Department of Economics and Finance

University of Guelph

Discussion Paper 2016-10

Forming a Majority Coalition for Carbon Taxes Under a State-Contingent Updating Rule

By:
Ross McKitrick and Jamie Lee

Forming a Majority Coalition for Carbon Taxes Under a State-Contingent Updating Rule

Ross McKitrick[1,*] and Jamie Lee[1]

University of Guelph
Department of Economics and Finance
Discussion Paper 2016-10

Abstract

Uncertainty and political polarization over global warming make it difficult to achieve a stable majority coalition supporting carbon taxes, especially since expectations about the future optimal values sharply diverge. We present an alternative approach in which the tax path is not announced in advance but is set to track observed future temperatures. Agents thus form expectations which imply the tax path will be correlated with their preferred price trajectory. Whereas greater variance in beliefs about future global warming undermines support for a compromise policy, the state-contingent proposal attracts majority support irrespective of the divergence of views, and even has robustness properties to strategic voting by dishonest agents.

[1] Department of Economics and Finance, University of Guelph. *corresponding author (ross.mckitrick@uoguelph.ca).

JEL Codes: Q54, Q58, H23, D72

Keywords: Carbon tax, State-contingent model, Majority voting, Climate change, Uncertainty.

The authors acknowledge with gratitude funding support from the Institute for New Economic Thinking.

1 Introduction

When damages due to pollutant emissions are observable and can be traced with certainty to specific emitting activity, use of an emissions tax can yield efficient incentives for abatement activity. But some externalities, such as global warming from greenhouse gas emissions, have complex intertemporal features that make it difficult to identify an optimal shadow price. The emissions (carbon dioxide or CO_2 in the case of global warming) do not directly affect welfare, instead they affect an environmental state variable *s*, normally thought of as some measure of atmospheric temperatures. Changes in *s* then give rise to damages, but *s* is also subject to natural variability, making it difficult to identify the effect due to emissions. Not only is the damage function uncertain, but the function relating emissions to the state variable is also unknown and must be estimated with considerable uncertainty. The effects of current emissions may also operate over a long time lag, the length of which is itself unknown, which also implies that the current state exhibits the effects of historical emissions to an unknown lag. These features lead to at least two major difficulties for devising policy responses.

First, it is effectively impossible to know whether a sequence of future emission tax rates derived from a computer model are optimal or not. In the case of CO_2, the analysis is typically done using Integrated Assessment Models (IAMs) which embed myriad assumptions about the many uncertain parameters describing the climate and economic systems (e.g. William Nordhaus, 2007, Robert Pindyck 2013, IWG 2010). Confidence intervals around key parameters are so wide as to yield an arbitrarily large range of marginal damage calculations. For instance, the IPCC (2007) Synthesis Report stated that peer-reviewed estimates of the social cost of carbon ranged from $3 to $95 per ton of CO2 emissions, and other recent surveys show even wider ranges (e.g. Richard Tol 2007; Mikhail Golosov et al., 2014, IWG 2013). These differences trace to divergent assumptions about climate sensitivity, response lags, discount factors etc. Ongoing attempts to tackle the issue through the use of IAMs require imposition of functional forms and parameter values that effectively assume away much of what makes the problem difficult in the first place (Pindyck 2013). Incorporating Bayesian learning into a model may correct the initial parameter values over time. But the learning routine for even two unknown variables can take millennia to reach a 5% critical value, thus making it irrelevant for climate policy (David Kelly and Charles Kolstad, 1999; Andrew Leach, 2007).

Second, the uncertainties about the effects of emissions and marginal damages imply individuals will form divergent preferences over the optimal policy response. If a tax instrument must receive majority support in a voting system, and people can vote against it because they perceive it either to be too high or too low, the formation of a majority coalition to support implementation can be effectively impossible. Illustrations of this are the recent rejection of a carbon tax proposal in Switzerland[1] and the repeal of the Australian carbon tax.[2] The dynamics of the problem interact with the political aspect, since voters know that the starting value will be subject to some form of adjustment over time. Someone who thinks the tax

[1] See http://www.wsj.com/articles/swiss-voters-reject-initiative-to-replace-vat-system-with-carbon-tax-1425822327.

[2] See http://www.news.com.au/national/australias-carbon-tax-has-been-axed-as-repeal-bills-clear-the-senate/story-fncynjr2-1226991948152.

initially too low may nonetheless support it if they expect it to rise, and vice versa. But someone who thinks the initial value acceptable may oppose it if they do not believe it will rise as quickly as it should in the future.

The purpose of this paper is to examine how a proposed solution to the first problem also addresses the second. Ross McKitrick (2010, herein M10) noted that while the function relating emissions to the climate state may be unknown, the state itself, s, is observable, and contains information that can be used to circumvent some of the major uncertainties in the computation of the optimal tax path. Specifically M10 suggested using observations on s to calibrate a dynamic pricing rule that, under certain assumptions, will closely approximate the unobservable optimal tax path. We extend this reasoning to show why voters would be more likely to support a tax based on this rule than one based on a pre-announced path.

In the M10 set-up, the value of s at time t (i.e. the current global average temperature) is a function of current and past emissions:

$$s_t = s(e_t, e_{t-1}, e_{t-2}, \dots, e_{t-k}) \quad (1)$$

out to lag length k, where k may be unknown. We assume that there are Q infinitesimally small emitters, so in any period t, total emissions are

$$e_t = \Sigma_{j=1}^{Q} e_t^j.$$

Damages at time t are a function of the state variable, i.e. $D(s_t)$, implying that the current value $V(t)$ of the externality is the discounted present value of damages from the present (time 0) out to the distant future at time T:

$$V(0) = \sum_{m=0}^{T} \beta^j D(s_{t+m}) \tag{2}$$

where β is the discount factor and T is the policy planning horizon. The socially optimal price $\tau^*(t)$ on emissions at time t is the change in the value of the externality as a result of marginal current emissions:

$$\tau^*(t) = \frac{\partial V(t)}{\partial e_t}. \tag{3}$$

A policy plan would consist of an announced sequence of current and future tax rates $\tau(t), \dots, \tau(t+T)$. Any attempt to derive such a path at time t would run into the computational problems noted above, and any attempt to secure majority agreement to implement such a path would run into the difficulty that most observers would expect the price path to be higher or lower than the one they would prefer based on their beliefs about the severity of the problem. As an alternative approach, M10 proposed a tax path that begins with an announced rate at time zero, then rather than the entire future path being announced at the same time, a rule is announced by which the rate will be updated in real time:

$$\tau_t^s = \gamma s_t \frac{e_t}{\bar{e}_t} \tag{4}$$

where $\bar{e}_t$ is a moving average of past emissions over the regulator's best estimate of k periods and γ is a parameter that must be chosen to determine an initial value of the tax sequence. The state-contingent mechanism is a direct analogue to monetary policy approaches based on committing to updating rules rather than interest rate paths determined long in advance. M10 showed that, over time, the tax path described by Equation (4) will be highly correlated with the unobservable optimal path based on equation (3) that would have been implemented if the planner had enough information to compute it. To implement (4) the regulator only needs to determine a value for γ and then use other information available contemporaneously with the rate revisions, but as a result, agents will not know its future levels, since they will rise or fall in step with s_t.[3] They will therefore have to make decisions based on expectations: those who expect rapid global warming, for instance, will expect τ_t^s to increase rapidly, whereas those who expect little warming will expect it to remain relatively unchanged from its initial value.

Shi-Ling Hsu (2011) and McKitrick (2011) also proposed supplementing the implementation of (4) with a futures market for certificates dated t, … , t+T defined so that each one exempts the holder from paying the tax on one tonne of CO_2 in the year indicated. In order to price such certificates, market participants would need to form expectations about the future path of s_t, and therefore of τ_t^s. The existence of such a market would allow agents to hedge against future policy costs, thus providing complete pricing certainty, and would also provide a visible indicator of

[3] Emissions in period t must also be known. As a practical matter the form in which they enter equation (4), as a ratio with the historical moving average could either be replaced with an estimate, or a 1-period lag.

the market's dominant forecast of the path of future temperatures. Experts whose beliefs about global warming deviate from the market consensus could then make investments based on their private information set, depending on how much confidence they have in their views. If their declared deviation from the consensus is merely ideological, for instance if they oppose the emissions tax for political reasons while privately believing the warming will occur, they will have no incentive to bet against the market.

There remains the second problem, namely whether this instrument can obtain majority support for implementation. We present herein a simple model of voting behaviour and show that when agents disagree about the climate issue, a conventional emissions tax based on political compromise will be less likely to get majority support, the higher is the variance of beliefs about the underlying issue. The problem is compounded by the fact that the tax rate is expected to change over time, so even if a majority can agree on a starting value of the tax, as beliefs about future global warming diverge the coalition will tend to break down, affecting the initial adoption decision. We show herein that the state-contingent approach addresses this problem: if the initial range of beliefs is sufficiently constrained that a starting value of the tax can obtain majority support, the future value of the tax will as well, even if a conventional approach would have failed.

We first show this in a case in which voters are honest, in the sense that they only care about implementing the socially optimal tax rates, but they differ in their beliefs about how the state variable will evolve over time. We then allow voters to be dishonest, such that they declare a preference for low or high tax rates irrespective of their actual beliefs about marginal damages, and we examine the incentives to support or reject the tax mechanism against an alternative in which the regulator implements a compromise tax rate. We find that the state-contingent rule may still obtain majority support, though it is not assured. If there is a combination of (polarized) dishonest as well as honest voters, as long as no one group has an outright majority it will be possible to form a majority coalition in support of the state-contingent tax.

The remainder of the paper is organized as follows. Section 2 briefly discusses earlier studies on voting mechanisms for public goods and externalities. Sections 3 and 4 develop the theoretical structure of our model and provides propositions and their proofs. Lastly, Section 5 presents conclusions.

2 Voting on taxes for externalities and public goods

Numerous authors have examined the way in which voting systems influence the adoption or rejection of proposed taxes. Experimental results of Simon Dresner et al (2006) show that the success of adopting a new tax policy depends on how well the voters understand the proposed policy. For example, some voters may not support a pollution tax because they do not fully understand how it is used to enhance efficiency. Similarly, Peter Clinch et al (2006) conclude that public trust in the government plays a key role in determining the support for new taxes. Several natural field experiments have shown that framing affects voting behavior. For instance, according to Edward McCaffery and Jonathan Baron (2003), some people may react negatively even to the use of the word “tax”. On the other hand, Rupert

Sausgruber and Jean-Robert Tyran (2005) showed experimentally that some people prefer indirect over direct taxes, an effect they call "fiscal illusion."

An earlier, related literature examined positive externalities such as publicly funded education. John Creedy and Patrick Francois (1990) showed that if education provides a positive externality to the economy by inducing economic growth, and if only certain (high) type of individuals can benefit from education, then under certain conditions a majority of uneducated individuals would be willing to pay taxes to subsidize education for high types in return for (higher) economic growth. Johnson (1984) draws the same conclusion, however, his model does not incorporate opportunity cost of education in terms of forgone wages.

Alberto Alesina and Francesco Passarelli (2013) analyze majority voting outcomes when the government has three environmental policy tools: a rule, which is an instrument that sets an upper limit to the activity; a *quota* that requires a proportional reduction of the activity; and an emissions tax. They show that majority voting may not yield a socially optimal outcome when there are several policy options and voters have divergent preferences. If the group responsible for the externality is in the minority, then the majority will choose a policy that puts the greatest compliance cost burden on the minority group, and vice versa. These results are in line with those of Friedrich Schneider and Juergen Volkert (1999) who show that when the voting community is composed of groups with differentiated interests, the voting outcome may not be socially optimal.

Per Fredriksson and Thomas Sterner (2005) incorporate differences in abatement technologies across firms and show that "clean" firms may lobby for higher tax rates if the revenue is used for rebates. Shinya Kawahara (2011) builds a model with assumptions that voters do not observe politician types and environmental damage. Under this model, a pooling equilibrium results in a sub-optimal tax rate, whereas in the separating equilibrium, pro-environmental politicians choose a tax rate that is too high in order to distinguish themselves from other types. Lastly, Helmuth Cremer et al. (2004) examine revenue recycling and voting outcomes. They show that if environmental tax revenue is used to subsidize income and capital taxes, then the majority will choose an environmental tax that is too low.

Overall, the literature finds that the voting outcome depends not only on preferences of the voters but also on the perceived distribution of expected benefits (and/or costs). Our analysis herein shows similar effects, but the potential popularity of the instrument is increased by replacing a specific future tax proposal with one that correlates with individuals' own expectations of what it ought to be.

3 Voting on a Static Carbon Tax

The voting environment is as follows. There are N voters indexed by $i = \{1,\ldots,N\}$ and one policy maker who proposes an emissions tax path τ_t at time zero. Each voter chooses either to support or oppose the tax policy. The proposal is implemented only if it obtains majority support. In any period, each voter's loss depends on the squared distance between the proposed tax rate and his or her privately-held belief about the ideal tax rate $\tilde{\tau}_{it}$. Throughout this analysis, the tilde always denotes an agent's privately-preferred tax rate. For now we assume that each agent wants the tax to be set at his or her estimate of marginal damages, which we denote $V_{it}' = E(V_t'|\Omega_{i0})$ where Ω_{i0} is person *i*'s information set at the time of the vote. We assume that everyone holds the same beliefs about the form of the damage function D but people have differing beliefs about how emissions will affect the future path of the state variable s_t. As noted in the US Interagency Working Group on the Social Cost of Carbon, the climate sensitivity to greenhouse gas emissions is a key parameter for determining the behaviour of IAMs and hence estimates of the marginal social cost of CO_2 emissions (see also Kevin Dayaratna et al. 2016).

Denote individual *i*'s loss in period t when the proposed tax rate is τ_t as

$$L_{it} = (\tilde{\tau}_{it} - \tau_t)^2 \,. \qquad (5)$$

A voter preferring $\tilde{\tau}_{it}$ will support the proposed tax if and only if the loss is less than or equal to a cut-off value d. We assume that d is the same for all agents. The following must therefore hold for a voter who supports the tax:

$$\tilde{\tau}_{it} \in \tau_t \pm \sqrt{d}.$$

A higher value of d thus increases the propensity to vote yes in each period.

The policy maker's objective is to choose a tax path that minimizes the summed losses of voters in each period, $L_t = \Sigma_{i=1}^{N} L_{it}$. We leave aside for now the problem that the regulator does not know the true values of $\tilde{\tau}_i$.[4] The first order condition of the single-period minimization problem implies that the tax should equal the mean of the preferred tax rates:

$$\bar{\tau}_{\mathrm{t}} = \frac{1}{N}\,\Sigma\tilde{\tau}_{it} \qquad (6)$$

Total losses will then be

$$L_t = (N-1)\sigma_t^2$$

where σ_t^2 is the variance of preferred tax rates among all voters in at time *t*. It is clear that total losses are increasing in the variance of beliefs. Hence for a given number of voters, when uncertainty grows and the variance of the preferred tax rate rises, we expect the mean loss to go up, implying an increased probability that the median voter will reject the tax.

This can be shown more formally as follows. If a majority consists of 50 percent of voters, we expect the tax will pass in a period if the probability of person *i* voting "No" is less than or equal to 0.5, so $P\left(|\tilde{\tau}_{it} - \bar{\tau}_t| > \sqrt{d}\right) \leq 0.5 \Rightarrow$Expected Majority Yes vote. Note that by Chebychev's inequality (John Rice 1988), as long as the distribution of $\tilde{\tau}_i$ has a finite second moment,

[4] Under the assumption that voters want the externality priced at marginal damages, this is no different, in principle, than any non-market elicitation problem which would require use of a technique like contingent valuation.

$$P\left(|\tilde{\tau}_{it} - \bar{\tau}_t| > \sqrt{d}\right) \leq \frac{\sigma_t^2}{d}.$$

Hence a sufficient condition for a majority to support the tax is $\sigma_t^2 \leq 0.5d$. For a given value of d, greater variance in the preferred tax rates reduces the likelihood of the vote passing.[5]

In sum, for a given value of d, in every period, the higher the variance of preferred tax rates, the less likely it is that a proposed tax rate will be supported by a majority of voters, even if the proposal is the optimal compromise computed by a planner who knows all the privately-preferred tax rate levels. A sufficient condition for a majority to support the tax is that the loss cut-off d is greater than or equal twice the variance of beliefs about the optimal tax.

In the next section we look at a simple intertemporal version of the problem and contrast the compromise (loss-minimizing) approach with the state-contingent approach, showing that as expectations about the future state diverge the compromise tax may eventually fail to obtain support, but the state-contingent tax is robust to this problem.

3 Two-Period Case Under a State-Contingent Tax

3.1 Pre-Announced Path

[5] If we know that the preferred tax rates are normally distributed we can weaken this condition somewhat, since in that case at least 50% of the distribution will vote yes if $0.68\sqrt{\sigma_t^2} \leq \sqrt{d}$ or $\sigma_t^2 \leq 2.13d$. However, it is unlikely the $\tilde{\tau}_{it}$'s are symmetric and Normally distributed, even if we allow them to take negative values.

Preferences over the optimal tax rate diverge because CO_2 emissions do not cause direct harm but have an indirect effect through an uncertain and potentially long term influence on the climate state, leading to disagreement about the ultimate severity of the externality. The most recent IPCC report (IPCC 2014, Figure SPM.7) shows computer projections of global warming over the coming century spanning half a degree in the current decade, diverging to a span of over 5 degrees by 2100, and of course individuals may privately have an even wider range of expectations. Hence while it might be possible to get a majority to support a compromise emissions tax rate now ($\bar{\tau}_0$), if the policy package also contains a commitment to a specific future values of the tax, the same voters may reject the policy on the grounds that the future rate is too high or too low, because the variance of preferred rates grows over time. In this section we consider this scenario, showing that the state-contingent approach is robust to the problem.

Dynamics are kept simple by reducing the planning horizon to two-periods. Modifying the notation from before, in period 0, the policymaker proposes a loss-minimizing tax rate $\bar{\tau}_0$ which will change to a preannounced value of $\bar{\tau}_1$ in period 1. Individual *i* has preferred tax rates $\tilde{\tau}_{i0}$ and $\tilde{\tau}_{i1}$ in the two periods, respectively, and the associated variances of preferred rates across *N* voters are denoted σ_0^2 and σ_1^2.

The set of preferred values each period is determined by $\tilde{\tau}_{i0} = V_{i0}' = E(V_0'|\Omega_{i0})$ and $\tilde{\tau}_{i1} = V_{i1}' = E(V_1'|\Omega_{i0})$. M10 shows that a reasonable approximation to the unobservable true value $V_t{}'$ is given by equation (4), hence the state-contingent tax is written as $\tau_t^s = \gamma s_t r_t$ where $r_t = e_t/\bar{e}_t$. This also implies the preferred tax rates at time 0 are (approximately) $\tilde{\tau}_{i0} = \gamma_i s_0 r_0$ where γ_i is agent i's preferred value of the scaling parameter, and the future preferred tax rate is

$$\tilde{\tau}_{i1} = \gamma_i E(s_1 r_1 | \Omega_{i0}). \tag{7}$$

M10 shows that γ is fully determined by the parameters of the damage function and the degree of homogeneity of the state function s_t. Since s_0 and r_0 (and all past values) are observable, this implies that γ is, in principle, identifiable even if future expectations of s_t differ across agents. The loss-minimizing solution thus yields $\gamma = s_0 r_0 / \bar{\tau}_0$. We assume henceforth that this value of γ is used by all agents. The loss-minimizing method implies that the policy maker should then announce the future tax rate as

$$\bar{\tau}_1 = \gamma \Sigma_i \tilde{\tau}_{i1} / N. \tag{8}$$

This requires elicitation of the $\tilde{\tau}_{i1}$'s, which are today's estimates of the preferred future tax rates, based on each individual's expectation of, among other things, how s_t will change. We denote the change between the periods in the preferred tax rates as $\Delta\tilde{\tau}_i$ and the change in the proposed tax rates as $\Delta\bar{\tau}$, and denote the difference between these as $\delta_i = \Delta\tilde{\tau}_i - \Delta\bar{\tau}$. Also denote $\varepsilon_{i0} = \tilde{\tau}_{i0} - \bar{\tau}_0$.

The variance of period 1 preferred rates from the perspective of time 0 is:

$$\sigma_1^2 = \left(\frac{1}{N-1}\right) \Sigma_i (\tilde{\tau}_{i1} - \bar{\tau}_1)^2$$

$$= \left(\frac{1}{N-1}\right) \Sigma_i (\tilde{\tau}_{i0} + \Delta\tilde{\tau}_i - \bar{\tau}_0 - \Delta\bar{\tau})^2$$

$$= \left(\frac{1}{N-1}\right) \Sigma_i (\varepsilon_{i0} + \delta_i)^2$$

$$= \left(\frac{1}{N-1}\right) \Sigma_i (\varepsilon_{i0}^2 + 2\varepsilon_{i0}\delta_i + \delta_i^2)$$

$$= \sigma_0^2 + \sigma_\delta^2 + 2\Sigma_i(\varepsilon_{i0}\delta_i)/(N-1)$$

where $\sigma_\delta^2 = (\Sigma_i \delta_i^2)/(N-1)$. The third term is twice the covariance between $\tilde{\tau}_{i0}$ and $\Delta\tilde{\tau}_i$. It may be zero if people perceive no connection between the current level of marginal damages and the likely growth rate in the future, and it might be positive if those who believe the marginal damages are currently high also anticipate relatively faster growth in the future. But it is unlikely to be negative. We do not expect those who believe marginal damages due to CO_2 to be low at present also tend to believe grow rapidly in the future, or vice versa. Hence the last two terms are positive and therefore $\sigma_1^2 > \sigma_0^2$.

In the two-period case we assume that the proposed policy is the pair $(\bar{\tau}_0, \bar{\tau}_1)$. Assume that $\sigma_0^2 < 0.5d$ so the period 0 tax rate on its own would pass a referendum, but that σ_1^2 grows sufficiently large that the period 1 tax $\bar{\tau}_1$ fails to get majority support at time 0, and therefore the entire proposal fails. This is a plausible representation of the climate case, since beliefs about the future path of warming, and thus preferences over preferred carbon tax rates, diverge sharply as the forecast horizon increases.

3.2 State-Contingent Path

The voting problem arises because the policy maker announces a specific value for $\bar{\tau}_1$ in period 0. Even if $\bar{\tau}_1$ is the optimal compromise, the divergence of views implies higher variance of $\tilde{\tau}_{i1}$, eventually making it impossible to hold a majority. The state-contingent mechanism gets around this by not specifying $\bar{\tau}_1$ directly. Instead the policy maker announces a rule that will be used to calculate it in period 1, which forces agents to form an expectation about the likely future tax rate and compare it against their preferred future tax rate. Specifically, the policymaker announces that, instead of applying equation (8), the future tax rate will be

$$\tau_1^s = \gamma s_1 r_1 \tag{9}$$

and that the actual rate will only be announced at the start of period 1 when s_1 and r_1 are known (or, for the latter, can be estimated). Each agent thus forms the expectation

$$E_i(\tau_1^s) = \gamma E(s_1 r_1 | \Omega_{i0}). \tag{10}$$

We then obtain the following result:

Proposition 1. If $\sigma_0^2 < 0.5d$ so the state-contingent tax (4) obtains majority support in period 0, it will in period 1 as well, regardless of the divergence of views on the evolution of s_t or the preferred tax rate over time.

Proof. From above we have $\sigma_1^2 = \sigma_0^2 + \sigma_\delta^2 + 2\Sigma_i(\varepsilon_{i0}\delta_i)/(N-1)$ and $\sigma_\delta^2 = (\Sigma_i\delta_i^2)/(N-1)$, where, in this case, $\delta_i = \Delta\tilde{\tau}_i - E(\Delta\tau^s)$. Using γ in equation (7), $\Delta\tilde{\tau}_i = \gamma(E(s_1 r_1|\Omega_{i0}) - r_0 s_0)$. Also $E(\Delta\tau^s) = \gamma(E(s_1 r_1|\Omega_{i0}) - r_0 s_0)$. Hence $\delta_i = 0$ for all *i*. It immediately follows that $\sigma_1^2 = \sigma_0^2 < 0.5d$ and the policy gets majority support. ∎

Proposition 1 works because everyone expects their preferred tax rate to be implemented in the future. This is a result of the tax rule (4) providing a correlated approximation to the unobservable socially optimal tax rate, so as long as a voter wants to see the socially optimal tax rate implemented at every time *t*, it does not matter whether that voter believes the state variable will increase or not over time. Those who believe in a rapid increase in s_t would both prefer and anticipate a steep increase in the tax rates, while those who believe s_t will not go up would expect and prefer low tax rates.

4. Dishonest Voting on a State-Contingent Tax

We now consider a case in which some voters do not actually care about correctly pricing the externality, but adopt one of two extreme views based on political or other exogenous considerations. There are n_1 voters in the first group who oppose the emissions tax under all circumstances, so we denote this as $\tilde{\tau}_t^1 = 0$. There are n_2 voters in the second group who prefer an emissions tax set to some maximum feasible level w under all circumstances, so we denote this as $\tilde{\tau}_t^2 = w$. If the government chooses a static compromise policy they will set the tax equal to the mean preferred rate, which in this case is:

$$\bar{\tau} = n_2 w / N \tag{11}$$

where $N = n_1 + n_2$. We denote the state-contingent option as $\tau^s = V'$, dropping the time subscripts for convenience but without losing generality.

If the compromise tax is implemented, from equation (5) the total losses for the first group will be $n_1 \bar{\tau}^2$ and the total losses for the second group will be $n_2(\bar{\tau} - w)^2$. Under the state-contingent option the total losses for the first group will be $n_1(V')^2$ and the total losses for the second group will be $n_2(V' - w)^2$.

Suppose the compromise option is currently in place and the policymaker proposes to change to the state-contingent option. The total increase in losses for group 1 will be

$$n_1(V'^2 - \bar{\tau}^2) \tag{12}$$

which may be a negative number. For group 2 the total increase in losses will be

$$\begin{aligned} n_2(V' - w)^2 - n_2(\bar{\tau} - w)^2 &= n_2(V'^2 - \bar{\tau}^2 - 2w(V' - \bar{\tau})) \\ &= n_2[(V' + \bar{\tau})(V' - \bar{\tau}) - 2w(V' - \bar{\tau})] \\ &= n_2(V' - \bar{\tau})(V' + \bar{\tau} - 2w). \end{aligned} \tag{13}$$

Since w is by definition the maximum possible tax rate we have $V' < w$ and $\bar{\tau} < w$, therefore $V' + \bar{\tau} < 2w$ so the second bracketed term is negative. It follows from (12) and (13) therefore that if one group experiences a total net benefit from the policy change, the other group must experience a total net loss, so they will always vote in opposite directions.

If there are only two groups the outcome must be one extreme or the other. If there is a third group n_3 consisting of honest voters who always prefer V', then there is an increased possibility for approval of a non-extreme tax rate. If no group has an outright majority it would not be able to impose its will, which rules out the extremes (0, w). Suppose therefore that the remaining options to be voted on are $\bar{\tau}$ and V'. As long as no group has an outright majority it must be the case that group 3 can combine with either of the other two groups to form a majority (so either $n_1 + n_3 > n_2$ or $n_2 + n_3 > n_1$). Recall that if group 1 prefers V' to $\bar{\tau}$ then group 2 prefers the opposite, and vice versa, so they will never form a coalition. By equation (12) if $\bar{\tau} > V'$ then group 1 will prefer V' and group 3 can combine with it to secure that outcome. If $\bar{\tau} < V'$ then group 2 will prefer V' and group 3 can combine with it to secure that outcome. Hence:

[Proposition 2] If $n_1 < 0.5$ and $n_2 < 0.5$ and $n_3 < 0.5$, then a majority coalition can always be formed that yields a tax rate of V'.

5 Conclusions

This paper has examined voting outcomes when the policy maker proposes an externality pricing instrument that is either based on a static political compromise or on a state-contingent updating rule. Firstly, in the case in which voters prefer the socially optimal price based on their honest expectation of marginal damages, a standard pricing mechanism aimed at minimizing political losses may never obtain majority support, depending on the variance of private expectations about the level of marginal damages. Even if the proposed rate is acceptable to a majority in the first period, a divergence of beliefs about the future severity of the externality implies that the majority coalition will tend to break down. But a state-contingent pricing rule can hold on to majority support regardless of the divergence of voters' beliefs about the future evolution of the state variable and its marginal damages, because everyone expects the tax path to follow their privately-preferred trajectory. We also show that if some of the voters are dishonest (they prefer either a zero tax or a maximum tax on a priori grounds, irrespective of marginal damages), but no group has an outright majority, then a coalition in favour of the state-contingent tax can always be formed.

References

Alesina, A.F., and Passarelli, F., 2014. Regulation versus taxation. *Journal of Public Economics*, 110, 147-156.

Baumol, W., 1972. On taxation and the control of externalities. *American Economics Review*, 307-332.

Clinch, J.P., Dunne, L., Dresner, S., 2006. Environmental and wider implications of political impediments to environmental tax reform. *Energy Policy*, 34(8), 960-970.

Creedy, J., Francois, P., 2010. Financing higher education and majority voting. *Journal of Public Economics*, 43(2), 181-200.

Cremer, H., Philippe, D.D., Gahvari, F., 2008. Political competition within and between parties: an application to environmental policy. *Journal of Public Economics*, 92(2), 532.547.

Dayaratna, K., R. McKitrick and D. Kreutzer (2016) Empirically-Constrained Climate Sensitivity and the Social Cost of Carbon. University of Guelph Department of Economics Discussion Paper 2016-08.

Damania, R., 1999. Political competition, rent seeking and the choice of environmental policy instruments. *Environmental and Resource Economics*, 13(4), 415-433.

Dresner, S., Dunne, L., Clinch, P., Beuermann C., 2006. Social and political responses to ecological tax reform in Europe: an introduction to the special issue. *Energy Policy*, 34(8), 895-904.

Fredriksson, P.G. and Sterner, T., 2005. The political economy of refunded emissions payment programs. *Economics Letters*, 87(1), 113-119.

Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1), 41-88.

Hsu, Shi-Ling. 2011 A Prediction Market for Climate Outcomes (February 26, 2011). Unpublished manuscript, available as SSRN Working paper http://ssrn.com/abstract=1770882

Interagency Working Group on the Social Cost of Carbon (IWG) (2010) "Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866." Washington: Environmental Protection Agency http://www.epa.gov/otaq/climate/regulations/scc-tsd.pdf

Interagency Working Group on the Social Cost of Carbon (IWG) (2013) "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866." Washington: Environmental Protection Agency https://www.whitehouse.gov/sites/default/files/omb/inforeg/social_cost_of_carbon_for_ria_2013_update.pdf

Intergovernmental Panel on Climate Change (IPCC) (2007) *Synthesis Report of the Fourth Assessment Report*. Cambridge: CUP.

Intergovernmental Panel on Climate Change (IPCC) (2013) *Summary for Policy Makers of the Working Group I Contribution to the Fifth Assessment Report*. Cambridge: CUP.

Johnson, G.E., 1984. Subsidies for higher education, *Journal of Labour Economics*, 2(3), 303-318.

Kallbekken, S., Kroll, S., 2011. Do you not like Pigou, or do you not understand him? Tax aversion and revenue recycling in the lab. *Journal of Environmental Economics and Management*, 62(1), 53-64.

Kawahara, S., 2011. Electoral competition with environmental policy as a second best transfer. *Resource and Energy Economics*, 33(3), 477–495.

Kelly, D.L., Kolstad, C.D., 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control*, 23(4), 491-518.

Leach, A.J., 2007. The climate change learning curve. *Journal of Economic Dynamics and Control*, 31(5), 1728-1752.

MacKenzie, I.A., Ohndorf, M., 2012. Cap-and-trade, taxes, and distributional Conflict., *Journal of Environmental Economics and Management*, 63(1), 51-65.

McCaffery, E.J., Baron, J., 2003. Heuristics and biases in thinking about tax. in: Proceedings : 96th Annual Conference on Taxation, Chicago, Illinois.

McKitrick, R., 2010. A simple state-contingent pricing rule for complex intertemporal externalities. *Energy Economics*, 33(1), 111-120.

McKitrick, Ross R. 2011 "State-Contingent Pricing as a Response to Uncertainty in Climate Policy" in Fouquet, Roger, ed. *Handbook on Energy and Climate Change*., Cheltenham: Edward Elgar.

Nordhaus, W.D., 2007. To tax or not to tax: alternative approaches to slowing global warming. *Review of Environmental Economics and Policy*, 1(1), 26-44.

Pachauri, R.K., 2008. Climate change 2007. *Synthesis report. Contribution of working Group I, II and III to The Fourth Assessment Report*.

Pigou, A.C., 1920. *The Economics of Welfare*. London: Macmillan.

Pindyck, R.S., 2013. Climate change policy: what do the models tell us? (No. w19244). National Bureau of Economic Research.

Ramsey, F.P., 1927. A contribution to the theory of taxation. *The Economic Journal*, 37(145), 47-61.

Rice, John A. 1988. *Mathematical Statistics and Data Analysis*. Belmont: Wadsworth and Brooks/Cole.

Sausgruber, R., Tyran, J.R., 2005. Testing the mill hypothesis of fiscal illusion. *Public Choice*, 122(1-2), 39-68.

Schneider, F., Volkert, J., 1999. No chance for incentive-oriented environmental policies in representative democracies? A public choice analysis. *Ecological Economics*, 31(1), 123-138.

Tol, R.S., 2008. The social cost of carbon: trends, outliers and catastrophes. *Economics: The Open-Access, Open-Assessment E-Journal*, 2(2008-25), 1-22.