimore, MD 21218; email: siddiqui@jhu.edu. Strand: Development Research Group, Environment and Energy Team, The World Bank, Washington DC 20433; email: Jstrand1@worldbank.org. We benefitted from comments by Mike Toman, Franz Wirl, and seminar participants at the World Bank and the International Energy Workshop in Paris, June 2013. This research was supported by a grant from the World Bank's Research Support Budget. The conclusions and viewpoints presented here are those of the authors and not necessarily of the World Bank, its management, directors or staff.

1 Introduction

A small group of countries account for most exports of fossil fuels, in particular oil on which this paper focuses. Another main group of countries, which includes most of the OECD, are fuel importers with only limited own oil production. The latter group of countries has already established, or may soon be expected to establish, policies for limiting the carbon emissions resulting from their fossil fuel consumption. The global consumption of fossil fuels increases stocks of greenhouse gasses (GHGs), likely altering the climate. Climate policy might serve as a coordination device, enabling a strategic bloc of importers to affect the price of fossil fuels, while also controlling carbon emissions. In a game involving strategic importers and exporters, the equilibrium policy levels depend on the choice of instrument, e.g. a trade tax or quota. The equilibrium may also be sensitive to the presence of nonstrategic (passive) countries with fixed trade policies. These nonstrategic countries are “innocent bystanders” in the game among the strategic countries, but their presence alters the equilibrium to the game. The strategic exporter in our setting represents OPEC, and the strategic importer represents a subset of developed countries that might at some point in the future agree on a unified climate policy. That policy would likely involve trade measures, and would therefore affect both terms of trade and climate-related outcomes. The nonstrategic “innocent bystander” represents the poorer countries. We have two research questions: How does the presence of these countries affect the equilibrium outcomes in the games involving the strategic importer and exporter, under various combinations of policy instruments? How do the equilibria in the different games affect the welfare of the poorer countries?

In order to address these questions, we study a model in which a monopsonistic importer and a monopolistic exporter exercise market power, using either a trade tax or quota. There are four policy combinations, leading to four different games. We call the nonstrategic agent “ R ”, for Rest-of-World. R ’s trade policy is fixed and exogenous: free trade in our setting. GHGs related to fossil fuel consumption accumulate in a stock variable that causes differing levels of damages to both the strategic importer and R ; the strategic exporter incurs no damages. This assumption captures the idea that climate-related damages differ across countries, both in their actual effects, and in the way in which countries take these effects into consideration when determining their climate-related policies. In particular, we assume that, while OPEC countries may suffer damages from adverse climate developments, such damages are not reflected in the fuel-related policies of these countries.

In order to understand the forces at work, it helps to begin with a static framework in which there are no climate damages. In general, a country’s optimal trade tax equals the

inverse of their trading partner's (tax or quota inclusive) elasticity of import demand or export supply. If a large importer and exporter are the sole agents in this market, and both use a trade tax, the Nash equilibrium taxes are positive. These taxes lower aggregate world welfare, but do not eliminate trade (Johnson 1953). In contrast, if both of these agents use quotas, there is zero trade in the Nash equilibrium (Tower 1975). For example, suppose that the importer takes the exporter's quota as given, as in a Nash equilibrium. For any export quota, the importer has an incentive to set a still lower import quota, in order to render the export quota non-binding and thereby capture all of the quota rents from the exporter. The exporter who takes the import quota as given has the same incentive. Because this incentive holds for any positive quota, the only Nash equilibrium in the quota-setting game involves zero quotas for both agents, and thus zero trade.

In a two stage game where the countries choose their policy instrument (either a tax or a quota) in the first stage and the level of that policy in the second stage, countries' first-stage dominant strategy is a tax. A country does not want to select a quota in the first stage, because it understands that if it does so, the second stage equilibrium involves zero trade if the rival also chooses a quota; if the rival chooses a tax, it extracts all of the quota rent by setting its tax at a level that makes the quota non-binding. This qualitative comparison survives in a dynamic setting where the level of a policy instrument (tax or quota) changes over time, as the stock of GHGs or some other state variable changes endogenously (Wirl 2012), (Rubio 2005).

The introduction of a nonstrategic third agent, R , into the static game qualitatively alters incentives, and therefore potentially alters the equilibrium (Karp 1988). Suppose for example that the exporter sets a quota and R has a downwardly sloping excess demand for the commodity. The combination of the export quota and R 's excess demand function causes the strategic importer to face a kinked, *but not perfectly inelastic* excess supply function. In this situation, the strategic importer does not, in general, want to set its trade policy to render the exporter's quota non-binding. Such a policy still eliminates the exporter's quota rents, but it (typically) is too costly for the importer, because most or all of the rents might be transferred to R . The exporter faces an analogous situation. It does not matter whether, in equilibrium, R is an importer or an exporter; its mere presence causes a country, whose rival uses a trade quota, to face a downwardly sloping excess demand function for some range of prices. Thus, the introduction of R eliminates both of the earlier results: the quota equilibrium does not drive trade to zero, and it may not be the case that choosing a tax is a dominant strategy in the first-stage game where countries choose their policy instrument. Even with R , the use of a quota causes the partner to face a less elastic excess supply or

demand curve (relative to the curves under a tax). Thus, using a quota encourages the trading partner to use an aggressive trade policy. For that reason, the forces that promote the adoption of taxes rather than quotas operate even with the presence of R , but they might no longer be determinative.

Our chief policy question concerns the equilibrium welfare effect, of different policy choices, on the poorer countries, R . In our setting, R is a net importer of fossil fuels. The strategic importer has two targets, its terms of trade and climate-related damages, and a single (state-dependent) instrument, the tax or quota. Both of its objectives encourage the importer to restrict trade. Its trade restriction lowers the equilibrium world fossil fuel price and slows the growth of the GHG stocks. R is a free rider; it benefits from both of these changes, because it is an importer and it suffers climate-related damages. R therefore prefers the strategic importer to use an aggressive trade restriction. The strategic exporter has a single target, improving its terms of trade, and a single instrument. A more aggressive export restriction raises the world price, harming R , and slows the accumulation of GHGs, benefiting R . Therefore, the welfare effect, on R , of the export policy is ambiguous. The presence of R causes carbon leakage. As the importer lowers the market price by restricting its own demand for fossil fuel imports, R increases its demand for those imports. As the exporter increases the market price by restricting its exports, R shifts from imports to domestic production.

Our calibration assumes that climate-related damages are small to moderate relative to the benefit of consuming fossil fuel. As a consequence, terms of trade considerations are more important than climate-related damages for both the strategic importer and R . Higher pollution stocks cause the strategic importer to use more aggressive equilibrium trade restrictions, in order to reduce future damages. As this importer's demand for fossil fuels diminishes, with higher pollution stocks, the strategic exporter also lowers its export quota. The higher stocks directly harm the strategic importer, but affect the strategic exporter only indirectly, via reduced importer demand. Consequently, the importer's policy is much more sensitive to the pollution stock, compared to the exporter's policy: higher stocks reduce both equilibrium (strategic) imports and exports, but the effect on the former is greater. Therefore, higher pollution stocks increase the supply of imports available for R . At least for low stock levels (and in some policy scenarios for all stock levels), climate-related damages actually increase R 's welfare, simply because these damages cause I to reduce its demand for imports.

We also find that R 's payoff is highest when the strategic importer uses a tax, and the strategic exporter uses a quota. The exporter's use of a quota encourages the strategic

importer to use a high tariff, in order to capture quota rents. The high tariff reduces the world fossil fuel price, benefitting R . If the strategic countries can choose the policy instrument (in addition to choosing the level of the policy), the unique Nash equilibrium in the policy selection game is for each to use a tax, just as in the simple static model without R . The tax is a dominant strategy for both players at every level of stock, so this equilibrium is subgame perfect, as is the level of every policy, conditional on the stock. Given our calibrations, the first-best stock trajectory under the social planner who uses a Pigouvian tax lies below the equilibrium trajectories in the four games corresponding to the four combinations of trade policy. The emissions reductions arising from strategic countries' desire to improve their terms of trade, exceed the reductions due to the Pigouvian tax. Under the Pigouvian tax, there are no terms of trade incentives.

To check robustness, we consider other calibration assumptions. Plausibly, the climate externality could dominate the consumption benefits for some countries, notably countries in the fringe R ; many of these could suffer substantial climate damages. In such cases, welfare to R would be higher in the long run, when bloc I uses quotas. I 's use of the quota lowers accumulation of atmospheric carbon, benefitting R in the long run.

Recent papers compare taxes and quotas in static models of the fossil fuel markets, with two strategic blocs and two fuels (Strand 2011), and with one fuel and a non-strategic bloc as in our paper (Strand 2013). Both papers find (as do we) that the strategic fuel importer prefers a tax policy over a quota policy. Earlier papers focus on using an import tax to capture a seller's resource rent, (Bergstrom 1982), (Brander and Djajic 1983), (Karp 1984), (Karp and Newbery 1991). Climate policy may also be a means of capturing resource rents (Wirl and Dockner 1995), (Wirl 1995), (Amundsen and Schöb 1999), (Liski and Tahvonen 2004), (Rubio 2005), (Kalkuhl and Edenhofer 2010), (Njopmouo 2010). Wirl (2012), the paper closest to ours, studies a dynamic model with only two (strategic) blocs and no third (passive) bloc. This simpler model can be solved analytically. In this setting also, tax policies are dominant for both importer and exporter. Dong and Whalley (2009)'s computable general equilibrium model suggests that a 20% ad valorem carbon tax could increase real income in the U.S., E.U. and China by 0.4 – 0.8%, while reducing OPEC real income by 5%. Jørgensen, Martín-Herrán, and Zaccour (2010) and Long (2010) survey applications of the type of dynamic game that we, and many of the other cited papers, use.

2 The Dynamic Game

There are three agents in the game, representing three regions: the strategic importer bloc (I), the strategic exporter bloc (E), and the nonstrategic rest of the world (R). The importer and exporter blocs, I and E , exercise market power, using either a trade tax or a quota. We take as given the combination of policy instruments and calculate the equilibria under the four policy combinations. By comparing payoffs, we determine the equilibrium policy choice. Region R is a price taker and can be either a net importer or exporter of fossil fuels, depending on the world price.

In period t , the strategic importer (I) incurs damages resulting from the stock of GHGs, x_t . In order to emphasize the situation where agent I is more concerned than agent E about GHG accumulations, we suppose that only I and R suffer stock-related damages. These stocks are the only source of dynamics. In particular, we assume that extraction costs are independent of cumulative extraction, and we ignore the fact that resource stocks are finite. These assumptions produce a model with a single state variable, x_t . In view of our functional assumptions and reliance on numerical methods, we could extend the model to include a second state variable, cumulative extraction, and thereby take into account the non-renewable resource aspect of the problem. However, in the one-state variable model we can present all important results graphically; those graphs would be less useful in a two-state model, and the results would be harder to interpret. Given the complexity of results in even the one state variable model, it is worth beginning there, despite the fact that such a model does not capture the real-world property that fossil resources are exhaustible.

The trajectory of the stock of GHGs is endogenous to the model. Our solution concept is a Markov Perfect equilibrium. In any period t , the current stock is predetermined, a function of past stocks and emissions. Both strategic players condition their period- t policy (level) on the period t stock level, the only “directly payoff relevant” state variable in this model. The equilibrium level of I ’s policy in period t is a function of x_t , which makes the importer’s problem dynamic. The GHG stock does not directly affect the exporter’s payoff, because by assumption E does not incur climate-related costs. However, I ’s equilibrium policy (level) is conditioned on the stock, and I ’s policies directly enter E ’s payoff. Therefore, E also has a dynamic problem, and its equilibrium policy level also depends on the stock of GHGs. Because E and I solve mutually related dynamic problems, they play a dynamic game.

The rest of the world, R , responds passively, taking the world price as given. R ’s presence in the model is essential for two reasons. First, we want to know how the strategic interaction of large buyers and sellers affects nonstrategic agents, in particular, developing countries.

Second, the presence of R 's net demand means that when either I or E use a quota, the other strategic agent does not face a perfectly inelastic demand or supply function. In the absence of R , there is 0 trade in the equilibrium when both strategic agents use a quota; if only one strategic agent uses a quota, the other strategic agent can capture all of the gains from trade by using a price policy, absent R . Matters are more complex in the more realistic situation where R is present in the market.

2.1 Flow payoffs

We assume that supply and demand curves are linear, the stock-related damage function quadratic, and that E and R 's average production costs increase in the rate of output (but are independent of the stock). The world fuel price, defined as the price that E receives and R pays is p . Consumers in I pay the price P . $P - p$ equals the (possibly implicit) unit tax or quota rent in I . We first state the single period payoffs of the three agents, and then use these to define the dynamic game.

Country I has no domestic production; its demand for imports equals $A - BP$. The tariff revenue or the quota rents equal $(P - p)(A - BP)$. The climate-related damages, conditional on x , is $\frac{d}{2}x^2$ where d is a constant. The stock x is an amalgam of all climate-related variables, e.g. carbon stocks and temperature changes, and thus does not have a simple physical interpretation. Merely for purpose of exposition, we refer to it as the pollution stock. I 's single period payoff equals consumer surplus plus tariff revenue (or quota rents) minus environmental damages:

$$\begin{aligned} I\text{'s flow payoff: } & \int_P^{\frac{A}{B}} (A - Bz) dz + (P - p)(A - BP) - \frac{d}{2}x^2 \\ & = \frac{1}{2} \frac{(A - BP)^2}{B} + (P - p)(A - BP) - \frac{d}{2}x^2. \end{aligned} \tag{1}$$

At price p , R 's domestic demand is $a - b_0p$ and its domestic supply is b_1p , so its net imports equal $a - bp$, with $b_0 + b_1 \equiv b$. R 's gains from trade minus its climate related damages $\frac{\kappa}{2}x^2$ equal its flow payoff:

$$R\text{'s flow payoff: } \int_p^{\frac{a}{b}} (a - bz) dz = \frac{1}{2} \frac{(a - bp)^2}{b} - \frac{\kappa}{2}x^2. \tag{2}$$

This payoff is not relevant to the solution to the game, because R is passive. However, the solution to the game determines the equilibrium trajectories of p and x . That information, together with R 's single period payoff, enables us to calculate the present value of the stream

of R 's payoff, and thereby enables us to see how different policies affect R 's welfare.

The exporter, E , has no domestic consumption and faces no stock-dependent costs. If the fuel export price is p and the (possibly implicit) export tax in region E is τ , E 's producers receive the price $p - \tau$. These producers' marginal cost function, equal to E 's supply function, is $g + f(p - \tau)$, where g and f are constants. The exporter's single period payoff equals its domestic profits plus the tax revenue or quota rents

$$E\text{'s flow payoff: } \int_{-\frac{g}{f}}^{p-\tau} (fs + g) ds + \tau(g + f(p - \tau)) = \frac{1}{2} \frac{2gpf + g^2 + f^2p^2 - f^2\tau^2}{f}. \quad (3)$$

Each agent has the same constant discount factor, β . Welfare for each agent equals the discounted stream of their single period payoff.

2.2 Single period equilibrium

We can express single period payoffs as functions of the state variable, x , and the control variables. The identity of the control variables depends on the policy scenario. A strategic player can either use a quota, Q for I and q for E , or they can use a unit tax, T for I or τ for E . It is also possible that one agent uses a quota and the other a tax, resulting in four scenarios.

If the agents both use quotas (Q and q), the equilibrium conditions in E and in the world at large are

$$\begin{aligned} g + f(p - \tau) &= q \text{ and } q - Q - (a - bp) = 0 \implies \\ p &= \frac{-q+Q+a}{b}, P = \frac{A-Q}{B}, \text{ and } \tau = \frac{g+fp-q}{f} = \frac{(-f-b)q+fQ+gb+fa}{bf}. \end{aligned} \quad (4)$$

If they both use taxes (T and τ) equilibrium requires

$$\begin{aligned} g + f(p - \tau) - (A - B(p + T) + a - bp) &= 0 \implies \\ p &= \frac{f\tau+a+A-BT-g}{B+f+b} \text{ and } P = \frac{f\tau+a+A-BT-g}{B+f+b} + T = \frac{(f+b)T+f\tau+a+A-g}{B+f+b}. \end{aligned} \quad (5)$$

If I uses the tax T and E uses the quota q , equilibrium requires

$$\begin{aligned} g + f(p - \tau) &= q \text{ and } q - (A - B(p + T) + a - bp) = 0 \implies \\ p &= \frac{a+A-q-BT}{B+b} \text{ and } P = \frac{bT+a+A-q}{B+b} \text{ and } \tau = \frac{g(B+b)+f(A+a)-(f+B+b)q-BfT}{(B+b)f}. \end{aligned} \quad (6)$$

If I uses the quota Q and E uses the tax τ , equilibrium requires

$$\begin{aligned} g + f(p - \tau) - (Q + a - bp) &= 0 \implies \\ p &= \frac{Q+a-g+f\tau}{f+b} \text{ and } P = \frac{A-Q}{B} \end{aligned} \tag{7}$$

The first two equations on the second lines of each of (4) – (7) give the equilibrium values of p and P as linear functions of the control variables (a combination of Q , q , T and τ) for the four scenarios. If E chooses the tax, τ , as in the scenarios that correspond to equations (5) and (7), the optimality condition for E 's problem determines τ . If E chooses a quota, q , the equilibrium condition $q = g + f(p - \tau)$, together with the requirement that aggregate demand equal aggregate supply, leads to an implicit tax, τ . This implicit tax is a linear function of the control variables, as shown by the third equation in the second lines of (4) and (6). The payoffs, presented in Section 2.1, are quadratic functions of the prices, controls and stock, T , τ , and x , and thus are quadratic functions of the control variables and x in each of the four scenarios. Given its rival's level of trade policy, a country is indifferent whether it supports its own trade restriction using a quota or a tax. However, the level of its rival's equilibrium trade restriction depends on both the level and form (a tax or quota) of its own policy instrument.

2.3 Dynamics

In a period, the current level of GHG stocks is predetermined, at level x . Region R 's domestic supply is b_1p and E supplies q , so total emissions are $b_1p + q$. The constant decay rate is δ so next period stock, denoted x' , is

$$x' = \delta x + q + b_1p. \tag{8}$$

We use the same procedure as above to write the right side of this equation in terms of the control variables. For example, if both countries use quotas, we replace p with $\frac{-q+Q+a}{b}$.

2.4 Calibration

This game is too complicated to easily produce analytic results, but too simple to provide an accurate empirical description of fossil fuel markets. We solve the model numerically and select parameter values to provide an economically meaningful context, so that the results are informative about world markets. We assume that, if I uses no trade restrictions, I

imports the fraction $\Lambda < 1$ of E 's exports, and R imports the remainder: for any price, I 's elasticity of demand equals R 's elasticity of net demand. We define $\Gamma = \frac{b_0}{b}$, the slope of R 's demand relative to the slope of its import demand. By varying Γ , we can change R 's fraction of world production in a competitive equilibrium. By choice of units, we set I 's demand intercept $A = 8$, and E 's supply slope $f = 1$. We assume that E 's production equals 0 at $p = 0$, implying $g = 0$. This assumption and the linearity of E 's supply implies that E 's elasticity of supply everywhere equals 1. Our second calibration assumption is that I 's elasticity of demand, evaluated at free trade, also equals 1. With these two calibration assumptions (and $g = 0$) and the normalizations $A = 8$ and $f = 1$, the choice of Γ and Λ determine the remaining supply and demand parameters.

Variable	$A = 8\Lambda$	$B = \Lambda$	$a = 8(1 - \Lambda)$	$b_0 = \Gamma(1 - \Lambda)$	$b_1 = (1 - \Gamma)(1 - \Lambda)$
Value	5.6	0.7	2.4	0.05	0.25

$b = b_1 + b_2 = 1 - \Lambda$	g	Λ	$\Gamma = 0.1667$	f	d	δ	β	$\kappa = 0.0000991$
0.3	0	0.7	$0.25 \frac{-3.0 + 4.0\Lambda}{-1.0 + \Lambda}$	1	2.31331×10^{-4}	.99	.95	$\frac{d(1-\Lambda)}{\Lambda}$

Table 1: Benchmark parameter values

For our baseline, we set $\Lambda = 0.7$, so (absent import restrictions) I accounts for 70% of E 's exports, and $\Gamma = 0.1667$, so that in a competitive equilibrium E accounts for 80% of world supply. Table 1 summarizes our parameter choices and the relation between Λ and Γ and the supply and demand parameters. Table 2 shows the formulae relating model parameters to the elasticities and Γ , Λ .

To assess the sensitivity of our results to parameters, we also considered the alternative $\Lambda = 0.3$. The choice $\Lambda = 0.3$ corresponds, roughly, to the situation where I represents Annex B countries under the Kyoto Protocol; $\Lambda = 0.7$ corresponds to a more aggressive policy scenario, where I includes the US and China; including all BRIC countries in I increases Λ above 0.8. The cumulative supply is about 10% higher with $\Lambda = 0.3$ compared to $\Lambda = 0.7$, because with the lower Λ , I has less market power and restricts its fuel demand less. But results are qualitatively unchanged. The results for this alternative calibration can be made available upon request.

I 's demand elasticity = R 's net demand elasticity	E 's supply elasticity	E 's production share	I 's consumption share
$\frac{8-g}{8f+g}$	$f \frac{8-g}{8f+g}$	$\frac{-8f-g}{-8f-8+8\Lambda-\Lambda g+8\Gamma-\Gamma g-8\Gamma\Lambda+\Gamma\Lambda g}$	Λ
1	1	0.8	0.7

Table 2: Economic interpretation of demand and supply parameters; all formulae evaluated at competitive equilibrium

We choose the unit of time equal to a year and set the discount factor $\beta = 0.95$, for an annual discount rate of about 5.3%. The persistence parameter $\delta = 0.99$ implies a half-life of the pollution stock of approximately 90 years. Despite the lack of a physical interpretation of the stock x (see above), it is important that there be an economic and physical interpretation of the parameter d , in order to give context to the model results. We obtain the parameter d as a function of previously chosen parameters and the level of a threshold stock above which it is optimal for I to consume nothing. We can choose the value of this threshold, and thereby choose the value of d , by answering the following question: How many years of consumption at the competitive level would it take to reach the threshold stock? Our choice of d is consistent with the answer “105 years”, implying a threshold value of $x = 900$, with an initial value $x_0 = 0$. Appendix A explains this calibration procedure, which we intend only as a means of providing context for a numerical value that would otherwise be hard to interpret. Our results imply that for this value of d the environmental objectives are low relative to the terms of trade objectives; in that respect, our calibration represents low to moderate levels of damages.

We set R ’s damage parameter $\kappa = \frac{d(1-\Lambda)}{\Lambda}$. With this choice, the ratio of I and R ’s damage, for any stock, equals the ratio of their import demand absent trade restrictions: I and R have the same relative benefit of consumption to cost of stock-related damage; they merely differ in size. As a second sensitivity experiment, we hold fixed other parameter values and double the value of κ , to represent a situation where R has much higher damages than I , taking into account their size difference.

3 Results

We study four scenarios, in which I chooses a sequence of either taxes or quantities, represented by T or Q , and E chooses a sequence of either taxes or quantities, τ or q . In each case, a player’s equilibrium control rule is a linear function of x , equal to $\rho + \sigma x$ for I and $\lambda + \mu x$ for E . For example, if I chooses T and E chooses q , we have $T = \rho + \sigma x$ and $q = \lambda + \mu x$. The values of the four endogenous parameters, $\rho, \sigma, \lambda, \mu$ are different in the different scenarios. R does not use a policy, so it has no control rule. The equilibrium payoff of each of the agents – the present discounted value of that agent’s future payoff stream – is a quadratic function of the current stock. The payoff for I is $V(x) = \chi + \psi x + \frac{\omega}{2}x^2$, for

E is $W(x) = \epsilon + \nu x + \frac{\phi}{2}x^2$, and for R is $Y(X) = \varsigma + \eta x + \frac{\gamma}{2}x^2$. Appendix B explains how we obtain the value of these parameters in the four scenarios. Table 3 lists the parameter names.

Parameter	Importer	Exporter	ROW
Coefficient of x^2	ω	ϕ	γ
Coefficient of x	ψ	ν	η
Constant in Value Function	χ	ϵ	ζ
Coefficient of x in Constant Rule	σ	μ	-
Constant in Control Rule	ρ	λ	-

Table 3: Definition of endogenous parameters

We use the model parameters from Section 2.4. We first discuss the parameter values for the endogenous value functions and control rules for the case where both strategic agents use quotas. We then compare the equilibrium stock trajectories, payoffs and prices in the four scenarios. We use information on the payoffs to determine the equilibrium to the game in which agents choose their policy instrument (a tax or quota).

With one exception, we compare results across different policy scenarios by comparing them to the corresponding result in the scenario where both strategic agents uses quotas, which we call the “reference scenario”. The graphs labeled ImpTExpT, ImpQExpT, and ImpTExpQ refer, respectively, to the graph of an outcome when both agents use tax policies, when I chooses a quota and E chooses a tax policy, and when I chooses a tax and E chooses a quota. In all cases but one, the outcome (e.g. a payoff, price, or quantity) is relative to the corresponding outcome in the reference scenario.

The exception is for I ’s payoff, $V(x)$, where the reference scenario trajectory passes through zero. For each of the four scenarios, these payoffs are initially positive, because the initial value of the stock is $x = 0$. However, as x increases, the payoffs become negative. The switch in sign occurs at a different time in each of the scenarios. Normalizing I ’s payoff in scenario ImpTExpT, for example, by dividing by the payoff when both agents use quotas (ImpQExpQ) would involve dividing by 0. To avoid this problem, we show the payoffs for I in the four scenarios as levels, rather than ratios.

3.1 Equilibrium parameters when both agents use quotas

Table 4 shows the equilibrium values of the endogenous parameters under our baseline calibration, when both E and I use quotas. For all x , the importer’s payoff decreases with x ,

so $\omega < 0$ and $\psi < 0$. I 's equilibrium imports decrease as the stock rises, so $\sigma < 0$.

Parameter	Importer	Exporter	ROW
Coefficient of x^2 in value function	$\omega = -3.3 \times 10^{-3}$	$\phi = 3.2156 \times 10^{-6}$	$\gamma = 2.3663 \times 10^{-6}$
Coefficient of x in value function	$\psi = -0.1435$	$\nu = -2.75 \times 10^{-2}$	$\eta = 9.5 \times 10^{-3}$
Constant in value Function	$\chi = 14.4634$	$\epsilon = 120.9244$	$\zeta = 19.6012$
Coefficient of x in control rule	$\sigma = -3.9290 \times 10^{-4}$	$\mu = -1.7066 \times 10^{-4}$	-
Constant in Control Rule	$\rho = 0.5049$	$\lambda = 1.2624$	-

Table 4: Equilibrium values of endogenous parameters when E and I both use quotas.

Over relevant state space, E 's value function also decreases in the stock: $\nu < 0$ and $|\nu|$ is large relative to ϕ . However, $\phi > 0$, so E 's value function is convex in x . The stock has no direct effect on E , but as x increases, E faces decreasing demand from I (because $\sigma < 0$). I eventually becomes a negligible part of the market, so further decreases in its demand have a negligible effect on E 's payoff; hence, the convexity of E 's payoff in x . As I 's demand falls with the increase in x , E 's exports also fall: $\mu < 0$. The fact that E suffers no direct loss in utility due to higher pollution stock means that its payoff is much less sensitive to x , compared to I 's payoff. (Compare the magnitudes of ω and ϕ and of ψ and ν .) I 's equilibrium quota is about twice as sensitive to the stock, compared to E 's equilibrium quota: $\frac{\sigma}{\mu} \approx 2$.

R 's payoff is a convex increasing function of the pollution stock (both γ and η are positive); over the relevant range of stocks, the relation is approximately linear ($\frac{\gamma}{\eta} \approx 0$). A higher stock has offsetting effects on R 's payoff. The higher stock increases R 's damages, lowering its net payoff. The higher stock also decreases E 's supply and I 's imports, but the second effect is approximately twice as large as the first ($\frac{\sigma}{\mu} \approx 2$), so on balance a higher stock increases the supply absorbed by R , increasing its gains from trade. With our calibration, the higher gains from trade dominate the higher damages, so on balance higher pollution stocks benefit R .

3.2 Equilibrium stock trajectories

Figure 1 (a) shows the pollution stock trajectories as functions of time in the quota-setting game, and in the first-best scenario where the social planner uses Pigouvian taxes. We defer discussion of the outcome under the social planner until Section 3.3 and here discuss the stock trajectories under the games corresponding to different combinations of trade policies.

Figure 1: (a) Stock Trajectory of the Quota Setting Game and under the first-best social planner; (b) Stock Trajectories Relative to the Quota Setting Game for other scenarios considered

After 150 years, the stock reaches only 22% of the threshold level ($x = 900$, at which it is optimal for I to cease imports). Recall that our calibration assumes that under unrestricted trade the stock reaches this threshold in 105 years. This comparison shows a very significant reduction (relative to free trade) in cumulative extraction, resulting from the quota-quota policy combination. The magnitude of that reduction is consistent with either high damages or a high incentive to exercise market power, or both. Our subsequent results show that our calibration actually implies rather low damages, and that the stock reduction is due primarily to agents' incentives to exercise market power.

Figure 1 (b) shows stock trajectories relative to the reference trajectory, beginning with the first period. The initial stock equals 0 and the graphs start at time $t = 1$. In the early periods, the graphs reflect primarily ratios of initial emissions, whereas later values of the graphs reflect ratios of cumulative emissions, adjusted for the stock decay. These graphs are quite flat, implying that relative flows, across policy scenarios, change little over time.

Cumulative stocks are 10-35% higher in the other policy scenarios, relative to the quota-setting game. The stocks are highest where both strategic agents use taxes, and are at intermediate levels where one agent uses a tax and the other uses a quota. For this comparison, it matters little which of the two agents uses a tax. We noted that in a static setting, equilibrium quotas tend to reduce trade to a much greater extent than equilibrium taxes.

When an agent uses a quota rather than a tax, its trading partner faces a less elastic excess supply or demand function, and therefore has an incentive to use a more aggressive trade restriction. Figure 1 (b) shows that this comparison also holds in our dynamic setting.

The steady state when both countries use taxes is $x = 329$, much lower than the assumed threshold of $x = 900$; after 150 years the stock reaches 80% of its steady state level. The steady state stocks in the other policy scenarios range from 254 to 280; by year 150 the stocks in these scenarios also equal about 80% of their respective steady states.

3.3 Payoffs and instrument selection

Figure 2 shows the importer and exporter continuation payoffs (value functions V and W) as functions of the stock. (Recall that the former is in levels, and the latter shows graphs relative to the ImpQExpQ levels, accounting for the difference in scale of the two figures.) The principal information from these figures is that the tax is a dominant strategy for both countries, at every stock level. If both countries believe that they can choose their policy instrument in perpetuity in the initial period, the unique Nash equilibrium is for both to choose a tax. If they have the opportunity to revisit this decision at any time in the future (i.e. at any stock level), the equilibrium policy choice does not change. Consider the more complex game in which, at each period, agents choose both their policy instrument and the level of the instrument. In the MPE to this game, both countries choose the tax, and the tax equals that of ImpTExpT.

Both countries' payoffs decrease with the stock. The importer suffers stock-related damages. As the stock increases, I tightens its trade restriction, reducing the aggregate demand that E faces, and reducing E 's flow payoff and its continuation payoff. Because the importer suffers direct damages, its payoff is more sensitive to the stock, relative to the exporter's payoff.

The dots on the vertical axis identify the payoffs in the static game, obtained by setting the damage parameter to 0. Comparison of the dots corresponding to the static games and the intercepts corresponding to the dynamic games shows two facts. First, the payoff ranking is the same in the static and in the dynamic settings. Second, the payoffs in the dynamic setting lie only slightly below the corresponding values in the static setting. The differences reflect the fact that I suffers from damages in the dynamic game (harming I) and therefore uses more aggressive trade restrictions (harming E). Our calibration is consistent with relatively small damages. The static terms of trade considerations are much more important to I 's payoff, compared to environmental damages. The largest difference

Figure 2: (a) Importer's value function as function of the stock; (b) Exporter's value function as a function of the state

between the static and dynamic counterparts corresponds to the game in which both agents use taxes. As Figure 1 (b) shows, the equilibrium stock is significantly higher in that policy scenario, so the welfare impact of damages is greatest there.

If the importer is constrained to use a quota, it does not (much) matter to it whether the exporter uses a tax or a quota. (The graphs corresponding to the importer's payoff under ImpQExpQ and ImpQExpT are nearly coincident.) In contrast, if the exporter is constrained to use a quota, it much prefers the importer to use a quota rather than a tax.

Although the payoff ranking (across policy combinations) does not change with the stock level (i.e. the graphs in Figure 2 do not cross), the payoff ranking for the importer does change as a function of time. After about 100 years, the importer's continuation value is lower under ImpTExpT than under the other policy scenarios. After a century, the stock is sufficiently higher when both strategic blocs use taxes, compared to other policy combinations. This higher stock reduces the importer's payoff. However, as noted above, if countries were able to reconsider their policy instrument after 100 years of the ImpTExpT equilibrium, the unique Nash equilibrium remains for both to continue using taxes.

3.4 R 's payoffs

Figure 3 (a) shows R 's payoffs over time in the dynamic games, and its corresponding payoffs in the static games (the dots on the vertical axis). As for I and E , R 's ranking of policy scenarios is the same in the static and dynamic settings; and for any policy scenario, the

Figure 3: (a) R 's payoff, $Y(x_t)$, in the four policy scenarios; (b) R 's payoff with twice the damages (2κ)

payoff level is similar in the static and dynamic settings. Again, these features reflect the fact that the static producer and consumer surplus are much more important to R 's payoff, compared to the dynamic pollution cost. R 's payoff is highest when I uses a tax and E uses a quota. I 's use of a tax rather than a quota reduces E 's incentive to restrict its supply, benefitting R , which in our calibration is an importer. From Figure 1 (b), the stock trajectory is highest when both I and E use taxes. But I consumes much of that additional supply. When I continues to use a tax and E switches from a tax to a quota, aggregate supply falls, lowering R 's gains from trade (and slightly lowering its damages). E 's switch to a quota causes I to face a less elastic excess supply function, inducing I to increase its tariff, and reduce its consumption. The net effect is to increase R 's supply, thus increasing its gains from trade, and (because damages are relatively small) increasing its payoff.

R 's payoff is higher in the dynamic setting (with damages) compared to the static setting without damages. In contrast, both I and E have lower payoffs in the dynamic setting. Section 3.1's discussion of endogenous parameters explains this relation: stock-related damages cause both I and E to impose tighter trade restrictions, lowering their equilibrium gains from trade; because I suffers directly from the higher stocks, and E suffers only indirectly (via the induced tightening of I 's trade restriction), I 's response to the higher stock is greater than E 's. Thus, the net effect of the higher stock is to increase supply available to R , increasing its gains from trade. That increased gain swamps the direct cost to R , arising from stock-related damages. In the quota-setting game, we noted that R 's payoff is monotonic in the stock, but this relation does not hold for all games.

The non-monotonicity is easiest to see in Figure 3 (b), where we double R 's damages by doubling κ . Higher damages do not alter the comparison of static and dynamic payoffs; in this respect, damages remain small relative to the gains from trade, even when κ doubles. However, for higher damages R 's payoff is non-monotonic in time. Because the stock is monotonically increasing over time, we conclude that R 's payoff is non-monotonic in the stock. As above, a higher stock decreases I 's demand more than it lowers E 's supply, thereby increasing the supply available to R and increasing its gains from trade; and the higher stock increases R 's damages. At low stocks, early in the program, the first effect dominates; at high stocks, later in the program, the second effect dominates when we double the damage parameter κ . In this case, the relation between R 's payoff is first increasing and then decreasing over both time and over stock levels. When the climate damages are sufficiently important for R , relative to fossil-fuel consumption, R 's payoff in the ImpTExpT game must eventually fall below its payoff in the ImpQExpT game, simply because the rate of accumulation of the carbon stock is greater in the former case.

3.5 Price and policy trajectories

Figure 4 (a) shows the equilibrium world price, p (the price that R pays and E receives) and Figure 4 (b) shows the importer's domestic price P . As the stock increases and I tightens its trade restriction, P rises. As the stock increases and I 's import demand falls, the world price falls. E is in the strongest position to exercise market power when it uses a tax and I uses a quota; therefore, this scenario leads to the highest market fuel world price. I is in the strongest position to exercise market power when it uses a tax and E uses a quota; therefore, this scenario leads to the lowest world price. The other scenarios, where both agents use taxes or both use quotas, result in intermediate levels of the world price.

Recall that absent R , the equilibrium when both agents use quotas implies that no fuel is traded. As discussed in the Introduction, the presence of R moderates this extreme result. With R , it is too costly for the strategic agents to try to capture all of their rival's quota rent. Nevertheless, trade between I and E is lowest in the quota setting game, so that scenario results in the highest domestic price for I . For similar reasons, trade between I and E is highest when both agents use tariffs, so that scenario leads to the lowest domestic price for I . The dots on the vertical axes show the equilibrium prices in the static games, where damage equal 0.

Figure 5 graphs the explicit or (in the case where an agent uses a quota) implicit trade tax. Consistent with our previous discussion, these figures show that an agent has the

Figure 4: (a) The world price, p , in the four scenarios; (b) The importer's domestic price in the four scenarios

greatest incentive to exercise market power, and therefore uses the most restrictive trade policy, when it uses a tax and its rival uses a quantity restriction. The agent's trade policy is aggressive when it uses a quantity restriction and its rival uses a tax. For all policy combinations, the importer's trade tax (or quota price) increases over time, i.e. it increases with the pollution stock. The exporter's implicit or explicit taxes fall slightly over time. E 's exports fall over time, with the fall in the price that E receives. As this price falls, a lower export tax supports reduced levels of exports. (In contrast, at a constant world price, the export tax would have to increase in order to support reduced exports.)

Figure 5 also shows the Pigouvian tax trajectory, for comparison with the equilibrium trade taxes in the different policy scenarios. The Pigouvian tax supports the first best outcome. Figure 1 shows that the stock trajectory under the social planner who uses a Pigouvian tax is lower than the trajectory under any of the four combinations of trade policy. The strategic countries want to improve their terms of trade and, in the case of I , to control the emissions-related future damages that they suffer. In pursuit of these objectives, the strategic countries reduce emissions. Those reductions exceed the reductions achieved by the social planner who uses a Pigouvian tax imposed on all units on fuel consumption. Under this tax, consumers in I and R and producers in E and R face the same prices; the difference between those prices equals the Pigouvian tax. In the absence of R , where one country (E) has production but no consumption, and the other (I) has consumption but no production, the first best output path can be supported with any combination of import and export tax that sum to the Pigouvian tax. The division of this sum between the import and export taxes determines the amount of tax revenue that each country collects, but has

Figure 5: (a) Exporter’s (explicit or implicit) trade tax; (b) Importer’s (implicit or explicit) trade tax

no effect on equilibrium sales, and therefore has no effect on efficiency.

In the presence of R , the first best outcome cannot be implemented using only trade policies for I and E (simply because the first best requires that all consumers face the same price, and all producers face the same price). Therefore, there is no direct way to compare the Pigouvian tax with the sum of the trade taxes in the different policy scenarios. However, we note that in all policy scenarios the sum of the equilibrium trade taxes exceeds the Pigouvian tax at least for the first 50 years (and, except for ImpTExpT, this comparison also holds for the entire 150 year period that we consider). In order to interpret this comparison, consider the case of a planner whose objective is to maximize the sum of world welfare, and who is constrained to use only an export tax for E and an import tax for I (or quota-equivalents to such taxes). This planner cannot achieve the first best. The trade taxes create a distortion in the process of achieving the desired reduction in the stock; therefore, in general, the sum of the optimal export and import tax for this planner is *less* than the Pigouvian tax. The fact that the sum of the equilibrium trade taxes *exceeds* the Pigouvian tax reflects the fact that the trade taxes are set (primarily) in order to improve a country’s terms of trade, rather than to correct the environmental distortion (which is the planner’s sole objective). Comparison of the two graphs in Figure 1 reinforces this interpretation.

4 Conclusion

This paper extends previous literature on dynamic games between a large bloc of fuel exporters and a large bloc of fuel importers by including a nonstrategic third bloc of countries, R , representing the group of developing countries with no climate policy nor strategic trade

policy. The presence of this nonstrategic bloc means that even if a strategic country uses a trade quota, the excess supply or demand function facing its trading partner is not perfectly inelastic. We find, under our preferred calibration assumptions, that a tax policy by both the strategic importer and exporter constitutes the Markov (or subgame) equilibrium to this game, at any value of the state variable. This result echos results from related models, especially the static three-bloc model in Strand (2013), and the dynamic two-bloc model (without the fringe) in Wirl (2012).

The strategic importer and exporter both use trade policies to improve their terms of trade. The strategic importer also uses trade policy to control the future stock-related damages, but does not internalize the damages facing R . The fact that the stock changes over time renders the importer’s problem dynamic. Although the exporter has no intrinsic concern about the stock, its equilibrium trade policy depends on the importer’s policy and therefore is also stock dependent. For our calibration, the terms of trade objectives dominate the environmental objective in explaining policy levels. OPEC countries appear to be concerned that a unified climate policy among OECD countries might provide both “green cover” and a coordinating device that would enable the OECD countries to exercise greater market power in the fuel markets. Our results indicate that OECD countries might indeed have an incentive to behave strategically in this way; although our model has little to say about whether unified climate policy would actually induce such behavior. In our calibration, the strategic countries’ terms of trade objectives and concern for country-specific damages, lead to smaller equilibrium pollution stocks than under the social planner who can use a Pigouvian tax.

The nonstrategic agent, R , also suffers stock-related damages. This country, a net fossil fuel importer, is a free rider, benefiting from the importer’s trade restriction; that restriction lowers the equilibrium price of fossil fuels and also reduces the equilibrium stock trajectory, lowering damages to R . R ’s equilibrium payoff is higher in the dynamic setting, where it incurs stock related damages, compared to the static setting where it incurs no damages. The explanation is that stock-related damages cause the strategic importer to use more aggressive trade restrictions, benefiting R . The reduced competition for fossil fuel imports more than offsets the stock-related damages. The social planner’s optimal solution is to set a Pigouvian tax applied to all fossil fuel consumption, including by the fringe. In our model, by contrast, the fringe faces lower fossil fuel prices than the strategic importer.

Our calibration assumes that, under free trade, the strategic importer accounts for 70% of fossil fuel imports. This scenario corresponds to a situation where most large countries cooperate on trade and environmental policy; those two policies are indistinguishable in our

setting, where the strategic importer consumes but does not produce fossil fuels. We have also considered an alternative calibration, where strategic importers account for only 30% of imports under free trade. The qualitative results in the two cases are similar, although the smaller importer obviously has less market power and therefore uses less aggressive trade restrictions.

References

- AMUNDSEN, E., AND R. SCHÖB (1999): “Environmental taxes on exhaustible resources,” *European Journal of Political Economy*, 15(2), 311–329.
- BERGSTROM, T. (1982): “On capturing oil rents with a national excise tax,” *The American Economic Review*, 72(1), 194–201.
- BRANDER, J., AND S. DJAJIC (1983): “Rent-extracting tariffs and the management of exhaustible resources,” *Canadian Journal of Economics*, pp. 288–298.
- DONG, Y., AND J. WHALLEY (2009): “A Third Benefit of Joint Non-OPEC Carbon Taxes: Transferring OPEC Monopoly Rent,” *CESifo Working Paper Series*.
- JOHNSON, H. (1953): “Optimum tariffs and retaliation,” *The Review of Economic Studies*, 21(2), 142–153.
- JØRGENSEN, S., G. MARTÍN-HERRÁN, AND G. ZACCOUR (2010): “Dynamic games in the economics and management of pollution,” *Environmental Modeling and Assessment*, 15(6), 433–467.
- KALKUHL, M., AND O. EDENHOFFER (2010): “Prices vs. quantities and the intertemporal dynamics of the climate rent,” *CESifo Working Paper Series No. 3044*.
- KARP, L. (1984): “Optimality and consistency in a differential game with non-renewable resources,” *Journal of Economic Dynamics and Control*, 8(1), 73–97.
- KARP, L. (1988): “A comparison of tariffs and quotas in a strategic setting,” Giannini Foundation Working 88-6.
- KARP, L., AND D. NEWBERY (1991): “Optimal tariffs on exhaustible resources,” *Journal of International Economics*, 30(3-4), 285–299.
- LISKI, M., AND O. TAHVONEN (2004): “Can carbon tax eat OPEC’s rents?,” *Journal of Environmental Economics and Management*, 47(1), 1–12.
- LONG, N. V. (2010): *A Survey of Dynamic Games in Economics*. World Scientific, Singapore.
- NJOPMOUO, O. (2010): “On Capturing Foreign Oil Rents,” *Working Paper, University of Montreal*.

- RUBIO, S. J. (2005): “Tariff Agreements and Non-Renewable Resource International Monopolies: Prices Versus Quantities,” *Department of Economic Analysis, University of Valencia*, Discussion paper no 2005-10.
- STRAND, J. (2011): “Taxes and Caps as Climate Policy Instruments With Domestic and Imported Fuels,” *Gilbert Metcalf (ed.): U.S. Energy Tax Policy, Cambridge University Press*.
- (2013): “Strategic Climate Policy with Offsets and Incomplete Abatement: Carbon Taxes Versus Cap-and-Trade,” *Journal of Environmental Economics and Management*, forthcoming, (Published online, April 28, 2013).
- TOWER, E. (1975): “The optimum quota and retaliation,” *The Review of Economic Studies*, 42(4), 623–630.
- WIRL, F. (1995): “The exploitation of fossil fuels under the threat of global warming and carbon taxes: A dynamic game approach,” *Environmental and Resource Economics*, 5(4), 333–352.
- (2012): “Global warming: prices versus quantities from a strategic point of view,” *Journal of Environmental Economics and Management*, 64, 217–229.
- WIRL, F., AND E. DOCKNER (1995): “Leviathan governments and carbon taxes: Costs and potential benefits,” *European Economic Review*, 39, 1215–1236.

A Appendix: The calibration of d

Suppose that I believes that if it were to drop out of the market (e.g. use a prohibitive tariff or set its import quota to 0), E would subsequently behave as a monopolist with respect to R 's import demand function. In that case (assuming $f = 1, g = 0$), E would set $q = \frac{a}{2+b}$, implying that $p = \frac{a+ab}{b^2+2b}$. The single period emissions in this case is the constant $y \equiv \frac{a}{2+b} + b_1 \frac{a+ab}{b^2+2b}$ and the equation of motion is $x_{t+1} = \delta x_t + y$. If I ceases consumption when the stock reaches z , the stock n periods later, denoted x_n , equals

$$x_n = \delta^n z + y \sum_{n=0}^{n-1} \delta^n = \delta^n z + y \frac{\delta^n - 1}{\delta - 1}.$$

The present discounted value of the stream of marginal damages, when the stock reaches z , is then

$$\begin{aligned} d \sum_{t=0}^{\infty} \beta^t x_t &= d \sum_{t=0}^{\infty} (\beta \delta)^t \left(z + \frac{y}{\delta - 1} \right) - \frac{dy}{\delta - 1} \sum_{t=0}^{\infty} \beta^t \\ &= d \frac{(1 - \beta) z + \beta y}{(\beta - 1)(\delta \beta - 1)}. \end{aligned}$$

The marginal value to I of consuming the first unit is the difference between its choke price and the monopoly price, $\frac{A}{B} - \frac{a+ab}{b^2+2b}$. If it is optimal for I to cease consumption, under the belief that subsequent emissions would be y in each period, then the marginal benefit of an additional unit of production equals the present discounted value of the stream of future marginal damages,

$$\frac{A}{B} - \frac{a+ab}{b^2+2b} = d \frac{(1 - \beta) z + \beta y}{(\beta - 1)(\delta \beta - 1)}. \quad (9)$$

This expression gives d as an implicit function of z , the threshold stock above which it is optimal for I to cease consumption.

Under perfect competition, let annual production equals s . Denote N as the number of years that it would take the stock to reach z units, starting from a zero stock level, given annual emissions s : N is the solution to $z = s \frac{\delta^N - 1}{\delta - 1}$. We can use this equation to eliminate z from equation (9), resulting in an implicit expression for d as a function of the previously defined parameters and the new parameter, T . Our choice $d = 3.3043 \times 10^{-4}$ is equivalent to setting $N = 105$. In summary, our choice of d is consistent with a circumstance where it would be optimal for I to stop consuming the carbon intensive good after approximately 105 years of world consumption at the competitive level, given I 's belief that subsequent

consumption would be at the monopoly price with respect to R excess demand.¹

B Appendix: The solution to the model

We first explain how we re-write the problem in order to unify the four scenarios. This procedure enables us to solve a single game, and then obtain each of the policy scenarios by appropriate choice of parameters. We then explain how to solve the unified model.

B.1 The unified model

In all four scenarios, corresponding to the different policy mixes, we can write the single period payoffs of E and I and their “perceived” equation of motion (defined below) as

$$\begin{aligned} I\text{'s payoff: } & f_I Q^2 + g_I Qx + h_I Q + r_I x + s_I - \frac{d_I}{2} x^2 \\ & \text{Equation of motion: } x' = k_I x + m_I Q + n_I. \end{aligned} \tag{10}$$

$$\begin{aligned} E\text{'s payoff: } & f_E q^2 + g_E qx + h_E q + r_E x + s_E - \frac{d_E}{2} x^2 \\ & \text{Equation of motion: } x' = k_E x + m_E q + n_E. \end{aligned} \tag{11}$$

We intentionally abuse notation here in order to obtain a unified (for all four policy scenarios) expression of the game, so that we can use a single program to obtain the equilibrium in all four cases. We now explain the relation between equations (10) and (11).

Consider first the case where both I and E choose quantities, Q and q . In a linear MPE both agents believe that their rival uses a linear control rule. Suppressing time subscripts, I believes that E sets $q = \lambda + \mu x$ and E believes that I sets $Q = \rho + \sigma x$, where the endogenous parameters $\lambda, \mu, \rho, \sigma$ are to be determined. The beliefs are confirmed in equilibrium. That is, given I 's belief about E 's policy, I 's optimal policy is $Q = \rho + \sigma x$, and given E 's belief about I 's policy, E 's optimal policy is $q = \lambda + \mu x$.

Using the price under quotas, and I 's belief, I expects the equilibrium price to be

$$p = \frac{Q + a - q}{b} = \frac{Q + a - (\lambda + \mu x)}{b}.$$

¹As noted above, this explanation is intended to provide context for an otherwise hard-to-interpret numerical value, not to represent a plausible outcome. In particular, the calibration described here implies $z = 900$. However, world equilibrium production under the monopoly price, when I has exited the market, would be too little to sustain the stock at that level.

Using this expression and $P = \frac{A-Q}{B}$ in I 's flow payoff, equation (1), we write that payoff as a quadratic function in q and x , as in the first line of equation (10). Equating coefficients of terms of the same power (e.g., equating the coefficient of x^2 in both equations), we obtain the formulae for f_I, g_I, h_I, r_I, s_I . Similarly, given its beliefs, I 's “perceived” equation of motion (i.e., its belief about the equation of motion) is

$$\begin{aligned} x' &= \delta x + (\lambda + \mu x) + b_1 \frac{Q + a - (\lambda + \mu x)}{b} \\ &= \left(\delta + \mu - b_1 \frac{\mu}{b} \right) x + \frac{b_1}{b} Q + \lambda + b_1 \frac{a - \lambda}{b}, \end{aligned}$$

which has the same form as the second line in equation (10). Again, equating coefficients of terms of the same power, we obtain the formulae for k_I, m_I, n_I . We obtain the formulae for the coefficients in equation (11) using the same procedure.

We use the same method to obtain formulae for the coefficients of the other three control problems.

B.2 Solution to the unified model

We now work with the control problems defined by equations (10) and (11). Each agent's equilibrium control rule, $q = \lambda + \mu x$ for E and $Q = \rho + \sigma x$ for I , appears in the other agent's control problem. Consider E 's control problem. Its dynamic programming equation (DPE) is

$$\begin{aligned} W(x) &= \max_q [f_E q^2 + g_E q x + h_E q + r_E x + s_E \\ &\quad - \frac{d_E}{2} x^2 + \beta W(k_E x + m_E q + n_E)], \end{aligned} \tag{12}$$

where the second line uses the second line in equation (11) to write $W(x')$ as a function of the current x and the current choice q . Because of our choice of a linear equilibrium, E solves a linear quadratic control problem, for which it is well known that the unique solution is a quadratic value function.² We write this function as $W(x) = \epsilon + \nu x + \frac{\phi}{2} x^2$, where the parameters ϵ, ν, ϕ are to be determined. Using this function to eliminate $W(x')$ on the right side of equation (12), we express the right side as a linear quadratic function of q, x and the unknown coefficients. We maximize this expression with respect to q to obtain the

²There are non-linear equilibria in this model. However, the linear equilibrium is an obvious choice to study, because it is the limit of the sequence of equilibria in the finite horizon game, as the time horizon goes to infinity. It is also the only equilibrium that is defined for all state space, ignoring inequality constraints.

The linearity of the equilibrium requires that inequality constraints, e.g. the non-negativity of prices and quantities, are never binding in equilibrium. For our calibration, these inequality constraints are in fact never binding for values of the state variable reached in equilibrium.

coefficients of E 's control rule $Q = \lambda + \mu x$:

$$\begin{aligned}\lambda &= -\frac{h_E + \beta \nu m_E + 2\beta \phi m_E n_E}{(2f_E + \beta \phi m_E^2)} \\ \mu &= -\frac{g_E + \beta \phi m_E k_E}{(2f_E + \beta \phi m_E^2)}.\end{aligned}\tag{13}$$

The maximized value of the right side of the DPE (12) is a quadratic function in x , as is the left side. The DPE holds identically in x if and only if the coefficients of terms of order of x are equal. We define

$$\Delta_E = (2f_E + 2g_E \beta m_E k_E + d_E \beta m_E^2 - 2\beta f_E k_E^2)^2 - 4\beta m_E^2 (g_E + 2d_E f_E) \tag{14}$$

and equate coefficients of terms of order of x on the two sides of the maximized DPE to obtain the following formula for the unknown parameters³

$$\begin{aligned}\phi &= \frac{1}{2\beta m_E^2} \left(-(2f_E + 2g_E \beta m_E k_E + d_E \beta m_E^2 - 2\beta f_E k_E^2) - \Delta_E \right) \\ \nu &= -\frac{h_E \beta \phi m_E k_E + g_E \beta \phi m_E n_E - 2\beta \phi f_E n_E k_E + g_E h_E - r_E (2f_E + \beta \phi m_E^2)}{2f_E + \beta \phi m_E^2 - 2\beta f_E k_E + g_E \beta m_E} \\ \epsilon &= \frac{1}{2} \frac{-2\beta \phi f_E n_E^2 + h_E^2 + 2h_E \beta \nu m_E + 2h_E \beta \phi m_E n_E + \beta^2 \nu^2 m_E^2 - 4\beta \nu n_E - 2s_E (2f_E + \beta \phi m_E^2)}{(2f_E + \beta \phi m_E^2)(\beta - 1)}.\end{aligned}\tag{15}$$

The importer I solves a similar control problem, where its single period payoff is the first line of equation (10) and its perceived equation of motion is the second line of that equation. Denoting I 's value function as $V(x)$, we write its DPE as

$$\begin{aligned}V(x) &= \max_Q [f_I Q^2 + g_I Qx + h_I Q + r_I x + s_I \\ &\quad - \frac{d_I}{2} x^2 + \beta V(k_I x + m_I Q + n_I)]\end{aligned}\tag{16}$$

Equations (16) has the same form as the exporter's DPE (12), except that the subscript I replaces the subscript E on parameter coefficients, the function V replaces W , and the control Q replaces q . Denote the quadratic value function as $V(x) = \chi + \psi x + \frac{\omega}{2} x^2$. Substituting

³The equations for ϕ and for ω are quadratics. For both of these equations we take the smaller root, leading to the first line of equation (15). The smaller root satisfies the transversality condition. In addition, when we repeat this procedure for the importer, the smaller root is the only negative root. The coefficient of x^2 in the importer's value function must be negative, as discussed in the text.

We confirmed that the choice of the smaller root for both quadratics is correct by solving these equations for the other three combinations of roots. For two of these combinations, there was no equilibrium candidate because there was no solution to the two equations given by the two roots. For the third combination, there was a solution to these two equations, but it resulted in negative stocks, and thus violates the requirement that stocks be non-negative.

this function into the DPE (16) we repeat the procedure above to obtain expressions for the endogenous parameters $\chi, \psi, \omega, \sigma, \rho$. These formulae are identical to those in equations (13) and (15), except that the subscript I replaces the subscript E , and the parameters $\chi, \psi, \omega, \sigma, \rho$ replace the parameters $\epsilon, \nu, \phi, \lambda, \mu$; we also define a function Δ_I using an equation analogous to (14).

The system consisting of (13) and (15) and the definition (14), together with the corresponding equations (not shown) for I can be solved recursively. We first solve the four equations that determine $\omega, \phi, \sigma, \mu$. This four dimensional system can be reduced to a two-dimensional system by noting that for all policy scenarios, g_E is a linear function of σ , and g_I is a linear function of μ . The second line of equation (13) shows that μ is a linear function of g_E , and hence a linear function of σ . Inspection of the analogous equation for I (not shown), shows that σ is a linear function of μ . We can solve this two dimensional linear system to obtain values of σ and μ as functions of ω and ϕ . Substituting these expressions into the equations that determine ω and ϕ (the first line of equation (15) for ω and the corresponding equation – not shown – for ϕ), we obtain two cubics in ω and ϕ . We can numerically solve these two cubics to find the correct values of ω and ϕ .

Given the values of ω and ϕ , we can then obtain σ and μ using the expressions described in the previous paragraph. With numerical values for $\omega, \phi, \sigma, \mu$, we then use the equations for λ and ν and the corresponding equations (not shown) for ρ and ψ to solve for these four parameters; this system is linear. We then solve the decoupled equations for τ and χ (again, the equation for χ is not shown).

We also need an expression for the present discounted value of the stream of R 's payoff. Equation (2) gives R 's single period payoff. Denote $p = \mu_R x + \lambda_R$ and $Q = \sigma_R x + \rho_R$ as the equilibrium values of p and Q . The parameters of these functions depend on the particular policy scenario, and their values are obtained from the solution to the different games. R 's flow payoff depends on p , which in equilibrium is a function of x , and the evolution of x depends on both p and Q , via equation (8). R 's continuation payoff is therefore a function of x , which we denote $Y(x)$. The value of the stream of R 's payoff equals its flow payoff plus its discounted continuation payoff. Therefore, $Y(x)$ must satisfy the functional equation

$$Y(x) = \frac{1}{2} \frac{(a - bp)^2}{b} + \beta Y(\delta x + Q + b_1 p). \quad (17)$$

Substituting the quadratic trial solution, $Y(x) = \frac{\gamma}{2} x^2 + \eta x + \varsigma$, into equation (17) and equating coefficients of terms in order of x provides the equations for the parameters of R 's

value function:

$$\begin{aligned}\gamma &= \frac{-b\mu_R^2}{\beta\delta^2+2\beta\delta\sigma_R+\beta\sigma_R^2+2\beta\delta\mu_Rb_1+2\beta\sigma_R\mu_R+\beta\mu_R^2b_1^2-1} \\ \eta &= \frac{-a\mu_R+b\mu_R\lambda_R+\beta\delta\gamma\rho_R+\beta\gamma\rho_R\sigma_R+\beta\gamma\mu_R\rho_Rb_1+\beta\gamma\sigma_R\lambda_Rb_1+\beta\gamma\mu_R\lambda_Rb_1^2}{(1-\beta\delta-\beta\sigma_R-\beta\sigma_R^2-\beta\mu_Rb_1)} \\ \zeta &= \frac{-a\lambda_R+\frac{1}{2}b\lambda_R+\eta\beta\rho_R+\frac{1}{2}\beta\gamma\rho_R^2+\eta\beta\lambda_Rb_1+\beta\gamma\lambda_Rb_1\rho_R+\frac{1}{2}\beta\gamma\lambda_R^2b_1^2}{1-\beta}.\end{aligned}$$

C Appendix: Calculation of a Pigouvian Tax

As in the text, the world price, defined as the price that E receives, is p . Consumers in I pay an additional Pigou Tax (Υ) added to the price: $p + \Upsilon$ and consumers in R face the same price.

Country I has no domestic production; its demand for imports equals $A - B(p + \Upsilon)$. The climate-related damages, conditional on x , is $\frac{d}{2}x^2$ where d is a constant. I 's single period payoff equals consumer surplus minus environmental damages:

$$\begin{aligned} I\text{'s flow payoff: } & \int_{p+\Upsilon}^{\frac{A}{B}} (A - Bz) dz - \frac{d}{2}x^2 \\ & = \frac{1}{2} \frac{(A - B(p + \Upsilon))^2}{B} - \frac{d}{2}x^2. \end{aligned} \quad (18)$$

At price $p + \Upsilon$, R 's domestic demand is $a - b_0(p + \Upsilon)$ and its domestic supply is b_1p , so its net imports equal $a - bp - b_0\Upsilon$, with $b_0 + b_1 \equiv b$. R 's gains from trade minus its climate related damages $\frac{\kappa}{2}x^2$ equal its flow payoff:

$$R\text{'s flow payoff: } \int_{p+\Upsilon}^{\frac{a}{b_0}} (a - b_0z) dz + \int_0^p (b_1z) dz = \frac{1}{2} \frac{(a - b_0(p + \Upsilon))^2}{b_0} + \frac{b_1p^2}{2} - \frac{\kappa}{2}x^2. \quad (19)$$

The exporter, E , has no domestic consumption. These producers' marginal cost function, equal to E 's supply function, is $f + gp$, where f and g are constants. The exporter's single period payoff equals its domestic profits

$$E\text{'s flow payoff: } \int_0^p (f + gz) dz. \quad (20)$$

Each agent has the same constant discount factor, β . Welfare for each agent equals the discounted stream of their single period payoff.

The social planner maximizes the sum of the payoffs plus rents collected through the tax.

$$\begin{aligned} \text{social payoff: } & \frac{1}{2} \frac{(A - B(p + \Upsilon))^2}{B} - \frac{d}{2}x^2 + \frac{1}{2} \frac{(a - b_0(p + \Upsilon))^2}{b_0} + \frac{b_1p^2}{2} - \frac{\kappa}{2}x^2 + fp + \frac{1}{2}gp^2 + \Upsilon(f + gp + b_1p) \end{aligned} \quad (21)$$

We can write the total demand equal to total supply (to get p in terms of Υ) and the

“perceived” equation of motion (defined below) as

$$\text{Equating Supply with Demand: } f + gp + b_1p = a - b_0(p + \Upsilon) + A - B(p + \Upsilon)$$

$$\text{which results in } p = \frac{a - b_0\Upsilon + A - B\Upsilon - f}{g + b_1 + b_0 + B} \quad (22)$$

$$\text{Equation of motion: } x' = \delta x + f + gp + b_1p$$

$$\text{which results in: } x' = \delta x + f + g\left(\frac{a - b_0\Upsilon + A - B\Upsilon - f}{g + b_1 + b_0 + B}\right) + b_1\left(\frac{a - b_0\Upsilon + A - B\Upsilon - f}{g + b_1 + b_0 + B}\right)$$

The social planner will choose a tax Υ which in equilibrium is a linear function of the state, $\Upsilon = \lambda + \mu x$. The social planner solves the following optimization problem

$$\begin{aligned} S(x) = \max_{\Upsilon} & \left[\frac{1}{2} \frac{(A - B(p + \Upsilon))^2}{B} + \frac{1}{2} \frac{(a - b_0(p + \Upsilon))^2}{b_0} + \frac{b_1 p^2}{2} + fp + \frac{1}{2} gp^2 + \Upsilon(f + gp + b_1p) \right. \\ & \left. + \left(\frac{-\kappa - d}{2}\right) x^2 + \beta S(\delta x + f + gp + b_1p) \right], \\ & s.t. \\ & p = \frac{a - b_0\Upsilon + A - B\Upsilon - f}{g + b_1 + b_0 + B} \end{aligned} \quad (23)$$

where the second line uses the equation of motion to write $S(x')$ as a function of the current x and the current choice Υ . The social planner solves a linear quadratic control problem, for which it is well known that the unique solution is a quadratic value function. We write this function as $S(x) = \epsilon + \nu x + \frac{\phi}{2} x^2$, where the parameters ϵ, ν, ϕ are to be determined. Using this function to eliminate $S(x')$ on the right side of equation (12), we express the right side as a linear quadratic function of Υ, x and the unknown coefficients. We maximize this expression with respect to Υ to obtain the coefficients of the control rule $\Upsilon = \lambda + \mu x$:

The maximized value of the right side of the DPE (12) is a quadratic function in x , as is the left side. The DPE holds identically in x if and only if the coefficients of terms of order of x are equal. We equate coefficients of terms of order of x on the two sides of the maximized DPE to obtain the unknown coefficients. Hence, we obtain $\Upsilon = \lambda + \mu x$, the optimal Pigou tax as determined by the social planner.