Dynamics of the market for corporate tax-avoidance advice

Kai A. Konrad*

May 24, 2017

Abstract

This paper studies the dynamics of the market for corporate tax avoidance products with tax accounting firms that may innovate new tax-avoidance products and other firms that may imitate these products with some time lag. The government can regulate and implement effective anti-avoidance regulation, also with some lag. A moderate regulatory time lag may lead to a stationary equilibrium with innovation and regulation in each period. An even larger regulatory time lag may lead to an equilibrium with fluctuations. Periods with much tax avoidance at low prices and no innovation interchange with periods with a moderate level of tax avoidance, much innovation, high profits in the accounting sector, and regulatory activity. The equilibrium analysis highlights the existence of synergies between highly innovative tax accounting firms and governmental tax regulation, with the innovative tax-accounting firms benefitting from governmental regulation.

Keywords: corporate taxation, tax avoidance, anti-tax-avoidance regulation, innovation, competition among tax accountants

JEL classification codes: M48, H26

*Max Planck Institute for Tax Law and Public Finance, Marstallplatz 1, D-80539 Munich, Germany. Email: kai.konrad@tax.mpg.de. I thank Michael P. Devereux and James R. Hines Jr. for helpful discussions at an early stage of this research project, and Tim Stolper for very helpful comments. The usual caveat applies.

1 Introduction

This paper considers the nexus between the tax accounting industry, governments, and multinational corporate firms (MNEs). It studies the dynamics of a market for tax-avoidance products with innovating and imitating tax-accounting firms, and with anti-avoidance tax legislation/regulation. Expert tax accounting firms offer tax optimization advice to MNEs. Some tax accounting firms may search for and innovate new tax optimization products, other firms may imitate such products, with some time delay. Governments react, possibly with some time lag, and implement anti-tax-avoidance regulation. Governments may be unable to detect and ban the newest tax avoidance models, particularly if they have never been used and have never been observed in operation. Governments may, however, choose to close the loopholes once they are innovated and used for some time. Anti-tax-avoidance legislation reduces the variety of applicable tax avoidance products and may reduce or prevent competition from accounting firms with lower innovative ability. The study highlights the crucial role of the speed by which governments can react to tax avoidance models that exist and are in use. If governments can react instantaneously, a stationary equilibrium may emerge that is characterized by full compliance. If the government needs some time lag, then an equilibrium with permanent innovation and governmental regulation may emerge. If the time lag is even longer, periods with much tax avoidance at low prices and no innovation may alternate with periods with a moderate level of tax avoidance, much innovation, high profits in the accounting sector, and high regulatory anti-tax-avoidance activity. The study also reveals a possible congruence of the interests of highly innovative tax accounting firms on the one side and of the governments on the other side.

The analysis makes a number of assumptions that have empirical counterparts and have been discussed in the existing literature. We discuss these briefly in what follows.

- Tax optimization by multinational firms (MNEs) is observed worldwide. To describe the demand for tax planning models, Slemrod (2004) coined the term "The economics of corporate tax selfishness" as a logical implication of shareholder value maximization: the desire and the obligation of company management to reduce the tax burden of the corporation and its shareholders. The various empirical observations suggest firm heterogeneity as regards their ability or willingness to, and their cost of using a tax sheltering/avoidance tool.[1]

- The literature documents an intimate role of tax professionals in the provision of tax optimization models to MNEs.[2] We take this as evidence of a market relationship between professional tax experts as sellers and corporate firms as customers in which tax-optimization solutions are the market product.

- Bankman (1999, p. 1790) highlights the importance of economies of scale which accounting firms may have: "Tax shelters are expensive to develop; these promoters are able to sell development costs in a way that in-house counsel cannot replicate." We take this as evidence of a market where the variable cost of making a given blueprint for tax optimization available to an additional customer is very low, and as the reason why MNEs do not innovate their tax avoidance models themselves, but purchase the advice of expert tax accountants.

- Governments in OECD countries may identify tax-avoidance schemes that are in use by MNEs and the tax loopholes they are based on. This may occur with some time lag. Depending on the avoidance

[1] For factors and firm characteristics that make it more or less likely for firms to use a tax shelter see Lisowsky (2010). Dharmapala (2014) surveys the empirical evidence on tax-optimizing behavior.

[2] One of the first writers highlighting the role of tax professionals for tax avoidance is Bankman (1999). Slemrod and Venkatesh (2002) describe the type and size of firms' expenditure for the use of tax professionals. They also report that a majority of the respondent firms were either approached by tax professionals that advertised tax shelters or approached tax professionals on this matter. Slemrod (2004) concludes that "undoubtedly, nearly all large corporations have been solicited by tax shelter promoters, and nearly all have considered pursuing them." Graham et al. (2014) can be interpreted along similar lines. In their paper they report survey-based evidence: "5.2 percent of our respondents answered that tax planning strategies are 'always' pitched as a way to increase earnings, 26.8 percent said it was 'often' the case, 50.7 percent responded that this was 'sometimes' the marketing strategy, and 17.4 percent said that this 'never' was part of the pitch."

scheme, they may change their national legislation, their international tax treaties or persuade other countries to change their legislation to eliminate existing avoidance schemes. Bankman (1999) discusses possible anti-avoidance actions, including some that address the role of accounting firms. Hines (2004) and Slemrod (2004) allude to how the market for tax-avoidance tools may react to governmental policy that bans existing tax-avoidance models: the innovation of replacement products for the ones that are put out-of-business by the governments.

- A tax rule thicket exists that is caused by the large number of countries, their complex national tax regulation, the variety of financial and ownership structures of MNEs, and the bilateral taxation rules that apply between different countries. Possible tax-avoidance vehicles for MNEs often draw on the national legislation of more than one country and on the bilateral agreements between the countries. The much quoted "double Irish with a Dutch sandwich" may illustrate this.[3] It is hard to imagine that this model is the design outcome of purposeful and intelligent collaboration between the tax administrations of a whole set of countries.[4]

- The thicket of international taxation is likely to provide a very large number of tax-optimization models. Professional expert firms can provide existing expertise and can also generate new expertise by investing in exploration and innovation. New tax-avoidance products must be discovered before they can be used. Leading accounting firms may have the skills and technological know-how to explore and find such new products, whereas many other, often smaller firms, may copy the models innovated by leading firms and cannot generate new innova-

[3] See, e.g., Kleinbard (2011). This tax optimization model typically has a mother company in a high-tax country, two companies in Ireland, one company in the Netherlands, and one offshore company located in a country with low or zero taxation. A specific ownership structure and different legal forms of incorporation are used to reduce make the sandwich an effective tax-avoidance model.

[4] One cannot rule out that some existing schemes are the outcome of a purposeful design and are appropriately described as the outcomes of tax competition between countries. Such tax-haven activity is not the focus of the analysis here.

tions. Imitation needs time. This protects the innovating firm from immediate competition and allows for some innovation rents.

These elements point at a dynamic model of innovation and imitation of tax avoidance products by expert tax accountants, competition between these experts, and tax regulation by governments that fight these products. Several writers discuss this relationship. Hines (2004), for instance, mentions the interdependence between the innovation of tax-avoidance tools and regular governments' legislative anti-tax-avoidance efforts. He speculates about possible drawbacks of anti-tax-avoidance efforts. Curry et al. (2014) consider taxpayers' incentives to invest in tax-planning opportunities if the government can close uncovered loopholes. Their analysis confirms the potentially welfare deteriorating interdependence between investments in tax planning opportunities and the government's choice. They also look at the search for tax planning opportunities as an innovation contest, and the role of patents for tax avoidance products. Slemrod (2004, p. 889) suggests that "adapting theoretical models of patent races to the case of tax shelters suggest caution regarding the success of a strategy of reforming tax laws and regulations to reduce the effectiveness of elaborate tax-avoidance techniques as soon as they are identified."

This paper takes up Slemrod's sugestion and considers a formal, fully dynamic innovation model. Distinguishing features of the model are the central role of the tax accounting firms as suppliers of tax avoidance products, the distinction between innovating and imitating tax accounting firms, and the multi-period structure with an infinite horizon. This structure allows us to highlight the role of a time lag needed to imitate newly innovated tax-avoidance products, and the role of the size of a time lag by which the government is able to react and implement anti-tax-avoidance regulation. It shows that the anti-avoidance regulation may, but need not, provide incentives for innovative activities of the accounting industry. It highlights the key role of the government's speed of reaction. Immediate reaction may eliminate the scope for tax avoidance tools. A moderate regulatory time lag may lead to a stationary equilibrium with innovation and regulation in each period. Regulation may still successfully increase the tax revenues that emerge in such a framework. An even larger regulatory time lag may lead to an equilibrium with an infinitely repeated cycle. The equilibrium then alternates between periods with innovation and regulation, and periods without such actions. Little tax avoidance is observed in the periods with innovation and regulation. In the other periods the price for tax avoidance products is very low and tax avoidance is widely used. The equilibrium analysis also finds that synergies exist between highly innovative tax accounting firms and governmental tax regulation, with the tax accounting firms benefitting from governmental regulation.[5]

The literature on innovation and competition provides fundamental concepts that are building blocks for the analysis in the main sections. A key consideration is that there is potentially a race for innovations that has been described in other contexts (e.g., Harris and Vickers, 1987, Reinganum, 1989, Bhattacharyya and Nanda, 2000). A particular new feature of the innovation competition in the tax-avoidance context is the role of the government. The government can change the tax law, thereby eliminating tax loopholes that enabled well-known tax-optimizing technologies. Rather than in a race for ever further improvement of the latest design, the government enters the game as a player, and its anti-avoidance regulation endogenously determines the obsolescence of existing products.

Our analysis is also related to other strands of literature. Some authors addressed the tax competition between jurisdictions, and the role of low-tax jurisdictions for location choices, investment decisions, production output, and, ultimately, consumer welfare. This literature identifies and quantifies different effects that lead to opposing views of the benefits of low-tax jurisdictions. Seminal theory contributions are by Slemrod and Wilson (2009) who show that haven-countries may be harmful, and Hong and Smart (2010) and Johannesen (2010) who identify reasons for why low-tax jurisdictions may have positive effects. For a discussion and overview see Keen and Konrad (2013). A further important literature analyses firms' ability to apply profit-shifting and other tax-avoidance means under various conditions (see, for a brief review and original results, Hopland et al., 2015). Much of this

[5]The dynamic nature and the asymmetries between innovators and imitators distinguish our approach further from the only formal approach of a competitive market for tax planning models in Curry et al. (2014).

literature focuses on the impact of tax havens on real investment, assuming that the MNEs can look freely through what appears to be a thicket of international tax law. The approach taken in this paper can be seen as complementary to this broad literature. It adds the sector of tax experts to it, considers a market for tax-avoidance products, and the role of innovation and regulation in this market.

The analysis is also related to the large economic and tax law literature and the lively discussion about "base erosion and profit shifting" (BEPS) that takes place both in the legal and in the economic profession (see, e.g., Ault, 2013). This policy debate is reflected in the recent report on the profit shifting project by the OEDC (2016). While this report addresses numerous institutional issues related to BEPS, it does not push the role of professional tax experts into the limelight. The role of tax professionals has become a matter of consideration as a side aspect of what is called the "Lux-Leaks" hinting at a role of one of the big four accounting companies in the context of Advance Tax Rulings by the Luxembourg tax authorities (see also, e.g., European Parliament 2015). The BEPS process can be seen as the "repair shop" of the community of non-haven countries that aims at closing the loopholes that have been discovered and are used for tax avoidance.

The analysis proceeds as follows. Section 2 outlines the general framework. Section 3 sequentially considers three types of Markov perfect equilibria that emerge for different lengths of the time lag of governmental anti-avoidance regulation. Section 4 concludes.

2 The dynamic framework

Players, actions, sequencing There are three players: two tax accounting firms, one labeled by L ("leader") and one labelled by F ("follower"), and the government G. They interact in an infinite series of periods $t = 1, 2, \ldots.$ Each period t has three consecutive decision phases: an innovation phase, a legislation phase and a market phase.

The innovation phase. Firm L can innovate a new tax-avoidance product and decides whether or not to innovate in period t. This decision is denoted by $\theta_t \in \{0, 1\}$, where $\theta_t = 0$ describes that L does not innovate, and $\theta_t = 1$

means that L expends a fixed effort $e > 0$ and as a result innovates a new tax-avoidance product successfully and with certainty. We keep a record of the innovation activity of firm L in the different periods up to period t by a vector $\boldsymbol{\theta}_t = (\theta_1, ..., \theta_t)$, and $\boldsymbol{\theta}_t$ is costlessly observed by all players. An innovated "product" is a blueprint for a tax-avoidance tool that uses an existing loophole in the national/international tax rules. The blueprint is under exclusive control of the innovating firm L in the period t in which the innovation is made, and firm F cannot imitate this blueprint in this period. The details of the blueprint become known to F in period $t+1$. Firm F can costlessly imitate the product that was innovated in period t in period $t+1$ and all further periods.

The legislation phase. The government G chooses between two possible regulatory actions, denoted by $\gamma_t \in \{0, 1\}$. The choice $\gamma_t = 0$ means that the government takes no action in this period and has no cost. The choice $\gamma_t = 1$ means that the government takes action and has a given cost of size $\lambda > 0$. The government's action in t bans all tax avoidance products that have ever been innovated in periods up to period $t - k$. The parameter k measures a time lag. We denote three regimes that differ by the length of the time lag $k \in \{0, 1, 2\}$. If $k = 0$, this allows the government to react immediately to any newly innovated tax avoidance product. The product can be banned before it can be used or imitated. If $k = 1$, the government can observe and learn about tax avoidance products and can take anti-avoidance measures in the subsequent period - the same period in which the product could be imitated for the first time. If $k = 2$, it takes two periods until a new innovation can be banned by anti-avoidance actions. Anti-avoidance legislation is slower than imitation. The three regimes will be discussed in more detail when we turn to an analysis of each regime separately in section 3. We record the series of government actions that emerged until period t by a vector $\boldsymbol{\gamma}_t = (\gamma_1, ..., \gamma_t)$ that is costlessly observed by all players.

The market phase. The history of innovation and regulatory legislation determines which are the tax-avoidance products that the accounting firms may legally offer to corporate firms in period t. We distinguish three market conditions that may prevail in a given period t.

Competition: Both firms F and L have at least one not-banned tax-avoidance product which they can offer to the corporate customers in period t. The products can be offered to any corporate firm for a marginal cost of zero. All products are assumed to be perfect substitutes from the perspective of customer firms. Firms L and F independently choose prices $p_{L,t} \in [0,1]$ and $p_{F,t} \in [0,1]$. If a firm has several products available, as they are perfect substitutes, we can assume without loss of generality that the firm offers them at the same price.

The buyers of the product are the corporate firms. These are price takers and not players in a game-theory sense. They are sorted uniformly along the unit interval, where the firm (of type) x is characterized by its location $x \in [0,1]$. The location x will determine the advantage of this firm from using a tax-avoidance product as follows. If firm x does not purchase a tax-avoidance product, it pays the corporate tax. The size of the corporate tax is normalized to $T = 1$ for each firm. If firm x purchases the product from accounting firm $J \in \{L, F\}$ then it pays zero taxes, but pays the product price $p_{J,t}$, and has a transaction cost of size x. This transaction cost may reflect that a specific financial or legal structure is required for the use of the tax avoidance product, and that this structure is not optimal, for instance, from the perspective of corporate governance. As the tax burden is normalized to 1, firm x's maximum willingness to pay for the tax avoidance model is $(1 - x)$. Depending on x and the prices, any firm $x \in [0,1]$ purchases the product from the least-price supplier firm or does not purchase a tax avoidance product.[6] The resulting market demand and firms' shares in it are,

$$D_{L,t} = \begin{cases} \max\{1 - p_{L,t}, 0\} & \text{if} \quad p_{L,t} < p_{F,t} \\ \max\{\frac{1-p_{L,t}}{2}, 0\} & \text{if} \quad p_{L,t} = p_{F,t} \\ 0 & \text{if} \quad p_{L,t} > p_{F,t} \end{cases} \tag{1}$$

[6] Note that corporate firms apply a static rule. This may be the case because the periods are defined as sufficiently long and the avoidance decision has to be made in each period. The formal set-up and the description for the market had to be adjusted if tax-optimization models were long-lasting durable products.

and

$$D_{F,t} = \begin{cases} \max\{1 - p_{F,t}, 0\} & \text{if} \quad p_{F,t} < p_{L,t} \\ \max\{\frac{1-p_{F,t}}{2}, 0\} & \text{if} \quad p_{F,t} = p_{L,t} \\ 0 & \text{if} \quad p_{F,t} > p_{L,t} \end{cases} \tag{2}$$

for the two tax-accounting firms.

Monopoly: let L be the monopolist supplier of tax-avoidance tools in period t. This applies if firm L innovated a new product in this period t that is not banned in this same period, and F has no product it can offer. The latter is the case if there were no innovations made prior to t, or if older innovations of tax-avoidance blueprints that could be imitated by F were banned by the government and cannot be used in period t. The firm L can choose the price $p_{L,t} \geq 0$ at which it offers the tax-avoidance tool to the whole set $[0,1]$ of corporate firms. Each of the corporate firms $x \in [0,1]$ pays one unit of corporate tax or it purchases the tax-avoidance product of firm L. Corporate firms act according to the same rules as previously stated for the competition phase. Accordingly, the turnover of the tax-accounting firms is

$$D_{L,t} = \max\{1 - p_{L,t}, 0\} \tag{3}$$

and

$$D_{F,t} = 0. \tag{4}$$

Note that F cannot be a monopolist supplier because any available blueprint in t was invented by L and is therefore also available to L.

No market: none of the firms can offer a tax avoidance product in period t. This applies if L never innovated a product, or if all innovations that have been made have been banned. Evidently, $D_{L,t} = D_{F,t} = 0$ in this case.

For notation, let us collect the price choices by the accounting firms in periods up to period t as $\mathbf{p}_t = ((p_{L,1}, p_{F,1}), ..., (p_{L,t}, p_{F,t}))$. For periods in which both firms have a tax-avoidance product which they can offer, the prices will be $(p_{L,t}, p_{F,t}) \in [0,1] \times [0,1]$. For periods in which only L can offer a product, the prices are $(p_{L,t}, p_{F,t}) \in [0,1] \times \{\varnothing\}$ with $p_{L,t} \in [0,1]$, and $p_{F,t} = \varnothing$ symbolizing that firm F has no product to offer. For periods in which none of the accounting firms can offer a tax-avoidance product, we have $(p_{L,t}, p_{F,t}) = (\varnothing, \varnothing)$.

Histories and strategies The complete history in a given period t is described by the combination $(\boldsymbol{\theta}_{t-1}, \boldsymbol{\gamma}_{t-1}, \mathbf{p}_{t-1})$ at the innovation phase, by $(\boldsymbol{\theta}_t, \boldsymbol{\gamma}_{t-1}, \mathbf{p}_{t-1})$ at the legislatory phase, and by $(\boldsymbol{\theta}_t, \boldsymbol{\gamma}_t, \mathbf{p}_{t-1})$ at the market phase of period t. Local strategies at the respective period t and phase are defined by functions that map a given history into an action, for each of the players, and the collection of these local strategies defines the strategy profile.

Period payoffs and overall payoffs The period payoffs of the tax-accounting firms are

$$\pi_{L,t} = \begin{cases} -e\theta_t + p_{L,t} D_{L,t} & \text{if} \quad p_{L,t} \in [0,1] \\ -e\theta_t & \text{if} \quad p_{L,t} = \varnothing \end{cases} \tag{5}$$

and

$$\pi_{F,t} = \begin{cases} p_{F,t} D_{F,t} & \text{if} \quad p_{F,t} \in [0,1] \\ 0 & \text{if} \quad p_{F,t} = \varnothing \end{cases} . \tag{6}$$

The government maximizes tax revenue net of the cost of legislation that eliminates tax loopholes to ban tax-avoidance products. It disregards the profits of the accounting firms and of the corporate firms.[7] The tax revenue is equal to the number of corporate firms that do not purchase a tax-avoidance product but rather pay the tax $T = 1$. Analytically, the period payoff of the government G is

$$g_t = -\gamma_t \lambda + (1 - D_{L,t} - D_{F,t}).$$

The overall payoff of a player is equal to the discounted sum of the player's period payoffs. Hence, the continuation payoff at a given period t is

$$\Pi_t^L = \sum_{z=t}^{z=\infty} \delta^{z-t} \pi_{L,z},$$

$$\Pi_t^F = \sum_{z=t}^{z=\infty} \delta^{z-t} \pi_{F,z},$$

[7] Tax revenue maximization is not uncommon as an objective in tax competition models. The reasoning behind this assumption can be that the government is a pure Leviathan. Alternatively, the reason here could be that the rents of the accounting firms and of the corporate customers of tax avoidance tools are fully discounted by the government for other reasons.

and

$$\Gamma_t = \sum_{z=t}^{z=\infty} \delta^{z-t} g_z.$$

Here, $\delta \in (0,1)$ is a time- and player-invariant discount factor. This completes the characterization of the components of the dynamic framework.

3 Equilibrium

In what follows let us search for Markov perfect equilibria (MPE) in stationary strategies or, if a stationary equilibrium does not exist, an equilibrium in stationary cycles for time-lag parameters $k = 0$, $k = 1$ and $k = 2$. We do this sequentially for each time-lag size in a separate subsection.

3.1 The MPE for the regime with lag $k = 0$

For a lag-parameter $k = 0$ the choice of $\gamma_t = 1$ bans all tax avoidance products that have been innovated up to this point of time, including a possible innovation ($\theta_t = 1$) in period t.

For the formal analysis and the characterization of the stationary MPE, we need to partition the sets of all histories of a given length $(\boldsymbol{\theta}_{t-1}, \boldsymbol{\gamma}_{t-1}, \mathbf{p}_{t-1})$ into two sets which we denote by H_t^0 and H_t^1, and which define two states of a Markov process. The set H_t^0 includes all histories $(\boldsymbol{\theta}_{t-1}, \boldsymbol{\gamma}_{t-1}, \mathbf{p}_{t-1})$ for which there is no previously innovated tax avoidance model available at the beginning of period t. This may result because no innovations ever occurred prior to period t, or because all previously innovated products have been banned by the government. The set H_t^1 includes all histories $(\boldsymbol{\theta}_{t-1}, \boldsymbol{\gamma}_{t-1}, \mathbf{p}_{t-1})$ for which there is at least one innovated tax avoidance model available at the beginning of period t.

Suppose the process starts with a history denoted by H_1^0, in which no previous innovation exists at the beginning of period $t = 1$.[8] The following result holds:

[8] As can be shown, if a previous innovation exists in $t = 1$, then $(\theta_1, \gamma_1) = (0, 1)$ in the equilibrium and the process reaches the stationary equilibrium in $t = 2$ with $(\theta_t, \gamma_t) = (0, 0)$ for all $t \geq 2$.

Proposition 1 *If $\lambda \in (0, \frac{1}{2})$ then a unique stationary Markov perfect equilibrium exists which has $(\theta_t, \gamma_t) = (\theta^*, \gamma^*) = (0, 0)$ for all t and payoffs $\Pi_t^L = \Pi_t^F = 0$ and $\Gamma_t = \frac{1}{1-\delta}$.*

The proof is in the Appendix. The proposition describes a very intuitive outcome: An MPE in which accounting firms never innovate emerges if the government can ban new innovations of tax-avoidance products before they can be offered to customers, if the cost of banning these new innovations is sufficiently low. A new product could be innovated in this equilibrium, but does not offer any reward for the innovating firm, because it is banned before it could be used. Accordingly, both accounting firms have zero profits, all corporate firms pay the tax of $T = 1$, and the present value of these taxes yields the government's payoff in the MPE. Uniqueness (among the set of stationary MPEs) can easily be established by showing that any other stationary candidate MPE invites profitable deviations. The no-tax-avoidance equilibrium in Proposition 1 is the benchmark for possible departures that emerge if the government has a larger regulatory time lag.

3.2 The MPE for the regime with lag $k = 1$

Let us now turn to the case with time-lag parameter $k = 1$. For $k = 1$ the choice of γ_t bans all tax avoidance products that have been innovated in period $t - 1$ or earlier. Again we define two sets H_t^0 and H_t^1 of histories at the beginning of period t which define two Markov states. Moreover, we start with a history in the set H_1^1; i.e., we assume that there is one blueprint already available at the beginning of period $t = 1$.[9] The following result holds:

Proposition 2 *If $e \in (0, \frac{1}{4})$ and $\lambda \in (0, \frac{1}{2})$, then a unique stationary Markov perfect equilibrium exists that is characterized by the path $(\theta_t, \gamma_t) = (1, 1)$.*

The proof is in the Appendix. Proposition 2 describes an MPE in which the innovative firm comes up with a newly innovated product in each period,

[9] As can be shown, if the process starts in a history from H_1^0, i.e., no previously innovated history exists at the beginning of period $t = 1$, then $(\theta_1, \gamma_1) = (1, 0)$ in $t = 1$ and the process reaches the stationary path with $(\theta_t, \gamma_t) = (1, 1)$ for all $t > 1$.

and in which the government chooses anti-tax-avoidance regulation and bans all older, previously innovated products (if there are any). As a result, this removes all competitor products that could be imitated and offered by F in the market in any given period. But due to the time lag, the tax avoidance product most recently innovated cannot be banned nor imitated. The innovative firm L earns monopoly profit from selling its most recent product, in each period. The government nevertheless benefits from banning old tax-avoidance products: this choice drives up the market price for tax avoidance goods from zero to the monopoly price. This high price prevents a share of the corporate firms from purchasing the tax avoidance product. They prefer to comply and pay their taxes.

The equilibrium outcome reveals a positive interaction between innovating and banning tax-avoidance products. Banning old products induces innovation, and at the same time it reduces tax avoidance. The reduction takes place via a market competition effect: due to the regulatory time lag, the government cannot make tax-avoidance products completely unavailable in this equilibrium. However, banning the old products provides the innovating accounting firms with market power. This results in a higher market price. This higher price reduces the set of corporate customers that purchase this product and makes more of them pay their taxes. Anti-avoidance regulation changes the market competition between accounting firms. It creates opportunities and market power for more advanced accounting firms that innovate and reap monopoly rents from their innovations. The government benefits from this as well, compared to Bertrand competition among the tax-accountant firms: firm L's rents are accompanied by higher prices for avoidance products, and these higher equilibrium prices may deter a number of corporate firms from making use of these products. These firms rather pay the regular tax rates. In conclusion, closing tax loopholes is beneficial for both the government and the innovative tax-accounting firm.

A comparison between the MPE in Proposition 1 and the MPE in Proposition 2 shows the important role of a regulatory time lag. If the government can instantaneously react to innovations and ban the most recent products at a low cost, it can and will induce full tax compliance. The credible threat of banning is sufficient to prevent accounting firms from innovating new products. The outcome is an MPE with full compliance. However, if it takes some time before a newly discovered tax avoidance product can be banned, then this induces a permanent race between the innovating firm and the government, unless innovating or banning are too expensive.

Let us also compare these results with standard innovation models. In regular product markets it may be the scope for product quality improvements and cost reductions that drives a continued innovation race. In the market for tax-avoidance products for MNEs the mechanism that drives innovation is very different. Quality improvements or process improvements probably play a minor role in this market. It may even hold that the most effective products that combine the simplicity of the required financial and the legal structure of the MNEs with a large reduction in the tax burden have been innovated early on. What causes repeated innovation is the regulatory action taken by governments of high-tax countries or the initiatives by the OECD. This action causes the obsolescence of existing products.

3.3 The MPE for the regime with lag $k = 2$

Let us now turn to the case with time-lag $k = 2$. This larger time-lag reduces the government's ability to timely ban existing tax avoidance products. A tax avoidance product which is innovated in period t can be offered to corporate customers by firm L (and only by this firm) in the innovation period t. And this tax avoidance product will also be available in period $t + 1$. In period $t + 1$ both firms L and F will be able to offer this product to corporate customers. The earliest moment when the government can regulate and ban this product is period $t+2$. The implication of lag $k = 2$ is that any product that is innovated may potentially yield innovation rents for the innovator firm at the time when it has just been innovated, but the product will also be available in the next period and will cause Bertrand competition between L and F in this next period.

Before we turn to the equilibrium results, we define a more fine-grained partition of the set of histories $(\boldsymbol{\theta}_{t-1}, \boldsymbol{\gamma}_{t-1}, \mathbf{p}_{t-1})$ at the beginning of each period t. This partition constitutes the Markov states. We distinguish between four subsets of histories at the beginning of any given period t which we denote by H_t^{00}, H_t^{10}, H_t^{01}, and H_t^{11}. The first superscript x in H_t^{xy} refers to whether there is a tax avoidance product that is at least two periods old and has not been banned previously. The second component y refers to $\theta_{t-1} \in \{0, 1\}$, i.e., whether L innovated in period $t-1$ or not. More in detail, the set H_t^{00} includes all histories with no blueprints that were innovated in period $t-2$ or earlier and are available at the beginning of period t, and with $\theta_{t-1} = 0$. The set H_t^{10} includes all histories for which at least one blueprint was innovated in $t-2$ or earlier, and is still available at the beginning of period t, and with $\theta_{t-1} = 0$. The set H_t^{01} refers to the set of histories for which no blueprints are available at the beginning of t that are older than one period, and $\theta_{t-1} = 1$. The set H_t^{11} refers to the set of histories at the beginning of period t such that at least one blueprint exists and is available at the beginning of t that is older than one period, and $\theta_{t-1} = 1$.

A first observation is a non-existence result of stationary equilibria.

Proposition 3 *If $e \in (0, \frac{1}{4})$ and $\lambda \in (0, \frac{1}{2})$, then a unique Markov perfect equilibrium in stationary strategies does not exist for a lag $k = 2$.*

The proof is in the appendix. A stationary equilibrium is characterized by $(\theta_t, \gamma_t) = (\theta^*, \gamma^*)$, at least from a certain t on. The proof shows that any such stationary equilibrium candidate offers profitable one-step deviations either for firm L or the government G or for both of them. This non-existence result hints at the possibility of more interesting equilibria with a more dynamic time structure. Indeed, we find that the equilibrium that emerges is characterized by cycles, in which periods with massive tax avoidance alternate with periods with innovation, regulatory intervention and more limited amounts of tax avoidance.

For the characterization of an MPE with such cycles we assume that the process starts from a history $h_1 \in H_1^{10}$, i.e., in period $t = 1$ there is at least one innovation available that is old enough such that it can be banned by regulatory action in period $t = 1$ and would be available to both firms if not banned, and there is no recent innovation that can be imitated by F in period t but cannot yet be banned in t through regulatory action, due to the time lag of regulatory measures. Given this starting condition[10] the following

[10] As can be shown and is implied by the reasoning in the proof, if the process starts in

holds:

Proposition 4 *Let $e \in (0, \frac{1}{4})$ and $\lambda \in (0, \frac{1}{2})$. An MPE with stationary cycles exists in which $(\theta_t, \gamma_t) = (1, 1)$ for all uneven periods $t = 1, 3, 5, ...$ and $(\theta_t, \gamma_t) = (0, 0)$ for all even periods $t = 2, 4, ...$*

The proof is in the appendix. The proposition characterizes an equilibrium in which the prevalence of tax avoidance, innovation, regulation, firm profits and tax revenues fluctuates over time. In one period tax avoidance products are widely available, their equilibrium market price is zero and all corporate customer firms use these products to avoid taxation. The tax revenue is very low, as none of the corporate firms pays taxes. However, the blueprints used in this period are old blueprints, and sufficiently well-known such that the government can use anti-avoidance regulation to ban them in the next period. This is what the government will do. In addition, the innovative accounting firms anticipate that these old and widely available tax-avoidance products will be banned and removed from the market. Hence, in the same period when this ban is enacted, the innovative firm among the tax accounting industry invents a new blueprint. As all the widely known and generally available products will be taken off the market, the innovator firm earns entrepreneurial rents on this newly innovated product. It takes yet another period until this new product is available for the tax accounting sector more generally, and one further period until this new product can again be banned, and so on.

4 Discussion and Conclusion

Many of the assumptions in the innovation model above are clearly for the sake of simplicity and to describe results in closed-form solutions. For instance, the assumption about the type of heterogeneity among customer firms and the linear demand function emerging from it is of this kind. The

H_1^{01}, then the behavior in even and odd periods is interchanged. If the process starts in a state other than H_1^{10} or H_1^{01}, then the actions in period $t = 1$ move the process to state H_2^{10} or H_2^{01} in period $t = 2$ and the stationary cycle begins there.

assumption that all tax-avoidance products are perfect substitutes is motivated similarly. If, instead, several tax-avoidance tools are on the market and some are better for some firms but others are better suited for other firms, then the market phase has competition with multiple products and a description of the precise market structure is cumbersome. Such a departure implies that the number of possible Markov states also becomes large, particularly if the government is selective in its choice of which loopholes to close now and which ones to keep for later decision-making. However, an extension in this direction should be complemented with a further element of realism that is likely to interact with this enhanced action/product space. In particular, product diversity should be considered only in a context with more than one innovator firm and with more than one imitator firm and with heterogeneity among them, etc. This would be a technically demanding exercise. It may be potentially valuable, but it clouds the picture and deflects the focus away from the key question and towards aspects that are not central to our main research question.

Departing from the assumption about deterministic effects of innovation and loophole-closing leads to a stochastic dynamic model. The forces that induce firms to innovate, the implications of innovation for the market outcome, and the role of governments' attempts to close down tax loopholes can be expected to remain qualitatively similar, but modeling the problem in a stochastic dynamic model becomes more cumbersome. These aspects are seemingly orthogonal to the issues studied in the main part of the paper. Departures from the assumptions about complete and perfect information go in a similar direction, raising complexity, potentially uncovering new aspects, but not systematically undoing the key insights of the analysis in the main sections.

A different direction is to depart from the Markov assumption and to ask if there are collusive arrangements that are sustainable by standard reasoning along the theory of infinitely repeated games. Even disregarding the corporate firms / customers who are not players in a game-theory sense in this set-up, there is a considerable rent that is "left on the table" in some of the stationary MPEs. Hence, one may consider a collusive outcome in which the two firms coordinate on sufficiently high prices such that the government earns sufficient tax revenue to prevent the government from closing tax loopholes and from moving the game toward one of the MPEs considered. Within the framework with two firms and one government the existence of such collusive equilibrium outcomes is seemingly possible if the discount factor is large. But one should note that a larger number of countries and accounting firms strongly reduces the scope for such collusive equilibria.

The key issue in the main analysis here is the size of the regulatory lag. The main section considered three lag sizes that yield qualitatively very different results. As shown in Proposition 1, if the government can react instantaneously, it will -unless the regulatory cost is excessive. This will lead to an elimination of tax avoidance. The government may like the equilibrium in Proposition 1, as this equilibrium has zero tax avoidance and zero regulation cost for the government.

An intermediate size of the time lag may cause an outcome in which the government always bans the old products, but they are replaced by a new product in each new period. A one-period lag of possible government action triggers a permanent innovation/regulation race. This race yields positive rents to the innovator firms. The government could abstain from regulation, as there will be a tax-avoidance product available in every period. But regulation nevertheless yields higher tax revenue, due to reactions on the market for avoidance products. From the leading accounting firms' perspective, the outcome for an intermediate time lag is the most attractive one. They have to innovate in each period, but they earn monopoly rents in each period as well.

The structure that is perhaps most interesting emerges for a larger time lag. A larger time lag causes a more complex equilibrium behavior that is characterized by cycles. Phases with cheap and abundant tax avoidance products and almost complete tax avoidance behavior alternate with phases with high regulatory effort, expensive and new tax avoidance products and reduced prevalence of tax avoidance activity.

Overall, the considerations hint at a partial congruence of the interests of a tax revenue-oriented government and of accounting firms with the highest innovative capabilities. Governmental tax regulation that removes the better-known and more widely available tax-avoidance products from the market for tax accounting services provides the tax accountants with innovation rents and makes these accounting firms the main beneficiaries of tax regulation.

We did not consider accounting firms that may be able to inform the government about existing loopholes, or can reduce the government's time lag. It is evident, however, that accounting firms which outperform their competitors in their innovation abilities have an incentive to inform the government and to shorten the time lag to a lag of intermediate size. Similarly, if the cost of regulation is prohibitive, these firms had an incentive to reduce the government's cost of regulatory action to implement an intermediate time lag. They dislike market outcomes in which imitating tax accountant firms drive down the price of tax-avoidance advice to zero. They need the government to remove tax avoidance products offered by their imitating competitors. If they succeed to move from the equilibrium in Proposition 4 to the one in Proposition 2, this eliminates the zero-profit periods with Bertrand competition.

5 Appendix

Proof of Proposition 1 Suppose that players follow the behavior that is described in the candidate equilibrium in periods after t. This behavior defines the continuation value (i.e., the present value in period $t+1$ of the sum of this period and future periods' payoffs) for L as

$$\Pi_{t+1}^{L} = 0, \tag{7}$$

the continuation value for F as

$$\Pi_{t+1}^{F} = 0, \tag{8}$$

and the continuation value for the government as

$$\Gamma_{t+1} = \frac{1}{1-\delta}. \tag{9}$$

Note next that the payoffs of all three players are continuous at infinity: any possible period payoffs for the accounting firms are within finite intervals which are bounded from below by $-e$ for L and 0 for F, and bounded from above by 1. The period payoff for the government is bounded from below by $-\lambda$ and from above by 1. Moreover, the discount factor is $\delta \in (0,1)$. We can therefore rely on the one-step deviation principle (Theorem 4.2 in Fudenberg and Tirole, 1991) such that we only need to check whether players can improve upon their payoffs in the candidate equilibrium by one-step deviations.

For the formal analysis we partition the sets of all histories $(\boldsymbol{\theta}_{t-1}, \boldsymbol{\gamma}_{t-1}, \mathbf{p}_{t-1})$ for given t into two sets which we denote by H_t^0 and H_t^1. Let h_t denote a representative element of the set $H_t^0 \cup H_t^1$ of all possible histories at the beginning of t. The set H_t^0 includes all histories h_t for which there is no previously innovated tax avoidance model available at the beginning of period t, either because no innovations occurred prior to period t, or because all previously innovated products are banned. The set H_t^1 includes all histories h_t for which there is at least one innovated tax avoidance model available at the beginning of period t. Further, we have chosen the starting condition at $t = 1$ such that $h_1 \in H_1^0$: there is no blueprint available at the beginning of period $t = 1$. Consider now one-step deviations in period t. The process is in $h_1 \in H_1^0$ by the starting condition, and at a given period t either $h_t \in H_t^0$ or $h_t \in H_t^1$ applies.

The market phase. The market game in period t has no impact on whether $h_{t+1} \in H_{t+1}^0$ or $h_{t+1} \in H_{t+1}^1$. Both firms maximize their period payoff by their price choices for given $(\boldsymbol{\theta}_t, \boldsymbol{\gamma}_t, \mathbf{p}_{t-1})$. If none of the firms have a tax-avoidance product to sell then $p_{L,t} = p_{F,t} = \varnothing$ and $\pi_{L,t} = \pi_{F,t} = 0$. All corporate firms pay taxes. If only L has a tax-avoidance product to sell then $\pi_{L,t} = \frac{1}{4}$ and $\pi_{F,t} = 0$ and half of the firms pay taxes. If both firms have a tax-avoidance product to sell then $p_{L,t} = p_{F,t} = 0$ and $\pi_{L,t} = \pi_{F,t} = 0$ and none of the firms pays taxes.

The regulation phase. Consider the payoff maximization of the government in period t anticipating the market phase in t as just described and that all players follow the candidate equilibrium in all future periods $t+1, \ldots$.

Let $h_t \in H_t^0$ and $\theta_t = 0$. In this case $h_{t+1} \in H_{t+1}^0$, independent of γ_t, and the government's payoff is higher by λ for $\gamma_t = 0$ than for $\gamma_t = 1$.

Let $h_t \in H_t^0$ and $\theta_t = 1$. In this case $\gamma_t = 1$ leads to $h_{t+1} \in H_{t+1}^0$. Hence, the government payoff is $1 - \lambda + \frac{\delta}{1-\delta}$. In comparison, assuming candidate equilibrium play in the continuation game, $\gamma_t = 0$ yields the government a payoff of $\frac{1}{2} + \delta(-\lambda + \frac{1}{1-\delta})$. The choice $\gamma_t = 1$ yields higher payoff than $\gamma_t = 0$ if $1 - \lambda > \frac{1}{2} - \delta\lambda$, which simplifies to $\lambda < \frac{1}{2(1-\delta)}$. This condition is weaker than $\lambda < \frac{1}{2}$ mentioned in the proposition.

Let $h_t \in H_t^1$ and $\theta_t = 0$. Then the government's payoff for $\gamma_t = 1$ as in the candidate equilibrium is $1 - \lambda + \frac{\delta}{1-\delta}$. The government's one-step payoff for $\gamma_t = 0$ is $0 - 0 - \delta\lambda + \frac{\delta}{1-\delta}$. So $\gamma_t = 1$ yields a (weakly) higher profit if $\lambda < \frac{1}{1-\delta}$. Again, this condition is weaker than $\lambda < \frac{1}{2}$ mentioned in the proposition.

Finally, let $h_t \in H_t^1$ and $\theta_t = 1$. Then the government's payoff in the candidate equilibrium for $\gamma_t = 1$ is $1 - \lambda + \frac{\delta}{1-\delta}$ and the government's one-step deviation payoff for $\gamma_t = 0$ is $0 - \delta\lambda + \frac{\delta}{1-\delta}$. Again, $\gamma_t = 1$ yields a higher payoff if $\lambda < \frac{1}{1-\delta}$.

The innovation phase. Consider the innovation behavior of firm L. Given the equilibrium continuation play and that $\gamma_t = 1$ if $\theta_t = 1$, not innovating yields a higher payoff for L than innovating.

We now turn to *uniqueness*. Any stationary equilibrium has $\theta_t = \theta^*$ and $\gamma_t = \gamma^*$ along the equilibrium path. If $\theta_t = \theta^* = 0$, then the stationary equilibrium either has $\gamma_t = \gamma^* = 0$, which is the equilibrium in Proposition 1, or it has $\gamma_t = \gamma^* = 1$. In this latter case a payoff increasing one-step deviation $\gamma_t = 0$ in period t exists, as there is no product that can be banned, and the cost of anti-avoidance regulation is pure waste. Hence, we can rule out the combination ($\theta_t = \theta^* = 0$ and $\gamma_t = \gamma^* = 1$) as a possible stationary MPE. Consider next $\theta_t = \theta^* = 1$. This may go along with $\gamma^* = 1$ or with $\gamma^* = 0$. The case $\theta^* = 1$ and $\gamma^* = 1$ can be ruled out, as $\theta_t = 0$ would be a payoff-increasing deviation for L: any innovation is banned in the same period, so the cost of innovating is purely wasteful. The case $\theta_t = \theta^* = 1$ combined with $\gamma^* = 0$ has $g_t = 0$ for all $t > 1$. A deviation to $\gamma_t = 1$ at $t > 1$ increases the government's payoff if $\lambda < \frac{1}{2}$. ■

Proof of Proposition 2 Any choices of actions (θ_t, γ_t) in period t either lead to a state $h_{t+1} \in H_{t+1}^1$ or to a state $h_{t+1} \in H_{t+1}^0$. Suppose that players follow the strategies as described in the candidate equilibrium. This defines continuation values (i.e., the period-$(t+1)$ present values of the sum of period

$t+1$'s and future periods' payoffs) as follows: firm L has a period profit of $\frac{1}{4} - e$ in every period. This sums up to

$$\Pi_t^L = (\frac{1}{4} - e)\frac{1}{1-\delta}. \tag{10}$$

Firm F's period payoff is zero in all periods, so F's continuation value is

$$\Pi_t^F = 0. \tag{11}$$

The government's period payoff depends on $h_t \in H_t^1$ or $h_t \in H_t^0$. Assuming candidate equilibrium play in all periods after t, the continuation value is

$$\Gamma_t = \begin{cases} (\frac{1}{2} - \lambda)\frac{1}{1-\delta} & \text{if} \quad h_t \in H_t^1 \\ \frac{1}{2} + (\frac{1}{2} - \lambda)\frac{\delta}{1-\delta} & \text{if} \quad h_t \in H_t^0 \end{cases}. \tag{12}$$

Note next that we can apply the one-stage deviation principle for the same reasons as in the context of the proof of Proposition 1. We now check whether players can improve upon their payoffs in the candidate equilibrium by one-step deviations.

The market phase. The market interaction follows the logic already described and has no impact on whether $h_{t+1} \in H_{t+1}^0$ or $h_{t+1} \in H_{t+1}^1$. Both firms maximize their period payoff by their price choices for given $(\boldsymbol{\theta}_t, \boldsymbol{\gamma}_t, \mathbf{p}_{t-1})$. If none of the firms have a tax-avoidance product to sell then $p_{L,t} = p_{F,t} = \varnothing$ and $\pi_{L,t} = \pi_{F,t} = 0$. All corporate firms pay taxes. If only L has a tax-avoidance product to sell then $\pi_{L,t} = \frac{1}{4}$ and $\pi_{F,t} = 0$ and half of the firms pay taxes. If both L and F have a tax-avoidance product to sell then $p_{L,t} = p_{F,t} = 0$ and $\pi_{L,t} = \pi_{F,t} = 0$ and none of the corporate firms pays taxes.

The legislation phase. Consider the optimizing behavior of the government in t anticipating that firms make pricing choices in period t as just described and all players follow the candidate equilibrium behavior from $t+1$ onwards. Note that equilibrium continuation play from $t+1$ onwards yields the continuation values as in (10), (11), and (12).

If $h_t \in H_t^0$, irrespective of whether $\theta_t = 1$ or $\theta_t = 0$, given the time lag $k = 1$, there are no products that could be banned in period t. The government prefers $\gamma_t = 0$ to $\gamma_t = 1$, as $\gamma_t = 1$ induces a cost of λ and does not ban any product and does not affect the market phase in period t and does not change the initial condition in the subsequent period $t+1$, which is in H^0_{t+1} if $\theta_t = 0$ and in H^1_{t+1} if $\theta_t = 1$.

If $h_t \in H^1_t$ and $\theta_t = 0$ then the equilibrium choice $\gamma_t = 1$ leads to a period payoff $(1-\lambda)$ in t and to $h_{t+1} \in H^0_{t+1}$. The government's payoff is $(1-\lambda) + \delta\frac{1}{2} + \delta^2(\frac{1}{2} - \lambda)\frac{1}{1-\delta}$. A deviation to $\gamma_t = 0$ leads to symmetric Bertrand competition between accounting firms with tax revenue equal to zero in period t and a history $h_{t+1} \in H^1_{t+1}$ in period $t+1$. Government's payoff is $0 + \delta(\frac{1}{2} - \lambda)\frac{1}{1-\delta}$. A comparison shows

$$(1-\lambda) + \delta\frac{1}{2} + \delta^2(\frac{1}{2} - \lambda)\frac{1}{1-\delta} > 0 + \delta(\frac{1}{2} - \lambda)\frac{1}{1-\delta}$$

if $\lambda < \frac{1}{1-\delta}$. This confirms the optimality of $\gamma_t = 1$ for $h_t \in H^1_t$ if $\theta_t = 0$.

If $h_t \in H^1_t$ and $\theta_t = 1$ then the equilibrium choice $\gamma_t = 1$ leads to a monopoly of firm L in this period t and to $h_{t+1} \in H^1_{t+1}$. The government's payoff is $\frac{1}{2} - \lambda + \delta(\frac{1}{2} - \lambda)\frac{1}{1-\delta}$. A deviation $\gamma_t = 0$ leads to symmetric Bertrand competition between accounting firms with tax revenue equal to zero in period t and to $h_{t+1} \in H^1_{t+1}$. The government's payoff is $0 + \delta(\frac{1}{2} - \lambda)\frac{1}{1-\delta}$. The payoff for $\gamma_t = 1$ exceeds the payoff for $\gamma_t = 0$ if $\frac{1}{2} > \lambda$.

The innovation phase. Turn now to the innovation choice of firm L in period t. Consider deviations from $\theta_t = 1$. If $h_t \in H^1_t$, then both $\theta_t = 1$ and $\theta_t = 0$ are followed by $\gamma_t = 1$. For $\theta_t = 1$ the payoff for L is $(\frac{1}{4} - e + \delta\Pi^L_t) = \Pi^L_t$. For $\theta_t = 0$ the payoff is $0 + \delta\Pi^L_t$. Hence, $\theta_t = 1$ leads to higher payoff for L than $\theta_t = 0$ if $\frac{1}{4} > e$. If $h_t \in H^0_t$, the choice of $\theta_t = 1$ is followed by $\gamma_t = 0$ and by $h_{t+1} \in H^1_{t+1}$, and yields firm L a payoff of $\frac{1}{4} - e + \delta\Pi^L_t = \Pi^L_t$. The choice of $\theta_t = 0$ is followed by $\gamma_t = 0$ and $h_{t+1} \in H^0_{t+1}$. It yields firm L a payoff of $0 + \delta\Pi^L_t$. Accordingly, the candidate equilibrium choice $\theta_t = 1$ yields higher payoff than $\theta_t = 0$ if $\frac{1}{4} > e$.

This completes the analysis of possible one-step deviations in period t and shows that deviations are unprofitable for the deviating player if $e < \frac{1}{4}$ and $\lambda < \frac{1}{2}$.

We now turn to *uniqueness*. Any stationary equilibrium must have $\theta_t = \theta^*$ and $\gamma_t = \gamma^*$ along the equilibrium path. If $\theta_t = \theta^* = 1$, then the stationary equilibrium either has $\gamma_t = \gamma^* = 1$, which is the equilibrium in Proposition 2. Or it has $\gamma_t = \gamma^* = 0$. In this case a one-step deviation $\gamma_t = 1$ in period

t exists that increases the period payoff for the government from 0 to $\frac{1}{2} - e$. Hence, we can rule out ($\theta_t = \theta^* = 1$ and $\gamma_t = \gamma^* = 0$) as a possible stationary MPE. Consider next $\theta_t = \theta^* = 0$. The case $\theta^* = 0$ and $\gamma^* = 1$ can be ruled out, as it allows for a one-step deviation towards $\gamma_t = 0$ which increases the government's period payoff from $-\lambda$ to 0 in the respective period. Finally, consider $\theta^* = 0$ and $\gamma^* = 0$. Given $h_1 \in H_1^1$, this yields period payoffs for F, for L, and for the government that are equal to zero in all periods. In comparison, a one-step deviation to $\gamma_t = 1$ increases the government's period payoff to $1 - \lambda$. Note that the suboptimality of $\theta^* = \gamma^* = 0$ does not hinge on the starting condition. Should $h_1 \in H_1^0$, then a deviation towards $\theta_t = 1$ increases the payoff of firm L in period t to $\frac{1}{4} - e > 0$, and the firm's continuation payoff cannot fall below zero. In sum this rules out any stationary MPE other than $(\theta^*, \gamma^*) = (1, 1)$ for $e < \frac{1}{4}$ and $\lambda < \frac{1}{2}$. ■

Proof of Proposition 3 The four candidates for a stationary MPE are $(\theta^*, \gamma^*) \in \{(0,0), (1,0), (0,1), (1,1)\}$. We rule them out one after the other.

(i) If $(\theta^*, \gamma^*) = (0, 0)$, then, after a finite number of periods and depending on the initial conditions that apply, in period t the process is either in the set of histories H_t^{00} (this happens if no tax-avoidance product was ever innovated prior to $t = 1$), or the process is in the set of histories H_t^{10} (this happens if a tax-avoidance product existed already in $t = 1$). For $h_t \in H_1^{00}$, the equilibrium continuation payoff of player L in the candidate equilibrium is $\Pi_t^L = 0$. A profitable one-stage-deviation for player L that yields at least a payoff of $\frac{1}{4} - e$ is, hence, $\theta_t = 1$. For $h_t \in H_1^{10}$, the equilibrium continuation payoff of player G in the candidate equilibrium is $\Gamma_t = 0$. A profitable one-step deviation for player G that increases the payoff by at least $\frac{1}{2} - \lambda$ is, hence, $\gamma_t = 1$.

(ii) If $(\theta^*, \gamma^*) = (1, 0)$, then, after a small number of periods, in period t the process is in H_t^{11} for t and all future periods. The equilibrium continuation values are $\Pi_t^L = \frac{-e}{1-\delta}$ and $\Gamma_t = 0$. A one-step deviation $\gamma_t = 1$ increases G's payoff by at least $\frac{1}{2} - \lambda$.

(iii) If $(\theta^*, \gamma^*) = (0, 1)$ then, after a small number of periods, the process is in H_t^{00} in t and in all further periods. For $h_t \in H_t^{00}$, the equilibrium continuation payoff of player L in the candidate equilibrium is $\Pi_t^L = 0$. A profitable one-stage deviation for player L is $\theta_t = 1$, if $\frac{1}{4} - e > 0$. Also player G has an incentive at $h_t \in H_t^{00}$ to deviate to $\gamma_t = 0$ and save the cost of banning non-existing products.

(iv) If $(\theta^*, \gamma^*) = (1, 1)$, then, after a small number of periods the process is in H_t^{11} in period t and for all further periods. For $h_t \in H_t^{11}$, the equilibrium continuation payoffs of players L and G in the candidate equilibrium are $\Pi_t^L = \frac{-e}{1-\delta}$ and $\Gamma_t = \frac{-\lambda}{1-\delta}$. Both players have payoff-increasing opportunities for one-step deviations if $\frac{1}{4} - e > 0$ and $\frac{1}{2} - \lambda > 0$. ■

Proof of Proposition 4 We show the existence of an equilibrium that alternates from histories in H_t^{01} to histories in H_{t+1}^{10} and back from histories in H_{t+1}^{10} to histories in H_{t+2}^{10}, due to the following local strategies:

- The government chooses $\gamma_t = 1$ in t if and only if $h_t \in H_t^{10} \cup H_t^{11}$. That is, if there are blueprints old enough so that they can be banned, then the government bans them.
- The firm L innovates if and only if the history $h_t \in H_t^{00} \cup H_t^{10}$. That is, L innovates in t if L did not innovate in $t-1$. Firm L does not innovate in t if it innovated in $t-1$.

This behavior moves the process according to the following pattern from the four sets of histories in t to two sets of histories in $t+1$:

$$\begin{aligned}
\text{If } h_t &\in H_t^{01}, \text{ then } (\theta_t, \gamma_t) = (0, 0) \text{ and } h_{t+1} \in H_{t+1}^{10}. \\
\text{If } h_t &\in H_t^{11}, \text{ then } (\theta_t, \gamma_t) = (0, 1) \text{ and } h_{t+1} \in H_{t+1}^{10}. \\
\text{If } h_t &\in H_t^{00}, \text{ then } (\theta_t, \gamma_t) = (1, 0) \text{ and } h_{t+1} \in H_{t+1}^{01}. \\
\text{If } h_t &\in H_t^{10}, \text{ then } (\theta_t, \gamma_t) = (1, 1) \text{ and } h_{t+1} \in H_{t+1}^{01}.
\end{aligned} \tag{13}$$

This mapping has the feature that the two sets of history that are visited in the future are either H_{t+i}^{10} or H_{t+i}^{01}. We can also associate the continuation payoffs of the two relevant players (L and G) given the mapping (13) for any initial state H_t^{xy}. Using the candidate equilibrium pricing behavior in the market stage which is fully determined by whether competition, monopoly or absence of market products prevails, we find:

$$\begin{aligned}\Pi_t^L(h_t \in H_t^{01}) &= 0+\delta\left(\frac{1}{4}-e\right)+0+\delta^3\left(\frac{1}{4}-e\right)+... \\ &= \delta\frac{1}{1-\delta^2}\left(\frac{1}{4}-e\right)\end{aligned} \tag{14}$$

$$\begin{aligned}\Pi_t^L(h_t \in H_t^{10}) &= \left(\frac{1}{4}-e\right)+\delta^2\left(\frac{1}{4}-e\right)+... \\ &= \frac{1}{1-\delta^2}\left(\frac{1}{4}-e\right)\end{aligned} \tag{15}$$

$$\Pi_t^L(h_t \in H_t^{11}) = 0+\delta\Pi_t^L(h_t \in H_t^{10}) \tag{16}$$

$$\Pi_t^L(h_t \in H_t^{00}) = \left(\frac{1}{4}-e\right)+\delta\Pi_t^L(h_t \in H_t^{01}). \tag{17}$$

Similarly, for the government, given the choices in (13),

$$\begin{aligned}\Gamma_t(h_t \in H_t^{01}) &= 0+\delta\left(\frac{1}{2}-\lambda\right)+0+\delta^3\left(\frac{1}{2}-\lambda\right)+... \\ &= \delta\frac{1}{1-\delta^2}\left(\frac{1}{2}-\lambda\right)\end{aligned} \tag{18}$$

$$\Gamma_t(h_t \in H_t^{10}) = \frac{1}{1-\delta^2}(\frac{1}{2}-\lambda) \tag{19}$$

$$\Gamma_t(h_t \in H_t^{11}) = 0+\delta\frac{1}{1-\delta^2}\left(\frac{1}{2}-\lambda\right) \tag{20}$$

$$\Gamma_t(h_t \in H_t^{00}) = \frac{1}{2}+0+\delta^2\frac{1}{1-\delta^2}\left(\frac{1}{2}-\lambda\right). \tag{21}$$

The payoffs of the players in each period are strictly bounded from above and below, and $\delta \in (0,1)$. Hence, we can apply the one-step deviation principle and consider only such deviations from the equilibrium path. Like in the cases $k=0$ and $k=1$, the optimizing behavior of L and F at the market phase has no dynamic implications for the state in the next period and is fully determined by whether competition, monopoly or absence of market products prevails in the way that has been described before. So we can take the equilibrium period payoffs from market interaction into consideration and turn directly to one-step deviations in states H_t^{01} and H_t^{10}, first in the legislation phase, then in the innovation phase.[11]

Let $h_t \in H_t^{01}$. Compare the payoffs for the equilibrium choice $\gamma_t = 0$ and a deviation to $\gamma_t = 1$. This deviation reduces the period payoff by λ in period t and has no consequences for the state in the next period, as there is no old product that is affected by this regulatory decision. Hence, the deviation reduces the government's payoff compared to the equilibrium choice. Note also that this suboptimality of $\gamma_t = 1$ is independent of $\theta_t \in \{0, 1\}$ in period t.

Let $h_t \in H_t^{10}$. The government's continuation payoff in the candidate equilibrium with $\gamma_t = 1$ is $\frac{1}{1-\delta^2}(\frac{1}{2} - \lambda)$. Consider a deviation to $\gamma_t = 0$. If $\theta_t = 1$ this leads to a state in H_{t+1}^{11}. Hence, government payoff is $0 + \delta\Gamma_t(h_t \in H_t^{11}) = \delta^2 \frac{1}{1-\delta^2}(\frac{1}{2} - \lambda) < \frac{1}{1-\delta^2}(\frac{1}{2} - \lambda)$. The deviation is not profitable if $\lambda \in (0, \frac{1}{2})$. If $\theta_t = 0$ this deviation leads to a state in H_{t+1}^{10}. The resulting government payoff is $0 + \delta\Gamma_t(h_t \in H_t^{10}) = \frac{\delta}{1-\delta^2}(\frac{1}{2} - \lambda) < \frac{1}{1-\delta^2}(\frac{1}{2} - \lambda)$ if $\frac{1}{2} - \lambda > 0$. The deviation is also not profitable.

We turn to the innovation phase. Let $h_t \in H_t^{01}$. The candidate equilibrium choice of L is $\theta_t = 0$, is followed by $\gamma_t = 0$ (as has just been shown) and moves the process to a state in H_{t+1}^{10}. It yields firm payoff $\Pi_t^L(h_t \in H_t^{01}) = \delta \frac{1}{1-\delta^2}(\frac{1}{4} - e)$. A deviation to $\theta_t = 1$ is also followed by $\gamma_t = 0$ (as has just been shown) and moves the process to $h_{t+1} \in H_{t+1}^{11}$. The resulting deviation payoff is $-e + \delta \frac{1}{1-\delta^2}(\frac{1}{4} - e)$. This shows that the equilibrium payoff exceeds the deviation payoff if $e \in (0, \frac{1}{4})$.

Let $h_t \in H_t^{10}$. The candidate equilibrium choice of L is $\theta_t = 1$, followed by $\gamma_t = 1$, moves the process to a state $h_{t+1} \in H_{t+1}^{01}$, and yields the firm L a payoff of size $\frac{1}{1-\delta^2}(\frac{1}{4} - e)$. A deviation to $\theta_t = 0$ is followed by $\gamma_t = 1$ (as we have seen above) but moves the process to a history $h_{t+1} \in H_{t+1}^{00}$. The payoff is $0 + \delta\left(\frac{1}{4} - e\right) + \delta^2\left(\delta\frac{1}{1-\delta^2}(\frac{1}{4} - e)\right) = \delta\frac{1}{1-\delta^2}(\frac{1}{4} - e)$. This deviation payoff is smaller than the payoff in the candidate equilibrium if $e \in (0, \frac{1}{4})$. ■

[11] Note that the same logic can be applied to show that one-step deviations from (13) for states H_t^{00} and H_t^{11} are also not profitable for L or G. We only consider one-step deviations along the equilibrium path here.

References

[1] Ault, Hugh J., 2013, Some reflections on the OECD and the sources of international tax principles, reprinted from *Tax Notes International*, 70(2), 2013, Working Paper of the Max Planck Institute for Tax Law and Public Finance No. 2013-03. Available at SSRN: http://ssrn.com/abstract=2287834 or http://dx.doi.org/10.2139/ssrn.2287834.

[2] Bankman, Joseph, 1999, The new market in corporate tax shelters, *Tax Notes*, 83(12), 1775-1795.

[3] Bhattacharyya, Sugato, and Vikram Nanda, 2000, Client discretion, switching costs, and financial innovation, *Review of Financial Studies*, 13(4), 1101-1127.

[4] Curry, Philip A., Claire A. Hill and Francesco Parisi, 2014, Tax avoidance and the optimal design of tax law, Minnesota Legal Studies Research Paper No. 09-16.

[5] Dharmapala, Dhammika, 2014, What do we know about base erosion and profit shifting? A review of the empirical literature, Illinois Public Law Research Paper No. 14-23; University of Chicago Coase-Sandor Institute for Law & Economics Research Paper No. 702. Available at SSRN: http://ssrn.com/abstract=2373549 or http://dx.doi.org/10.2139/ssrn.2373549.

[6] European Parliament, 2015, Special Committee on tax rulings and other measures similar in nature or effect, Draft Report on tax rulings and other measures similar in nature or effect, 2015/2066(INI).

[7] Fudenberg, Drew, and Jean Tirole, 1991, Game Theory, MIT Press, Cambridge, MA.

[8] The Group of 20 (G20), 2009, London summit communiqué: global plan for recovery and reform. Available at http://www.g20ys.org/upload/auto

/9c0ee8439921b1032fba1b8ca960ba8be8ca4e0d.pdf, accessed on September 21, 2015.

[9] Graham, John R., Michelle Hanlon, Terry Shevlin and Nemit Shroff, 2014, Incentives for tax planning and avoidance: evidence from the field, *Accounting Review*, 89(3), 991-1023.

[10] Harris, Christopher, and John Vickers, 1987, Racing with uncertainty, *Review of Economic Studies*, 54(1), 1-21.

[11] Hines, James R. Jr., 2004, On the timeliness of tax reform, *Journal of Public Economics*, 88(5), 1043-1059.

[12] Hong, Qing, and Michael Smart, 2010, In praise of tax havens: international tax planning and foreign direct investment, *European Economic Review*, 54(1), 82-95.

[13] Hopland, Arnt O., Petro Lisowsky, Mohammed Mardan, and Dirk Schindler, 2015, Income shifting under losses, NHH Discussion Paper FOR 21/2015.

[14] Johannesen, Niels, 2010, Imperfect tax competition for profits, asymmetric equilibrium and beneficial tax havens, *Journal of International Economics* 81(2), 253-264.

[15] Keen, Michael, and Kai A. Konrad, 2013, The theory of international tax competition and coordination, in: Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez (eds.), *Handbook of Public Economics*, Vol. 5. Elsevier, Amsterdam, 257-328.

[16] Kleinbard, Edward D., 2011, Stateless income, Florida Tax Review, 11(9), 699-774.

[17] Lisowsky, Petro, 2010, Seeking shelter: Empirically modeling tax shelters using financial statement information, *Accounting Review*, 85(5), 1693-1720.

[18] Organisation for Economic Co-operation and Development (OECD), 2016, OECD secretary-general report to the G20 finance ministers. Available at http://www.oecd.org/tax/oecd-secretary-general-tax-report-g20-finance-ministers-july-2016.pdf (accessed on August 10, 2016).

[19] Reinganum, Jennifer F., 1989, The timing of innovation: Research, development, and diffusion, in: Richard Schmalensee and Robert Willig (eds.), Handbook of Industrial Organization, Volume 1, 1989, 849-908.

[20] Slemrod, Joel, 2004, The economics of corporate tax selfishness, National Tax Journal, 57(4), 877-899.

[21] Slemrod, Joel, and Varsha Venkatesh, 2002, The income tax compliance cost of large and mid-size businesses, Working Paper No. 914, Ross School of Business Working Paper Series.

[22] Slemrod, Joel, and John D. Wilson, 2009, Tax competition with parasitic tax havens, *Journal of Public Economics,* 93(11-12), 1261-1270.