How robust is the Strategic Tax Competition model? An Experimental Study

Saileshsingh Gunessee*

School of Economics, University of Nottingham, Nottingham NG7 2RD, UK

06 April 2006
JOB MARKET PAPER

Abstract

The strategic tax competition model predicts that in equilibrium jurisdictions set taxes non-cooperatively below the efficient tax rates most of the time. In this paper, we provide a direct test of the model based on the experimental method to assess whether this claim is robust or not. Though there is some evidence to support the prediction that jurisdictions choose inefficiently low tax rates, the model is less successful in explaining our data when we have asymmetric jurisdictions and repeated interactions.

Keywords: Tax Competition, experimental economics.
JEL Classification Numbers: H21, H73, C92.

* *Telephone*: +44 115 846 6131; *Fax*: +44 115 951 4159; *E-mail*: lexsg3@nottingham.ac.uk. The author would like to thank Hölger Gorg, Steve Humphrey, Ingrid Kubin, Martin Sefton and participants at the GEP conference on ‘Foreign Direct Investment and Taxation’ for helpful comments and is grateful to the Leverhulme Trust for financial assistance under Programme grant F114/BF. The usual disclaimer applies.

1. Introduction

Strategic capital tax competition describes the situation where independent jurisdictions compete noncooperatively over the taxation of mobile capital to influence the location of that mobile tax base. Rooted in the work of Zodrow and Mieszkowski 1986), Wilson (1986), Wildasin (1988), Bucovetsky (1991) and Hoyt (1991)[1], a conventional result of the strategic tax competition (hereafter STC) model is that welfare-maximising jurisdictions choose to set taxes on capital, to finance the public good, inefficiently below the cooperative equilibrium level of taxation.

The reason behind this result is as follows. Knowing that a lower capital tax rate implies a higher net return on capital and what flows out of one jurisdiction must be flowing in another, creates a temptation on the part of each jurisdiction to act independently in using a lower tax rate to attract the mobile capital. Therefore, when capital is free to move between jurisdictions, each time a jurisdiction lowers its tax rate another jurisdiction will follow suit, so that this simultaneous undercutting causes an under-taxation of the mobile resource in equilibrium.

In short, the model predicts that jurisdictions should select inefficiently low tax rates in equilibrium relative to the efficient tax rates most of the time. This echoes the conventional view of most standard tax competition models that tax competition is 'inefficient', when a sourced-based tax is used to finance the supply of the public good.

The theoretical prediction of the STC model is interesting from a policy standpoint for two main reasons. First, one lesson that can be drawn from this prediction is that any taxing region should not use a sourced-based tax on mobile capital because it is a distortionary tax instrument which leaves each jurisdiction expecting that a tax rise will cause capital to flee abroad, hence making them unwilling to increase their tax rates.

Secondly, the natural policy inference drawn from the above prediction is that tax coordination or harmonisation is essential to guard against 'inefficient' tax competition. The 'tax competition versus tax coordination' theme is a hotly debated topic in policy

[1] The Zodrow-Mieszkowski-Wilson (ZMW) model is regarded as the 'standard' tax competition model. It differs from the Wildasin-Hoyt (WH) model in that it does not model any form of 'strategic interaction' between jurisdictions, while Bucovetsky (1991) allows for an asymmetry in population. In this paper, we focus on the strategic version of the symmetric tax competition model but allow for asymmetry in preferences in one treatment.

circles. Critics argue, that with capital being mobile across borders, governments become increasingly constrained to rely less and less on their capital tax policy to provide for the public good (Catenaro and Vidal, 2003). Giving credence to the inefficiency result, the European Union has been campaigning for coordinated tax setting to ensure that tax competition does not undermine the tax revenue of European countries derived from taxed capital income (Zodrow, 2003).

While the above insights from these standard models are interesting, several extensions have shown that the welfare-worsening prediction of the standard models may not be that trivial and in some cases have shown that tax competition can be welfare-improving (Wilson and Wildasin, 2004). In addition, recent works have used documentary evidence on the effective marginal tax rates to point out the absence of any sharp decline in tax rates as predicted by the standard models (Baldwin and Krugman, 2004).

Although these various extensions add new insights to the debate on tax competition, we feel that to judge the standard strategic model one should test the core prediction of the theory and hence provide useful inference as to its empirical validity.

In short we believe that the question as to how robust is the 'inefficient tax competition' postulate is an empirical one. Existing empirical approaches tend to consider patterns of falling tax rates on capital as 'evidence of tax competition' (Slemrod, 2004). One problem in interpreting this as a means to validate the STC model or even calling such patterns a direct test of the model has to do with the difficulty at identifying the exact source of these stylised facts and whether tax competition is the main driving force behind these patterns. Slemrod acknowledges this difficulty when he puts the following query: 'Is this [declining tax rate] due to changes in the domestic determinants of corporate taxation or increases in international pressures for tax competition?'(p.1169). Given the above difficulty the evidence can at best be suggestive but not 'definitive' about the role of tax competition (Slemrod, p. 1183).

This should not come as a surprise as using field data it is difficult to disentangle the influence of tax competition per se and other factors on capital taxes. In this optic we choose to conduct a laboratory experiment which allows us to *control* for these external factors. By using the experimental method we can reproduce the setting where the model works best. Put differently we look to design an experiment that is based on the simple framework of the STC model. For instance, with the facts discussed above we are looking at the evolution of tax rates over time, but the model itself is static by nature. With the aid of the experimental method we can reproduce a design that captures this static feature of the model. Hence, our experimental test constitutes a direct test of the theory.

Another reason why the experimental approach can be used to test the above prediction has to do with the fact that in the field it is difficult to obtain information on the coordinated efficient tax rates (and hence the corresponding optimal provision of the public good), as jurisdictions rarely set taxes cooperatively. Therefore, it is difficult to tell, using field data analysis, whether the observed equilibrium tax rate is inefficient or not. In employing the experimental tool we can actually compare both inefficient and efficient outcome to tell whether actual behaviour indeed matches predicted behaviour.

Using the experimental method also offers the opportunity to uncover new insights in the workings of the STC model, as such we try to decipher the conditions under which the model does or does not work. With certain degree of control achieved, we can identify the set-up where the theoretical prediction organise the data with some success and alternatively the set-up where it is less successful.

To test for the robustness of the model, the treatment variable used is the 'preference for the public good', which is varied to check observed jurisdictional behaviour against predicted jurisdictional behaviour.

As a second stress test, we investigate the robustness of the model with respect to repeated interactions. One problem with the inefficiency postulate and the subsequent tax coordination policy inference is that it is derived from a static framework, where jurisdictions interact only once. However, in reality regions do interact repeatedly. The question repeated interactions raise is whether there is any possibility for cooperative tax setting between jurisdictions[2]. With the added feature of repeated interactions, it is worth asking and investigating whether cooperative arrangements emerge and for how long they can be sustained.

[2] Benoit and Krishna (1985) among others have shown that cooperation can emerge in a finitely repeated setting.

This paper contributes to the literature on tax competition in that we provide a *direct* test of the STC model. Using a simple tax competition model that has a unique Nash equilibrium we construct a design that employs a reduced form representation of the full model, while still preserving the incentive structure and the key property of inefficiency. Our aim is to try and ascertain whether jurisdictions actually choose the Nash tax rates most of the time (and that relative to the efficient tax rates). We give the theory its best shot by testing it in the controlled environment of the laboratory, using the experimental methodology. Under such circumstances, if the model fails to track jurisdictional behaviour in our data then there would be some cause for concern.

To our knowledge, this is the first study that directly tests the model. Though we find some support for the model, it is less robust when we have asymmetric jurisdictions and repeated interactions.

The paper is organised as follows. The second section presents the theoretical model, while we present the experimental design and tell the reader how we implement the design in section 3. In section 4, we report our results, leaving it in section 5 to conclude.

2. A simple STC model

In this section we solve for the model's equilibrium tax rates and demonstrate why tax competition is inefficient if jurisdictions set taxes non-cooperatively. This prediction will form the basis of our refutable hypothesis for our experimental test. To keep the analysis tractable we use specific functional forms for the production function and utility function[3]. Since our interest is to perform a robustness check of the model with respect to the preference for the public good parameter γ, the model is thus parameterised only in terms of γ which in turn allows us to derive an equilibrium tax function solely dependent on γ.

2.1 Theoretical set-up

[3] It is a well known result that existence of equilibrium is not that straightforward for a general tax competition model (Bucovetsky, 2003). So the advantage of using specific functional forms not only helps keep the analysis tractable but has the added bonus of allowing for uniqueness of equilibrium on top of existence.

The model considered is a simplified version of the STC model with 2 jurisdictions *i* and *j*, labelled as i = 1,2 such that $i \neq j$, with the production function[4] of each jurisdiction being quadratic in capital per head k_i

$$f(k_i) = k_i - \frac{1}{4}k_i^2 \tag{1}$$

and $f'(k_i)$ is the marginal product of capital or gross return, derived from the quadratic production function as

$$f'(k_i) = 1 - \frac{1}{2}k_i, \qquad f''(k_i) = -\frac{1}{2} < 0 \tag{2}$$

The representative household's utility function[5] is linear in private good consumption, C_i and public good consumption, G_i

$$U(C_i, G_i) = C_i + \gamma_i G_i, \qquad \gamma_i > 1 \tag{3}$$

where γ_i denotes the preference for the public good of the representative household. A high γ implies the representative consumer has a strong preference for the public good. The parameter γ will be used in section 3 to set-up different treatments for the experiment.

Our representative household owns immobile labour and mobile capital which earns labour income, $f(k_i) - f^l(k_i).k_i$ and capital income, $\frac{1}{2}r\bar{k}$, which is used to finance private consumption

$$C_i = f(k_i) - f'(k_i)k_i + \frac{1}{2}r\bar{k} \tag{4}$$

where $\bar{k}$ stands for the capital stock per head owned by the two jurisdictions, which under the capital market equilibrium is assumed to be equal to the amount of capital employed in each jurisdiction

$$\bar{k} = k_i + k_j = 2 \tag{5}$$

---

4 The typical production function takes the form $f(k_i) = ak_i - bk_i^2$ where a > b. Here a =1 and b = $\frac{1}{4}$.

5 Cardarelli, Taugourdeau and Vidal (2002) use a similar functional form to define their representative consumer's utility function, where γ_i captures 'preferences for the public good'. However, their repeated interactions model differs from the mainstream model in that they assume the existence of a sunk cost to investing in capital abroad.

Given capital is freely mobile across jurisdictions this implies the net-of-tax return, *r*, must be equal in all jurisdictions

$$f'(k_i) - t_i = f'(k_j) - t_j = r \tag{6}$$

The public good in each jurisdiction is assumed to be solely financed by a per unit source-based tax on capital

$$G_i = t_i k_i \tag{7}$$

where t_i is the per unit tax in jurisdiction *i*. Equation (7) also interpreted as the budget constraint of each jurisdiction's government, defines the level of public good provided which is directly dependent on the tax rate and the capital invested in jurisdiction *i*.

2.2 Equilibrium

The objective of each jurisdiction is to set a tax rate that maximises their respective welfare function, *U(t)*. In Appendix A we solve for $t^* = \arg\max_t U(t)$ in more details to arrive at each jurisdiction's equilibrium tax rate given as

$$t_i^* = \frac{\left[\left(\frac{\gamma_i - 1}{2\gamma_i - \frac{1}{2}}\right) + \left(\frac{\gamma_i - \frac{1}{2}}{2\gamma_i - \frac{1}{2}}\right)\left(\frac{\gamma_j - 1}{2\gamma_j - \frac{1}{2}}\right)\right]}{\left[1 - \left(\frac{\gamma_i - \frac{1}{2}}{2\gamma_i - \frac{1}{2}}\right)\left(\frac{\gamma_j - \frac{1}{2}}{2\gamma_j - \frac{1}{2}}\right)\right]} \tag{8}$$

The expression defines the Nash equilibrium tax rate when jurisdictions set taxes non-cooperatively and they share the same production and objectives.

2.3 Inefficient tax competition

The proposition of inefficient tax competition can be stated as follows,

Proposition. *Tax competition leads jurisdictions to set an inefficiently low source-based capital tax rate in equilibrium, relative to the coordinated tax rate.*

Proof. See Appendix B.

The intuition behind this proposition can be explained by the fact that a decrease in the tax rate in one of the jurisdiction will be met by a reduction in the tax rate of the other region in an attempt to influence the destination of capital by influencing the return on capital, for the simple reason that if a tax cutting is not met by a reduction then capital will fly from one region to another. Therefore, in the presence of this simultaneous under-cutting, tax rates will eventually be set below unity in equilibrium. In turn the welfare levels corresponding to the non-cooperative Nash equilibrium tax rates will be lower than the levels under cooperative tax setting.

3. Experimental Design and Implementation

To render the theoretical prediction of 'inefficient tax competition' readily testable in the laboratory we need to set up a design that gives the model its best shot. We have three options: use a combination of parameters and payoff functions to present the model, use a continuous strategy space or use a discretised strategy space. We choose the latter. In what follows we explain our decision.

We simplified our experimental set-up considerably rather than presenting the problem in terms of the model's parameters and welfare functions or with continuous strategies where subjects choose a tax rate between $[t^{min},1]$. The reason we decided to discretise our tax strategies rather than adopt the other two options has to do with the fact we wanted to keep the design as simple as possible while still preserving the incentive structure and the inefficiency postulate of the full model.

So how we went about the designing exercise? A look at the equilibrium tax function and welfare function reveal that these depend on the parameter γ, the preference for the public good. In making the journey from theory to the lab we studied three cases where we assigned specific values to γ. This is an important part of converting our model into an experiment. In actual fact γ is our treatment variable and based on different values it takes we can readily compute the equilibrium tax rates (with the corresponding public good supplies) and the welfare levels.

Bearing in mind that $\gamma > 1$, our choice of γ is led by the consideration that we want to test the robustness of the prediction with respect to γ and with our focus being on whether jurisdictional behaviour changes with a change in γ, we allowed for the three cases: $\gamma_1 = \gamma_2 = 2$; $\gamma_1 = 3$ and $\gamma_2 = 2$; and $\gamma_1 = \gamma_2 = 3$. Appendix C details the relationship between γ, the equilibrium tax rate t^*, the equilibrium welfare levels U^*, and the efficient welfare levels U^{Coop}. It shows the equilibrium tax rates yield inefficiently low welfare levels relative to cooperative tax setting in both symmetric and asymmetric cases, irrespective of the values γ takes.

Discretisation involves simplifying the STC model to a reduced-form tax competition game with only a few tax options to select from. So equipped with the Nash tax rates and other individual tax rates we construct payoff matrices for our non-cooperative tax competition game based on these three specific cases that we then use to conduct our experiment. The rationale for going with a 4x4 strategy space is outlined ahead when we explain the implementation process.

Case 1: Symmetric Tax Competition [$\gamma_1 = \gamma_2 = 2$]. First we look at the case where the preference for the public good is set such that $\gamma_1 = \gamma_2 = 2$. Table 1 (the first matrix) reproduces the welfare levels for the two jurisdictions for the sixteen possible pair of tax rates 0.4, 0.5, 0.667 and 1. By examining the action pairs in the payoff matrix we can see that the pair (0.5, 0.5) satisfies the condition for a unique Nash equilibrium (which also survives the iterated elimination of dominated strategies). If the two regions were competing and under-cutting each other in tax rates, they would not coordinate at the Pareto efficient coordinated equilibrium outcome. A look at the payoff matrix illustrates that the welfare levels for the Nash tax rates are below the welfare levels under the coordinated tax pair, $U^* = 1.25 < U^{Coop} = 1.75$. In addition, Table 2 shows the corresponding public good supply, computed using $G = t_1 - t_1^2 + t_1 t_2$. It shows that in correspondence to capital being under-taxed, the public good is under-provided (0.5<1).

Case 2. Asymmetric Tax Competition [$\gamma_1 = 3$ and $\gamma_2 = 2$]. We next consider the case where the two jurisdictions have asymmetric preferences for the public good, such that $\gamma_1 = 3$ and $\gamma_2 = 2$. This time using the asymmetric equilibrium tax rates $t_1^* = 0.613$ and $t_2^* = 0.548$ and the tax rates 0.667 and 1, the calculations for the utility and public good levels are shown in the second matrix of Tables 1 and 2. By analysing the various action pairs the Nash equilibrium boils down to being sub-optimal relative to the coordinated equilibrium again, as the predicted welfare level of $U_i^* = 1.858$ and $U_j^* = 1.37$ for regions 1 and 2 happen to be lower than $U_i^{Coop} = 2.75$ and $U_j^{Coop} = 1.75$ respectively. Table 2 reports the public good supply to be undersupplied.

Case 3. Symmetric Tax Competition [$\gamma_1 = \gamma_2 = 3$]. The third case stipulates that regions are symmetric with preferences given by $\gamma_1 = \gamma_2 = 3$. The case $\gamma_1 = \gamma_2 = 3$ as reported in Table 1 in the third matrix, shows that under non-cooperative tax setting jurisdictions would set taxes at (0.667, 0.667) in equilibrium which is sub-optimally lower than the cooperative pair, as welfare is lower $U^* = 2.084 < U^{Coop} = 2.75$. The corresponding public good provision, reported in Table 2, reveals an under-provision of the public good.

So the model has a clear message: the Nash equilibrium is unique and delivers inefficiently low tax rates (and under-provision of the public good) in all three cases considered. Tax competition is 'inefficient' in our simple set-up.

How we implemented our test of the model? Based on these 3 cases we modelled our experimental set-up[6] to characterise three treatments based on the 'one-shot' nature of the model: Treatment 1 (referred to as T1 hereafter) with $\gamma_1 = \gamma_2 = 2$; Treatment 2 (T2) with $\gamma_1 = 3$ and $\gamma_2 = 2$; and Treatment 3 (T3) with $\gamma_1 = \gamma_2 = 3$. We call these one-shot treatments (OST hereafter). We also had a further three treatments to stress test the model with respect to repeated interactions: Treatment 4 (referred to as T4 hereafter) with $\gamma_1 = \gamma_2 = 2$; Treatment 5 (T5) with $\gamma_1 = 3$ and $\gamma_2 = 2$; and Treatment 6 (T6) with $\gamma_1 = \gamma_2 = 3$. These are referred to as repeated interactions treatments (RIT hereafter).

To minimise presentation effects, our experimental design employed a completely context free situation[7]. This was done with the consideration that the terms 'jurisdictions' and 'tax rates' may be interpreted differently by some people. So we opted for a neutral set-up where the situation was presented to subjects such that jurisdictions 1 and 2 were respectively called 'row player' and 'column player' and tax rates were referred to as 'options'.

---

[6] We designed our experiment in the spirit of the designs used by Huck, Müller and Normann (2004) and Engelmann and Normann (2005).

[7] Alm, McClelland and Schulze (1992), in a tax evasion experiment, find that behaviour is unaffected as to whether neutral or with-context instructions are used. Abbink and Henning-Schmidt (2006) obtain similar results in a bribery experiment, prompting them to conjecture that a neutral frame already transmit the key features of a bribery situation. However, there are studies that found that framing effects do matter (See Abbink and Henning-Schmidt).

Furthermore, as shown in Appendix D, the actual payoff matrices used in the experiment differed from the payoff tables 1-3 in the following ways. First, the payoff matrix used in the experiment had strategies labelled {1,2,3,4} instead of the actual strategies {0.4,0.5,0.667,1} for the first and third treatments and {0.548,0.613,0.667,1} for the second treatment and the payoffs were expressed in terms of dollars. Secondly, to make sure the best replies to any given strategy were quite distinct, we decided to manipulate the payoffs by first rounding all entries to one decimal place, then changing the 'best reply payoffs' to the higher decimal value while rounding the other payoffs close to their lower decimal value.

So the strategies are discrete rather than continuous as one would expect in practice. However, as explained earlier with the aim of giving the model a fair chance to succeed we went for the simpler of these two. Worth noting is that we went for a 4x4 space rather than a prisoner's dilemma, which T1 and T3 could have been reduced to, for the simple reason that T2 being asymmetric by nature rules out a 2x2 non-cooperative tax competition game of only the equilibrium and efficient rates. This would have left us with a 3x3 matrix, with three tax options. However, with the idea of making all three treatments comparable we added a fourth strategy which had the defection tax rate as common denominator (the tax rate 0.666 here). What is also interesting to note is that T1 and T3 have the same rates and differ only with respect to γ.

Another reason motivating our choice of a design based on discretised taxes can be explained by stressing the research question under study: whether jurisdictions choose the equilibrium tax rates most of the time. With our interest being firmly on outcome, we felt few options would suffice to pick behaviour. This differs from a design based on continuous tax options where the likelihood of choices being more spread around and the interest being more on deciphering any convergence pattern towards equilibrium.

Hence, these explain our choice of having a 4x4 discretise tax competition game which to us was complex-free and more suited to answering the hypotheses under investigation.

The data on the OST was collected at the University of Nottingham where a total of 9 sessions were conducted, with 3 sessions for each one-shot treatment. Instead, for the RIT data was collected for one session from each treatment. Subjects, who were preregistered undergraduate students from the entire university, were recruited through e-mail shots. Eight subjects took part in each session, and no subject participated in more than one session. A total of 96 subjects took part in the experiment with 72 subjects in the OST and 24 subjects in the RIT. Since each session of the OST comprises one independent observation, there were three independent observations per treatment and a total of 9 independent observations. In the RIT with one session of 8 subjects for each treatment, we have 12 independent observations, since in these treatments each pair of subjects represents one independent observation.

Upon arrival, participants were seated at a computer terminal, where they were given a set of instructions[8] and assigned the role of either a row player or a column player. After the instructions were read to them, subjects were asked to complete a question form (reproduced in Appendix F) to ascertain they understood the experiment.

The session consisted of 20 rounds. In each round, of the OST, subjects were paired randomly with a different participant, while in the RIT subject were paired with the same subject throughout the whole session. These were made clear to them.

Subjects were not allowed to communicate between themselves, such that all choices and feedback were made and received through the computer terminal. So, whom they were paired with remained unknown to the subjects. In each round, a subject chose from the four options 1, 2, 3 and 4. At the end of each round, when all subjects had made their choice, they received information about their choice, their opponent's choice, their own earnings for that round and their total accumulated earnings. A sample of the Subject's decision screen is given in Figure 1.

Then at the end of the session, subjects were paid in cash based on their accumulated total earnings for all 20 rounds, using the exchange rate £1 = $4. Each session lasted less than an hour and participants averaged earnings of £8.50 in the OST and £9.30 in the RIT[9].

All decisions regarding our experimental design described in this section were taken with the intention to give the model the 'best shot' possible.

[8] A sample copy of the instructions is placed in Appendix E.
[9] Average hourly earnings in the experiment were much higher than the money that can be earned by a student.

4. Experimental results

Keeping in mind the static nature of the key prediction being tested, we report the results for the 'ones-shot' treatments and 'repeated interactions' treatments separately, such that our analysis is mainly focussed on the first three treatments with the RIT offering a further robustness check. We start with the OST.

4.1 Results from the 'one-shot' treatments

The main variables we look at are the frequency of the Nash tax rates and the frequency of individual tax rates observed. The hypotheses in question-form that were under investigation are:

Q1. Do jurisdictions choose the Nash pair of tax rates most of the time?

Q2. Do jurisdictions choose the Nash pair of tax rates relative to the Pareto pair of tax rates (and the other pair of tax rates) most of the time and how often are the individual jurisdiction equilibrium tax rate selected relative to the other individual tax rates?

Q3. Do changes in the preference for the public good have any effect on the jurisdictional choice of the Nash tax rates and the individual tax rates?

We first analyse whether the observed behaviour is in line with the static Nash equilibrium predictions. To do so we precisely ask:

Q1. Do jurisdictions choose the Nash pair of tax rates most of the time? Table 3 reports the Nash choices of tax rates for the three treatments individually and aggregated, over the 20 rounds. The STC model predicts that the Nash pair of tax rates is chosen most of the time. The total number of Nash tax rates was 344 out of 720. This means about 47.78% of the observed choices are Nash tax rates, which the binomial test confirms to be equally probable to be observed (p-value = 0.124). This can be summed up in the following result:

Result 1. *Jurisdictions choose the Nash tax rates around 47.78% of the time for all treatments.*

Looking at the treatments individually reveals a different picture. Figure 2(a) offers a graphic portrayal of these three treatments. Consider Treatment 1. Around 39.58% of the observed tax rates were Nash, implying the binomial test rejects the hypothesis that it is equiprobable to observe Nash and non-Nash tax rates (*p*-value = 0.001). Treatment 2 reveals similar results. With 41.25% of the decisions being the equilibrium tax rate and the binomial test to confirm it (*p*-value = 0.004), we again have to accept that there were less than half equilibrium jurisdictional decisions. In contrast to the other two treatments, reported Nash observations were around 62.5% in T3. The hypothesis that Nash tax rates are more likely to be observed than non-Nash tax rates cannot be rejected (*p*-value = 0.228 for H_0: p = 0.65 against H_1: p < 0.65).

Observation 1.1. *Jurisdictions select the Nash equilibrium rates 39.58%, 41.25% and 61.25% of the time in T1, T2 and T3 respectively.*

So far, the analysis offer some support for the prediction that jurisdictions do choose the equilibrium tax rates. Now that the static behaviour of jurisdictions has been deciphered it is worth finding out about behaviour over time. It is possible that experience may have had an impact on the equilibrium selections, such that the observed Nash frequencies may increase over time.

To find out we present empirically observed Nash frequencies in Table 4. Close scrutiny of the table highlights the fact that jurisdictions behave in an uncoordinated fashion when selecting their taxes over time. The support for this observation comes from first comparing the overall round 1 and round 20 frequencies. Regions significantly choose more Nash tax rates (binomial test with *p*-value = 0.001 for round 1 and *p*-value = 0.566 for round 20). The second support comes by observing the rise in Nash percentages for the four 5-rounds intervals (the *p*-values of the binomial test for the first three intervals are 0.000, 0.301, 0.206 and 0.239 for the last interval under H_0: p = 0.65).

The same pattern can be observed for the individual treatments. Figure 2(b) depicts these frequencies when the data is split into first ten rounds and last ten rounds. There appears to be an unambiguous increase in the equilibrium selections observed between the first ten rounds and last ten rounds. Thus, the same story of rising Nash proportion is revealed.

The binomial test rejects that it is equiprobable to observe an equilibrium tax rate for the first ten rounds but cannot reject that regions equilibrium options are more likely in the last ten rounds (binomial test respectively reports a *p*-value = 0.000 as opposed to a *p*-value = 0.191 under H_0: p = 0.65). An OLS regression of 'Nash Tax Rate' on 'Rounds' generates a positive estimate which provides some support for the hypothesis of unintended learning to play Nash. However, a further segmentation of the dataset into quarters (1-5, 6-10, 11-15, 16-20) points to one exception to the rising proportion of regional equilibrium selection. There was a slight decline in the Nash proportion in T1 from the second quarter (6-10) to the third quarter (11-15). Despite this episode, overall it seems that experience improves conformance between data and model predictions.

The findings can be summarised in Observation 1.2

Observation 1.2. *It appears that jurisdictions chose the Nash tax rates regularly more over time.*

With 47.78% of the data being consistent with Nash behaviour, it is clear that some jurisdictions were choosing non-equilibrium choices (the remaining 52.22%).

One question regarding these non-equilibrium choices is whether they are random in the sense they are spread about and possibly clustered around the optimal choice. If not random, then the second query amounts to finding whether jurisdictions are selecting some other tax option more systematically. For instance, regions may be coordinating on taxes. So, one might wonder whether the coordinated tax rates are being chosen with the same regularity or not and also how spread are the non-equilibrium choices. To find out the following question was under investigation.

Q2. Do jurisdictions choose the Nash pair of tax rates relative to the Pareto pair of tax rates (and the other pair of tax rates) most of the time and how often are the individual jurisdiction equilibrium tax rate selected relative to the other individual tax rates? The STC model predicts that jurisdictions set taxes below the Pareto pair of tax rates, where the Nash pair of tax rates is chosen most of the time, relative to other tax rates.

Table 5 presents the payoffs for the three OST, with the choice frequencies for the *individual* tax rates and the tax *pairs* shown in bold. So equilibrium prediction in T1 is for regions to choose the tax pair (0.5, 0.5), tax pair (0.613, 0.548) to be chosen by regions in T2 and tax pair (0.667, 0.667) to be chosen in T3. In T1, only 39.58% of the regions choices are equilibrium decisions, while it is 41.25% in T2 and much higher at 62.5% in T3.

Figure 3 confirms the difference between the Nash and efficient frequencies in all three treatments. When regions have symmetrically low preferences, in T1, it can be observed that the equilibrium tax rates are the most commonly chosen options. Jurisdictional non-equilibrium choices appear to be clustered around the equilibrium pair of inefficiently low tax rates. Interestingly we can observe that jurisdictions chose their coordinated tax rates only 1.25% of the time. More interesting is the significantly higher *individual* equilibrium frequencies of the two regions, 61.67% and 59.58% for region 1 and region 2 respectively. This point a tendency for both jurisdictions to *individually* select taxes non-cooperatively at the equilibrium level most of the time.

Examining the other game where preferences are symmetrically higher (T3), a picture with similar features of clustering close to equilibrium emerges. Indeed, the choice of other tax pairs was clearly dwarfed by the success rate at which regions competed at equilibrium. The only observation worth reporting here is that of jurisdiction 1 choosing the equilibrium strategy significantly more than jurisdiction 2 (Fisher's test *p*-value = 0.000). In terms of individual choices one can observe the relatively equilibrium frequencies of 71.66% for jurisdiction 1 and 87.5% for jurisdiction 2.

However, the choice frequencies for T2 differ from the observed frequencies in T1 and T3. The table reports that action pair (0.667, 0.548) accounts for 34.58% of the empirical frequencies and is as significant as the equilibrium tax choice frequencies of 41.25%. Although region 2's behaviour conforms to equilibrium behaviour, region 1 was equally likely to choose tax rates 0.613 or 0.667 (binomial tests for the two taxes reveal *p*-value = 0.2807 and *p*-value = 0.0876 respectively and Fisher's test reveals a *p*-value = 0.3237). In effect jurisdiction 1 was taxing above equilibrium. As to the efficient outcome, it was never chosen.

One explanation for this anomalous behaviour is that jurisdiction 1 may be choosing the particular tax option that yields the highest expected welfare level under the belief the other region is equally likely to choose one of the four options. This then can explain why jurisdictions may be choosing the tax rate 0.667 as significantly as the equilibrium tax rate 0.613, as the expected welfare of the latter rate is lower than the former. This explanation also fits with behaviour of region 2 in T2 and jurisdictional behaviour in T1 and T3, where the expected payoff is highest under the equilibrium rates. Observed behaviour matches predicted behaviour.

While jurisdiction 2 appeared to be behaving as a ‘best reply’ player setting taxes 81.67% of the time at equilibrium level, the significantly lower proportion of Nash pair of 41.25% can be explained by some non-equilibrium behaviour on the part of jurisdiction 1 in opting for around 47.92% of the time for the predicted tax rate of 0.613 and 45.42% for the higher tax rate of 0.667.

We summarise the above results in Result 2.

Result 2. *Jurisdictions tend to set tax rates non-cooperatively below the cooperative tax rates, as the model predicts, but jurisdictions do also tend to quite significantly tax above the Nash equilibrium rates when preferences for the public good are asymmetric*.

Figures 4(a)-(c) show the *individual* jurisdictional frequencies split into the first ten rounds and the last ten rounds. Save in T2 where region 1’s individual choices for the higher tax rate 0.667 appear to fall slowly over time, the rising individual region proportion for the equilibrium was balanced by the falling choice of non-equilibrium individual rates in all treatments.

So, what do we make of these non-equilibrium observations in our data? In T3, these non-equilibrium observations seem to be minimal, while in T1 it exhibits patterns of dying out over time, with clustering close to equilibrium. In T2 these are systematic and appear to persist. What is clear from the experimental data is that there are some systematic non-Nash observations (though not one of coordination but one of competition still) when jurisdictions have asymmetric preferences.

With the Nash equilibrium being chosen around half of the time, coupled with individual frequencies which were even higher and the fact that the coordinated Pareto superior tax rates were rarely chosen, point towards some support for the theoretical prediction of the STC model. However, this has also to be balanced against the fact that jurisdictional behaviour remained unexplained by theory in T2, where jurisdictions are asymmetric.

We summarise these observations in

Observation 2.1. *Jurisdictions rarely set taxes cooperatively, indicating some features of non-cooperative tax setting that is present in our data.*
Observation 2.2. *There is significant difference in jurisdictional behaviour in the choice of the individual equilibrium strategy, when preferences are asymmetrically different.*

So far we have looked at results that portrayed the Nash choices in the different treatments individually. Another way to check on the robustness of the inefficiency claim is to examine the choices and percentages of tax rates across treatments by investigating the following question.

Q3. Do changes in the preference for the public good have any effect on the jurisdictional choice of the Nash tax rates and the individual tax rates? The hypothesis being tested is whether the proportion of Nash tax rates differs across treatments. The STC model predicts that changes in the preference for the public good have no effect on the Nash tax rates and the individual tax rates chosen by jurisdictions.

A look at Figure 2(a) and Table 3 can tell us whether the proportion of equilibrium tax rates across treatments is comparably similar or quite distinct. An answer to this question is valuable to the robustness exercise. What emerges from the graph and the table is that the percentages are similar between T1 (39.58%) and T2 (41.25%) but quite distinct between T1 (39.58%) and T3 (62.5%) and T2 (41.25%) and T3 (62.5%).

To determine whether there are any differences in jurisdictional behaviour across treatments we use Fisher's exact test. We find no significant difference in Nash proportions when comparing T1 and T2 (p-value = 0.355) but do find that an increased preference for the public good can matter from our comparisons of T1 to T3 (p-value = 0.000) and T2 to T3 (p-value = 0.000). The finding that the equilibrium proportion are significantly higher in T3 compared to T1, seem to suggest an increase in the intensity of tax competition when jurisdictional preference for the public services are raised.

With evidence to the fact that the Nash proportion is not the same across treatments this leads us to Result 3 and Observation 3.1

Result 3. *Jurisdictional behaviour, with respect to their choice of Nash equilibrium tax rates and individual tax rates, appear different when the preference for the public good changes.*

Observation 3.1. *A symmetric increase in the preference for the public good leads to an unpredicted rise in the intensity of tax competition.*

How to interpret this result in terms of our tests of the model? Since the parameter change in principle should leave the unique Nash postulate of inefficient tax competition unchanged, this *may* be interpreted as a lack of support for the model in a strict sense. However, we have to bear in mind that in spite of the difference between T1 and T3 there is still a certain degree of non-cooperative tax setting present in our observations on T1 highlighted by the clustering close to equilibrium and insignificant choice of the efficient outcome.

But what is of concern is the difference between T2 and T3 which is being driven only by the non-equilibrium behaviour of region 1. This was already highlighted before.

Next, we offer some further analysis on the difference across treatments using a probit regression approach which takes into account demographic characteristics. Table 6 reports two random-effects probit regressions, with Treatment 3 as the baseline treatment. The dependent variable 'action' is 1 if the subject chooses the individual equilibrium strategy (strategy corresponding to the Nash tax rates) or 0 if another strategy is chosen. The independent variables used include two treatment variables, 'treatment 1' and 'treatment 2'. Both specifications use a constant term to capture the idea of Treatment 3 as the baseline treatment, to examine whether the treatment variables have any effect on the 'best-reply-cum-Nash' action. We are principally interested in the treatment variables. They should tell us whether there is any significant difference between treatments. The first specification reports estimated coefficients when only the treatment variables are used. To consider any potential effect of gender, the second specification adds the demographic variable 'Female' to the probit regression, along with the relevant interactive dummies, 'FemaleT1' and 'FemaleT2'.

Table 7 shows the number of subjects and observations across gender. Of the total number of observations 68.26% was a best reply option (983 out of 1440), of which 65.35% (562 out of 860) were male choices and 72.59% (421 out of 580) were female. A Fisher's exact test confirms that there is a statistically significant difference between choices of the two genders (p-value = 0.0022). Looking at the treatments individually reveals a different picture. There is a difference between gender in T1, while there appears to be no difference in T2 and T3. So gender difference in T1 is driving the overall diversity.

Consistent with what we found from Table 7 on gender, the results from the second probit model show that the interactive demographic variable 'FemaleT1' seems to increase the chance of observing equilibrium choices (the variable is significant at the 1% level). This seems to suggest that females are more competitive than males in T1. To summarise

Observation 3.2. *When public good preferences are symmetrically lower, gender differences seem to matter with females being more selfish than their male counterpart in T1.*

This result is a little bit surprising given that most existing findings on gender in the experimental literature tend to suggest the contrary, that women are less competitive and selfish than men (Eckel and Grossman, 1998; Croson and Gneezy, 2004). One exception to this can be found in the study by Dufwenberg, Gneezy and Rustichini (2004) who find no significant difference between the two genders.

4.2 Results from the 'Repeated Interactions' treatments

So far we have looked at out data from the OST. Now we consider our results from the RIT. The question under scrutiny is

Q4. Do repeated interactions matter for jurisdictional behaviour? The hypothesis we are interested in here is whether jurisdictions set taxes differently when they interact repeatedly. Table 8 reports the Nash frequencies for the 'repeated interactions' treatments. We find that in T4, T5 and T6 jurisdictions set taxes at the equilibrium level 26.25%, 30% and 15% of the time. Unlike in the one-shot treatments, there is no clear pattern as to a rise in Nash tax choices over time, save for T5 where jurisdictional behaviour has been to choose the Nash rates significantly more over time.

By a comparison of the OST (T1, T2, T3) and RIT (T4, T5, T6), as reported in Table 9, we can discern a stark difference across treatments (confirmed by Fisher tests). There are two interesting points worth making. First, by comparing the Nash proportion across the OST and RIT we can find a significant drop in equilibrium choices: 39.58% to

26.25% in the T1-T4 comparison; 41.25% to 30% in the T2-T5 comparison; and 62.5% to 15% in the T3-T6 comparison. Figure 5(a) paints a picture of this difference, confirming our findings.

An alternative comparison of the efficient outcomes reveals significantly higher proportion for the T1-T4 and T3-T6 comparison. This is depicted in Figure 5(b). The T2-T5 comparison shows no such disparity as the efficient outcome is rarely chosen. It appears repeated interactions lead to an increased cooperative tax setting when jurisdictions have symmetric preferences or reduced intensity of tax competition overall.

Our findings seem to fit in line with what Cardarelli et al. find in their paper[10]. They show that tax cooperation can be sustained when countries are *symmetric* and are sufficiently patient. Instead when countries are dissimilar, in that they have asymmetric preferences, cooperation may become unsustainable.

The above findings can be put on record as

Result 4. *Under repeated interactions jurisdictions choose significantly less of the Nash tax rates and cooperate with varying degree of success*.

Looking at previous experiments on non-cooperative games that have contrasted behaviour in a one-shot set-up to one with a repeated framework seems to point towards similar findings of increased cooperation rates in finitely repeated games compared to one-shot games (Montet and Serra, 2003, p.393).

4.3 Welfare Analysis

As a further analysis of the data we look at the welfare levels as predicted by theory and compare it to the actual welfare levels in both OST (T1 to T3) and RIT (T4 to T6). Table 10 shows the welfare for jurisdictions 1 and 2, reporting the theoretical equilibrium and efficient levels alongside the actually observed welfare. Welfare levels in the OST seem relatively closer to the Nash predicted welfare than the efficient welfare levels when compared to the welfare levels in the RIT. This seems to confirm our earlier findings that the model appears fairly robust in its static framework compared to a repeated framework.

5. Discussion and Conclusion

In this paper, we provided a direct test to examine the robustness of the STC model. There is some evidence to support the claim that jurisdictions choose the equilibrium tax rates most of the time, relative to the coordinated outcome, epitomised by a significant choice of the equilibrium pair of tax strategies and backed by even significantly more individual equilibrium choices. In addition, it appears that over time the model's predictions fit the data better.

However, there is also evidence that observed jurisdictional behaviour departs from predicted jurisdictional behaviour when jurisdictions have asymmetric preferences for the public good. Indeed, the results in our experiment suggest that asymmetry in terms of preferences for the public good appears to influence regions tax decisions in such a way that jurisdictions with higher preferences taxed above their equilibrium tax rate as opposed to their low preference counterparts (in Treatment 2).

Furthermore, though more of a stress test rather than a direct test, with repeated interactions the model appears less successful in organising the data. It emerges repeated interactions can induce some cooperative tax setting when regions have symmetric preferences or reduce the degree of non-cooperative tax setting when preferences are asymmetric. These findings from the 'repeated interactions' treatments were explained by drawing from the work by Cardarelli et al.

How do we explain the results from the 'one-shot' treatments? We offer two explanations which overlap to a certain degree with one another[11]. One explanation for the overall results and mainly based on the findings of some support for Nash behaviour and clustering close to the equilibrium pair of strategies in all three treatments follows from a *one-shot version* of the ε-equilibrium[12] (Radner, 1980), where each region may get close to the best response of the other region's tax rate, without actually reaching a best reply. In effect, due to bounded rationality and uncertainty about the choice of the tax rate

[10] Note that Cardarelli et al. study tax competition with an infinite horizon, while we implement a finite horizon. However, as shown by Friedman (1985) the backward induction argument in finitely repeated games need not rule out trigger strategies for these types games.

[11] In Appendix G we provide estimates from a Quantal response equilibrium (QRE) estimation exercise we conducted to see whether QRE fits behaviour better than the Nash equilibrium. The results are not illuminating.

the other jurisdiction would opt for, the relationship between the tax rates and welfare levels is rendered complex and uncertain to the other region.

Thus, what this explanation is suggesting is we should observe some Nash choices but also choices that cluster around the Nash outcome, because tax options (corresponding to the Nash choices and other choices) yield payoffs that are only ε from one another. Then the anomalous result in T2 can be explained by appealing to the additional feature that asymmetry may play a significant role in enhancing the uncertainty surrounding the choice of the tax rate. In a nutshell, we are suggesting that differences between treatments can be explained away by different uncertainty levels that underpins the link between tax options and welfare levels in the three treatments and an added complexity of asymmetry between regions (not sharing the same taste for the public good make them different in preferences for tax options) can explain behaviour in T2.

A second explanation at an attempt to rationalise our findings can draw from a mixed strategy type of equilibrium jurisdictional behaviour such that if there are beliefs on the part of jurisdictions as to what action the other jurisdiction would be taking, then regions tend to choose their own strategy believing the other region is equally likely to choose one of the four strategies. Therefore, given this belief about other players, the player then chooses that strategy which gives the highest expected payoff. A simple computation of the expected payoffs, for jurisdiction 1, in T2 reveals that the expected welfare under tax rate 0.667 is greater than that under equilibrium tax rate 0.613, hence this could explain why region 1 taxes significantly at the non-equilibrium rate of 0.667. This explanation is also consistent with jurisdiction 2's behaviour in T2 and overall regional behaviour in T1 and T3. The high proportion of individual choices that corresponds to the predicted rates can be attributed to regions getting the highest expected welfare under the equilibrium rates.

What insights do our results provide? These results are interesting in that they highlight the importance of symmetry as one pre-condition for the theory to work. In addition, they highlight the importance of asymmetry and repeated interactions from a modelling perspective. These two elements are absent from most models of tax

[12] 'An ε-equilibrium is a pair of strategies, one for each player, such that each player's strategy is within ε of being the best response to the other player's strategy' (Radner, 1986).

competition. Careful modelling of asymmetry and repeated interactions with some dose of uncertainty *may* go a long way in bringing some much needed insights to the debate on tax competition.

Tables

Table 1: Jurisdictions payoff matrices

		0.4	0.5	0.667	1
Case 1	**0.4**	1.15,1.15	1.233,1.153	1.381,1.079	1.72,0.64
	0.5	1.153,1.233	1.25,1.25	1.424,1.201	1.813,0.813
	0.667	1.079,1.381	1.201,1.424	1.417,1.417	1.889,1.112
	1	0.64,1.72	0.813,1.813	1.112,1.889	1.75,1.75

		0.548	0.613	0.667	1
Case 2	**0.548**	1.846,1.298	1.949,1.284	2.045,1.262	2.64,0.897
	0.613	1.858,1.37	1.976,1.363	2.076,1.346	2.725,1.013
	0.667	1.849,1.432	1.977,1.429	2.084,1.417	2.778,1.112
	1	1.445,1.844	1.626,1.875	1.779,1.889	2.75,1.75

		0.4	0.5	0.667	1
Case 3	**0.4**	1.55,1.55	1.673,1.603	1.888,1.568	2.36,1.04
	0.5	1.603,1.673	1.75,1.75	2.007,1.757	2.563,1.313
	0.667	1.568,1.888	1.757,2.007	2.084,2.084	2.778,1.779
	1	1.04,2.36	1.313,2.563	1.779,2.778	2.75,2.75

Table 2: Jurisdictions public good provision

		0.4	0.5	0.667	1
Case 1	**0.4**	0.4,0.4	0.44,0.45	0.507,0.489	0.64,0.4
	0.5	0.45,0.44	0.5,0.5	0.584,0.556	0.75,0.5
	0.667	0.489,0.507	0.556,0.584	0.667,0.667	0.889,0.667
	1	0.4,0.64	0.5,0.75	0.667,0.889	1.00,1.00
		0.548	**0.613**	**0.667**	**1**
Case 2	**0.548**	0.548,0.548	0.584,0.573	0.613,0.588	0.796,0548
	0.613	0.573,0.584	0.613,0.613	0.646,0.631	0.85,0.613
	0.667	0.588,0.613	0.631,0.646	0.667,0.667	0.889,0.667
	1	0.548,0.796	0.613,0.85	0.667,0.889	1.00,1.00
		0.4	**0.5**	**0.667**	**1**
Case 3	**0.4**	0.4,0.4	0.44,0.45	0.507,0.489	0.64,0.4
	0.5	0.45,0.44	0.5,0.5	0.584,0.556	0.75,0.5
	0.667	0.489,0.507	0.556,0.584	0.667,0.667	0.889,0.667
	1	0.4,0.64	0.5,0.75	0.667,0.889	1.00,1.00

Table 3: Number of Nash tax rates

Round	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*	*13*	*14*	*15*	*16*	*17*	*18*	*19*	*20*	*Total*
Treatment 1	2	3	1	3	3	7	3	6	6	5	6	4	5	3	8	5	4	7	7	7	95 (240)
Treatment 2	2	1	2	1	4	3	3	7	6	5	2	7	5	8	9	5	9	9	7	4	99 (240)
Treatment 3	4	8	5	5	6	10	6	7	6	6	7	9	8	7	8	10	11	10	10	7	150 (240)
Total	8	12	8	9	13	20	12	20	18	16	14	20	18	18	25	20	24	26	24	18	344 (720)

Note: Maximum number of Nash choices per round is 12.

Table 4: Nash Tax Rate Frequencies

Treatments	*Rounds*								
	1	*20*	*1-5*	*6-10*	*11-15*	*16-20*	*1-10*	*11-20*	*1-20*
1	16.67%	58.33%	20%	45%	43.33%	50%	32.5%	46.67%	39.58%
2	16.67%	33.33%	16.67%	40%	51.67%	56.67%	28.33%	54.17%	41.25%
3	33.33%	58.33%	46.67%	58.33%	65%	80%	52.5%	72.5%	62.5%
Overall	22.22%	50%	27.78%	47.78%	52.77%	62.22%	37.78%	57.5%	47.78%

Table 5. Tax Rate Frequencies

Treatment 1

	0.4 **3.75%**	0.5 **59.58%**	0.667 **27.09%**	1 **9.58%**
0.4 **2.5%**	1.10,1.10 **0.42%**	1.20,1.20 **1.25%**	1.30,1.00 **0.42%**	1.70,0.60 **0.42%**
0.5 **61.67%**	1.20,1.20 **1.25%**	1.30,1.30 **39.58%**	1.50,1.20 **15.42%**	1.80,0.80 **5.42%**
0.667 **23.75%**	1.00,1.30 **1.25%**	1.20,1.50 **13.33%**	1.40,1.40 **6.67%**	1.90,1.10 **2.5%**
1 **12.08%**	0.60,1.70 **0.83%**	0.80,1.80 **5.42%**	1.10,1.90 **4.58%**	1.70,1.70 **1.25%**

Treatment 2

	0.548 **81.67%**	0.613 **1.25%**	0.667 **13.75%**	1 **3.33%**
0.548 **2.08%**	1.80,1.30 **1.25%**	1.90,1.20 **0%**	2.00,1.20 **0.83%**	2.60,0.80 **0%**
0.613 **47.92%**	1.90,1.40 **41.25%**	1.90,1.30 **0.83%**	2.00,1.30 **5.42%**	2.70,1.00 **0.42%**
0.667 **45.42%**	1.80,1.50 **34.58%**	2.00,1.40 **0.42%**	2.10,1.40 **7.5%**	2.80,1.10 **2.92%**
1 **4.58%**	1.40,1.80 **4.58%**	1.60,1.80 **0%**	1.70,1.90 **0%**	2.70,1.70 **0%**

Treatment 3

	0.4 **1.667%**	0.5 **1.25%**	0.667 **87.5%**	1 **9.583%**
0.4 **10.42%**	1.50,1.50 **0.42%**	1.60,1.70 **0%**	1.80,1.50 **10%**	2.30,1.00 **0%**
0.5 **4.167%**	1.70,1.60 **0%**	1.70,1.70 **0%**	2.00,1.80 **2.917%**	2.50,1.30 **1.25%**
0.667 **71.663%**	1.50,1.80 **1.25%**	1.80,2.00 **0.83%**	2.10,2.10 **62.5%**	2.80,1.70 **7.083%**
1 **13.75%**	1.00,2.30 **0%**	1.30,2.50 **0.42%**	1.70,2.80 **12.083%**	2.70,2.70 **1.25%**

Table 6. Random Effects Probit regressions results

Independent Variables	Dependent Variable: *action* (T3 as baseline)	
	(1)	(2)
constant	1.039***	1.086***
	(0.130)	(0.269)
treatment 1	-0.419**	-1.143***
	(0.205)	(0.303)
treatment 2	-0.541***	-0.513*
	(0.179)	(0.301)
Female	-	-0.026
		(0.312)
FemaleT1	-	-0.827**
		(0.397)
FemaleT2	-	-0.154
		(0.405)
N	1440	1440
Log-likelihood	-723.74	-722.39

Standard errors in parentheses.
*, **, *** indicate significance at the 10%, 5% and 1% levels respectively

Table 7. Number of subjects and observations across gender

	T1		T2		T3		Overall	
	No. of Subjects	**No. of Observations**	**No. of Subjects**	**No. of Observations**	**No. of Subjects**	**No. of Observations**	**No. of Subjects**	**No. of Observations**
Male	16	320	14	280	13	260	43	860
Female	8	160	10	200	11	220	29	580
Total	24	480	24	480	24	480	72	1440
***p*-value of Fisher's test of gender differences**	0.000 Different		0.000 No Difference		0.000 No Difference		0.002 Different	

Table 8: Nash Frequencies in Repeated Interactions Treatments

Treatments	*Rounds*								
	1	*20*	*1-5*	*6-10*	*11-15*	*16-20*	*1-10*	*11-20*	*1-20*
4	0%	25%	15%	45%	25%	20%	30%	22.5%	26.25%
5	0%	75%	15%	20%	30%	55%	17.5%	42.5%	30%
6	0%	50%	15%	5%	10%	30%	10%	20%	15%
Overall	0%	50%	15%	23.33%	21.67%	35%	19.17%	28.33%	23.75%

Table 9. Differences across treatments and outcome

Repeated Interactions			
T1 and T4	*T1*	*T4*	*Fisher's Exact Test (T1 Vs T2) p-value*
Nash *Efficient*	39.58% 1.25%	26.25% 21.25%	0.021 0.000
Fisher's Exact Test (Nash Vs Efficient) p-value	0.000	0.289	
T2 and T5	*T2*	*T5*	*Fisher's Exact Test (T3 Vs T4) p-value*
Nash *Efficient*	41.25% 0%	30% 1.25%	0.048 0.250
Fisher's Exact Test (Nash Vs Efficient) p-value	0.000	0.000	
T3 and T6	*T3*	*T6*	*Fisher's Exact Test (T1 Vs T3) p-value*
Nash *Efficient*	62.5% 1.25%	15% 32.5%	0.000 0.000
Fisher's Exact Test (Nash Vs Efficient) p-value	0.000	0.008	

Table 10. Jurisdictions Welfare Levels

	Predicted		Actual
	Nash	*Efficient*	
Treatment 1:			
Jurisdiction 1	1.30	1.70	1.33
Jurisdiction 2	1.30	1.70	1.34
Treatment 2:			
Jurisdiction 1	1.90	2.70	1.89
Jurisdiction 2	1.40	1.70	1.44
Treatment 3:			
Jurisdiction 1	2.10	2.70	2.06
Jurisdiction 2	2.10	2.70	2.08
Treatment 4:			
Jurisdiction 1	1.30	1.70	1.48
Jurisdiction 2	1.30	1.70	1.39
Treatment 5:			
Jurisdiction 1	1.90	2.70	2.00
Jurisdiction 2	1.40	1.70	1.39
Treatment 6:			
Jurisdiction 1	2.10	2.70	2.23
Jurisdiction 2	2.10	2.70	2.23

Figures

Payoff Matrix (Row, Column)

	1	2	3	4
1	**\$1.10, \$1.10**	**\$1.20, \$1.20**	**\$1.30, \$1.00**	**\$1.70, \$0.60**
2	**\$1.20, \$1.20**	**\$1.30, \$1.30**	**\$1.50, \$1.20**	**\$1.80, \$0.80**
3	**\$1.00, \$1.30**	**\$1.20, \$1.50**	**\$1.40, \$1.40**	**\$1.90, \$1.10**
4	**\$0.60, \$1.70**	**\$0.80, \$1.80**	**\$1.10, \$1.90**	**\$1.70, \$1.70**

Submit Decision

History with the Most Recent Round Listed First
(current part only)

Prior Round	Your Decision	Other's Decision	Your Earnings	Total Earnings
1	4	3	\$1.10	\$1.10

Figure 1. Subject's Decision Screen

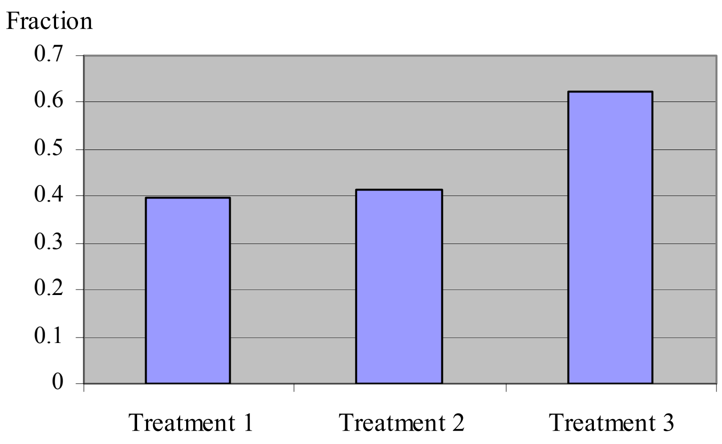

Figure 2(a). Nash proportion across treatments

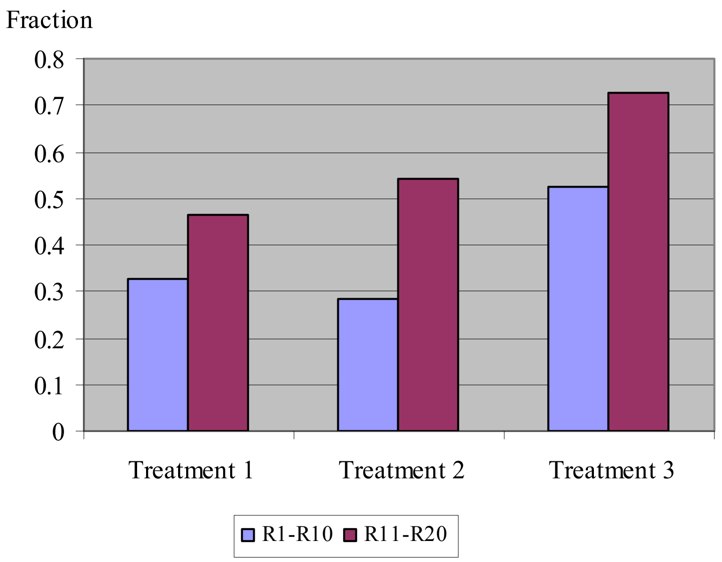

Figure 2(b). Nash Proportion for rounds 1-10 and 11-20

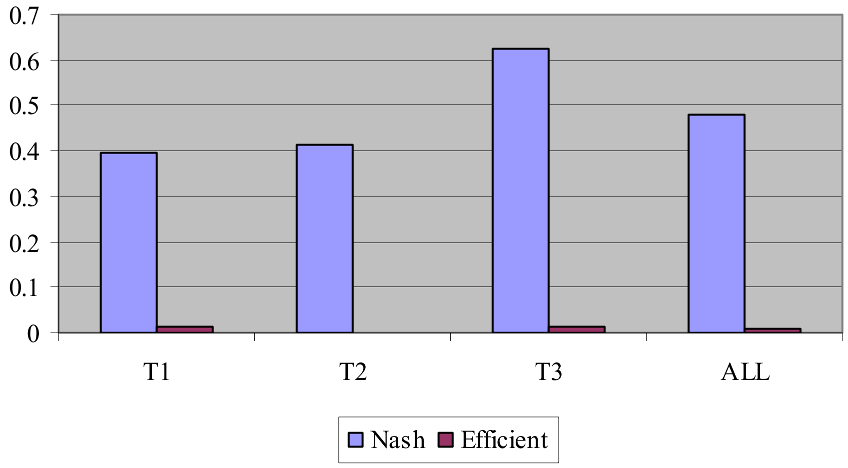

Figure 3. Nash and efficient frequencies

Jurisdiction 1

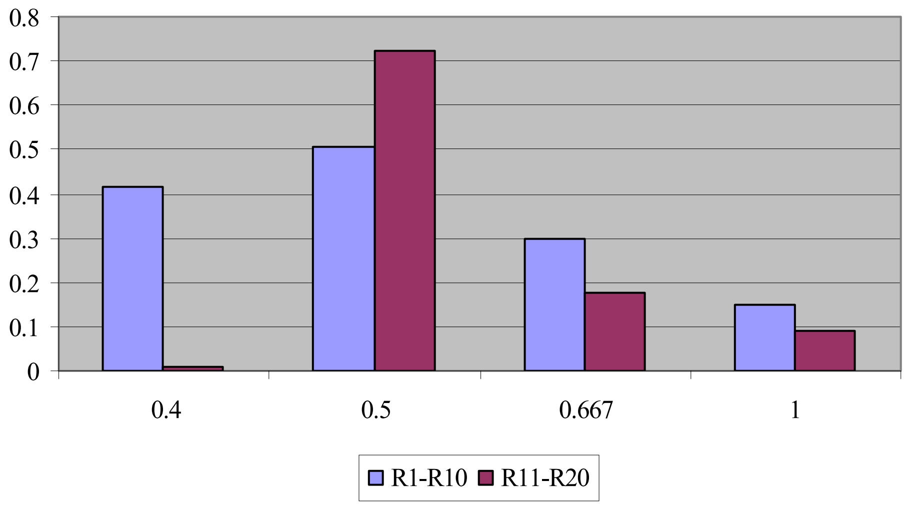

Jurisdiction 2

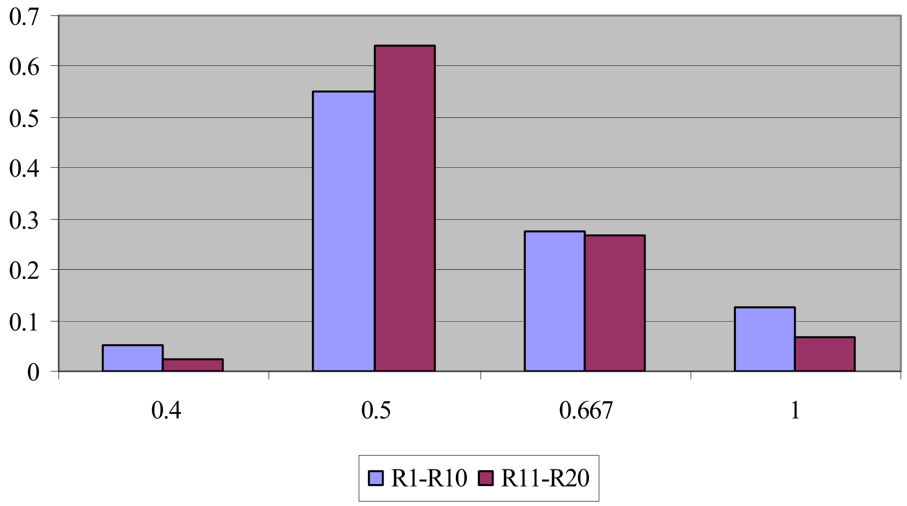

Figure 4(a). Jurisdictional Frequencies for rounds 1-10 and 11-20 in Treatment 1

Jurisdiction 1

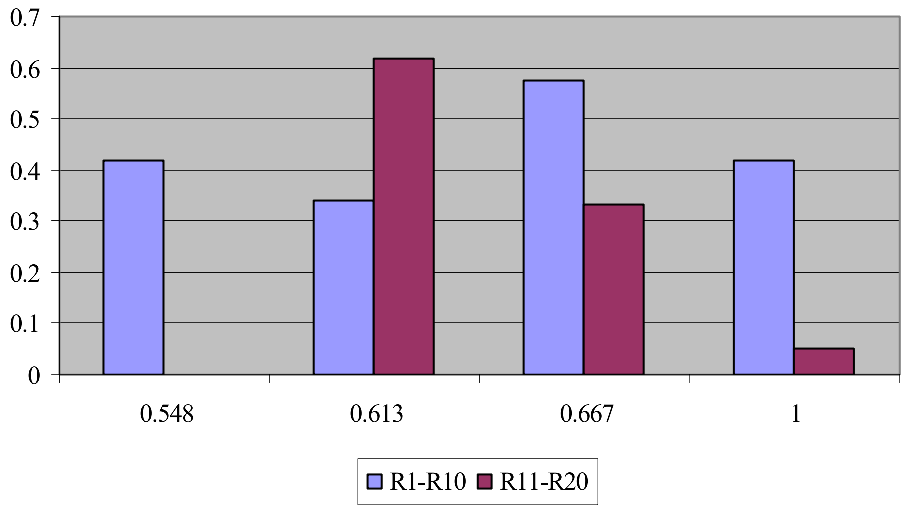

Jurisdiction 2

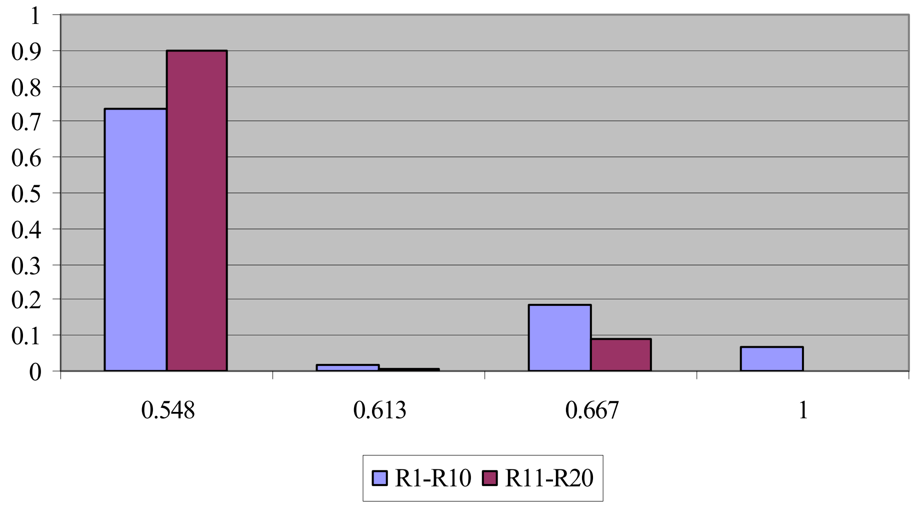

Figure 4(b). Jurisdictional Frequencies for rounds 1-10 and 11-20 in Treatment 2

Jurisdiction 1

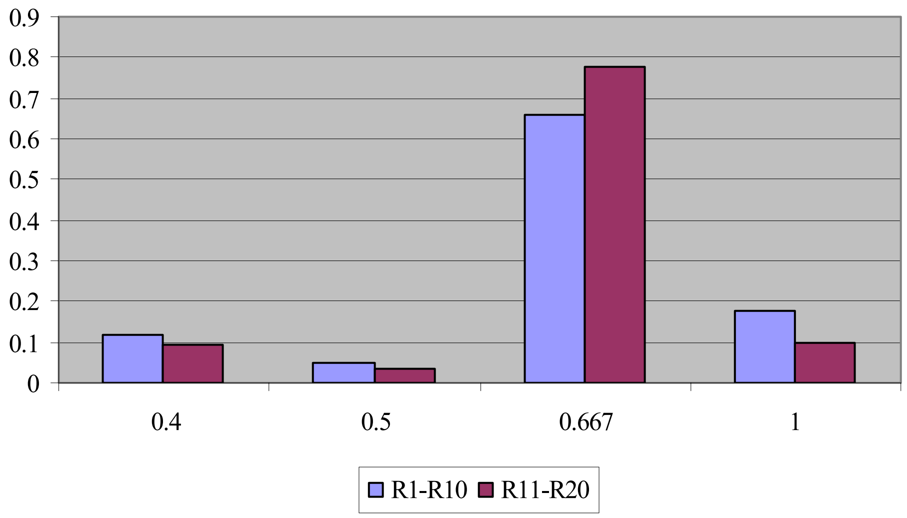

Jurisdiction 2

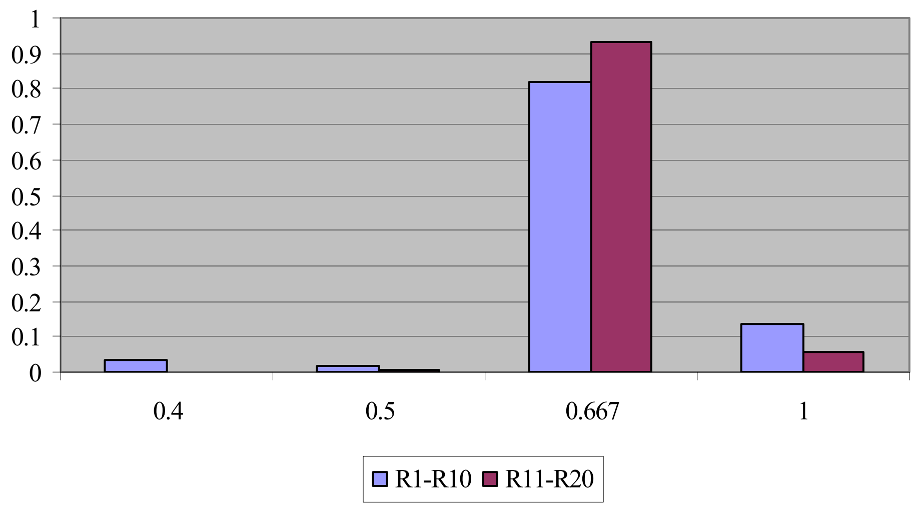

Figure 4(c). Jurisdictional Frequencies for rounds 1-10 and 11-20 in Treatment 3

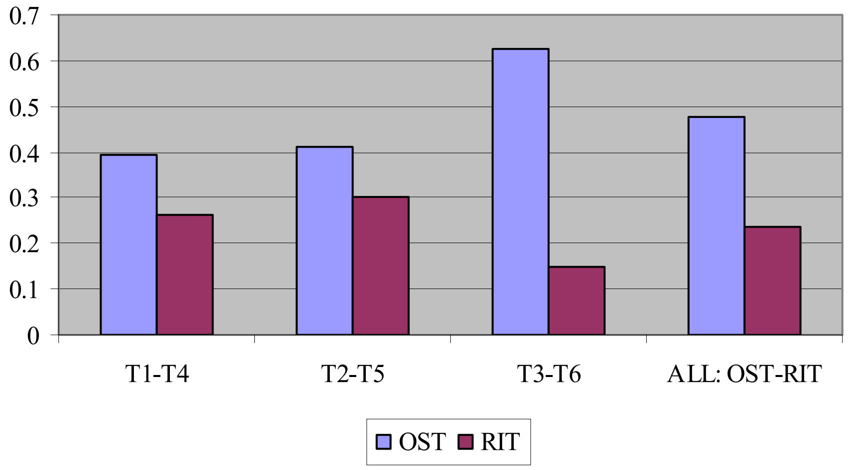

Figure 5(a). Nash Proportion across OST and RIT

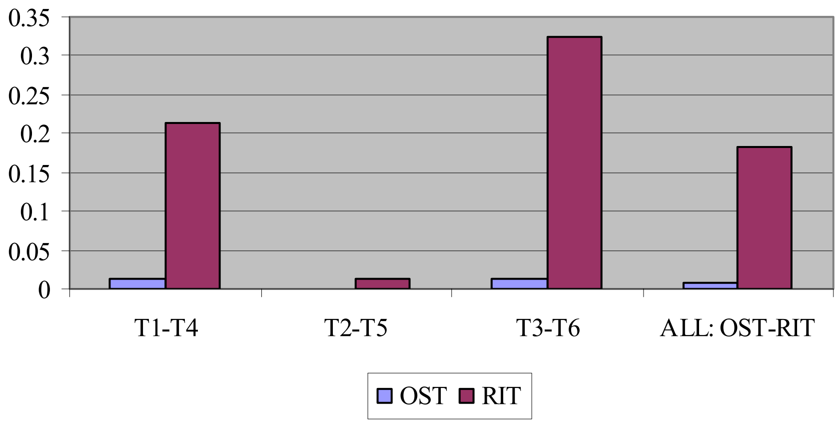

Figure 5(b). Efficient Proportion across OST and RIT

Appendix A: Equilibrium Tax Rates

Solving for the equilibrium tax rates involves the following steps. As a first step we need to get an expression of the utility function in terms of tax rates only. This means finding equations for *k(t)* and *r(t)* which we can then use to get an expression for private consumption solely in terms of tax rates, *C(t)*. Similarly using *k(t)* into the public good function should yield *G(t)*. This readily provides an expression for utility of the representative citizen only as a function of *t*. This exercise is carried out for jurisdiction *i*.

Where *k(t)* comes from? Through a process of substituting for $f'(k)$ and $\bar{k}$ as defined by equations (4) and (5) in equation (6), the expression for the net return on capital *r*, we arrive at

$$k_i = 1 - t_i + t_j \qquad \frac{\partial k_i}{\partial t_i} < 0, \frac{\partial k_i}{\partial t_j} > 0$$

for jurisdiction *i*.

To find *C(t)*, the expression for private consumption, we need to plug the equation of the capital employed as a function of tax rates, *k(t)*, into $f(k_i), f'(k_i)$, r which should leave us with

$$C_i = \frac{3}{4} - t_i - \frac{1}{2}t_i t_j + \frac{1}{4}t_i^2 + \frac{1}{4}t_j^2$$

Next, we again make use *k(t)* into the equation for the public consumption, (7), to yield

$$G_i = t_i - t_i^2 + t_i . t_j$$

Then, combining *C(t)* and *G(t)* produce the following utility function

$$U[C_i(t), G_i(t)] = \frac{3}{4} + (\gamma_i - 1)t_i + (\gamma_i - \frac{1}{2})t_i t_j - (\gamma_i - \frac{1}{4})t_i^2 + \frac{1}{4}t_j^2$$

The objective of each jurisdiction is to solve for $t^* = \arg\max_t U(t)$ which requires setting $\frac{\partial U_i(.)}{\partial t_i} = 0$. This yields the best response function for jurisdiction *i*

$$t_i = \left(\frac{\gamma_i - 1}{2\gamma_i - \frac{1}{2}}\right) + \left(\frac{\gamma_i - \frac{1}{2}}{2\gamma_i - \frac{1}{2}}\right).t_j \qquad ; t_i \geq \frac{\gamma_i - 1}{2\gamma_i - \frac{1}{2}} \text{ , } \gamma_i > 1$$

which can be used in conjunction with jurisdiction *j*'s best response function to solve for the equilibrium tax rate

$$t_i^* = \frac{\left[\left(\frac{\gamma_i - 1}{2\gamma_i - \frac{1}{2}}\right) + \left(\frac{\gamma_i - \frac{1}{2}}{2\gamma_i - \frac{1}{2}}\right)\left(\frac{\gamma_j - 1}{2\gamma_j - \frac{1}{2}}\right)\right]}{\left[1 - \left(\frac{\gamma_i - \frac{1}{2}}{2\gamma_i - \frac{1}{2}}\right)\left(\frac{\gamma_j - \frac{1}{2}}{2\gamma_j - \frac{1}{2}}\right)\right]}$$

Two interesting properties worth noting are the existence and uniqueness of equilibrium. The second-order condition can be solved as

$$\frac{\partial^2 U_i(.)}{\partial t_i^2} = -(2\gamma_i - \frac{1}{2}) < 0 \text{ for } \gamma_i > 1$$

revealing it holds. With respect to the uniqueness of equilibrium rates one can draw from the fact that the reaction functions are both continuous and strictly linearly increasing in the other jurisdiction's tax rate and hence establish the premise for single crossing, which imply uniqueness.

Appendix B: Proof

Proposition. *Tax competition leads jurisdictions to set an inefficiently low source-based capital tax rate in equilibrium, relative to the coordinated tax rate.*

Proof. Differentiating the utility of the representative consumer in jurisdiction i

$$U[C_i(t),\, G_i(t)] = \frac{3}{4} + (\gamma_i - 1)t_i + (\gamma_i - \frac{1}{2})t_i t_j - (\gamma_i - \frac{1}{4})t_i^2 + \frac{1}{4}t_j^2$$

gives

$$dU(t_i, t_j) = (\gamma_i - 1)\,dt_i + (\gamma_i - \frac{1}{2})\,dt_i t_j - (2\gamma_i - \frac{1}{2})\,dt_i t_i$$

For the symmetric case $\gamma_i = \gamma_j$ we have $t_i = t_j$ which reduces the above equation to

$$dU(.) = (\gamma_i - 1)\,dt_i + (\gamma_i - \frac{1}{2})\,dt_i t_i - (2\gamma_i - \frac{1}{2})\,dt_i t_i$$
$$= \gamma_i\, dt_i (1 - t_i) - dt_i$$

We can further simplify this equation under three cases: when symmetric jurisdictions set equilibrium taxes $t_i = t_i^*$; given $\gamma_i > \gamma_j$ asymmetric jurisdictions set taxes non-cooperatively $t_i > t_j$; regions prefer to cooperate over taxes $t_i = t_j = 1$.

Plugging for the equilibrium tax rates $t_i = t_i^* < 1$ we arrive at

$$dU(.)\big|_{\text{Sym}} = \gamma_i\, dt_i (1 - t_i^*) - dt_i$$

Alternatively for the asymmetric case $\gamma_i > \gamma_j$ we have $t_i > t_j$ which we can re-define as $t_j = t_i - \Delta$, where $0 < \Delta < 1$, which imply an alternative expression in the asymmetric case

$$dU(.)\big|_{\text{Asym}} = \gamma_i\, dt_i (1 - t_i^* - \Delta) - (1 - \frac{1}{2}\Delta)\, dt_i$$

Finally, if jurisdictions were to set taxes cooperatively such that $t_i = t_j = 1$ it follows that the equation for the change in welfare can be rewritten as

$$dU(.)\big|_{\text{Coop}} = -dt_1$$

How to tell that the pair of Nash tax rates (t_i^*, t_j^*) are inefficiently low? Observe that if tax rates are set below unity then for any pair of Nash tax rates less than unity a comparison of equations $dU(.)\big|_{\text{Sym}}$ and $dU(.)\big|_{\text{Asym}}$ to equation $dU(.)\big|_{\text{Coop}}$ reveals that both jurisdictions can raise welfare if they set tax rates cooperatively at unity above the equilibrium tax rate, t_i^* . Given the first term in both $dU(.)\big|_{\text{Sym}}$ and $dU(.)\big|_{\text{Asym}}$ are positive it shows the unexploited gains from coordination when taxes are set below unity. Then what remains in $dU(.)\big|_{\text{Coop}}$ is the temptation payoff for undercutting the other jurisdiction.

This in turn means if regions set tax rates non-cooperatively below unity they should be inefficiently low as there is a potential gain in welfare from raising tax rates simultaneously.

Appendix C: Choice of Parameters

The following tables show how t^*, U^* and U^{Coop} change with γ, when preferences are symmetric $\gamma_i = \gamma_j$ and asymmetric $\gamma_i > \gamma_j$.

γ	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
t^*	0.500	0.667	0.750	0.800	0.833	0.857	0.875	0.889	0.900
U^*	1.25	2.083	3.000	3.950	4.917	5.893	6.875	7.861	8.85
U^{Coop}	1.750	2.750	3.750	4.750	5.750	6.750	7.750	8.750	9.750

$\gamma_i = 3 > \gamma_j = 2$	$\gamma_i = 4 > \gamma_j = 3$	$\gamma_i = 5 > \gamma_j = 4$	$\gamma_i = 6 > \gamma_j = 5$	$\gamma_i = 7 > \gamma_j = 6$	$\gamma_i = 8 > \gamma_j = 7$	$\gamma_i = 9 > \gamma_j = 8$	$\gamma_i = 10 > \gamma_j = 9$
$t_i^* = 0.613$ $t_j^* = 0.548$	$t_i^* = 0.723$ $t_j^* = 0.692$	$t_i^* = 0.784$ $t_j^* = 0.766$	$t_i^* = 0.823$ $t_j^* = 0.811$	$t_i^* = 0.849$ $t_j^* = 0.841$	$t_i^* = 0.869$ $t_j^* = 0.863$	$t_i^* = 0.884$ $t_j^* = 0.880$	$t_i^* = 0.896$ $t_j^* = 0.893$
$U_i^* = 1.858$ $U_j^* = 1.370$	$U_i^* = 2.830$ $U_j^* = 2.198$	$U_i^* = 3.815$ $U_j^* = 2.558$	$U_i^* = 4.804$ $U_j^* = 4.226$	$U_i^* = 5.796$ $U_j^* = 4.997$	$U_i^* = 6.791$ $U_j^* = 5.965$	$U_i^* = 7.7\ 86$ $U_j^* = 6.940$	$U_i^* = 8.783$ $U_j^* = 7.921$
$U_i^{Coop} = 2.75$ $U_j^{Coop} = 1.75$	$U_i^{Coop} = 3.75$ $U_j^{Coop} = 2.75$	$U_i^{Coop} = 4.75$ $U_j^{Coop} = 3.75$	$U_i^{Coop} = 5.75$ $U_j^{Coop} = 4.75$	$U_i^{Coop} = 6.75$ $U_j^{Coop} = 5.75$	$U_i^{Coop} = 7.75$ $U_j^{Coop} = 6.75$	$U_i^{Coop} = 8.75$ $U_j^{Coop} = 7.75$	$U_i^{Coop} = 9.75$ $U_j^{Coop} = 8.75$

Appendix D: Payoff Matrices used in the experiment

During the experiment the payoff matrices had the heading **Payoff Matrix (Row, Column)** at the top with the payoffs shown in dollars. The row player's strategies and payoffs were shown in blue in the left corner and the column player's strategies and payoffs were shown in red on the top. The following payoff matrices are tables used for the three treatments.

	1	2	3	4
1	$1.10,$1.10	$1.20,$1.20	$1.30,$1.00	$1.70,$0.6 0
2	$1.20,$1.20	$1.30,$1.30	$1.50,$1.20	$1.80,$0.80
3	$1.00,$1.30	$1.20,$1.50	$1.40,$1.40	$1.90,$1.10
4	$0.60,$1.70	$0.80,$1.80	$1.10,$1.90	$1.70,$1.70

Treatments 1 & 4

	1	2	3	4
1	$1.80,$1.30	$1.90,$1.20	$2.00,$1.20	$2.60,$0.80
2	$1.90,$1.40	$1.90,$1.30	$2.00,$1.30	$2.70,$1.00
3	$1.80,$1.50	$2.00,$1.40	$2.10,$1.40	$2.80,$1.10
4	$1.40,$1.80	$1.60,$1.80	$1.70,$1.90	$2.70,$1.70

Treatments 2 & 5

	1	2	3	4
1	$1.50,$1.50	$1.60,$1.70	$1.80,$1.50	$2.30,$1.00
2	$1.70,$1.60	$1.70,$1.70	$2.00,$1.80	$2.50,$1.30
3	$1.50,$1.80	$1.80,$2.00	$2.10,$2.10	$2.80,$1.70
4	$1.00,$2.30	$1.30,$2.50	$1.70,$2.80	$2.70,$2.70

Treatments 3 & 6

Appendix E: Sample Instructions

Instructions

Welcome to the experiment. Please read and follow the instructions carefully. A good understanding of the task you will be asked to perform will increase your chances of earning some money. During the experiment your earnings will be given in dollars, which will be converted into pounds at the end of the experiment at a rate of £1=$4. This money will be paid in cash and in private at the end of the session. Please do not talk or communicate with the other participants in this room during the session. If you have any questions just raise your hand and I will come to where you are seated to answer any query that you may have. Please do not touch the computer until the experiment starts.

Introduction

This is an experiment on decision making. There are 8 people in this room who are participating in this experiment. The experiment consists of 20 rounds. In each round of this session, the computer will randomly pair you with someone else in the room. However, you will not learn which of the people in the room you are paired with. How much you earn during the session depends on your decisions and those of the people you are paired with, during these 20 rounds.

What you have to do?

In the experiment you will play the role of either a 'row' player or a 'column' player. [**You will be a row player in all rounds**]. You will face the payoff matrix on the next page for the 20 rounds. The payoff matrix has 4 rows and 4 columns. Since there are four possible options for each player, there are sixteen possible combinations or 'cells' in the table. The numbers shown in those 16 cells are simply the payoffs, with two numbers reported for each cell. The row player's options and payoffs are shown in blue in the left corner and the column player's options and payoffs are shown in red on the top. The intersection of the row selected by one person and the column selected by the other will determine the earnings for each person.

Here is an example on how to read the payoff matrix. If the row player chooses 1 and the column player chooses 3 the payoff shown is $1.30,$1.00: the row player gets a payoff of $1.30, the column player gets $1.00.

For each round, the row player will choose one of the rows from the left corner and the column player will choose one of the columns on the top of the payoff matrix. This means in each round you have to choose between the four possible options labelled as: 1, 2, 3, and 4. You make your decision on your computer by choosing one of these 4 numbers, using the computer mouse. At the end of each round, after taking your decision, you will be prompted to confirm or change your decision. Once you confirm your decision, you will either get a screen telling you to wait, if the other person's decision has not yet been received, and that you can use the 'check to see if others have finished' button. Otherwise, your choice and the choice of the other player will be displayed. Your terminal will also display your earnings for that round and your accumulated total earnings from all rounds. At the beginning of each round, along with telling you the round number, the computer will display a history for each round, showing your decision you had taken, the other person's decision, your earnings and total earnings, with the most recent round shown first.

Payoff Matrix (Row, Column)

	1	2	3	4
1	\$1.10,\$1.10	\$1.20,\$1.20	\$1.30,\$1.00	\$1.70,\$0.60
2	\$1.20,\$1.20	\$1.30,\$1.30	\$1.50,\$1.20	\$1.80,\$0.80
3	\$1.00,\$1.30	\$1.20,\$1.50	\$1.40,\$1.40	\$1.90,\$1.10
4	\$0.60,\$1.70	\$0.80,\$1.80	\$1.10,\$1.90	\$1.70,\$1.70

Appendix F: Pre-experiment Question Form

Question Form

In answering the questions, please circle one option and where appropriate fill in the blanks. Raise your hand, if you need any help and if you finish earlier, to have your answers checked.

1. You will play the role of which type of player during the experiment?

   Row Column

2. In what colour will your options and payoffs be during the experiment?

   Red Blue

3. What is the available number of options you can choose from during the experiment?

   1 4

4. You will be randomly matched with another person in each round of the experiment.

   True False

5. If you choose 2 and the person you are matched with chooses 4, you earn $ ______ and the other person gets $ ______.

6. If you choose 3 and the person you are matched with chooses 1, you earn $ ______ and the other person gets $ ______.

7. The cash you will be paid at the end of the experiment will depend on your **total accumulated earnings**.

   True False

8. The cash you will be paid at the end of the experiment will be converted at what rate?

   £1= $1 £1= $4

Appendix G: Behavioural analysis

In light of the fact that the Nash prediction fails to track behaviour across treatments and hence leaves unexplained part of our dataset unexplained we explored the possibility that the QRE might offer some useful clues about jurisdictional behaviour.

The QRE approach, developed by McKelvey and Palfrey (1995), consists of the introduction of limited rationality, often termed random elements, in the way players solve a game. The model assumes that players' choices are related to expected payoffs such that decisions with higher expected payoffs have a higher likelihood of being selected. Furthermore, in the QRE model the parameter λ is the measure of bounded rationality which describes how equilibrium probabilities converge to the Nash equilibrium as λ approaches infinity. The logit equilibrium for λ such that $\lambda \in [0,\infty]$ produces a logit equilibrium correspondence. When λ = 0 all options are equally likely, while when λ increases more weight is given to the those strategies that yield a higher expected payoff and when λ tends to infinity the probability of observing the payoff maximising strategy approaches unity. Using a probabilistic choice rules one can estimate the probabilities for each individual strategy and thus compute the log-likelihood for the model.

Figure G.1 shows the QRE correspondences for the three treatments. The probabilities from the Logit equilibrium are displayed on the vertical axis and the horizontal axis depicts λ. The first and third panels show the plots of choosing equilibrium strategies 0.5 and 0.67 (for Treatments 1 and 3) respectively. The two curves in the second panel shows the probability of choosing options 0.613 and 0.548 for region 1 and region 2 respectively. Interestingly the QRE predicts the equilibrium strategy to be overplayed by region 2 relative to region 1's equilibrium strategy play, for intermediate values of λ. The actual data does show something along those lines. Table G.1 reports the QRE estimates of the data for the three treatments. The data column and row show the actual frequencies (in bold), while the Nash column and row reports the probability based on equilibrium predictions and the last column and row gives the estimated choice probabilities using the QRE model. What emerges is that neither model seems to fit the data well, with the QRE undershooting the actual choice probabilities most of the time. The one occasion where the QRE model seems to pick some patterns of behaviour is in T2 for region 1's choice of the equilibrium strategy, but overall QRE is not very informative as an alternative explanation to the Nash equilibrium.

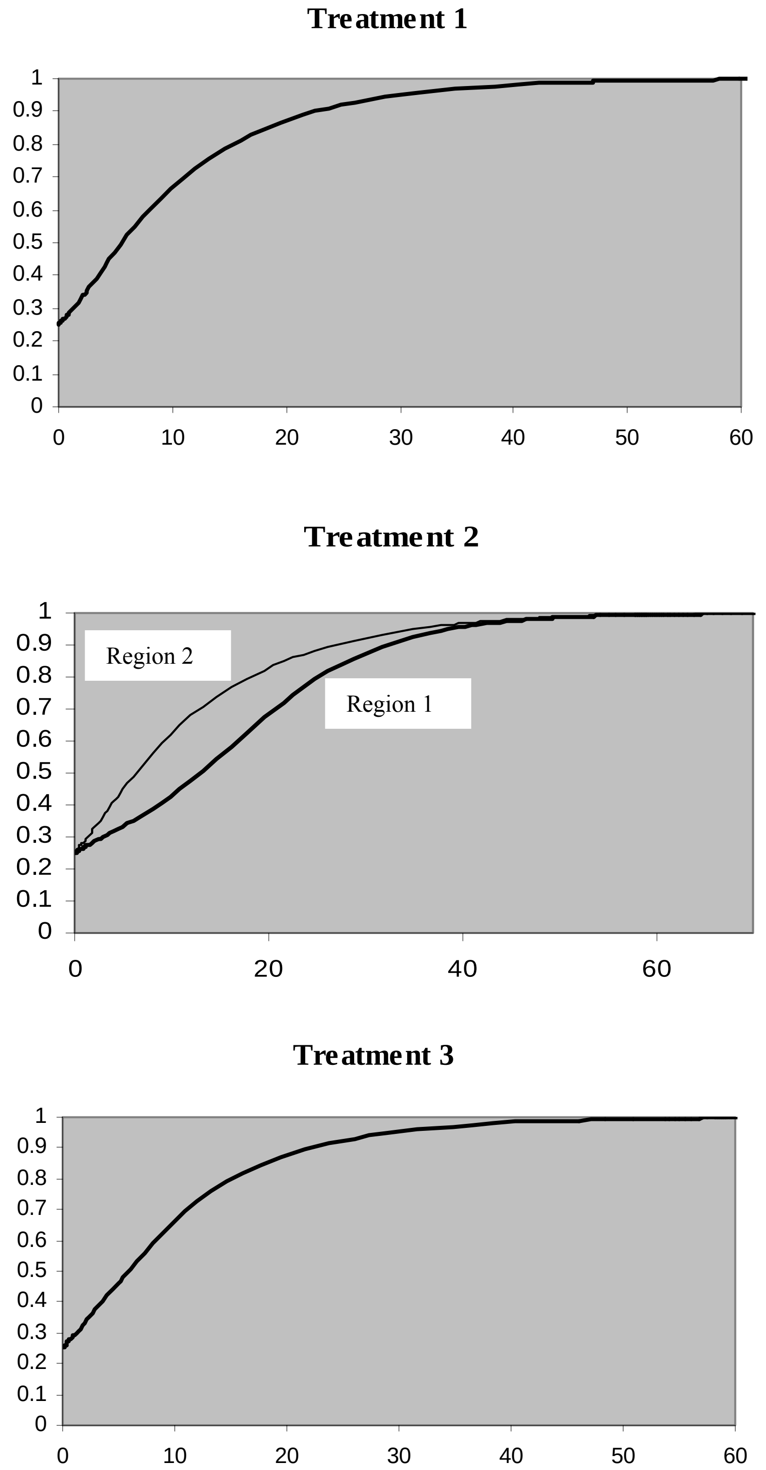

Figure G.1. QRE Correspondence

Table G.1. QRE estimation

Treatment 1

	0.4	0.5	0.667	1	*Data*	*Nash*	*QRE*
0.4	1.10,1.10	1.20,1.20	1.30,1.00	1.70,0.60	**0.025**	**0**	**0.269**
0.5	1.20,1.20	1.30,1.30	1.50,1.20	1.80,0.80	**0.617**	**1**	**0.441**
0.667	1.00,1.30	1.20,1.50	1.40,1.40	1.90,1.10	**0.238**	**0**	**0.225**
1	0.60,1.70	0.80,1.80	1.10,1.90	1.70,1.70	**0.121**	**0**	**0.065**
Data	**0.038**	**0.596**	**0.271**	**0.096**			
Nash	**0**	**1**	**0**	**0**			
QRE	**0.269**	**0.441**	**0.225**	**0.065**			**-ℓ = 582.27**

Treatment 2

	0.548	0.61	0.667	1	*Data*	*Nash*	*QRE*
0.548	1.80,1.30	1.90,1.20	2.00,1.20	2.60,0.80	**0.021**	**0**	**0.225**
0.613	1.90,1.40	1.90,1.30	2.00,1.30	2.70,1.00	**0.479**	**1**	**0.438**
0.667	1.80,1.50	2.00,1.40	2.10,1.40	2.80,1.10	**0.454**	**0**	**0.331**
1	1.40,1.80	1.60,1.80	1.70,1.90	2.70,1.70	**0.046**	**0**	**0.005**
Data	**0.817**	**0.013**	**0.138**	**0.033**			
Nash	**1**	**0**	**0**	**0**			
QRE	**0.635**	**0.179**	**0.180**	**0.006**			**-ℓ = 471.68**

Treatment 3

	0.4	0.5	0.667	1	*Data*	*Nash*	*QRE*
0.4	1.50,1.50	1.60,1.70	1.80,1.50	2.30,1.00	**0.104**	**0**	**0.090**
0.5	1.70,1.60	1.70,1.70	2.00,1.80	2.50,1.30	**0.042**	**0**	**0.317**
0.667	1.50,1.80	1.80,2.00	2.10,2.10	2.80,1.70	**0.717**	**1**	**0.570**
1	1.00,2.30	1.30,2.50	1.70,2.80	2.70,2.70	**0.138**	**0**	**0.023**
Data	**0.017**	**0.013**	**0.875**	**0.096**			
Nash	**0**	**0**	**1**	**0**			
QRE	**0.090**	**0.317**	**0.570**	**0.023**			**-ℓ = 511.03**

References

Abbink, K. and Hennig-Schmidt, H. (2006). 'Neutral versus Loaded Instructions in a Bribery Experiment', *Experimental Economics*, forthcoming.

Alm, J., McClelland, G.H. and Schulze, W.D. (1992). 'Why do people pay taxes', *Journal of Public Economics* 48, 21-38.

Baldwin, R. and Krugman, P. (2004). 'Agglomeration, integration and tax harmonisation', *European Economic Review* 48(1), 1-23.

Benoit and Krishna (1985). 'Finitely Repeated Games', *Econometrica* 53, 905-922.

Bucovetsky, S. (1991). 'Asymmetric Tax Competition', *Journal of Urban Economics*, 30, 167-181.

Bucovetsky, S. (2003). 'Existence of Nash Equilibrium for Tax Competition among a Large Number of Jurisdictions', mimeo, York University.

Cardarelli, R., Taugourdeau, E. and Vidal, J-P. (2002). 'A Repeated Interactions model of Tax Competition', *Journal of Public Economic Theory* 4(1), 19-38.

Catenaro, M. and Vidal, J-P. (2003). 'Implicit Tax Co-ordination under repeated policy interactions', *European Central Bank Working Paper* No. 259, September.

Croson, R. and Gneezy, U. (2004). 'Gender Differences in Preferences', mimeo, University of Chicago Graduate School of Business.

Dufwenberg, M., Gneezy, U. and Rustichini, A. (2004). 'Price Competition: The role of of gender and education' in Huck, S. (ed.) *Advances in understanding strategic behaviour: Game theory, Experiments and Bounded Rationality*. Palgrave Macmillan, Hampshire.

Eckel, C. and Grossman, P.J. (1998). 'Are women less selfish than men?: Evidence from dictator experiments', *Economic Journal* 108, 726-735.

Engelmann, D. and Normann, H.T. (2005). 'An Experimental Test of Strategic Trade Policy', *Journal of Economic Behavior and Organization*, forthcoming.

Friedman, J. (1985). 'Cooperative Equilibria in Finite Horizon Supergames', *Journal of Economic Theory* 35, 390-398.

Hoyt, W.H. (1991). 'Property taxation, Nash equilibrium and market power', *Journal of Urban Economics* 30, 123-131.

Huck, S., Müller, W. and Normann, H.T. (2004). 'Strategic Delegation in Experimental Markets', *International Journal of Industrial Organization* 22, 561-574.

McKelvey, R.D. and Palfrey, T.R. (1995). 'Quantal Response Equilibrium for Normal Form Games', *Games and Economic Behavior* 10, 6-38.

Montet, C. and Serra, D. (2003). *Game Theory and Economics*. Palgrave Macmillan, Hampshire.

Radner, R. (1980). 'Collusive Behavior in Oligopolies with Long but Finite Lives', *Journal of Economic Theory*, 22, 136-156.

Radner. R. (1986). ‘Can Bounded Rationality resolve the Prisoner’s Dilemma?’ in Hildenbrand W. and Mas-Colell A. (eds.) *Contributions to Mathematical Economics in Honor of Gerard Debreu*. North Holland, Amsterdam.

Slemrod, J. (2004). ‘Are corporate tax rates, or countries, converging?’, *Journal of Public Economics* 88, 1169-1186

Wildasin, D.E. (1988). ‘Nash equilibria in models of fiscal competition’, *Journal of Public Economics* 35, 299-240.

Wildasin, D.E. (1991). ‘Some rudimentary duopolity theory’, *Regional Science and Urban Economcis* 21, 393-421.

Wilson, J.D. (1986). ‘A theory of interregional tax competition’, *Journal of Urban Economics* 19, 296-315.

Wilson, J.D. and Wildasin, D. (2004). ‘Capital Tax Competition: Bane or Boon?’, *Journal of Public Economics* 88, 1065-1091.

Zodrow, G. (2003). ‘Tax Competition and Tax Coordination in the European Union’, *International Tax and Public Finance* 10, 651-671.

Zodrow, G. and Mieszkowski, P. (1986). ‘Pigou, Tiebout, property taxation and the underprovision of local public goods’, *Journal of Urban Economics* 19, 356-370.