

Becker, Johannes; Davies, Ronald B.

Working Paper

Learning to Tax - Interjurisdictional Tax Competition under Incomplete Information

CESifo Working Paper, No. 6699

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Becker, Johannes; Davies, Ronald B. (2017) : Learning to Tax - Interjurisdictional Tax Competition under Incomplete Information, CESifo Working Paper, No. 6699, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/172975>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Learning to Tax - Interjurisdictional Tax Competition under Incomplete Information

Johannes Becker, Ronald B. Davies

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Learning to Tax – Interjurisdictional Tax Competition under Incomplete Information

Abstract

How do countries compete for mobile tax base when they lack precise information on how tax rates affect the tax base? We present a multi-period version of a classic tax competition model in which countries set source-based taxes under incomplete information on the tax base elasticity. This information, however, improves as they observe both their own and their neighbours' experiences. In contrast to the existing work on policy learning, we focus on learning in the presence of (fiscal) externalities. We show that, because learning can exacerbate this externality, the value of learning can be negative and, thus, learning may be too fast. Given that variance in tax policies enhances learning, this implies that, in the sequence of Markov perfect equilibria, tax rates can be too heterogeneous. Furthermore, we contribute to the empirical tax competition literature by showing that learning generates tax patterns that look as if countries react to each other even if there are no fiscal externalities. We conclude that the existing results typically taken as evidence of tax competition may be more nuanced than heretofore recognized.

JEL-Codes: H250, H320, H870.

Keywords: social learning, policy diffusion, tax competition.

Johannes Becker
University of Münster
Institute of Public Economics
Wilmergasse 6-8
Germany – 48143 Münster
johannes.becker@wiwi.uni-muenster.de

Ronald B. Davies
University College Dublin
G215 Newman Building, Belfield
Ireland – Dublin 4
ronbdavies@gmail.com

September 19, 2017

We thank Karolien De Bruyne, Andreas Haufler, Vilen Lipatov, Ben Lockwood and seminar and conference participants at the Universities of Ghent and Siegen, the 2013 Oxford Tax Conference, and the Trade and Networks Conference in Leuven. All errors are our own.

1 Introduction

Optimal tax policy in open economies is a trade-off between the benefits of taxation (e.g. public goods, redistribution) and its costs (e.g. loss in tax base, distortion of economic activity). The latter crucially depend on the elasticity of the tax base.² The existing literature generally approaches tax competition by assuming that the elasticity of the tax base is common knowledge. However, it seems more than plausible that policy-makers have no better access to information than researchers do and – at best – have to measure elasticity as they must: based on past experience, both their own and their neighbours'. In this paper, we extend the literature on optimal taxation of mobile resources by embedding it in a setting of policy learning – an innovation that has non-trivial impacts both on the welfare implications of tax competition *and* learning.³

Under incomplete information, a government can learn not just from its own experience, but also by observing the actions and outcomes in other jurisdictions. These observations allow it to refine its beliefs about the state of nature – the tax base elasticity in our model. As has been shown by Bala and Goyal (1998) and Gale and Kariv (2003) among others, learning in networks increases the speed of convergence in beliefs about the elasticity towards the actual value. Whenever information is the only factor that spills across borders, this faster evolution of beliefs is welfare improving. The positive welfare effect of new observations gives rise to a positive information externality and, thus, a potentially slower-than-optimal speed of learning in equilibrium (a point emphasized, e.g., by Rose-Ackerman 1980).

However, as we show in our model, with a cross-border resource externality (such as the fiscal externality generated by tax competition for mobile capital) this is not generally the case. Learning about the true tax base elasticity may, in instances to be specified, imply more aggressive tax com-

²See Feldstein (1999) and Saez (2001) for the notion that the elasticity of the tax base in a sufficient statistic for optimal tax policy.

³Please note that the literature's nomenclature of 'social experimentation' means that agents learn from others' actions and payoffs, not that they experiment with their actions to generate additional information.

petition, exacerbating the welfare losses in the short- and medium-term. As such, an ‘ignorance is bliss’ scenario emerges in which learning can be inefficiently fast. Further, since heterogeneity in taxes improves the accuracy of beliefs, tax rate heterogeneity may be excessively high in equilibrium.

Applied to the classical symmetric tax competition framework (Zodrow and Mieszkowski 1986), the curious case of a non-learning equilibrium may arise. Since the symmetry assumption implies that tax bases across jurisdictions do not depend on the degree of tax competition, learning may be prevented and – in principle – the equilibrium may achieve the first-best outcome. This knife-edge case is, however, more a matter of theoretical curiosity than of practical relevance.

Our second insight builds on the observation that, with policy learning, convergence of beliefs results in convergence in tax rates (up to variation driven by other differences across locations, such as size). Importantly, this is true even if there are no resource externalities of tax policy. This finding allows us to reinterpret the empirical tax competition literature which finds that tax rates in other nations affects one’s own tax. Observable interdependence of tax rate setting is typically attributed to competition for mobile resources or yardstick competition. In a world with policy learning, however, such interdependence also arises because taxes reflect beliefs about the (supposedly common) state of nature.⁴ This suggests that the empirical evidence on tax interdependence may be more nuanced than previously recognized, as such pattern may occur even in the absence of fiscal externalities and political features inherent to the yardstick competition literature. Distinguishing between these explanations is important, though, since their policy implications differ considerably.

Our model combines a baseline tax competition model (e.g. Zodrow and Mieszkowski, 1986) with standard policy learning. This combination is designed to highlight the interactions of the two externalities in the most transparent fashion. To this end, we use an infinite horizon model with mul-

⁴For instance, the common state of nature could refer to the elasticity of entrepreneurial labor supply which, by assumption, is perfectly immobile across borders and, therefore, does not give rise to resource externalities.

multiple countries whose governments gradually learn about the true elasticity of the tax base. Each country sets its source-based tax rate on business profits in each period. While our focus will be on a setting with fiscal externalities (i.e. a tax base that leaves a given country and to be taxed by some other country), our model can also encompass their absence (e.g. tax base reductions due to decreasing entrepreneurial effort). Starting from an *a priori* belief on the true state of nature, a country observes its own payoff and neighbouring countries' taxes and outcomes and then updates its belief function.⁵ To mimic the one-shot setting of the standard tax competition models, we restrict our analysis to Markov perfect equilibria with the information set being the only aspect that is transferred from one period to the next. For the purpose of reinterpreting the empirical evidence, we consider two variants of the model: one in which information spreads with a lag (which is rationalized by a social learning framework as per Gale 1996, Gale and Kariv 2003, and Vives 1996); another one in which observations of close or more similar neighbors contain more information (which rationalizes the use of weighted spatial lags in empirical measurement).

The international interdependence of tax policies has been documented in a number of empirical studies, including Devereux et al. (2008), Overesch and Rincke (2009, 2011), and Heinemann, et al. (2010).⁶ There are three dominant theories that may explain this interdependence: competition for mobile resources (tax competition), yardstick competition and policy learning. The literature on 'tax competition' for mobile factors in the tradition of Zodrow and Mieszkowski (1986) and Wilson (1986), assumes that tax policies across jurisdictions interact because mobile resources respond to tax rate differentials. Specifically, because taxes abroad affect the domestic

⁵In contrast to the social *experimentation* literature which assumes that agents learn from each other by observing their neighbours' actions *and* their outcomes (e.g. Bala and Goyal (1998)), the social *learning* literature (Gale (1996), Gale and Kariv (2003), Vives (1996)) assumes that agents can observe their neighbours' actions but not their outcomes. We consider social learning as an extension.

⁶Although our discussion focuses on international policy diffusion, our theory equally applies to the literature on policy competition between jurisdictions within a nation. Empirical work here includes Fredricksson, List, and Millimet (2003), Levinson (2003), Heyndels and Vuchelen (1998), Besley and Case (1995), and Mintz and Smart (2004).

endowment with mobile resources, a change in tax in a neighbouring country triggers a change in domestic policy. A well-known prediction of this approach is the race to the bottom, i.e. a sharp reduction in source-based taxes on mobile entities due to international competition. In equilibrium, the competing jurisdictions are in a prisoners' dilemma-like situation in which coordination of tax policies may yield a Pareto improvement.⁷

The second explanatory approach is the 'yardstick competition' of Besley and Case (1995) in which voters judge the performance of domestic policy makers by comparing their choices to the policies implemented elsewhere.⁸ Whereas politicians have knowledge about the underlying state of nature, voters do not but can learn about it by observing policy choices abroad. By mimicking other countries' policies, politicians can improve their chances of reelection (since this is a signal for being an able policy maker). Competition for office therefore forces politicians to adopt 'good' policies in order to stay in office. Therefore, competition has beneficial effects (and policy coordination could endanger these benefits from competition).

The third theory to explain interdependent policy choices, policy learning, has – so far – almost been neglected in the literature on international taxation. In the classical policy learning framework (see e.g. Rose-Ackerman 1980, Shipan and Volden 2008, or Callander and Harstad 2015), countries are part of a network which is defined by information flows.⁹ In such a network, beliefs converge over time to the true state of nature and thus actions converge to those under full information. As such, learning will almost always be perfect in the long run (Bala and Goyal, 1998). Set in the context of international tax policy, beliefs on the true tax base elasticity are predicted to converge over time and, eventually, the full-information Nash equilibrium tax rates are implemented. Since actions (tax policy choices) affect the size and quality of information, they have information externalities and the pace

⁷Other models, including Davies (2005), provide additional settings where inter-jurisdictional tax competition can be welfare improving.

⁸See Salmon (1987) for an initial application to taxes and Brueckner (2003) for an overview.

⁹When all agents in a group can observe each others' actions and outcomes, as in our case, the network is 'fully connected'.

of learning is, at least in many cases, inefficiently slow. This finding has been emphasized in the literature on policy experimentation in federal structures (e.g. Rose-Ackerman 1980, Strumpf 2002, Kotsogiannis and Schwager 2006, Cai and Treisman 2009) which shows that, due to a positive information externality, federations yield inefficiently low levels of experimentation.¹⁰ Buera et al. (2011) consider the policy learning of countries which decide between market-oriented and state-oriented economic policies. Callander and Harstad (2015) show that the federal authority may implement a mechanism that compensates for the insufficient incentives to experiment.

The remainder of the paper is organized as follows. Section 2 presents the model and derives the main results. Section 3 turns to the empirical literature and reinterprets the existing evidence using the learning model and two extensions. Section 4 concludes.

2 Model

2.1 Setup

We consider an infinite horizon model with $N \geq 2$ countries, indexed by i (where N is finite). In each country, there is a representative firm which uses an input good k to produce output $g^s(k_{i,t}, \mathbf{x}_i)$ where t denotes the period and \mathbf{x}_i is a vector of country-specific, time-invariant characteristics. The input good may be tradable across borders (e.g. mobile capital) or may be completely immobile internationally (e.g. entrepreneurial effort). We assume that $g_k^s > 0$ and $g_{kk}^s < 0$.¹¹ The exact shape of the output function depends on the state of nature s which is initially unobserved. It is drawn in period 0 from a cumulated distribution function $F_s(s)$ where $f_s(s) = F'_s(s)$

¹⁰The yardstick competition literature following Besley and Case (1996) points out that this information externality may be used by other countries to control politicians. Mukand and Rodrik (2005) demonstrate that there is a downside to this if adopting policies is a signal for not being corrupt (as adopted policies may not be suitable for the jurisdiction under consideration).

¹¹ $g^s(\cdot)$ can be interpreted, for example, as either a decreasing returns production function with a constant price or a constant returns production with a declining price of output. For simplicity, we adopt the former nomenclature.

denotes the density function. Let s^* be the actual state of nature and $g^*(.)$ the actual output function.

The input is purchased at a cost of r_i (in the case of freely tradable input goods, the law of one price would result in $r_i = r$). Let τ_i be the source-based tax in country i . Net-of-tax profits are given by:

$$g_i^s(k_i, \mathbf{x}_i) - r_i k_i - \tau_i B_i \quad (1)$$

where B_i is the tax base which may depend on inputs k_i and output $g_i^s(k_i, \mathbf{x}_i)$ (for instance, $B_i = k_i$ in the classical tax competition framework, i.e. τ_i is a unit tax on capital use in country i).

Assuming that the firm takes r_i as given, profit maximization implies $g_k^s(k_i, \mathbf{x}_i) = r_i + \tau_i B_{ik}$. The equilibrium input price r_i may be affected by demand in other countries (e.g. the interest rate for capital could be determined on the world capital market with the capital supply being fixed as in Zodrow & Mieszkowski (1986)), therefore r_i can be expressed as a function of all tax rates and the state of nature, i.e. $r_i^s = r_i^s(\boldsymbol{\tau})$ where $\boldsymbol{\tau} = (\tau_1, \dots, \tau_i, \dots, \tau_N)$ denotes the vector of tax rates. Input demand can thus be expressed as $k_i^s(\boldsymbol{\tau})$. Accordingly, the tax base depends on both, s and $\boldsymbol{\tau}$, i.e. $B_i = B_i^s(\boldsymbol{\tau})$

Each country is ruled by a government which maximizes tax revenue. We assume that the government cannot directly observe s^* , nor can it observe the level of input demand $k_i^s(\boldsymbol{\tau})$ or the input price $r_i^s = r_i^s(\boldsymbol{\tau})$.¹² It does, however, observe tax revenues, which for i are given by:

$$\pi_{i,t} = \tau_{i,t} (B_i^s(\boldsymbol{\tau}) + \varepsilon_{i,t}) \quad (2)$$

where $\varepsilon_{i,t}$ is an idiosyncratic shock (a measurement error) with mean of zero which is drawn from a cumulated distribution function $F_\varepsilon(\varepsilon)$ where $f_\varepsilon(\varepsilon) = F'_\varepsilon(\varepsilon)$ denotes the density function. $F_\varepsilon(\varepsilon)$ is common knowledge as

¹²Note that this complete unobservability is not necessary. What is important here is that the government is not able to perfectly deduce s^* from firm choices. By making this completely unobservable, only tax revenues and tax rates provide information, streamlining the presentation of the model.

are the \mathbf{x}_i s. By observing tax revenue and the tax rate, the government can calculate the tax base (both its own and its neighbours'). This does not, however, directly identify s^* due to the white noise arising from the shocks $\varepsilon_{i,t}$. Examples of such shocks could be errors in tax collection, randomness in tax evasion, and the like. Note further that the variance of the tax base, $B_i^s(\boldsymbol{\tau}) + \varepsilon_{i,t}$, is only affected by the tax rate via the input use.

In each period $t \geq 1$, the government chooses its own tax rate $\tau_{i,t}$. Define $\boldsymbol{\tau}_{-i,t}$ as the vector of tax rates of all countries except for i in period t . The government in i estimates its expected revenue in period t , for a given tax environment $\boldsymbol{\tau}_t$, to be

$$T_{i,t}(\boldsymbol{\tau}_t) = \tau_{i,t} \int B_i^{\tilde{s}}(\boldsymbol{\tau}_t) f(\tilde{s}|I_{i,t}) d\tilde{s} \quad (3)$$

where \tilde{s} denotes potential states of nature, $I_{i,t}$ the set of information in period t , and $f(\tilde{s}|I_{i,t})$ the density of \tilde{s} conditional on $I_{i,t}$.

In line with the literature on social experimentation (Bala and Goyal 1998, and Buera et al. 2011), we will assume that the government sets $\tau_{i,t}$ in each period to maximize (3). That is, it does not willingly experiment in order to gain information. This assumption is crucial and can be justified as follows. First, political cycles and short election terms (not explicitly modeled here) may force the incumbents towards a certain degree of myopia. Second, governments may not have the analytical power to solve for optimal experimentation strategies. Third, allowing for forward-looking experimentation makes the solution incredibly difficult and, in many instances, intractable.

For expositional reasons and simplicity, we limit our analysis to cases in which the following holds.

Assumption 1 For each environment s and $\boldsymbol{\tau}_{-i,t}$, there is a unique optimal tax rate, denoted by $\tau_{i,t}^s(\boldsymbol{\tau}_{-i,t})$, that strictly decreases in s .

Since the equilibrium tax rate is mainly driven by the tax base elasticity (a typical feature of revenue maximization behavior), the state of the world

s indicates how elastic the tax base is. Depending on the shape of the tax base $B_i^s(\tau_t)$, the above assumption defines conditions for the shape of the production function. For instance, if $B_i^s(\tau_t) = k_i(\tau_t)$, Ass. 1 implies that g_{kk}^s needs to decrease in s (i.e. become larger in absolute terms). Then, an example for an output function that is compatible with Assumption 1 is $g^s(k, \mathbf{x}_i) = (k)^s$ where $s \in (0, 1)$.¹³

Expected revenue maximization in each period implies the optimality condition

$$\int B_i^{\tilde{s}}(\tau_t) f(\tilde{s}|I_{i,t}) d\tilde{s} + \tau_{i,t} \int \frac{dB_i^{\tilde{s}}(\tau_t)}{d\tau_{i,t}} f(\tilde{s}|I_{i,t}) d\tilde{s} = 0 \quad (4)$$

in each t . The first term on the left hand side represents the expected gain in tax revenue (sometimes called the 'mechanical effect' as it abstracts from behavioral adjustment). The second term depicts the expected cost of taxation, i.e. the response of the tax base to a tax rate change, over all potential states of the world \tilde{s} . The available information $I_{i,t}$ affects both, the expected gain and the expected cost of taxation. Thus, actual policy depends on the information set in each period, which is the only difference to the typical tax competition result.

The timing is as follows. In period 0, Nature draws s^* from the distribution function $F_s(s)$. In period 1, the governments simultaneously choose tax rates, $\tau_{i,1}$, and first period payoffs, $\pi_{i,1}$, are received. In period 2, each government refines its belief about the true state of nature, based on the observation of $\pi_1 = (\pi_{1,1}, \dots, \pi_{i,1}, \dots, \pi_{N,1})$ and sets a new tax rates. All subsequent periods are equal, the only difference being that the information set (i.e. the amount of information available) becomes larger over time.

¹³If $g^s(\cdot)$ is interpreted as a revenue function, i.e. the product of price and quantity, an example that is compatible with Assumption 1 is $g^s(k, \mathbf{x}_i) = (A - \frac{1}{s}k)k$ where the price of output is given by $(A - \frac{1}{s}k)$, with $A > 0$ a constant, and the quantity by k (assuming a linear production function where one unit of k produces one unit of output good).

2.2 Learning

While s^* is unobservable, the functions $F_s(\cdot)$ and $F_\varepsilon(\cdot)$ as well as the country-specific vectors \mathbf{x}_i are common knowledge. Moreover, the history of tax rates and realized tax revenues for all i are known by all players (in the terms of the learning theory, we consider a model of social experimentation with a fully connected network). As a result, the information sets are equal across countries.

By observing one's own and the neighbours' tax bases a country i will update its beliefs as information arrives. In period t , the information set has $2Nt$ elements and is given by $I_{i,t} = \{\boldsymbol{\pi}_l, \boldsymbol{\tau}_l\}_{l=1, \dots, t-1}$. The updated density is

$$f_s(\tilde{s}|I_{i,t}) = \frac{\prod_{l=1}^{t-1} \prod_i f_\pi(\pi_{i,l}(\boldsymbol{\tau}_l) | \tilde{s})}{\prod_{l=1}^{t-1} \prod_i f_\pi(\pi_{i,l}(\boldsymbol{\tau}_l))} \cdot f_s(\tilde{s}) \quad (5)$$

for $t \geq 2$, where $f_s(\tilde{s})$ is the unconditioned (a priori) density of \tilde{s} and $f_\pi(\pi_{i,t}|s) = f_\varepsilon(\pi_{i,t}/\tau_{i,t} - B_i^s(\boldsymbol{\tau}_t))$ is the derived density of $\pi_{i,t}$.

For further use, it is helpful to define the *precision of beliefs* which measures the accuracy of current beliefs about the true state of nature s^* . We define the precision of beliefs as

$$\left(\int (\tilde{s} - s^*)^2 f_s(\tilde{s}|I_{i,t}) d\tilde{s} \right)^{-1}. \quad (6)$$

With perfect learning, the precision of beliefs approaches infinity.¹⁴

2.3 Equilibrium and welfare

In the following, we will describe the sequence of equilibria over time. As is well-known, in an infinitely repeated game a broad range of allocations can be sustained as an equilibrium. We will therefore constrain ourselves to Markov perfect equilibria, where the only feature that is transferred from

¹⁴Since s^* is unknown, the expected precision may be calculated by replacing s^* with $\int \tilde{s} f_s(\tilde{s}|I_{i,t}) d\tilde{s}$. With unbiased learning, we have $\lim_{t \rightarrow \infty} \int \tilde{s} f_s(\tilde{s}|I_{i,t}) d\tilde{s} = s^*$ and the expected precision is an unbiased estimator of the precision.

period to period is the increasing size of the information set.¹⁵ With \mathbf{x}_i and s^* being time-invariant, the only feature that changes across time is the beliefs expressed in $f(\tilde{s}|I_{i,t})$ due to changes in the size of the information set $I_{i,t}$.

Definition 1 *An equilibrium is defined as a vector $\boldsymbol{\tau}_t$ which satisfies*

$$\tau_{i,t} = \arg \max T_{i,t}(\boldsymbol{\tau}_t) \quad \text{given } \boldsymbol{\tau}_{-i,t}$$

for all i .

In the following, we will assume that there exists a unique Markov-perfect equilibrium.¹⁶ As a first step, we confirm the established result that more information can only increase the precision of beliefs.

Lemma 1 *In expected terms, an increase in the number of observations (either over time or through an increase in observed countries) weakly increases the precision of beliefs.*

Proof. In expected terms, new information reduces the first part on the right hand side of (5) if $\tilde{s} \neq s^*$ and increases it if $\tilde{s} = s^*$. ■

Thus, the general finding in Bala and Goyal (1998) can be applied showing that, in fully connected networks, beliefs converge to the true state of nature.

Proposition 1 *For $t \rightarrow \infty$ and/or $N \rightarrow \infty$, learning is perfect, except for in the case below in Lemma 3.*

Proof. In line with the proof of Lemma 1 the first part on the right hand side of (5) approaches zero if $\tilde{s} \neq s^*$ and the number of periods or countries

¹⁵Note that this is also most in line with the Zodrow & Mieszkowski (1986) class of static models where, by definition, nothing carries over from one period to the next.

¹⁶Proving the existence and the uniqueness of tax competition equilibria is inherently difficult. It is usually presumed that, if the objective function (here, tax revenue) is sufficiently concave, that a unique equilibrium exists. See Bayindir-Upmann and Ziad (2005) for a discussion.

approaches infinity. In contrast, if $\tilde{s} = s^*$, this term approaches infinity.

■

As beliefs converge, so too will actions, resulting in ‘social conformism’ (Bikhchandani et al. 1992, 1998). In the tax competition framework, this means that, in the long run, our model approaches the solution described in Zodrow & Mieszkowski’s (1986) full information setting. With symmetric countries, social conformism implies complete convergence in taxes; however when the \mathbf{x}_i s differ, tax rates will generally also differ in the full-information Nash equilibrium.

We now examine the efficiency of learning in our framework. The precision of beliefs depends on tax policy itself. Moreover, the effectiveness of learning depends on the signal-to-noise ratio (SNR) which is defined as the ratio of the variance of the signal over the variance of the shock ε ,

$$SNR = \frac{\int (B^{\tilde{s}}(\boldsymbol{\tau}) - B^*(\boldsymbol{\tau}))^2 f_s(\tilde{s}|I_{i,t}) d\tilde{s}}{\sigma_\varepsilon} + 1$$

where σ_ε denotes the variance of ε . Taxes affect the SNR by changing the variance of the signal. The higher the SNR, the larger the increase in the precision of beliefs. Note that countries ignore the effect of their tax policy on learning since they only maximize current tax revenue. Moreover, it ignores the information externality on all other countries.

Lemma 2 *In equilibrium, a small tax rate change in country i that increases the signal-to-noise ratio has a positive externality on the belief precision in all countries $-i$.*

Proof. See the considerations above.

■

This is the well established positive information externality discussed by Rose-Ackerman (1980), Strumpf (2002), Kotsogiannis and Schwager (2006), Cai and Treisman (2009), and Callander and Harstad (2015), among others. It follows that, if the value of learning is positive, the pace of learning may be inefficiently slow in equilibrium because each government ignores the

beneficial effect of (marginally) increasing the SNR by adjusting its tax rate.

In the following, we will show that accounting for the resource externality inherent to tax competition can offset the benefit to learning established in that work.

2.4 Learning with fiscal externalities

The results so far are derived for both interpretations of the input k : internationally immobile entrepreneurial effort (no resource externalities) or internationally mobile capital (fiscal externalities). Now, we tighten the set of assumptions by assuming that the input good k is internationally mobile capital. Further, let the global capital supply be exogenously given as in Zodrow and Mieszkowski (1986). As is well established, a tax rate reduction by an individual country attracts capital from others, lowering their tax bases and their revenues. As such, a tax cut by one country imposes a negative fiscal externality on others.

Our first finding in this framework is that, in the special case of perfect symmetry (implying identical \mathbf{x}_i s) and fixed world supply of capital, learning may break down entirely. To see this, consider the case of $B_i^s(\boldsymbol{\tau}) = k_i(\boldsymbol{\tau})$, i.e. as in Zodrow and Mieszkowski (1986), the tax is a unit tax on capital use.

As is well-known in this model framework, with uniform tax rates across countries, the firm's capital demand is independent of the common tax rate applied in all countries.¹⁷ In such a situation, regardless of whether all countries believe that the tax base elasticity is high and apply low tax rates or that the elasticity is low and set high tax rates, the input choice k_i is the same. In other words, for given identical tax rates, all states of the world are observationally equivalent. This is summarized in Lemma 3 which provides the caveat in Proposition 1.

¹⁷A coordinated increase in tax rates drives down the interest rate one-by-one, leaving the cost of capital and thus capital demand unaffected.

Lemma 3 (Non-learning equilibrium) *In a tax competition model with given capital supply, identical countries and $B_i^s(\tau) = k_i(\tau)$, the equilibrium tax rates depend on (common) initial beliefs and no learning takes place.*

Proof. In period 1, all countries set their tax rates based on the initial belief function $F_s(s)$. The first period equilibrium yields identical tax rates and, thus, identical capital stocks in all countries. This observation does not reveal any information that allows for updating the belief function since it is in line with all potential states of the world. As a consequence, period 2 tax policies are based on the same information as period 1 policies.

■

Since learning only takes place by observing tax bases, observational equivalence implies that the true state of nature cannot be identified. False beliefs are ‘confirmed’ in the sense that the evidence does not contradict it.¹⁸ So, curiously, the classical tax competition model in its symmetric version fails in reaching the equilibrium associated with the true state of nature once one adds uncertainty about capital demand elasticity and common priors. In this case, learning requires some heterogeneity in \mathbf{x}_i or differences in the belief function across countries.

Outside of a non-learning equilibrium, however, learning takes place. In the following, we will show that, in the presence of fiscal externalities, learning about tax base elasticities may have an expected negative impact on welfare. To be precise, countries may expect their welfare to deteriorate the more information on the true state of the world s^* is revealed. For a better understanding, recall that countries set their tax rates based on expected tax bases and tax base elasticities. As a consequence, current tax revenue equals some (weighted) average tax rate times the expected tax base. In contrast, the expected future tax revenue depends on expected tax base elasticities and tax bases *multiplied* by the tax rate tailored to the then-revealed state of the world.

¹⁸Observational equivalence does not imply that an individual government is indifferent between all potential equilibria. Of course, governments will prefer an equilibrium in which all countries believe that tax base elasticity is (close to) zero.

To formalize this argument, from (4) we see that the current equilibrium optimal tax rate can be expressed as:

$$\tau_{i,t}^* = \frac{\mathbf{E} [B_i^{\tilde{s}}(\tau_t^*)]}{-\mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} \right]} \quad (7)$$

where $\mathbf{E} [B_i^{\tilde{s}}(\tau_t^*)] = \int B_i^{\tilde{s}}(\tau_t^*) f(\tilde{s}|I_{i,t}) d\tilde{s}$ and $\mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} \right] = \int \frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} f(\tilde{s}|I_{i,t}) d\tilde{s}$. Expected tax revenue without learning is thus given by $\tau_{i,t}^* \mathbf{E} [B_i^{\tilde{s}}(\tau_t^*)]$. In contrast, expected tax revenue with learning is given by $\mathbf{E} [\tau_{i,t}^{\tilde{s}*} B_i^{\tilde{s}}(\tau_t^*)]$ where $\tau_{i,t}^{\tilde{s}*}$ denotes the optimal tax rate for the case in which \tilde{s} is revealed as the true state of the world. We may therefore define the expected value of perfect learning as follows.

Definition 2 *The expected value of perfect learning is given by the difference*

$$\mathbf{E} [\tau_{i,t}^{\tilde{s}*} B_i^{\tilde{s}}(\tau_t^*)] - \tau_{i,t}^* \mathbf{E} [B_i^{\tilde{s}}(\tau_t^*)].$$

The value of learning can be positive and negative, as the following examples show. To start with, consider the case of Zodrow and Mieszkowski (1986) where $B_i = k_i$ and countries are perfectly symmetric. Then, the tax base B is always equal to \bar{K}/N , i.e. does not depend on the actual state of nature.¹⁹ The value of learning thus can be expressed as $(\mathbf{E} [\tau_{i,t}^{\tilde{s}*}] - \tau_{i,t}^*) \bar{K}/N$ – which can be shown to be positive as follows. If \tilde{s} has turned out to be the true state of nature, the equilibrium tax rate satisfies $B_i^{\tilde{s}}(\tau_t^{\tilde{s}*}) + \tau_{i,t}^{\tilde{s}*} \frac{dB_{ik}^{\tilde{s}}(\tau_t^{\tilde{s}*})}{d\tau_{i,t}^{\tilde{s}*}} = 0$. Integrating over all \tilde{s} gives $\mathbf{E} [B_i^{\tilde{s}}(\tau_t^*)] + \mathbf{E} \left[\tau_{i,t}^{\tilde{s}*} \frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^{\tilde{s}*}} \right] = 0$. Eq. (7) implies $\mathbf{E} [B_i^{\tilde{s}}(\tau_t^*)] + \tau_{i,t}^* \mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} \right] = 0$. Therefore, $\tau_{i,t}^* \mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} \right] = \mathbf{E} \left[\tau_{i,t}^{\tilde{s}*} \frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^{\tilde{s}*}} \right]$. With $\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} = \frac{dk(\tau_t^*)}{d\tau_{i,t}^*} = \frac{1}{g_{kk}^{\tilde{s}}}$ which does not depend on the specific level of $\tau_{i,t}^{\tilde{s}*}$, we can write $\tau_{i,t}^* \mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} \right] =$

¹⁹ The reason is that for every potential value of s , a symmetric Nash equilibrium emerges with uniform tax rates across countries and, thus, identical capital stocks everywhere.

$$\mathbf{E} \left[\tau_{i,t}^{\tilde{s}*} \right] \mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^{\tilde{s}*})}{d\tau_{i,t}^{\tilde{s}*}} \right] + \mathbf{cov} \left[\tau_{i,t}^{\tilde{s}*}, \frac{dB_{ik}^{\tilde{s}}(\tau_t^{\tilde{s}*})}{d\tau_{i,t}^{\tilde{s}*}} \right] \text{ and}$$

$$(\tau_{i,t}^* - \mathbf{E} [\tau_{i,t}^{\tilde{s}*}]) \mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} \right] = \mathbf{cov} \left[\tau_{i,t}^{\tilde{s}*}, \frac{dB_{ik}^{\tilde{s}}(\tau_t^{\tilde{s}*})}{d\tau_{i,t}^{\tilde{s}*}} \right]. \quad (8)$$

With $\frac{dB_{ik}^{\tilde{s}}(\tau_t^{\tilde{s}*})}{d\tau_{i,t}^{\tilde{s}*}} = \frac{1}{g_{kk}^{\tilde{s}}}$, $\mathbf{E} \left[\frac{dB_{ik}^{\tilde{s}}(\tau_t^*)}{d\tau_{i,t}^*} \right] < 0$ and $\mathbf{cov} \left[\tau_{i,t}^{\tilde{s}*}, \frac{dB_{ik}^{\tilde{s}}(\tau_t^{\tilde{s}*})}{d\tau_{i,t}^{\tilde{s}*}} \right] > 0$ follows that $\tau_{i,t}^* < \mathbf{E} [\tau_{i,t}^{\tilde{s}*}]$.²⁰ That is, the expected value of learning in the classical tax competition model is positive.

Now, consider a class of models in which $B_i = g_i(k_i)$, i.e. the tax base equals output. By assumption, input cost cannot be deducted from the tax base (which is in line with the usual tax treatment of equity financing).²¹ A positive value of learning can be shown for $g_i(k_i) = k^s$ or $g_i(k_i) = (A - \frac{1}{s}k_i)k_i$. In contrast, a negative value of learning can be demonstrated for $g_i(k_i) = D - ds + k^s$ for appropriate levels of D and d (see the Appendix for detailed calculations). Thus, the value of learning can be positive and negative, depending on the production function and the dimension of heterogeneity.

Moreover, as is well-known from the literature on asymmetric tax competition (Bucovetsky 1991), certain countries may gain from tax competition while others lose. Thus, the above condition may hold for some countries and not for others. Specifically, small countries may gain from high capital mobility whereas large countries may be expected to lose.

Proposition 2 *Assume that the economy is not in a non-learning equilibrium. Then, uncoordinated tax rate setting implies inefficiently fast (slow) learning if the value of learning (as defined in Def. 2) is negative (positive). Specifically, learning in the classical tax competition framework is inefficiently slow.*

²⁰For simplicity, we assumed that countries ignore their tax policy's effect on the interest rate. The results, however, do not change if this effect is accounted for.

²¹It is straightforward to introduce tax deductibility of input cost into the model as long as we rule out full deductibility (which would eliminate any tax effects); however, it adds no further insights.

Proof. By assumption, each country sets its tax rate to maximize current tax revenue given the available information. If learning has an expected negative effect on equilibrium tax revenue, a small tax rate change that reduces the SNR, starting from the uncoordinated equilibrium has a zero expected impact on own revenue and a positive effect on future revenue of all countries by preventing learning. Such a tax change would therefore be welfare-enhancing (i.e. revenue-increasing compared to a situation with faster learning) in expected terms.

■

Further, the pace of learning depends on the heterogeneity of tax rates. As discussed above, uniform tax rates do not allow for learning as shifts in the common level of tax rates do not affect the tax bases. Just as learning is faster with more countries, a more heterogeneous set of tax rates leads to faster learning.

Corollary 1 *Assume that the economy is not in a non-learning equilibrium and the value of learning (as defined in Def. 2) is negative. Then, tax rates are inefficiently heterogeneous.*

Proof. See the considerations above.

■

The above Corollary is in direct contrast to the experimentation in a federation literature where there is generally too little policy heterogeneity. This is because, in our setting, the conditions laid out in Proposition 2 define an ‘ignorance is bliss’ scenario in which learning only pushes the equilibrium towards a less efficient full-information outcome at a faster rate.

Under revenue maximization, equilibrium tax rates are almost always too low (due to the negative externality of tax cuts). The information externality may provide another reason why tax rates are too low if the SNR decreases in tax rates. That is, higher tax rates would slow down harmful learning.

Corollary 2 *Assume that the economy is not in a non-learning equilibrium and the value of learning (as defined in Def. 2) is negative. Then, tax rates are always and unambiguously too low if the SNR decreases in the tax rates (for a given degree of heterogeneity).*

Proof. Omitted.

■

The above Corollary points to equilibrium tax rates which are possibly too high in the learning transition. This can be illustrated with the following thought experiment. Assume that the countries' governments have a one-time opportunity for coordinated increase or decrease in tax rates. Afterwards, tax competition would go on uncoordinatedly. Then, the following trade-off emerges: Governments would have to weigh the one-time gain in tax revenue against a permanent increase in expectedly harmful information (i.e. persistently lower tax revenues in the future). If the latter is large enough, it may be optimal to decrease tax rates.

3 A reinterpretation of the empirical evidence

As indicated in the introduction, the evidence on international tax rate setting is usually interpreted as supporting the classical tax competition theory (countries compete for mobile resources) or the yardstick competition theory (politicians adjust policies to get re-elected). In this section, we argue that the policy learning model can rationalize the existing evidence on tax competition and, thus, provides a complementary perspective to the conventional interpretations of the empirical findings.²²

To see this, consider the following variant of the above model. Assume that the input good k is entrepreneurial effort which is immobile across borders. Now, there are no resource externalities since tax bases only depend on the entrepreneurial effort (we assume here that if the entrepreneur reduces her effort, her leisure increases which has no externality on other individuals). If s^* were known, tax policy in country i would be completely independent of the tax policy in other countries. With an unknown state of nature, however, country i has an incentive to observe tax policy and revenues in other countries, since they reveal information on the common state of nature.

²²It needs to be noted that, in other policy areas, a wider variety of policy diffusion channels are considered, see Shipan and Volden (2008).

Now, due to the absence of fiscal externalities, there are no cross-border tax effects, i.e. $k_i^s(\tau_t) = k_i^s(\tau_{i,t})$, and the value of learning is strictly positive. The reason is that learning reduces, in expectation, the difference between the current tax rate and the optimal tax rate. In line with the existing literature, learning in this situation is too slow since tax policy has a positive information externality (see Kotsogiannis and Schwager 2006), i.e. an increase or a decrease in a country's tax rate increases the precision of beliefs in all other countries. Thus learning has an unambiguously positive welfare impact.

The crucial point we wish to make, however, is that in such a model tax changes are correlated across borders although there are no resource externalities. To see this, note that information sets are equal across countries. This implies that changes in the information set will have similar consequences in each country (or equal consequences if the \mathbf{x}_i are identical). Since there is a unique optimal tax for each s which monotonously increases in s , news will trigger either tax increases or tax decreases in all countries.

Lemma 4 (Correlated tax changes) *Tax changes are correlated internationally, even in the absence of resource externalities. In the case of perfect symmetry, correlation is perfect, i.e. tax rates are equal in each period.*

Proof. Omitted. ■

Thus, countries set policies as if they react to each other in a strategic sense, even though no such motive exists.²³ Put differently, although each country's payoff is completely independent of the taxes set elsewhere, those taxes alter the information set it has and therefore the tax it chooses. In what follows, we will reconsider the evidence in favor of tax competition and analyze whether the existing evidence allows for a clear-cut identification of the underlying model.

Before we do so, we emphasize a methodological point. The empirical literature is, of course, interested in *causation*. The model above, in the

²³More subtly, an indirect implication of Lemma 1 is that the learning curve is concave, i.e. the largest gains in belief precision occur at the beginning. Thus, to the degree that beliefs become more precise over time, one should expect tax changes to become less frequent and less pronounced over time.

version with observable outcomes, derives *correlation*: news incorporated in lagged outcomes cause both tax rates to change. Still, with the methods used in the empirical studies covered here, the pattern of tax rates will look as if neighbouring tax changes *cause* tax changes at home. In econometric terms, there is a omitted variable (observation of outcomes in the preceding period) not controlled for in the empirical design. Indeed this is the point of the criticism of Gibbons and Overman (2012) who take the empirical tax competition literature to task over identification issues.

For instance, Devereux, et al. (2008) regress statutory tax rates on a weighted average of neighbouring tax rates from the same period. In their preferred estimation, a one percentage point increase in the neighbours' tax rate increases the domestic tax rate by 0.67 percentage points. This finding could be interpreted as being in line with our model. By observing outcomes, all countries adjust their tax rates in the same direction albeit not at the same rate (as country characteristics may differ and states of nature may be heterogeneous).²⁴

Before we propose two extensions of the model framework in order to deal with specific features of the empirical literature, we summarize our argument and its implications. The tax competition model and the learning-without-fiscal-externalities model are, in many aspects, observationally equivalent. However, the policy implications differ diametrically. The theory of tax competition suggests that, under some circumstances, tax harmonization may yield a Pareto improvement. From the viewpoint of the model outlined above, however, policy harmonization may decrease efficiency as it prevents countries from learning.

3.1 Time lags

Overesch and Rincke (2011) regress statutory tax rates on their own lag (since tax rates are sticky) and the neighbours' lagged tax rate measure

²⁴In order to identify evidence for tax competition instead of evidence for other theories, e.g. yardstick competition, Devereux, et al. (2008) control for openness. They find some evidence that tax rate interactions are stronger in the absence of capital controls. However, Overesch and Rincke (2011) apply a different method and cannot find an impact of openness on tax rate interactions.

(an average weighted by distance). Their estimation results imply that the long-run results are quantitatively in line with those found by Devereux et al. (2008). Interestingly, though, it is the lagged variables that show up to have significant impact. In order to reconcile this finding with the above policy learning model, we modify the model as follows.

So far, the model is based on the assumption that countries can observe their neighbours' tax revenues. This may be considered a strong assumption given that tax data is hard to collect, difficult to interpret and is usually only available after a lag of two or more years. We therefore change the above model by assuming that a country's revenues cannot be immediately observed by its neighbours; tax rates, however, can be observed as before. Thus, we now consider a social learning model (see e.g. Gale and Kariv (2003)).

Now, the timing of the model is as follows. In period 1, the government decides on a tax rate $\tau_{i,1}$ after which the first period tax base, $\pi_{i,1}$, is realized. In period 2, the government refines its beliefs based on the observation of $\pi_{i,1}$. In period 3, beliefs are updated again, but now the neighbours' tax rate choices of the preceding period contain information on the experience made two periods before, i.e. $\pi_{-i,1}$. In all subsequent periods, country i can observe and may perfectly infer their tax bases from two periods before.

The crucial difference to the case of observable outcomes is that information sets are no longer equal across countries since each country has one piece of information (its own current outcome) that other countries lack. Because countries base their tax rate choice on all available information, tax rates can only be equal by chance. As a consequence, a non-learning equilibrium cannot exist. Once tax rates are set, other countries can infer the missing information from the tax rate choice.²⁵ That is, all information available from $t - 2$ is reflected in tax rates in t .

²⁵To be precise, updating works as follows. In period 1, the information set consists of a single element, $I_{i,1} = \{\pi_{i,0}(\boldsymbol{\tau})\}$. Given this, $f_s(\tilde{s}|I_{i,1}) = \frac{f_\pi(\pi_{i,0}^s(\boldsymbol{\tau}_0)|\tilde{s})}{f_\pi(\pi_{i,0}^s(\boldsymbol{\tau}_0))} \cdot f_s(\tilde{s})$. In period 2 and all subsequent periods, country i infers its neighbours' payoffs from two periods before, i.e. $\pi_{-i,0}$ for $t = 2$ and, generally, $\pi_{-i,t-2}$ for period t . In period t , the information set has $N(t - 2) + 1$ elements and is given by $I_{i,t} = \{\pi_{i,t-1}(\boldsymbol{\tau}_{t-1}), \pi_{j,l}(\boldsymbol{\tau}_l)\}_{j=1,\dots,N}^{l=0,\dots,t-2}$. The

As a consequence, a tax rate increase (decrease) in i in $t - 1$ causes a tax rate increase (decrease) in all other countries $-i$ in t . There is an important but subtle distinction in this between the social experimentation and social learning settings. With social experimentation, tax rates in one period are set based on information gleaned from outcomes which occurred in prior periods. Since information sets are the same across countries, taxes in a given period are *correlated*. With social learning, the tax rates themselves convey information. Therefore, tax rate changes in $t - 1$ *cause* neighbouring tax rate changes in t – which is reflected in the findings in Overesch and Rincke (2011).

3.2 Spatial and similarity lags

Our model may also account for spatial lags using weighted averages.²⁶ When the states of nature in neighbouring countries are more strongly correlated than in countries which are more distant from each other, a country's tax policy would react more to its close neighbours' tax rates than to more distant ones.

Therefore, we relax the assumption that the state of nature is equal across countries. Instead, let $\mathbf{s} = (s_1, \dots, s_N)$ denote the vector of states of natures where, although country i 's and country j 's state of nature may differ, there may be correlation. In line with the baseline, let the true states of nature vector, \mathbf{s}^* , be a realization of a common distribution $H(\mathbf{s})$ with density function $h_{\mathbf{s}}(\mathbf{s})$.

updated density is

$$f_{s,i}(\tilde{s}|I_{i,t}) = \frac{\prod_{l=0}^{t-2} \prod_i f_{\pi}(\pi_{i,l}(\boldsymbol{\tau}_l) | \tilde{s})}{\prod_{l=0}^{t-2} \prod_i f_{\pi}(\pi_{i,l}(\boldsymbol{\tau}_l))} \cdot \frac{f_{\pi}(\pi_{i,t-1}(\boldsymbol{\tau}_{t-1}) | \tilde{s})}{f_{\pi}(\pi_{i,t-1}(\boldsymbol{\tau}_{t-1}))} \cdot f_s(\tilde{s})$$

for $t \geq 2$, with the second term being country specific and the first one equal across country.

²⁶In this empirical literature, the tax rates of other countries are typically aggregated into a weighted average that is used as an (endogenous) explanatory variable called a spatial lag. See Devereux, et al. (2008) for a presentation of the methodology in the context of international taxation.

As before, countries observe their own and their neighbours' policy outcomes and update their beliefs. With $\tilde{\mathbf{s}}$ denoting a potential realization of \mathbf{s} , the updated density is given by

$$h_{\mathbf{s}}(\tilde{\mathbf{s}}|I_{i,t}) = \frac{\prod_{l=0}^{t-1} h_{\boldsymbol{\pi}}(\boldsymbol{\pi}_l(\boldsymbol{\tau}_l) | \tilde{\mathbf{s}})}{\prod_{l=0}^{t-1} h_{\boldsymbol{\pi}}(\boldsymbol{\pi}_l(\boldsymbol{\tau}_l))} \cdot h_{\mathbf{s}}(\tilde{\mathbf{s}}) \quad (9)$$

where $h_{\boldsymbol{\pi}}(\boldsymbol{\pi}_l(\boldsymbol{\tau}_l) | \tilde{\mathbf{s}})$ is the derived density of the outcome variable vector $\boldsymbol{\pi}_l(\boldsymbol{\tau}_l)$.

Since an individual country is ultimately interested in its own payoffs, it sets payoffs to maximize expected revenue given by

$$T_{i,t}(\boldsymbol{\tau}) = \tau_{i,t} \int B_i^{\tilde{\mathbf{s}}}(\boldsymbol{\tau}_t) h_{s,i}(\tilde{\mathbf{s}}|I_{i,t}) d\tilde{\mathbf{s}} \quad (10)$$

with $h_{s,i}(\tilde{\mathbf{s}}|I_{i,t}) = \int_{s_1} \int_{s_2} \dots \int_{s_n} h_{\mathbf{s}}(\tilde{\mathbf{s}}|I_{i,t}) d\tilde{s}_n \dots d\tilde{s}_2 d\tilde{s}_1$, i.e. the integral over all states of nature expect for s_i .

How much country i may increase its precision of beliefs by observing other countries' policy outcomes depends on the degree to which updates in those countries' beliefs yield information on country i 's state of nature. If states of natures are completely independent variables, updates do not include useful information. However, if the pairwise covariances in $H(\mathbf{s})$ is not strictly zero, observations from other countries yield valuable information for country i . The stronger the covariance, the more information is implied.

As a consequence, an observation of a neighbour's policy outcome only increases the precision of beliefs if the covariance of the neighbour's and one's own state of nature is non-zero. Moreover, the increase in the belief precision increases in the covariance.

In terms of measurement, this can give rise to spatial lags that include only the policies of neighboring countries, as done by Altshuler and Goodspeed (2007), or those weighting by inverse distance (e.g. Overesch and Rinke (2011)). An alternative is the use of political similarity weights, as

utilized in Davies and Klasen (forthcoming), who use a measure of political affinity based on United Nations voting as a weighting scheme for their study of overseas development assistance donations. Taking this idea a step farther, one can imagine a setting in which a subset of nations, for instance the EU, closely observe each other, but non-members are ignored. In this case, 'clubs' occur with different learning pattern inside and outside the EU. This would then fit the pattern found by Redoano (2014) and Davies and Voget (2009), who find that while non-EU members respond equally to the corporate taxes of both EU and non-EU countries, members respond less to the taxes of non-members than to those of members.

4 Conclusion

This paper combines a standard model of tax competition with one of social experimentation in order to investigate the interaction of the former's fiscal externality with the latter's information externality. We do so in a model of tax competition in which the tax base elasticity is an *a priori* unknown variable which can be learned over time. In line with other policy learning models, learning is gradual, i.e. countries do not jump to a new equilibrium once there is a change in the underlying environment. Instead, they gradually approach the new equilibrium with learning from one's neighbours' experiences increasing the convergence of beliefs, and thus taxes, to the full-information equilibrium.

We contribute to the literature by showing that, in the presence of resource externalities, learning may have a negative value and may thus be inefficiently fast. This occurs because, if the tax base is more elastic than is initially believed, in the short run before learning is perfect, higher taxes are used, resulting in higher welfare. However, as countries learn of the true tax base elasticity, they revise their beliefs upwards and their tax rates downwards, lowering welfare. Thus, by shorting the convergence process to the well-known race to the bottom equilibrium, learning is inefficiently fast. This is in contrast to models where information externalities are the only coordination failure and thus learning is inefficiently slow.

Moreover, if states of nature are similar across countries, tax policies interact even in the absence of cross-border externalities. This would give rise to the patterns the empirical tax literature generally ascribed to tax competition for mobile resources or to yardstick competition. Policy learning provides another interpretation of the empirical literature's findings.

Although our model uses the elasticity of firm decisions as the unknown variable, the results here are easily applicable to a broad variety of situations, including the productivity of research and development, the ease of profit shifting, or the costs associated with relocation. Therefore while we are not suggesting that traditional motives for the correlation of taxes across borders (such as tax competition) are not important, we hope that our model provides insights that can further the debate on international taxation.

Appendix

In this Appendix, we consider three examples with $B_i = g(k_i)$, i.e. the tax base is output (capital inputs are not deductible which is the typical treatment for equity). We focus on symmetric equilibria in which the capital stock is equal across countries and independent of s . Then, the following variant of eq. (8) holds:

$$\left(\frac{\tau_{i,t}^*}{1 - \tau_{i,t}^*} - \mathbf{E} \left[\frac{\tau_i^{\tilde{s}*}}{1 - \tau_i^{\tilde{s}*}} \right] \right) \mathbf{E} \left[\frac{(g_k)^2}{g_{kk}} \right] = \mathbf{cov} \left[\frac{\tau_i^{\tilde{s}*}}{1 - \tau_i^{\tilde{s}*}}, \frac{(g_k)^2}{g_{kk}} \right] \quad (11)$$

With $\mathbf{E} \left[\frac{(g_k)^2}{g_{kk}} \right] < 0$, it follows that $\mathbf{E} \left[\tau_{i,t}^{\tilde{s}*} \right] > \tau_{i,t}^*$ (i.e. the value of learning is positive) if $\mathbf{cov} \left[\tau_{i,t}^{\tilde{s}*}, \frac{(g_k)^2}{g_{kk}} \right] > 0$ and vice versa.

Example 1: $g(k_i) = k_i^s$ with $s \in (0, 1)$

The first order condition of the revenue maximization problem is, for a given \tilde{s} , given by $k_i^{\tilde{s}} + \frac{\tau_i^{\tilde{s}*}}{1 - \tau_i^{\tilde{s}*}} \frac{\tilde{s}}{\tilde{s} - 1} k_i^{\tilde{s}} = 0$. It follows that $\frac{\tau_i^{\tilde{s}*}}{1 - \tau_i^{\tilde{s}*}} = \frac{1 - \tilde{s}}{\tilde{s}}$. With $\frac{(g_k)^2}{g_{kk}} = \frac{\tilde{s}}{\tilde{s} - 1} k_i^{\tilde{s}}$, it follows that $\mathbf{cov} \left[\frac{\tau_i^{\tilde{s}*}}{1 - \tau_i^{\tilde{s}*}}, \frac{(g_k)^2}{g_{kk}} \right] > 0$. Thus, in this example,

the value of learning is positive.

Example 2: $g(k_i) = (A - \frac{1}{s}k_i) k_i$ **with** $s > 0$

The first order condition of the revenue maximization problem is, for a given \tilde{s} , given by $(A - \frac{1}{s}k) k + \frac{\tau_i^{\tilde{s}*}}{1-\tau_i^{\tilde{s}*}} \frac{(A - \frac{2}{s}k)^2}{-\frac{2}{s}} = 0$. It follows that $\frac{\tau_i^{\tilde{s}*}}{1-\tau_i^{\tilde{s}*}} = \frac{\frac{2}{s}(A - \frac{1}{s}k)k}{(A - \frac{2}{s}k)^2}$. With $\frac{(g_k)^2}{g_{kk}} = \frac{(A - \frac{2}{s}k)^2}{-\frac{2}{s}}$, it follows that $\mathbf{cov} \left[\frac{\tau_i^{\tilde{s}*}}{1-\tau_i^{\tilde{s}*}}, \frac{(g_k)^2}{g_{kk}} \right] > 0$. Thus, in this example, the value of learning is positive.

Example 3: $g(k_i) = D - ds + k_i^{1/s}$ **with** $s > 1$ **and** $D, d > 0$

The first order condition of the revenue maximization problem is, for a given \tilde{s} , given by $D - ds + k_i^{1/s} + \frac{\tau_i^{\tilde{s}*}}{1-\tau_i^{\tilde{s}*}} \frac{1}{1-s} k_i^{1/s} = 0$. It follows that $\frac{\tau_i^{\tilde{s}*}}{1-\tau_i^{\tilde{s}*}} = (s-1) \frac{D-ds+k_i^{1/s}}{k_i^{1/s}}$. With $\frac{(g_k)^2}{g_{kk}} = \frac{1}{1-s} k_i^{1/s}$, there are certain parameter ranges for D and d , in which $\mathbf{cov} \left[\frac{\tau_i^{\tilde{s}*}}{1-\tau_i^{\tilde{s}*}}, \frac{1}{1-s} k_i^{1/s} \right] < 0$. Thus, in this example, the value of learning is negative.

References

- [1] Altshuler, R. and Goodspeed, T. (2007). Follow the Leader: Evidence on European and U.S. Tax Competition, Working Paper.
- [2] Bala, V. and Goyal, S. (1998). Learning from Neighbours, Review of Economic Studies 65(3): 595-621.
- [3] Bayindir-Upmann, T., Ziad, A. (2005). Existence of equilibria in a basic tax competition model. Regional Science and Urban Economics 35 (1), 1-22.
- [4] Besley, T. and Case, A. (1995). Incumbent Behavior: Vote-Seeking, Tax-Setting, and Yardstick Competition, American Economic Review 85(1): 25-45.
- [5] Bikhchandani, S. Hirshleifer, D. and Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades, Journal of Political Economy 100(5): 992-1026.

- [6] Bikhchandani, S. Hirshleifer, D. and Welch, I. (1998). Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades, *Journal of Economic Perspectives* 12: 151–170.
- [7] Brueckner, J.K. (2003). Strategic Interaction Among Governments: An Overview of Empirical Studies, *International Regional Science Review* 26(2): 175-188.
- [8] Bucovetsky, S. (1991). Asymmetric Tax Competition, *Journal of Urban Economics* 30(2): 167-181.
- [9] Buera, F. J.; Monge-Naranjo, A.; Primiceri, G.E. (2011). Learning the Wealth of Nations, *Econometrica* 79(1): 1–45.
- [10] Cai, H. and Treisman, D. (2009). Political Decentralization and Policy Experimentation, *Quarterly Journal of Political Science* 4(1): 35-58.
- [11] Callander, S. and Harstad, B. (2015). Experimentation in Federal Systems, *Quarterly Journal of Economics* 130(2): 951-1002.
- [12] Davies, R.B. (2005). State Tax Competition for Foreign Direct Investment: A Winnable War?, *Journal of International Economics*, 67(2), pp. 498-512.
- [13] Davies Ronald B. and Stephan Klasen (forthcoming). Of Donor Coordination, Free-Riding, Darlings, and Orphans: The Dependence of Bilateral Aid on Other Bilateral Giving, *Scandinavian Journal of Economics*.
- [14] Davies, R.B. and Voget, J. (2009). Tax Competition in an Expanding European Union, *The Institute for International Integration Studies Discussion Paper Series No. 276*.
- [15] Devereux, M.P., B. Lockwood, and M. Redoano (2008). Do countries compete over corporate tax rates?, *Journal of Public Economics* 92(5–6): 1210–1235.
- [16] Feldstein, M.S. (1999). Tax Avoidance and the Deadweight Loss of the Income Tax, *Review of Economics and Statistics* 81(4): 674-680.
- [17] Fredriksson, P.G., List, J.A., and Millimet, D.L. (2003) Bureaucratic Corruption, Environmental Policy and Inbound US FDI: Theory and Evidence, *Journal of Public Economics*, 87, 1407-1430.

- [18] Gale, D. (1996). What Have We Learned from Social Learning, *European Economic Review* 40(3-5): 617-628.
- [19] Gale, D. and Kariv, S. (2003). Bayesian Learning in Social Networks, *Games and Economic Behavior* 45(2): 329-46.
- [20] Gibbons, S. and Overman, H. (2012). Mostly Pointless Spatial Econometrics. *Journal of Regional Science*, 52(2), 172-191.
- [21] Heinemann, F.; Overesch, M. and Rincke, J. (2010). Rate-Cutting Tax Reforms and Corporate Tax Competition in Europe, *Economics and Politics* 22(3): 498-518.
- [22] Heyndels, Bruno and Jef Vuchelen (1998). Tax Mimicking among Belgian Municipalities, *National Tax Journal*, 51: 89-101.
- [23] Kotsogiannis, C. and Schwager, R. (2006). On the Incentives to Experiment in Federations, *Journal of Urban Economics* 60(3): 484-97.
- [24] Levinson, Arik. (2003). Environmental Regulatory Competition: A Status Report and Some New Evidence, *National Tax Journal*, 56(1), 91-106.
- [25] Mintz, Jack and Michael Smart (2004). Income shifting, investment, and tax competition: theory and evidence from provincial taxation in Canada, *Journal of Public Economics*, 88: 1149-1168.
- [26] Mukand, S. W. and Rodrik, D. (2005). In Search of the Holy Grail: Policy Convergence, Experimentation, and Economic Performance, *American Economic Review* 95(1): 374-83.
- [27] Overesch, M. and Rincke, J. (2009). Competition from low-wage countries and the decline of corporate tax rates: evidence from European integration, *World Economy* 32(9):1348–1364.
- [28] Overesch, M. and Rincke, J. (2011). What Drives Corporate Tax Rates Down? A Reassessment of Globalization, Tax Competition, and Dynamic Adjustment to Shocks, *Scandinavian Journal of Economics* 113(3): 579-602.
- [29] Redoano, M. (2014). Tax Competition among European Countries. Does the EU Matter? *European Journal of Political Economy* 34: 353–371.

- [30] Rose-Ackerman, S. (1980). Risk taking and reelection: Does federalism promote innovation? *Journal of Legal Studies* 9: 593–616.
- [31] Saez, E. (2001). Using Elasticities to Derive Optimal Income Tax Rates, *Review of Economic Studies* 68: 205-229.
- [32] Salmon, P. (1987). Decentralization as an Incentive Scheme, *Oxford Review of Economic Policy* 3: 24–43.
- [33] Shipan, C. R. and Volden, C. (2008). The Mechanisms of Policy Diffusion, *American Journal of Political Science* 52(4): 840-857.
- [34] Strumpf, K. S. (2002). Does Government Decentralization Increase Policy Innovation?, *Journal of Public Economic Theory* 4(2): 207-41.
- [35] Vives, X. (1996). Social Learning and Rational Expectations, *European Economic Review* 40(3-5): 589-601.
- [36] Wilson, J.D. (1986). A Theory of Interregional Tax Competition, *Journal of Urban Economics* 19(3): 296-315.
- [37] Zodrow, G.R. and Mieszkowski, P. (1986). Pigou, Tiebout, Property Taxation, and the Underprovision of Local Public Goods, *Journal of Urban Economics* 19(3): 356-371.