The B.E. Journal of Theoretical Economics

Contributions

Volume 9, *Issue* 1 | 2009 | *Article* 25

A Positive Theory of Income Taxation

Oriol Carbonell-Nicolau*

*Rutgers University and CODE, Universitat Autònoma de Barcelona, carbonell@econ.rutgers.edu

Recommended Citation

Oriol Carbonell-Nicolau (2009) "A Positive Theory of Income Taxation," *The B.E. Journal of Theoretical Economics*: Vol. 9: Iss. 1 (Contributions), Article 25.

Available at: http://www.bepress.com/bejte/vol9/iss1/art25

A Positive Theory of Income Taxation*

Oriol Carbonell-Nicolau

Abstract

We explore the consequences of electoral competition for nonlinear income taxation. Our model is a dynamic version of the standard two-party electoral competition model adapted to nonlinear income taxation. The theory has a number of desirable features. First, equilibria always exist, even though the set of admissible tax policies is multidimensional. Second, the Nash set can be characterized generically, and its components give sharp predictions. Third, the features of equilibrium tax policies depend only on empirically meaningful fundamentals.

Equilibrium tax schedules benefit the more numerous income groups and place the burden of taxation on income groups with fewer voters. For empirical income distributions, the features of an equilibrium tax schedule are reminiscent of Director's law of public income redistribution (Stigler [39]).

KEYWORDS: nonlinear income taxation, electoral competition, Director's law, extensive zero-sum game

---

*I am indebted to Colin Campbell and Rich McLean for very useful discussions. Thanks also to Marcus Berliant, Alessandro Lizzeri, Daijiro Okada, Sjaak Hurkens, the editor, an anonymous referee, and seminar participants at Rutgers University, Universitat Autònoma de Barcelona, Universidad Carlos III de Madrid, Universidad de Málaga, Universitat Pompeu Fabra, and the 2008 North American Winter Meetings of the Econometric Society; John Bowblis for research assistance; and the Spanish Ministry of Education and Science and the Research Council at Rutgers University for financial assistance. Part of this research was conducted during the author's visit to CODE at Universitat Autònoma de Barcelona. The author is grateful to CODE for its support.

1. Introduction

Director's law of income redistribution argues that 'public expenditures are made for the primary benefit of the middle classes, and financed with taxes which are borne in considerable part by the poor and rich' (Stigler [39], p. 1). For instance, publicly financed schools and universities, social security benefits, public housing, and mortgage interest tax deductions for homeowners are often argued to disproportionately benefit the middle class.

Stigler [39] sketches a theoretical argument whereby the middle class benefited from (1) being in coalition with the rich in the nineteenth century and (2) entering in coalition with the poor in the twentieth century, as the increase in the flexibility of taxes and expenditure programs redistributed income increasingly toward lower income classes. This argument is hard to formalize, for it entails redistribution of income among (at least) three income groups, and it is well-known that collective decision-making processes modeled as strategic games with multidimensional action spaces (such as spaces of nonlinear tax schemes) typically have no (pure-strategy) equilibrium.

The essence of the problem that arises when tax policies are nonlinear can be grasped in the context of the standard model of two-party competition, where voters switch support from one candidate to the other if promised a more favorable policy. When the set of admissible policies is sufficiently rich, this creates incentives for a bidding war between the parties, which leads to cycling over alternative platforms. This is a fundamental problem that is not specific to taxation settings, but rather rooted in Arrow's impossibility theorem and intrinsic to environments of collective choice over many dimensions.[1] This argument (or some variation of it) can be used to explain why the constraints on the set of admissible tax schemes cannot be dispensed with in most of the literature on positive income taxation.

To overcome this problem, an important part of the existing literature on voting over income taxes assumes policy spaces that are artificially constrained. For instance, to be able to make use of the median voter theorem, the seminal papers of Romer [36], Roberts [34], and Meltzer and Richard [30] consider only linear tax schemes.[2] These studies do not fully account for Director's law in

[1]See Austen-Smith and Banks [3] for a general treatment.

[2]In general, resort to various forms of constraints imposed on the set of admissible tax schemes for the sole purpose of obtaining a coherent model, namely one for which an equilibrium can be shown to exist, is pervasive. Despite the constraints, the field has produced studies that are useful to understand various aspects of the political economy of income taxation (cf. Romer [36], Roberts [34], and Meltzer and Richard [30], Cukierman and Meltzer [14], Gouveia and Oliver [19], Snyder and Kramer [38], Marhuenda and Ortuño-

that redistribution of income via linear tax schedules benefits not only the middle class (or the median voter) but also the poor.

Models of pork-barrel politics (e.g., Lindbeck and Weibull [25] and Dixit and Londregan [17]), which build on theories of probabilistic voting (cf. Austen-Smith and Banks [4] and references therein), can handle nonlinear tax schemes and relate to Director's law. In these models, voters have attachments to political parties that can be loosened through the offer of private consumption in the form of pork-barrel spending. Lindbeck and Weibull [25] derives a version of Director's law under the assumption that the middle class cares less about the candidates' personal attributes (and more about the policy implemented) than the poor and the rich. While Lindbeck and Weibull [25] requires that voter preferences on ideological aspects of public policy be independent of the distribution of income, Dixit and Londregan [17] removes this separation between ideology and income redistribution: as long as the rich and the poor have a strong ideological affinity for the rightist and the leftist parties, respectively, party loyalties are most likely to be loosened for middle income voters, who consequently receive the largest strategic transfers.

This paper shows that Director's law can emerge in a dynamic version of the standard Downsian model of electoral competition, independently of the distribution of voter preferences beyond private consumption.

We use Downs' [18] view that electoral competition is 'a mechanism whereby political parties that are engaged in what Schumpeter called a "competitive struggle for the people's vote" are obliged to take account of the preference of the electors for one policy rather than another' (Barry [5]). In this regard, we model electoral competition as a standard non-cooperative game played by two candidates who strive to maximize their vote shares (or their probability of holding office). Specifically, we extend the static two-party electoral competition model studied in Carbonell-Nicolau and Ok [11] by allowing the parties to reveal their tax platforms gradually in more than one period. The candidates reveal—when it is their turn to do so—small pieces of information concerning their platform and must commit to any current and past announcements.[3,4] We assume a discrete money unit and formulate this scenario as a finite extensive game. Imposing a smallest money unit $\varepsilon > 0$ means that all money amounts (tax liabilities, pre-tax income levels, etc.) must be multiples of ε. This, together with the assumption that, when it is their turn to speak,

Ortín ([28],[29]), Roemer [35], Benabou [6], Berliant and Gouveia [7], Austen-Smith [2], Hindriks [21], Kranich [22], and De Donder and Hindriks [15].)

[3]This assumption is discussed in Subsection 2.1.

[4]The dynamical process may be interpreted as a stylized instance of political campaigning.

the candidates must provide some information (however minimal this information may be) about their prospective tax policy, implies that tax functions are described in finite time. Finiteness of the strategy spaces guarantees the existence of a subgame perfect equilibrium.

The theory has a number of desirable features. First, equilibria always exist, even though the set of admissible tax policies is multidimensional. Second, the Nash set can be characterized generically, and its components give sharp predictions. Third, the features of equilibrium tax policies depend only on empirically meaningful fundamentals—the shape of the income distribution and the government's target revenue.

We show that, at each (generic) component of the Nash set, equilibrium tax schemes lie within a small set of admissible policies that benefit the more populous voter groups and place the burden of taxation on income groups with fewer voters. When the income distribution resembles a log-normal density function, the features of an equilibrium tax schedule are reminiscent of Director's law of income redistribution.[5]

Finally, our model allows for the introduction of sources of voter heterogeneity other than pre-tax income, such as marital status, immigration status, etc., according to which tax structures may discriminate between taxpayers. Results are obtained for any given partition of the population consisting on various groups of 'similar' individuals (where the relation of similarity is defined in terms of the individuals' pre-tax income and possibly in terms of other attributes that may be relevant for tax purposes).

A number of research avenues have been explored to study collective decision-making under nonlinear tax schemes.[6] Nonetheless, the forces at work in our model—and their implications for equilibrium outcomes and Director's law—have (to the best of our knowledge) not been highlighted. We argue that equilibrium outcomes can be viewed as consequence of the fact that the candidates have a desire to render their policies ambiguous (or flexible), and contrast this idea with a related (albeit different) notion: delay of electoral commitment. The candidates tend to make ambiguous announcements in order to minimize the amount of tax that is precisely assigned to particular voter groups (or, in other words, in order to maximize the amount of tax whose incidence is

[5] On the other hand, the results suggest a u-shaped pattern of effective marginal tax rates, which is observed in the data on US effective marginal tax rates (see, for instance, [9]).

[6] See, for instance, Aumann and Kurz [1], Hettich and Winer [20], Lindbeck and Weibull [25], Chen [13], Myerson [33], Lizzeri and Persico [26], Laslier and Picard [23], Carbonell-Nicolau and Klor [10], Carbonell-Nicolau and Ok [11], Dekel, Jackson, and Wolinsky [16], and Ledyard [24].

left unspecified). Intuitively, obscuring future moves gives flexibility to shape policy through time. We believe that our analysis emphasizes effects that are likely to play a role in more general settings (encompassing, e.g., endogenous labor supply and party ideology).

The convenience of introducing a discrete money unit in a dynamic model of electoral competition was first exploited by Dekel, Jackson, and Wolinsky [16]. Our analysis differs from that in [16] in the following respects: in [16], voters' heterogeneity arises from differences in (observable) preference biases towards one candidate (or fixed platform), not from differences in initial endowments, and candidates try to offset the biases by offering money payments (taken from candidate-specific endowments) to voters, rather than collecting taxes from different groups. Absent any voter preference bias, the setting in [16] features N identical voters, and in this case our exercise is not meaningful. Consequently, our characterization of the effects of income taxation on pre-tax income distributions is not possible within the setting of [16]. In addition, unlike [16] we provide a full analysis of the entire set of Nash equilibria of (a slight perturbation of) the electoral game.

The paper is organized as follows. Section 2 introduces the setup and discusses the modeling strategy. The results appear in Section 3. Subsection 3.1 contains an example, and the general results are provided in Subsections 3.2 and 3.3. Subsection 3.2.1 furnishes intuition for the general results and discusses the effects of population grouping on equilibrium outcomes. Section 4 concludes. The proofs are relegated to Section 5.

2. The model

Society consists of a continuum of individuals and two political parties, denoted as A and B. Let $\overline{X}$ be a large positive real, and, for $+\infty > \varepsilon > 0$, define

$$X_\varepsilon := \{0, \varepsilon, 2\varepsilon, ...\} \cap \left[0, \overline{X}\right].$$

The set X_ε represents the universe of possible pre-tax income levels (multiples of ε). We refer to ε as the *money unit* for X_ε.

Fix a nonempty finite set $\mathfrak{A}$. The set $\mathfrak{A}_\varepsilon := X_\varepsilon \times \mathfrak{A}$ represents a set of individual attributes. Each individual is characterized by an element (x, a) of $\mathfrak{A}_\varepsilon$, which is a description of the individual's pre-tax income x along with other attributes a that may be relevant for tax purposes (e.g., single/married, homeowner/renter, etc.).[7]

[7]Imagine a situation where society has identified a partition of the set of all individuals

A pre-tax ***income distribution*** is defined as an element of

$$\mathcal{D} := \left\{ d \in \bigcup_{\varepsilon>0} [0,1]^{\mathfrak{A}_\varepsilon} : \sum_{(x,a)\in\mathfrak{A}_\varepsilon} d(x,a) = 1 \right\}.$$

A distribution $d \in \mathcal{D}$ determines the measure $d(x,a)$ of individuals with characteristic (x,a). For $d \in \mathcal{D}$, let $\varepsilon(d)$ denote the money unit corresponding to the domain of d. In the remainder of the paper, $\{d > 0\}$ shall be used to designate the set of characteristics (x,a) that are assigned positive mass:

$$\left\{(x,a) \in \mathfrak{A}_{\varepsilon(d)} : d(x,a) > 0\right\} =: \{d > 0\}.$$

Let

$$\mathcal{M} := \left\{ (d,r) : d \in \mathcal{D}, 0 \le r \le \sum_{(x,a)\in\mathfrak{A}_{\varepsilon(d)}} d(x,a)x \right\}.$$

Each tuple (d,r) in $\mathcal{M}$ consists of an income distribution d and a target revenue r. A ***tax policy*** in $(d,r) \in \mathcal{M}$ is a map $t : \{d > 0\} \to X_{\varepsilon(d)}$ that assigns to each vector of attributes (x,a) a total tax liability $t(x,a)$ with the property that $0 \le t(x,a) \le x$ for all $x \in \{d > 0\}$. The first inequality rules out negative taxation, that is, subsidies.[8] The second inequality says that an individual can never be required to pay more than her endowment. Let $\mathcal{P}_{(d,r)}$ represent the set of all tax policies.

A tax policy t is ***admissible*** for $(d,r) \in \mathcal{M}$ if

$$\sum_{(x,a)} t(x,a)d(x,a) \ge r.$$

That is, t is admissible if it collects at least the target revenue r. The set of all tax policies that are admissible for (d,r) is designated by $\mathcal{T}_{(d,r)}$.

---

such that each element of the partition contains individuals that are identical with respect to a number of characteristics (pre-tax income, marital status, immigrant status, etc.). Tax structures cannot discriminate between people belonging to the same element of the partition, and may discriminate between members of different elements of the partition. Thus, the partition is a specification of the relation of 'similarity' between individuals that is necessary to objectify the notion of horizontal equity (here we are referring to the traditional public finance concept of horizontal equity; see Berliant and Strauss [8]). In this paper, we take this partition as given. In terms of our notation, the population is partitioned into as many groups as there are elements in $\mathfrak{A}_\varepsilon$, and each $(x,a) \in \mathfrak{A}_\varepsilon$ can be interpreted as the list of characteristics (including pre-tax income) shared by the members of group (x,a). (If $\mathfrak{A}$ is a singleton, then income is the only source of discrimination.)

[8] Allowing for subsidies would not change the essence of our results.

Before the election, each candidate advocates an admissible tax policy, possibly revealed gradually as follows. There is a number of rounds $1, 2, ...$ The candidates take actions in each round as indicated next. Choose $(d, r) \in \mathcal{M}$. In round 1, candidate A announces a mapping $f_1 \in \mathcal{P}_{(d,r)}$. Any such mapping is called an ***announcement***, and may be interpreted as a way of raising part (or all) of the required revenue r. By proposing f_1, candidate A commits to levying (at least) $f_1(x, a)$ on voter group (x, a). If $\sum_{(x,a)} f_1(x, a)d(x, a) < r$, then f_1 falls short of collecting the target revenue. In this case, by announcing f_1, candidate A reveals only part of her proposed policy. In subsequent rounds, the candidate will indicate how the remainder of the required revenue, $r - \sum_{(x,a)} f_1(x, a)d(x, a)$, will be collected. Also in round 1, candidate B makes an announcement $g_1 \in \mathcal{P}_{(d,r)}$, with a similar interpretation. The announcements are revealed sequentially. Some candidate moves first and then the opponent takes an action having observed the other player's move. Nature determines the order of moves. To avoid difficulties generated by an asymmetric treatment of the players (as will become clear our game features a second-mover advantage), we shall assume that each candidate has a 50% chance of moving first.

Again in round 2, nature determines whether A moves first or B does. Candidate A's second announcement, $f_2 \in \mathcal{P}_{(d,r)}$, is made public in round 2, after A's observation of B's first announcement, g_1, and possibly B's second announcement (if A moves second in round 2); f_2 must be consistent with previous announcements made by A in the sense that $f_2 \geq f_1$.[9] After observing A's first proposal and possibly A's second move (if B moves second in round 2), candidate B makes a second announcement, g_2, also in round 2. This announcement must be consistent with B's first proposal, g_1, as specified above. The parties make proposals according to this time frame, each proposal being consistent with previous proposals as indicated. In each round, each candidate has a 50% chance of moving first.[10]

In any given period, a candidate's announcement f is ***final*** if $f \in \mathcal{T}_{(d,r)}$. With a final announcement, a candidate discloses all information about its advocated tax policy and commits to its implementation, conditional on winning the election. The sequence of campaign promises reaches an end when both parties have made a final announcement.[11]

[9] This assumption is discussed in Subsection 2.1.

[10] We know that most of our results would prevail if actions were taken simultaneously in each round, or if nature chose whether the moves are sequential or simultaneous at the beginning of each round. In the latter case, the probability of sequential moves could be history-dependent.

[11] Observe that an announcement $f \in \mathcal{P}_{(d,r)}$ could also be interpreted as a "promise" that

We assume that, when it is one candidate's turn to make an announcement, this candidate must give new information about its prospective policy. More precisely, given two successive announcements f and g of the same candidate such that f is not final, $g \neq f$.[12]

After each candidate has fully specified a final proposal, the election takes place. Each voter casts a ballot for one of the two candidates. The candidate that receives the most votes wins the election and implements her proposed tax policy. Ties are broken via an equal probability rule. Voters vote for the candidate whose final announcement most favors their economic interests. In other words, each voter chooses the candidate who will enact, if elected, a tax policy under which the voter's disposable income is maximal. In case of indifference, voters toss a fair coin to determine their choice.

The candidates are opportunistic; they wish to maximize their net plurality, which is defined as the difference between their vote share and the vote share of the opponent. Formally, if (f, g) represents the observed pair of final announcements, a candidate i receives a payoff of

$$u^i_{(d,r)}(f,g) := \sum_{(x,a):f(x,a)<g(x,a)} d(x,a) - \sum_{(x,a):f(x,a)>g(x,a)} d(x,a), \tag{1}$$

if $i = A$, and $u^i_{(d,r)}(f,g) := -u^A_{(d,r)}(f,g)$ if $i = B$.[13]

The above scenario can be embedded in the formal definition of a two-player zero-sum extensive game $G_{(d,r)}$ parameterized by an income distribution d and a revenue requirement r.[14] We focus on the notion of Nash equilibrium and subgame perfect equilibrium.

---

each group of individuals (x, a) will pay at most $f(x, a)$ plus the maximum additional tax this group could face given what is left to be collected.

[12]This assumption is discussed in Subsection 2.1.

[13]Other standard candidate objectives include the vote share and the probability of winning. Assuming that the candidates maximize the vote share would not change any of the results of the paper. Moreover, if $u^A_{(d,r)}$ were defined as a continuous, strictly increasing, and symmetric around zero transformation of the expression in (1), and similarly for $u^B_{(d,r)}$, all the results would remain unaltered. Observe that this transformation permits a pointwise approximation of the candidates' objective to the probability of winning. Finally, Theorem 1 is also true when the candidates' objective is exactly the probability of winning (and not just a pointwise approximation to it).

[14]Observe that the set of all possible histories is finite. Further, the game has finite horizon (i.e., all histories are finite).

2.1. Remarks on the modeling strategy

The model proposed here is in the tradition of the standard two-party electoral competition model (Downs [18]). There is, however, an important difference between the standard model and our model: the latter is richer than the former in the sense that in the current model candidates have more flexibility in the strategies they can use. For example, in the standard model, which is static, the candidates must completely reveal their policy in one shot; by contrast, in our model, the candidates could completely reveal their policy in the first period if they wished, but may decide to wait to do so. It turns out that this extra flexibility matters in that, in equilibrium, the players choose to wait (more on this in Subsection 2.2).

In light of the above comparison between models, we can view our assumption that the candidates must give new information about their prospective tax policy (each time it is their turn to make an announcement) as a weakening of the standard assumption that the candidates must completely reveal their policy in one period.[15]

On the other hand, the assumption that the candidates must commit to past announcements is also made in the standard model, where the candidates are not allowed to change their actions once a policy has been chosen. It is natural to assume that platform adjustments are costly in that they entail reversing previous promises. In this paper (as in virtually all the literature on electoral competition with commitment), we assume that it is too costly for the candidates to rectify past moves.[16]

Finally, our model is one possible extension of the static model towards models of gradual commitment, but there are obviously alternative formulations. While the analysis of some of these alternatives lies outside the scope of this paper, we discuss possible variants in Sections 3.2.1 and 4.

2.2. Second-mover advantage

The static (one-period) version of our model studied in Carbonell-Nicolau and Ok [11] lacks a pure-strategy equilibrium. In fact, given any admissible tax

[15] Alternatively, one may assume that candidates may remain silent, each time it is their turn to speak, at a cost. If the total cost incurred by a candidate is convex in the number of times the candidate fails to provide new information, then the game can be shown to possess an equilibrium. We conjecture that in this new game Theorem 1 would remain intact.

[16] One could envisage a game where, in each round, each candidate must either respect foregoing announcements or incur a cost to amend them. This is related to the variant proposed in footnote 15. We conjecture that the new game would not affect Theorem 1. A thorough analysis is left for future research.

function, it is always possible to find another tax policy meeting the revenue requirement and defeating the original tax function under pairwise majority voting. This would also be true here (at least for sufficiently small money units) if the parties were constrained to fully reveal their tax policy in one period. This means that each candidate i would like her opponent to completely reveal their tax policy in the first round, since then i could commit to some tax platform that defeats the opponent's policy, after observing the opponent's first move.[17] Obviously, in equilibrium, the opponent would never find it optimal to make a final announcement in the first round. Thus, there is an incentive for the candidates to reveal little information in each round, thereby gaining leverage to react to the opponent's announcements in future rounds.[18]

The intuition behind our main results is related to the second-mover advantage. This is discussed in Section 3.2.1.

3. Results

In the present model, the pre-tax income distribution is exogenously given, and therefore individuals cannot escape excessive tax burdens by reducing their labor supply. For this reason, in the equilibria described here, and absent any limits on the extent to which individuals may be taxed, smaller voter groups tend to be expropriated. This is obviously unrealistic and would not occur in a model à la Mirrlees [32], with endogenous labor supply. Since the introduction of distortionary taxation lies outside the scope of this paper, one might for now be content with the assumption that, for each group (x, a), there is a maximum (exogenously given) tax liability $\lambda_{(x,a)} \in X_{\varepsilon(d)}$ that may be imposed on group (x, a). This assumption would not change the results of the paper, yet we have omitted it to ease notation.[19]

Our first result states, roughly speaking, that $G_{(d,r)}$ possesses a subgame perfect equilibrium whose corresponding tax function is such that taxes are borne by less populous voter groups. Before stating the general result, we present a special case, with three income groups, which illustrates the idea behind the proof of Theorem 1.

[17]Even if the second mover made a final announcement in the first round, it would be possible for the first mover (for sufficiently small money units) to reveal little information in the first round and then choose, in the second round (and after observing the opponent's final policy), some tax scheme ensuring victory.

[18]This feature of the model contrasts with the first-mover advantage exhibited by other extensive forms, such as the bargaining game of alternating offers and the Stackelberg game.

[19]The results of Section 3, stated (with the obvious modifications) in terms of the upper bounds $\lambda_{(x,a)}$, remain valid.

3.1. An example

In this example we assume that the set $\mathfrak{A}$ of individual attributes other than income is a singleton, so that the tax system does not discriminate between individuals whose pre-tax endowment is identical.

Consider a population consisting of nine individuals, two of them endowed with (pre-tax) income \$3, three of them endowed with \$1 each, and the remaining four in possession of \$2 each.[20] Let d denote the corresponding pre-tax income distribution. Suppose that the money unit is set at $\varepsilon(d) = \$1$.

Say $r = 7$. Observe that in this case no admissible tax policy leaves the two poorest groups untaxed. We shall first construct one equilibrium profile for this case and then look at the associated equilibrium tax policy. Let us suppose that candidate B plays only strategies that increase taxes by \$1 in each round. This simplifies the description of the following strategy of A, but a similar treatment is possible if B plays any kind of strategy. Throughout the sequel, we consider paths of play in which player A is always the first mover (the worst-case scenario for this player). By determining A's payoffs along these paths at a given strategy profile, we can find a lower bound for this player's payoff in the game.

Suppose that A starts announcing that \$1 will be collected from the two richest individuals. Consider a path of play where B's first-round action does not coincide with A's announcement. In this case A can imitate B's first-round action and proceed as follows in subsequent rounds:

- If B keeps collecting revenue from groups other than the richest group, A can keep imitating B's moves until A falls $\$m$ short of meeting the revenue requirement, where $m \leq 4$. At this point, A can finish collecting the revenue from the richest group. This clearly secures a payoff of at least 0 against B's moves.

- If B chooses to collect \$1 from the richest group in some round t (and, in previous rounds, A has imitated B's actions) then at the end of round t both announcements coincide, and both candidates' collected revenue is at least \$5, so A can meet the revenue constraint by incrementing taxes on the richest group, and this guarantees a payoff of at least 0.

Now consider a path of play where B's first-round action coincides with A's initial announcement. Suppose that in the second round A levies an additional dollar on the richest group. If B's second-round action does not coincide with

[20]While this distribution is not exactly a member of $\mathcal{D}$, it can be transformed into a member of $\mathcal{D}$ without altering the essence of the example.

A's second announcement, A can proceed as before to secure a payoff of at least 0. If, on the other hand, B's second-round action coincides with A's second announcement, A can meet target revenue by levying an additional dollar on the poorest group. This gives A a minimum payoff of 0.

We have argued that an action plan specifying, in each round, an announcement for A contingent on B's action in the preceding round may be obtained that secures a payoff of at least 0 against any strategy of B. While this action plan does not constitute a full contingent plan for A, one can prove that an equilibrium strategy profile μ in $G_{(d,r)}$ may be constructed in which both players' moves along the equilibrium path are consistent with the said action plan. This results in the implementation of a tax policy t satisfying

$$t(x) = \begin{cases} 3 & \text{if } x = 3, \\ 1 & \text{if } x = 1, \\ 0 & \text{if } x = 2. \end{cases}$$

Observe that t exempts the more populated group from taxation.

To see how this argument interacts with the choice of a money unit, observe that if $\varepsilon(d) = \$0.01$ the previous argument gives an equilibrium tax policy

$$t(x) = \begin{cases} 3.5 & \text{if } x = 3, \\ 0 & \text{if } x = 1, \\ 0 & \text{if } x = 2, \end{cases}$$

which exempts the two more populated groups from taxation. Thus, smaller money units lead to equilibrium tax policies that are closer to the tax function that maximizes the number of exemptions on the more populous income groups. This idea appears in Theorem 1 (which generalizes this example to arbitrary income distributions) in terms of an error "band" around the said tax function: given an (arbitrarily small) error margin, there exists a sufficiently small money unit such that there is an equilibrium tax policy within the corresponding neighborhood of the tax function that maximizes exemptions on the more populous income groups.

3.2. Characterizing an equilibrium

Let $\widetilde{\mathcal{E}}_{(d,r)}$ be the set of all admissible tax policies $t \in \mathcal{T}_{(d,r)}$ such that for all (x, a),

$$t(x, a) > 0 \Rightarrow t|_{\{(y,b):d(y,b)<d(x,a)\}} = \mathfrak{i}|_{\{(y,b):d(y,b)<d(x,a)\}},$$

where $\mathfrak{i} : \{d > 0\} \to \mathbb{R}$ is defined by $\mathfrak{i}(x, a) := x$.[21] The set $\widetilde{\mathcal{E}}_{(d,r)}$ contains those admissible tax policies that tax more populous groups only if less numerous groups have been taxed to the fullest extent. Define

$$\mathcal{E}_{(d,r)} := \left\{ t \in \widetilde{\mathcal{E}}_{(d,r)} : \text{there is no } \tau \in \mathcal{T}_{(d,r)} \text{ with } \tau \lneqq t \right\}.$$

Thus, $\mathcal{E}_{(d,r)}$ is the set of admissible tax policies that levy r on the less numerous groups, leaving the more numerous groups untaxed.

The result below depends on a parameter $+\infty > \eta > 0$, which may be interpreted as an error margin for the graph of an equilibrium tax policy. Given η and a model $(d, r) \in \mathcal{M}$, consider the following statement: the tax policy implemented at a subgame perfect equilibrium of $G_{(d,r)}$ lies in $\mathcal{E}_{(d,r)}$ with an error margin of η or, more precisely, it lies in $N_\eta\left(\mathcal{E}_{(d,r)}\right)$.[22] Obviously, if η is very large, then the assertion is vacuous. If, on the other hand, η is small, then the equilibrium policy lies (approximately) in $\mathcal{E}_{(d,r)}$, and this characterizes the equilibrium policy quite sharply, given the 'smallness' of $\mathcal{E}_{(d,r)}$ within the set of all admissible tax policies.

We state our result for all the members of a sub-class of models in $\mathcal{M}$, which depends on η. Roughly speaking, we require that the money unit $\varepsilon(d)$ be sufficiently small relative to the error margin η. This imposes an upper bound on $\varepsilon(d)$ that decreases with the error margin. If, for example, the error margin η is 100 times the money unit $\varepsilon(d)$, and $\varepsilon(d)$ is taken to be one cent of a dollar, then Theorem 1 says that the graph of an equilibrium tax schedule lies within a neighborhood of radius one dollar of an element of $\mathcal{E}_{(d,r)}$.

Theorem 1. *Suppose that $+\infty > \eta > 0$. There exists $+\infty > \varepsilon_\eta > 0$ such that for every $(d, r) \in \mathcal{M}$ with $\varepsilon(d) \leq \varepsilon_\eta$, $G_{(d,r)}$ has a subgame perfect equilibrium whose corresponding tax policy lies in $N_\eta\left(\mathcal{E}_{(d,r)}\right)$.*

Clearly, the content of Theorem 1 is meaningful only if η is a small number, and the informativeness of the theorem is inversely related to the size of η. One can ensure a precise statement by choosing a small η. The following example illustrates the relationship between the error margin and the magnitude of the money unit using real data.

Example 1. Figure 1 depicts the US household income distribution for the year 2004. The data depicted can be presented as an element of $\mathcal{D}$, for some

[21] Here, $t|_{\{(y,b):d(y,b)<d(x,a)\}}$ stands for the restriction of t to the domain $\{(y, b) : d(y, b) < d(x, a)\}$, etc.

[22] Here $N_\eta\left(\mathcal{E}_{(d,r)}\right)$ denotes the set $\bigcup_{t\in\mathcal{E}_{(d,r)}} N_\eta(t)$, where $N_\eta(t)$ stands for the η-neighborhood of t in $\mathcal{T}_{(d,r)}$ (relative to the sup metric).

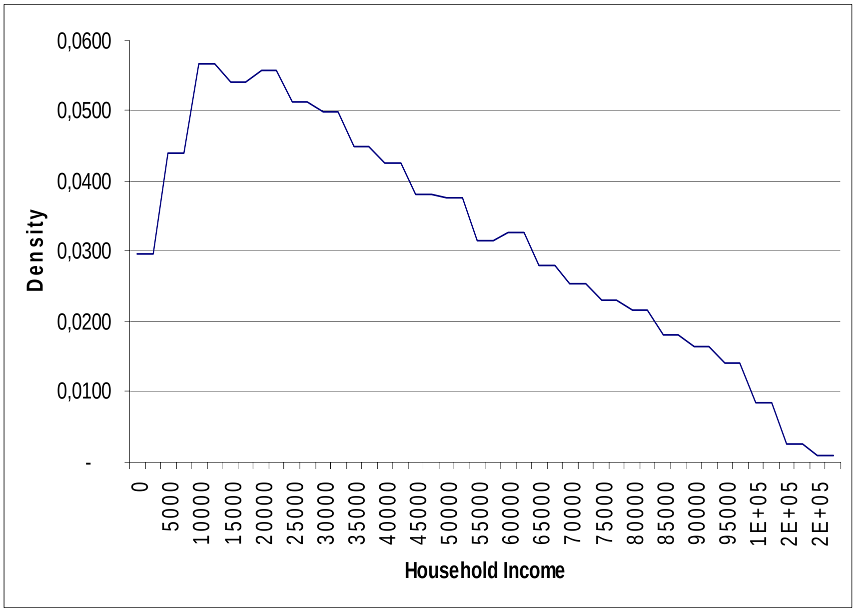

Source: U.S. Census Bureau, Current Population Survey, 2005 Annual Social and Economic Supplement

Figure 1. US Income Distribution for Households: 2004

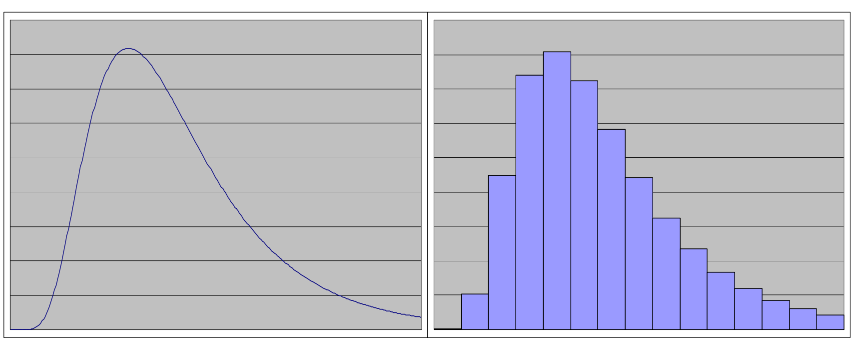

Figure 2. A log-normal density and its discrete analogue

choice of a money unit. (In this example we assume that the set $\mathfrak{A}$ of individual attributes other than income is a singleton, so that the tax system does not discriminate between individuals whose pre-tax endowment is identical.) Let this distribution be denoted by d_{US}, where $\varepsilon(d_{\mathrm{US}}) = \$0.01$ (i.e., say that the money unit is one cent of a dollar). Let r_{US} be the total amount of taxes collected by the Internal Revenue Service in 2004.[23] If one sets $\eta = \$1.74$, then the model $(d_{\mathrm{US}}, r_{\mathrm{US}})$ is such that the tax policy implemented at some subgame perfect equilibrium of $G_{(d_{\mathrm{US}}, r_{\mathrm{US}})}$ lies in $N_\eta\left(\mathcal{E}_{(d_{\mathrm{US}}, r_{\mathrm{US}})}\right)$. Thus, if one takes the money unit to be one cent of a dollar, Theorem 1 gives, with an error margin of at most \$1.74, an equilibrium tax policy in $\mathcal{E}_{(d_{\mathrm{US}}, r_{\mathrm{US}})}$.

Observe the implications of Theorem 1 for the features of the equilibrium tax policy in a society where the income distribution is of a log-normal type (Figure 2).[24] For this type of distribution, at the equilibrium of Theorem 1, the tax revenue is collected from the tails of the distribution. This is consistent with Director's law of public income redistribution, which states that 'public expenditures are made for the primary benefit of the middle classes, and financed with taxes which are borne in considerable part by the poor and rich' (Stigler [39]).[25]

3.2.1. Discussion

It is useful to outline some intuition for Theorem 1. Given our discussion in Subsection 2.2, it may appear that incrementing taxes for smaller groups

[23]Net collections for individual income tax, IRS Data Book 2004, Table 1.

[24]This type of distribution is empirically relevant: the fact that income obeys a log-normal distribution is widespread. The ***log-normal distribution*** has the probability density function

$$f(x;\mu,\sigma) = \frac{\exp\left(\frac{-(\ln x-\mu)^2}{2\sigma^2}\right)}{x\sigma\sqrt{2\pi}},$$

for $x > 0$, where μ and σ are the mean and standard deviation of the variable's logarithm, respectively. Discrete analogues of f can be defined as follows. Given a partition $\mathcal{I} = \{(0,\delta), [\delta, 2\delta), ...\}$ of the positive real line into intervals of length $\delta > 0$, the discrete version $f_{\mathcal{I}}$ of f given $\mathcal{I}$ is

$$f_{\mathcal{I}}(x;\mu,\sigma) = \begin{cases} \frac{1}{\delta}\int_0^\delta f(y;\mu,\sigma)dy & \text{if } 0 < x < \delta, \\ \frac{1}{\delta}\int_\delta^{2\delta} f(y;\mu,\sigma)dy & \text{if } \delta \leq x < 2\delta, \\ \quad\vdots & \qquad\vdots \end{cases}$$

[25]On the other hand, the locus of equilibrium marginal tax rates against income is reminiscent of a u-shaped pattern. A first look at some US data (see, for instance, [9]) reveals that effective marginal tax rates (e.g., for single taxpayers) first decrease and then increase.

will serve the candidates' purpose: In each round, the candidates harden their platforms by incrementing taxes for some voter group, and individuals within this group are treated symmetrically. Because the size of commitment (measured in terms of tax revenue) is proportional to the size of the group on which the tax rate is levied, and the candidates prefer seeing the opponent's move before taking an action, incrementing taxes for the smaller groups (as in the equilibrium of Theorem 1) implies a lesser commitment.

In light of this argument, one may be tempted to conclude that the candidates' desire to delay electoral commitment is the main driving force behind Theorem 1. It turns out that this effect alone does not explain the result. Indeed, there are situations in which incrementing taxes for small groups implies a *higher* amount of commitment than incrementing taxes for larger groups. This is illustrated in the following example.

Suppose that there are three income groups $x_1 = 1$, $x_2 = 2$, and $x_3 = 3$ (x_i denotes group i's endowment) with sizes .35, .4, and .25, respectively. Suppose that the revenue requirement is $r = .78$. It is clear that the revenue cannot be collected by taxing group x_3 only, while it is possible to meet the revenue requirement by taxing group x_2 only. Therefore, if a candidate's first move is to increment group x_3's tax liability by ε, this candidate is committing to taxing *not only* group x_3, *but also* either group x_1 or x_2. By contrast, incrementing group x_2 tax liability entails no commitment as to whether other groups will be taxed. Moreover, note that the size of group x_3 plus the size of either x_1 or x_2 exceeds the size of group x_2.

Theorem 1 tells us that, in equilibrium, the candidates start by incrementing taxes for the smallest group, in spite of the fact that they could choose actions that entail less commitment. To see that this strategy cannot be beaten by a strategy whereby the opponent taxes group x_2 only, suppose that candidate A chooses to increment group x_2's taxes by ε in each round, while candidate B starts levying taxes on group x_3. After the first round, B can imitate A's move in the previous round and keep doing this until a point is reached in which A has collected all revenue, say $k\varepsilon$, from group x_2, while B has committed to levying ε on group x_3 and $(k-1)\varepsilon$ on group x_2. At this point, if ε is sufficiently small, B can collect what is left to be collected from group x_3, thereby ensuring that the members of group x_2 will vote for B. Given A's position, this gives electoral victory to B.

Rather than delaying commitment, the candidates tend to make ambiguous announcements in order to minimize the amount of tax that is precisely assigned to particular voter groups (or, in other words, in order to maximize the amount of tax whose incidence is left unspecified). Intuitively, obscuring future moves gives flexibility to shape policy through time.[26]

We conclude this section with a discussion of the effects of population grouping on equilibrium outcomes.

Even though our model can handle partitions of the population that depend not only on income but also on an arbitrary set of attributes other than income, in the sequel we shall confine our attention to partitions based solely on income, since these partitions correspond to a stylized case in which our theory relates to Director's law.[27]

Suppose that the true (pre-tax) distribution of income is represented by a (Borel) probability measure p over some subset S of $\mathbb{R}_+$.[28] One can evaluate the robustness of equilibrium tax policies to the choice of discrete grouping by posing (and answering) the following question: are there drastic changes in equilibrium tax policies for perturbations of p with finite support that are "close" to p? (Evidently, our model can only handle probability measures with finite support.)

If p itself has finite support and S (interpreted as the universe of income levels) is finite, then any slight perturbation of p relative to standard topologies for the set of (Borel) probability measures on S will give equilibria that are close to equilibria of the game induced by p.[29]

To be a bit more precise, consider any sequence (p^n) of probability measures on S converging weakly to p, and let r be target revenue, so that each (p^n, r) can be identified with a member of $\mathcal{M}$ (looking at sequences (p^n, r^n) with $r^n \to r$ would leave our argument intact). Then the equilibrium outcomes characterized by Theorem 1 are similar along the sequence (p^n, r) for n large, and the corresponding sequence of equilibrium tax policies in $\mathcal{T}_{(p^n,r)}$ converges (when $\mathcal{T}_{(p^n,r)}$ is endowed, for example, with the sup metric) to an equilibrium tax policy for the game induced by (p, r).

[26] The reader may wonder what would happen in a different game whereby the candidates promise upper bounds to the voter groups. That is, suppose that, starting from an initial situation where everybody is taxed to the fullest extent, the candidates reduce taxes decrementally up to the point where the required amount of revenue is just barely collected. It can be shown, at least in the context of an example, that Theorem 1 survives if one changes the rules of the game according to this story.

[27] Note however that more generally our theory would be consistent with the tax treatment of mortgage interest paid by homeowners, as long as the group of beneficiaries of the said policy is sufficiently numerous. In this case the partition of the population would need to accommodate attributes other than income.

[28] Here p can be thought of as the "true" income distribution as (commonly) perceived by the two candidates.

[29] One could use, for instance, strong topologies such as the one induced by the total variation metric ρ^s (where $\rho^s(\mu, \nu) := \sup\{|\mu(B) - \nu(B)| : B \text{ measurable}\}$), or the weak* topology.

To illustrate, suppose that the initial income distribution is uniform, say with three income groups, and consider a slight perturbation of the uniform distribution resulting in one group's size being slightly above the other two sizes. For the perturbation, our results predict an outcome in which the smaller groups are expropriated, and the larger group is residually taxed to meet the revenue (Theorem 1 says that there is at least one subgame-perfect equilibrium with these features, and Theorem 2 says that for slightly perturbed versions of $G_{(d,r)}$ any Nash equilibrium has these features). This equilibrium outcome is still an equilibrium outcome for the game in which the distribution is exactly uniform, but in this case other equilibria emerge.[30,31]

So far we have assumed that p has finite support. What if one insists upon modeling the true distribution as one having non-finite support? For concreteness, let p be represented by a probability density function $f : [0, \overline{X}] \to \mathbb{R}$ with $f(x) > 0$ for all x. In this case the previous robustness exercise entails evaluating equilibrium outcomes induced by a sequence of discrete probability mass functions converging (in some sense) to f. Here the weak* topology will not give the stability of equilibrium outcomes obtained above.[32] In fact, even when f is, say, log normal, the said topology is weak enough to permit approximations of f (or p) via a sequence of measures with finite support for which the tails (of the corresponding probability mass functions) contain groups of larger size than those located around the mode of f. However, such approximations to f will not generally respect the moments of f. Furthermore, there are alternative discrete versions of a continuous distribution that approximate the moments of f more accurately (such approximations rely on on the Gaussian quadrature method of numerical integration; see, e.g., Miller

[30] If the income distribution is uniform, Theorem 1 still applies, but its conclusion loses precision, since in this case the set $N_\eta(\mathcal{E}_{(d,r)})$ is "large." But distributions for which several income groups have exactly the same size (such as the uniform distribution) are nongeneric. Our results give sharp predictions for a large family of generic income distributions (more precisely, for a subset of models $(d, r) \in \mathcal{M}$ that is open and dense in $\mathcal{M}$ (Theorem 2)). By contrast, lack of precision appears to be a feature of the equilibrium set in the standard static model: the results in [11] suggest that the set of mixed equilibria is large even in generic games.

[31] If p has finite support but one insists on viewing p as a member of the class of (Borel) probability measures over $[0, \overline{X}]$ (where $\overline{X}$ is a large positive real), the above conclusion remains intact if one considers the topology induced by ρ^s (footnote 29) or other strong topologies (such as those induced by the Radon metric or the Wasserstein distances). For reasons explained below, in this case the weak* topology is too weak to sustain the sought robustness.

[32] On the other hand, under the topology induced by ρ^s (footnote 29) (and under other strong topologies) it is not possible to approach a nonatomic probability measure that assigns positive mass to every open set by a probability measure with finite support.

and Rice [31]). If one considers discrete approximations to f such that the moments of the discrete distributions converge to those of f, then robustness of equilibrium outcomes is restored.

3.3. On the analysis of other equilibria

Theorem 1 states that the game $G_{(d,r)}$ has a subgame perfect equilibrium whose corresponding tax policy lies, approximately, in $\mathcal{E}_{(d,r)}$. However, the theorem does not say anything about the features of other Nash equilibria in $G_{(d,r)}$. We can provide a complete description of the Nash set for perturbed versions of $G_{(d,r)}$. We show that there are perturbations of $G_{(d,r)}$ such that at *each* component of the Nash set, *any* equilibrium tax policy lies, approximately, in $\mathcal{E}_{(d,r)}$.

We think of $G_{(d,r)}$ as a member of an enriched class of games where the players may not have perfect information about the order of moves, and where the second mover may receive distorted information about the first mover's action. Consider the following extension of the game analyzed in the previous section. In each round, the players do not observe nature's choice of the order of moves. Rather, they observe signals that contain information on the order of moves. Moreover, the first mover's actions are only indirectly observable by the second mover through a (not necessarily perfect) signal.

Formally, let h be any history of announcements in $G_{(d,r)}$. In the round that follows h, nature determines who will be the first mover (in that round) and then sends a private message to each player. Each message is an element of $\{0, 1\}$. If 0 is observed by player i, i's signal is interpreted as saying that i is the first mover in the round that follows h. If player i is chosen as the first mover, nature sends a message $(m_A, m_B) \in \{0, 1\}^2$ with probability $\chi^{(h,i)}_{(d,r)}(m_A, m_B)$, where m_A is the message sent to A and m_B is the message sent to B. Each m_i is private information of player i. If player i is chosen to be the first mover and nature sends the message $m = (m_A, m_B)$, the sequence of actions occurs as follows:

- If $m_i = 0$, then i chooses an announcement. If $m_{-i} = 0$, where $-i \neq i$, then i's move is followed by $-i$'s move. Before making a choice, the second mover does not have any information on the first mover's action. If, on the other hand, $m_{-i} = 1$, then i's move is followed by nature's choice of a message to $-i$. The content of the message is an announcement feasible for i in the round that follows h. This message signals i's move (which is not directly observed by $-i$) and need not be completely accurate. If i chooses g, the message received by $-i$ is f with probability

$\vartheta^{-i}_{(d,r)}(h,i,m,g)(f)$ (thus, $\vartheta^{-i}_{(d,r)}(h,i,m,g)$ is a probability measure on the set of all the announcements that are feasible for i in the round that follows h). After receiving the signal, $-i$ chooses an announcement.

- If $m_i = 1$ (i.e., if i receives information indicating that i is the second mover, even though i is the actual first mover), then nature sends a message to i. The content of the message is an announcement feasible for $-i$ ($-i \neq i$) in the round that follows h. The message received by i is f with probability $\vartheta^{i}_{(d,r)}(h,i,m)(f)$. After receiving the message, i chooses an announcement. To describe what happens next, we distinguish two cases. If $m_{-i} = 0$, then i's move is followed by $-i$'s choice of an announcement. In this case, $-i$ takes action without receiving any signal on i's move. If $m_{-i} = 1$, then i's move is followed by nature's choice of a message to $-i$. The content of the message is an announcement feasible for i in the round that follows h. This message signals i's previous move (which is not directly observed by $-i$). If i chooses g, then the content of the message is f with probability $\vartheta^{-i}_{(d,r)}(h,i,m,g)(f)$. After receiving the signal, $-i$ chooses an announcement.

If player i receives message 0 and $\psi^{(h,i)}_{(d,r)}$ is a probability measure on $\{0,1\}$ representing i's prior beliefs on the identity of the first mover at the beginning of the round that follows h, then, by Bayes' rule, i believes that she is the first mover with probability

$$\frac{\psi^{(h,i)}_{(d,r)}(0)\left(\chi^{(h,i)}_{(d,r)}(0,0)+\chi^{(h,i)}_{(d,r)}(0,1)\right)}{\psi^{(h,i)}_{(d,r)}(0)\left(\chi^{(h,i)}_{(d,r)}(0,0)+\chi^{(h,i)}_{(d,r)}(0,1)\right)+\psi^{(h,i)}_{(d,r)}(1)\left(\chi^{(h,-i)}_{(d,r)}(0,0)+\chi^{(h,-i)}_{(d,r)}(0,1)\right)},$$

where $-i \neq i$. Posterior beliefs about the first mover's action given the messages received by a player can be formed similarly.[33]

At the beginning of each round, any uncertainty concerning previous rounds vanishes, and the players observe the choices made by their opponent (as well as their own) in the preceding round.

In order to emphasize the dependence of the new game on the signals sent by nature in each round, we shall designate our game by $G_{(d,r)}(\boldsymbol{\lambda}_{(d,r)})$, where

$$\begin{aligned}\boldsymbol{\lambda}_{(d,r)} &= \left(\boldsymbol{\chi}_{(d,r)}, \boldsymbol{\vartheta}_{(d,r)}\right) \\ &= \left(\left\{\chi^{(h,i)}_{(d,r)}\right\}_{(h,i)}, \left\{\vartheta^{i}_{(d,r)}(h,i,m)\right\}_{(h,i,m)}, \left\{\vartheta^{-i}_{(d,r)}(h,i,m,g)\right\}_{\substack{(h,i,m,g)\\ i\neq -i}}\right).\end{aligned}$$

[33] The details are cumbersome. We omit the exact derivation in the general case, which is not needed for our purposes.

We shall often omit the second subscript and simply write $G_{(d,r)}(\boldsymbol{\lambda})$. Note that $G_{(d,r)}(\boldsymbol{\lambda})$ is now an extensive zero-sum game with imperfect information. We study Nash equilibria of this game.

Let $\boldsymbol{\Lambda}_{(d,r)}$ be the set of all maps like $\boldsymbol{\lambda}_{(d,r)}$. Each member $\boldsymbol{\lambda}$ of $\boldsymbol{\Lambda}_{(d,r)}$ describes the signals received by the players in each round of $G_{(d,r)}(\boldsymbol{\lambda})$. A member $\boldsymbol{\lambda}$ of $\boldsymbol{\Lambda}_{(d,r)}$ is ***perfect*** if the signals are completely accurate. Formally, $\boldsymbol{\Lambda}_{(d,r)}$ is perfect if the following is satisfied:

- For each (h, i) and every (m_A, m_B),

$$\chi_{(d,r)}^{(h,i)}(m_A, m_B) = \begin{cases} 1 & \text{if } (m_A, m_B) = (0,1) \text{ and } i = A, \\ 1 & \text{if } (m_A, m_B) = (1,0) \text{ and } i = B, \\ 0 & \text{elsewhere.} \end{cases}$$

- For each (h, i, m, g), $\vartheta_{(d,r)}^{-i}(h, i, m, g)$ assigns full support to g, where $-i \neq i$.

Figure 3 illustrates the extensive form of $G_{(d,r)}(\boldsymbol{\lambda})$, in a given round, when $\boldsymbol{\lambda} = (\boldsymbol{\chi}, \boldsymbol{\vartheta})$ satisfies

$$\operatorname{supp}\left(\chi^{(h,i)}\right) \in \begin{cases} \{(0,0),(0,1)\} & \text{if } i = A, \\ \{(0,0),(1,0)\} & \text{if } i = B, \end{cases}$$

for each (h, i), and $(\boldsymbol{\chi}, \boldsymbol{\vartheta})$ is otherwise identical to a perfect signal.[34] At the beginning of the round, nature (N in Figure 3) chooses the first mover (either A or B). Then nature sends signals about the order of moves to the players. If it chooses $(m_A, m_B) = (0, 0)$, each player is told that she is the first mover. If $(0, 1)$ is chosen, then player A is told that she is the first mover and player B is told that she is the second mover, and so on. If each m_i is 0, both players choose an announcement without receiving any information about the first mover's action. If $(m_A, m_B) = (0, 1)$, then A moves first and nature sends an accurate signal to B about A's action, and so on. In Figure 3, decision nodes belonging to the same information set are linked using dashed lines.

Note that if $\boldsymbol{\lambda}$ is perfect, the game $G_{(d,r)}(\boldsymbol{\lambda})$ coincides with $G_{(d,r)}$. We shall consider versions $G_{(d,r)}(\boldsymbol{\lambda})$ of $G_{(d,r)}$ that are close to $G_{(d,r)}$. To make this statement precise, we need the following notation.

[34] This case will be particularly relevant for our purposes; it describes a situation where the actual first mover always receives an accurate signal, but the second mover may sometimes receive a signal indicating that she is moving first. In all other instances, $\boldsymbol{\lambda}$ is like a perfect signal.

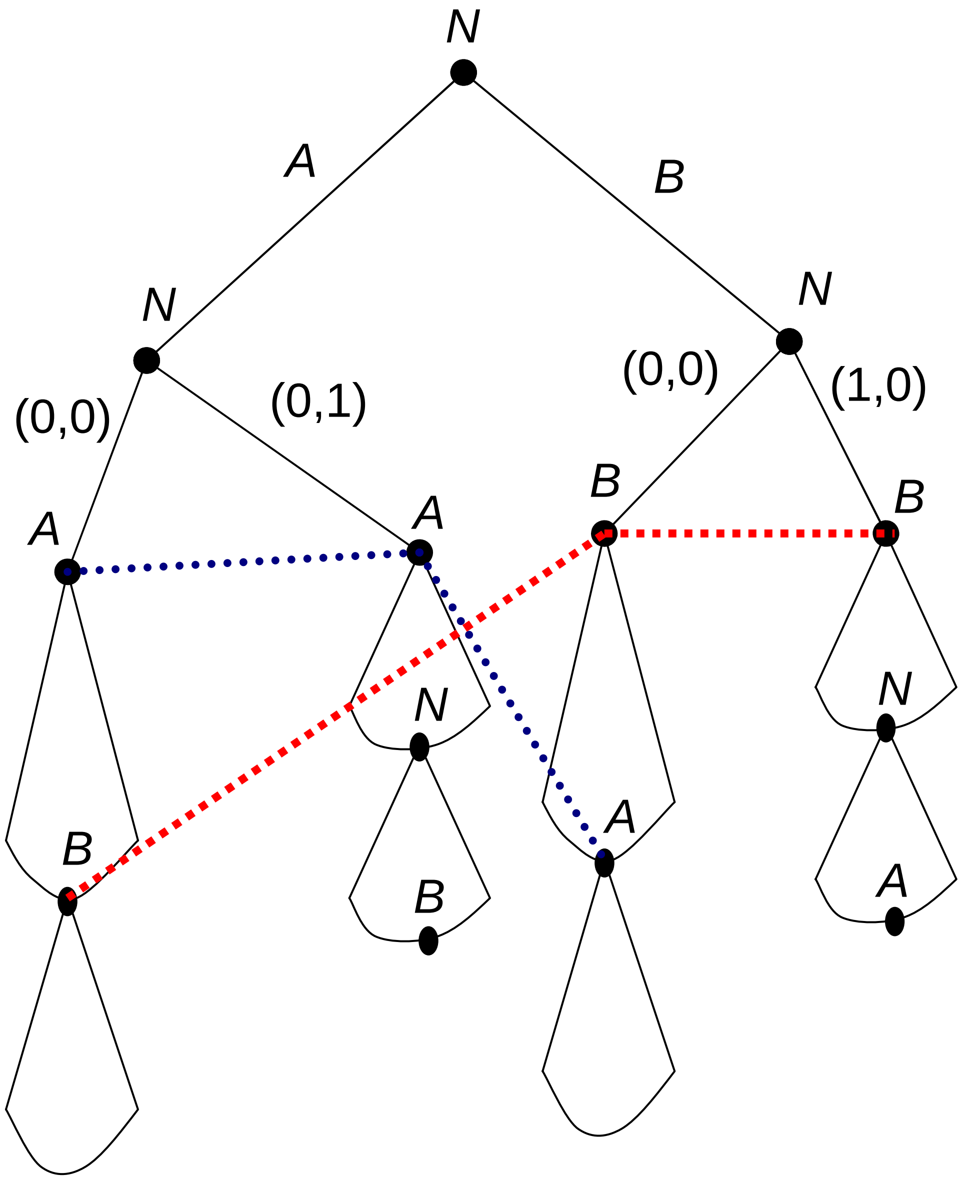

Figure 3

We view $\mathcal{D}$ as a metric space with metric $\varrho_{\mathcal{D}}$, where $\varrho_{\mathcal{D}}$ is defined as follows. Let ϱ_0 be the discrete metric on $\mathfrak{A}$ (i.e., $\varrho_0 : \mathfrak{A}^2 \to \mathbb{R}$, $\varrho_0(a,b) := 1$ if $a \neq b$, $\varrho_0(a,b) := 0$ if $a = b$). Define $\varrho : (\mathbb{R} \times \mathfrak{A} \times [0,1])^2 \to \mathbb{R}$ by

$$\varrho\left((x,a,d),(y,b,\delta)\right) := \max\left\{|x-y|, \varrho_0(a,b), |d-\delta|\right\}.$$

Given $d \in \mathcal{D}$, we define the graph of d as

$$\mathrm{gr}(d) := \left\{(x,a,d(x,a)) : (x,a) \in \{d > 0\}\right\}.$$

Observe that each $\mathrm{gr}(d)$ is a subset of $\mathbb{R} \times \mathfrak{A} \times [0,1]$. Identify the members of $\mathcal{D}$ with their graphs and define the distance between distributions in $\mathcal{D}$ to be the Hausdorff distance between their graphs in $\mathbb{R} \times \mathfrak{A} \times [0,1]$, where $\mathbb{R} \times \mathfrak{A} \times [0,1]$ is viewed as a metric space with associated metric ϱ. That is, given $d, \delta \in \mathcal{D}$, let

$$\varrho_{\mathcal{D}}(d,\delta) := \mathfrak{h}\left(\mathrm{gr}(d), \mathrm{gr}(\delta)\right),$$

where $\mathfrak{h}\left(\mathrm{gr}(d), \mathrm{gr}(\delta)\right)$ stands for the Hausdorff distance between $\mathrm{gr}(d)$ and $\mathrm{gr}(\delta)$ as induced by ϱ.[35] Take $\mathcal{M}$ as a metric space with associated metric

$$\varrho_{\mathcal{M}}\left((d,r),(\delta,\mathfrak{r})\right) := \max\left\{\varrho_{\mathcal{D}}(d,\delta), |r-\mathfrak{r}|\right\}.$$

View any subset $\mathcal{M}'$ of $\mathcal{M}$ as a metric subspace of $\mathcal{M}$ (with its relative topology). If a statement is true for all members of an open and dense subset of $\mathcal{M}'$, we say that it is ***generically true in*** $\mathcal{M}'$.

Let $\varrho_{\boldsymbol{\Lambda}_{(d,r)}}$ be a metric on $\boldsymbol{\Lambda}_{(d,r)}$ such that

$$\begin{aligned}\varrho_{\boldsymbol{\Lambda}_{(d,r)}}\left((\boldsymbol{\chi},\boldsymbol{\vartheta}),(\boldsymbol{\psi},\boldsymbol{\theta})\right) = \max\Bigg\{ & \max_{(h,i,m)}\left|\chi^{(h,i)}_{(d,r)}(m) - \psi^{(h,i)}_{(d,r)}(m)\right|, \\ & \max_{(h,i,m,f)}\left|\vartheta^{i}_{(d,r)}(h,i,m)(f) - \theta^{i}_{(d,r)}(h,i,m)(f)\right|, \\ & \max_{\substack{(h,i,m,g,f)\\ i\neq -i}}\left|\vartheta^{-i}_{(d,r)}(h,i,m,g)(f) - \theta^{-i}_{(d,r)}(h,i,m,g)(f)\right|\Bigg\}.\end{aligned}$$

Let $\mathfrak{G}_{(d,r)} := \left\{G_{(d,r)}(\boldsymbol{\lambda}) : \boldsymbol{\lambda} \in \boldsymbol{\Lambda}_{(d,r)}\right\}$ be a metric space with associated metric $\varrho_{\mathfrak{G}_{(d,r)}}$, where

$$\varrho_{\mathfrak{G}_{(d,r)}}\left(G_{(d,r)}(\boldsymbol{\lambda}), G_{(d,r)}(\boldsymbol{\lambda}')\right) := \varrho_{\boldsymbol{\Lambda}_{(d,r)}}(\boldsymbol{\lambda},\boldsymbol{\lambda}').$$

[35] One could also require that the distance between the money unit associated to each distribution be factored in. The following (more stringent) alternative definition of $\varrho_{\mathcal{D}}$ could also be adopted without altering our results: $\varrho_{\mathcal{D}}(d,\delta) := \max\left\{\mathfrak{h}(\mathrm{gr}(d),\mathrm{gr}(\delta)), |\varepsilon(d)-\varepsilon(\delta)|\right\}$.

By a ***neighborhood*** of $G_{(d,r)}(\boldsymbol{\lambda})$ we mean an open subset of $\mathfrak{G}_{(d,r)}$ containing $G_{(d,r)}(\boldsymbol{\lambda})$. We may think of the game $G_{(d,r)}(\boldsymbol{\lambda})$ as a perturbed version of $G_{(d,r)}$ (recall that $G_{(d,r)} = G_{(d,r)}(\boldsymbol{\theta})$ if $\boldsymbol{\theta}$ is perfect) if $G_{(d,r)}(\boldsymbol{\lambda})$ lies within a small neighborhood of $G_{(d,r)}$.

Note that, for each member of $\widetilde{\mathcal{E}}_{(d,r)}$ not in $\mathcal{E}_{(d,r)}$, there exists another tax policy that meets the revenue constraint and makes some group of voters strictly better off without making any other voter group worse off. Some of these policies cannot, in principle, be ruled out as equilibrium outcomes of $G_{(d,r)}(\boldsymbol{\lambda})$. In fact, in equilibrium, the last move of a player could be optimal and, at the same time, overtax some voter group (in the sense that one could reduce the tax burden and still meet the revenue requirement). This could happen if suppressing the tax excess did not switch the group's support from one candidate to the other.[36] To avoid situations where some voter group is soaked excessively (in the sense that some candidate keeps taxing them even when the revenue requirement has been met), we assume that the candidates are forced to choose final announcements in

$$\mathcal{O}_{(d,r)} := \left\{t \in \mathcal{T}_{(d,r)} : \text{there is no } \tau \in \mathcal{T}_{(d,r)} \text{ with } \tau \lneqq t\right\}.^{37}$$

The statement below refers to the version of $G_{(d,r)}(\boldsymbol{\lambda})$ in which the candidates' final announcements must be elements of $\mathcal{O}_{(d,r)}$. We slightly abuse notation and denote this game again by $G_{(d,r)}(\boldsymbol{\lambda})$. It is worth noting that $G_{(d,r)}(\boldsymbol{\lambda})$ has a subgame perfect equilibrium. In fact, this game has finite horizon, and each subgame has a Nash equilibrium (e.g., Mamer and Schilling [27]).

Theorem 2 says that (for a sufficiently small money unit) the components of the Nash set in perturbations of $G_{(d,r)}$ have the same features as the equilibrium of Theorem 1, in the sense that, for all these components, any equilibrium tax policy lies, approximately, in $\mathcal{E}_{(d,r)}$.

For $+\infty > \epsilon > 0$ and $1 > \theta > 0$, define the set

$$\mathcal{M}(\epsilon, \theta) := \left\{(d, r) \in \mathcal{M} : \varepsilon(d) < \epsilon, \tfrac{r}{\sum_{(x,a)} d(x,a)x} < \theta\right\}.$$

This set contains all the pairs (d, r) for which the money unit $\varepsilon(d)$ is smaller than ϵ and the revenue requirement as a fraction of total income lies below θ.

Theorem 2. *Suppose that $+\infty > \eta > 0$ and $1 > \theta > 0$. There exists $+\infty > \varepsilon_{(\eta,\theta)} > 0$ such that, generically in $\mathcal{M}\left(\varepsilon_{(\eta,\theta)}, \theta\right)$, any neighborhood of $G_{(d,r)}$ contains a game $G_{(d,r)}(\boldsymbol{\lambda})$ such that the tax policy implemented at any Nash equilibrium of $G_{(d,r)}(\boldsymbol{\lambda})$ lies in $N_\eta\left(\mathcal{E}_{(d,r)}\right)$.*

The proof is relegated to Section 5.

[36] Observe that this feature of some equilibria is particularly stark in a model with exogenous labor supply.

[37] Without this assumption, Theorem 2 is also valid with $\widetilde{\mathcal{E}}_{(d,r)}$ replacing $\mathcal{E}_{(d,r)}$.

4. Concluding remarks

We have modeled electoral competition between two candidates as a dynamical process in which each candidate gradually commits to a tax schedule that is to be implemented if she is elected by a majority of voters. We have characterized the Nash set of the associated game. At each component of this set, equilibrium tax schedules take a very particular form, which is reminiscent of Director's law of public income redistribution (Stigler [39]).

This paper has proposed a particular extensive form as a representation of the dynamical campaigning process, but other rules of the game could be envisaged. For example, the candidates could announce an arbitrary subset of the set of all tax policies in the first round, and then refine this subset in subsequent rounds. Another possibility would be for the players to announce intervals of possible tax liabilities for each voter group, each interval being a refinement of intervals announced in previous rounds. Since our analysis does not readily extend to these alternative formulations, the study of these variations is left for future research.

Finally, extending the analysis to the case of endogenous labor supply is desirable for at least two reasons. First, in a model à la Mirrlees [32], the limit on the extent to which a voter group may be soaked would be endogenous, and would depend on elasticities. Second, the predictions of the current model, augmented to allow for distortionary taxation, would give new insight on the interplay of the candidates' incentives to favor certain voter groups and the distortion of labor supply embedded in income redistribution.[38]

5. Proofs

The proof of Theorem 1 is omitted and available from the author upon request. The argument for this proof was illustrated in Section 3.1. Moreover, Theorem 1 can be viewed as a special case of Theorem 2.[39]

[38]If one could interpret equilibrium tax schedules as optimal ones, for some (endogenously determined) social welfare weights, then the machinery developed within the framework of optimal income taxation could be used to understand the features of equilibrium outcomes.

[39]Note however that Theorem 2 refers to a perturbation of $G_{(d,r)}$, while Theorem 1 is about $G_{(d,r)}$.

5.1. Theorem 2: preliminaries

In the interest of brevity, we shall confine attention here to a variant of $G_{(d,r)}$ in which the players are constrained to incrementing taxes by exactly $\varepsilon(d)$ on exactly one income group in each round (this is the game $G^*_{(d,r)} = G^*_{(d,r)}(\boldsymbol{\lambda})$, where $\boldsymbol{\lambda}$ is perfect, defined below). This game is easier to handle, and permits a shorter proof. Given our discussion of the second-mover advantage in Section 2.2, in principle it is intuitive that forcing the players to disclose minimal information in each round does not really affect the optimality of their strategies in the full-fledged version $G_{(d,r)}$. The variant of Theorem 2 proven here is stated at the end of this subsection. For a more detailed treatment of the general game $G_{(d,r)}$, the reader is referred to Carbonell-Nicolau [12].

Define $\mathfrak{F}^*_{(d,r)} : \mathcal{P}_{(d,r)} \rightrightarrows \mathcal{P}_{(d,r)}$ by

$$\mathfrak{F}^*_{(d,r)}(f) := \begin{cases} \left\{ g \in \mathcal{P}_{(d,r)} : \begin{pmatrix} g(x,a) = f(x,a) + \varepsilon(d), \\ \text{some } (x,a) \text{ with } f(x,a) < x, \\ g = f \text{ elsewhere} \end{pmatrix} \right\} & \text{if } f \text{ is not final,} \\ \{f\} & \text{if } f \text{ is final.} \end{cases}$$

Define $\mathbf{T}_{(d,r)} : \mathcal{P}_{(d,r)} \rightrightarrows \mathcal{T}_{(d,r)}$ by $\mathbf{T}_{(d,r)}(f) := \left\{ g \in \mathcal{O}_{(d,r)} : g \geq f \right\}$ and $Z_{(d,r)} : \mathcal{P}_{(d,r)} \rightrightarrows \mathfrak{A}_{\varepsilon(d)}$ by

$$Z_{(d,r)}(f) := \arg \min_{\substack{(x,a) \in \{d>0\} \\ (x - f(x,a) - \varepsilon(d)) d(x,a) \geq \max_{(y,b)} d(y,b)\varepsilon(d) \\ + \sum_{(y,b): d(y,b) < d(x,a)} (y - f(y,b)) d(y,b)}} d(x,a).$$

Given $(x,a) \in \{d > 0\}$, define $\varphi^{(x,a)}_{(d,r)} : \mathcal{P}_{(d,r)} \to \mathcal{P}_{(d,r)}$ by

$$\varphi^{(x,a)}_{(d,r)}(f)(y,b) := \begin{cases} f(y,b) + \varepsilon(d) & \text{if } (y,b) = (x,a), \\ f(y,b) & \text{otherwise.} \end{cases}$$

Let $G^*_{(d,r)}(\boldsymbol{\lambda})$ be a game exactly like $G_{(d,r)}(\boldsymbol{\lambda})$ with the following constraints on the actions available to the players: (1) at the beginning of the game, both players are forced to choose an announcement $f^A_1 := 0 =: f^B_1$; and (2) in round $k = 2, 3, ...$, each player i is forced to choose an announcement $f^i_k \in \mathfrak{F}^*_{(d,r)}(f^i_{k-1})$. When $\boldsymbol{\lambda}$ is perfect, we often denote $G^*_{(d,r)}(\boldsymbol{\lambda})$ as $G^*_{(d,r)}$. Let $\boldsymbol{\Lambda}^*_{(d,r)}$ be the analogue of $\boldsymbol{\Lambda}_{(d,r)}$ (that is, the set of all possible $\boldsymbol{\lambda}$ describing the information structure in $G^*_{(d,r)}(\boldsymbol{\lambda})$). The set $\boldsymbol{\Lambda}^*_{(d,r)}$ can be viewed as a metric space with corresponding metric $\varrho_{\mathbf{\Lambda}^*_{(d,r)}}$, where the definition of $\varrho_{\mathbf{\Lambda}^*_{(d,r)}}$ is analogous to that of $\varrho_{\mathbf{\Lambda}_{(d,r)}}$ (Subsection 3.3). One can then view $\mathfrak{G}^*_{(d,r)} := \left\{G^*_{(d,r)}(\boldsymbol{\lambda}) : \boldsymbol{\lambda} \in \mathbf{\Lambda}^*_{(d,r)}\right\}$ as a metric space with associated metric $\varrho_{\mathfrak{G}^*_{(d,r)}}$, where $\varrho_{\mathfrak{G}^*_{(d,r)}}\left(G^*_{(d,r)}(\boldsymbol{\lambda}), G^*_{(d,r)}(\boldsymbol{\lambda}')\right) := \varrho_{\mathbf{\Lambda}^*_{(d,r)}}(\boldsymbol{\lambda}, \boldsymbol{\lambda}')$. By a ***neighborhood*** of $G^*_{(d,r)}(\boldsymbol{\lambda})$ we mean an open subset of $\mathfrak{G}^*_{(d,r)}$ containing $G^*_{(d,r)}(\boldsymbol{\lambda})$.

By the ***subgame of*** $G^*_{(d,r)}(\boldsymbol{\lambda})$ ***induced by*** h we mean the subgame of $G^*_{(d,r)}(\boldsymbol{\lambda})$ that starts immediately after the history of announcements h in $G^*_{(d,r)}(\boldsymbol{\lambda})$, before nature chooses the order of moves in the round that follows h. We denote this subgame by $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$, and represent its value as $v^{*h}_{(d,r)}(\boldsymbol{\lambda}) = \left(v^{*(h,A)}_{(d,r)}(\boldsymbol{\lambda}), v^{*(h,B)}_{(d,r)}(\boldsymbol{\lambda})\right)$.[40]

A ***perfect*** $\boldsymbol{\lambda}$ in $\mathbf{\Lambda}^*_{(d,r)}$ is the analogue of a perfect signal in $\mathbf{\Lambda}_{(d,r)}$ (see Section 3.3). A member $\boldsymbol{\lambda} = (\boldsymbol{\chi}, \boldsymbol{\vartheta})$ of $\mathbf{\Lambda}^*_{(d,r)}$ is ***symmetric*** in $\mathbf{\Lambda}^*_{(d,r)}$ if the following holds:

- For each h whose last two announcements are identical, the following holds: $\chi^{(h,A)} = \chi^{(h,B)}$ for all h; $\vartheta^A(h, A, m) = \vartheta^B(h, B, m)$ for all m; and $\vartheta^A(h, B, m, g) = \vartheta^B(h, A, m, g)$ for all (m, g).

- For each h whose last two announcements are identical and each super-history $(h, f_1, g_1, ..., f_k, g_k)$ of h in $G^*_{(d,r)}(\boldsymbol{\lambda})$, the following holds:

  - $\chi^{((h,f_1,g_1,...,f_k,g_k),A)} = \chi^{((h,g_1,f_1,...,g_k,f_k),B)}$.
  - $\vartheta^A((h, f_1, g_1, ..., f_k, g_k), A, m) = \vartheta^B((h, g_1, f_1, ..., g_k, f_k), B, m)$ for all m.
  - $\vartheta^A((h, f_1, g_1, ..., f_k, g_k), B, m, g) = \vartheta^B((h, g_1, f_1, ..., g_k, f_k), A, m, g)$ for all (m, g).

A ***symmetric*** $\boldsymbol{\lambda}$ in $\mathbf{\Lambda}_{(d,r)}$ is defined analogously.

Let $\mathfrak{H}_{(d,r)}$ be a map that assigns to each history of announcements $h = (f_1, g_1, ..., f_k, g_k)$ in $G^*_{(d,r)}$ the set $\mathfrak{H}_{(d,r)}(h)$ of all histories of announcements $(t_1, \tau_1, ..., t_k, \tau_k)$ in $G^*_{(d,r)}$ such that for some κ for which $f_\kappa = g_\kappa$, $(t_l, \tau_l) = (f_l, g_l)$ for $l = 1, ..., \kappa$, and (if $\kappa < k$) $(t_{\kappa+1}, \tau_{\kappa+1}, ..., t_k, \tau_k) = (g_{\kappa+1}, f_{\kappa+1}, ..., g_k, f_k)$.

Let $\mathcal{M}^\circ$ be the set of all $(d, r) \in \mathcal{M}$ for which $\sum_{(x,a)\in X} d(x, a) \neq \sum_{(x,a)\in Y} d(x, a)$ whenever $\{d > 0\} \supseteq X \neq Y \subseteq \{d > 0\}$.

[40] Each $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ is a special case of the zero-sum game of incomplete information studied by Mamer and Schilling [27]. By Sion [37], this game has a value. Similar statements are true for $\Gamma^h_{(d,r)}(\boldsymbol{\lambda})$ (to be defined next).

Let $\mathcal{T}^{\circ}_{(d,r)}$ be the set of all $f \in \mathcal{O}_{(d,r)}$ such that there exist $f_1, ..., f_k$ and $(x_2, a_2), ..., (x_k, a_k)$ satisfying the following three conditions: (1) $f_1 = 0$ and $f_k = f$; (2) for all $\kappa = 2, ..., k$, $(x_\kappa, a_\kappa) \in Z_{(d,r)}(f_{\kappa-1})$ if $Z_{(d,r)}(f_{\kappa-1}) \neq \emptyset$; and (3) for all $\kappa = 2, ..., k$,

$$f_\kappa = \begin{cases} \varphi^{(x_\kappa,a_\kappa)}_{(d,r)}(f_{\kappa-1}) & \text{if } f_{\kappa-1} \text{ is not final,} \\ f_{\kappa-1} & \text{if } f_{\kappa-1} \text{ is final.} \end{cases}$$

(If $r = 0$ then $\mathcal{T}^{\circ}_{(d,r)}$ is simply $\{0\}$.) Let $\mathcal{E}^{*}_{(d,r)}$ be the set of all $t \in \mathcal{O}_{(d,r)}$ satisfying $t \leq g_\tau$ for some $\tau \in \mathcal{T}^{\circ}_{(d,r)}$, where

$$g_\tau(x,a) := \begin{cases} x & \text{if } d(x,a) < \max\limits_{(y,b):\tau(y,b)>0} d(y,b), \\ \tau(x,a) & \text{if } d(x,a) = \max\limits_{(y,b):\tau(y,b)>0} d(y,b). \end{cases}$$

For $+\infty > \eta > 0$, let $\mathcal{M}^{*}_{\eta}$ be the set of all $(d,r) \in \mathcal{M}$ satisfying the following:

- For any non-final $f \in \mathcal{P}_{(d,r)}$ for which $f \leq g$, some $g \in \mathcal{O}_{(d,r)}$, $Z_{(d,r)}(f) \neq \emptyset$.
- For all $f \in \mathcal{T}^{\circ}_{(d,r)}$ and all (x,a) with $f(x,a) > 0$,

$$\begin{aligned} \eta - \xi_{(x,a)}\varepsilon(d) \geq & \sum_{\substack{(y,b): \\ d(y,b)<d(x,a)}} \tfrac{20}{d(y,b)} \\ & \times \left(\max_{(z,c)} d(z,c)\varepsilon(d) + \sum_{\substack{(z,c): \\ d(z,c)<d(x,a)}} (z - f(z,c))d(z,c) \right), \end{aligned} \tag{2}$$

  where $\xi_{(x,a)} := 2 + \frac{1}{d(x,a)} \sum_{(y,b):d(y,b)\leq d(x,a)} d(y,b)$.

Let $\boldsymbol{\Lambda}^{\bullet}_{(d,r)}$ be the set of all symmetric $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^{*}_{(d,r)}$ satisfying the following property: Suppose that $h = (f^A_1, f^B_1, ..., f^A_k, f^B_k)$ is a history of announcements in $G^{*}_{(d,r)}(\boldsymbol{\lambda})$ such that $f^A_k \neq f^B_k$. Suppose that there is a player i such that for any announcement f^i that is feasible for i in the round that follows h (i.e., any member of $\mathfrak{F}^{*}_{(d,r)}(f^i_k)$), there exists an announcement f^j that the opponent j may choose (in the same round) in $\mathfrak{F}^{*}_{(d,r)}(f^j_k)$ such that i's value in $\Gamma^{*(h,f^A,f^B)}_{(d,r)}(\boldsymbol{\lambda})$ is negative. Then the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero.

Let $\mathbf{\Lambda}^{\circ}_{(d,r)}$ be the set of all $(\boldsymbol{\chi}, \boldsymbol{\vartheta})$ in $\mathbf{\Lambda}^{*}_{(d,r)}$ such that the following is satisfied: for each (h,i), $\operatorname{supp}\left(\chi^{(h,i)}\right) \neq \{(0,0)\}$,

$$\operatorname{supp}\left(\chi^{(h,i)}\right) \in \begin{cases} \{(0,0),(0,1)\} & \text{if } i = A, \\ \{(0,0),(1,0)\} & \text{if } i = B, \end{cases}$$

and $(\boldsymbol{\chi}, \boldsymbol{\vartheta})$ is otherwise identical to a perfect signal in $\mathbf{\Lambda}^{*}_{(d,r)}$. Each member of $\mathbf{\Lambda}^{\circ}_{(d,r)}$ has the property that the actual first mover always receives an accurate signal, but the second mover may sometimes receive a signal indicating that she is moving first. In all other instances, the members of $\mathbf{\Lambda}^{\circ}_{(d,r)}$ are like perfect signals.

The next section is devoted to proving the following variant of Theorem 2:

Theorem 2'. *Suppose that $+\infty > \eta > 0$ and $1 > \theta > 0$. There exists $+\infty > \varepsilon_{(\eta,\theta)} > 0$ such that, generically in $\mathcal{M}\left(\varepsilon_{(\eta,\theta)}, \theta\right)$, any neighborhood of $G^{*}_{(d,r)}$ contains a game $G^{*}_{(d,r)}(\boldsymbol{\lambda})$ such that the tax policy implemented at any Nash equilibrium of $G^{*}_{(d,r)}(\boldsymbol{\lambda})$ lies in $N_{\eta}\left(\mathcal{E}_{(d,r)}\right)$.*

5.2. Proof of Theorem 2'

Lemma 1. *Suppose that $+\infty > \eta > 0$ and $(d,r) \in \mathcal{M}^{*}_{\eta} \cap \mathcal{M}^{\circ}$. Let $\boldsymbol{\theta}$ be a perfect signal in $\mathbf{\Lambda}^{*}_{(d,r)}$. Fix any open subset $\mathfrak{O}$ of $\mathbf{\Lambda}^{*}_{(d,r)}$ containing $\boldsymbol{\theta}$. Then $\mathfrak{O} \cap \mathbf{\Lambda}^{\bullet}_{(d,r)} \cap \mathbf{\Lambda}^{\circ}_{(d,r)} \neq \emptyset$.*

Proof. Suppose that $+\infty > \eta > 0$ and $(d,r) \in \mathcal{M}^{*}_{\eta} \cap \mathcal{M}^{\circ}$. Let $\boldsymbol{\lambda}_0 = (\boldsymbol{\chi}_0, \boldsymbol{\vartheta}_0)$ be a perfect signal in $\mathbf{\Lambda}^{*}_{(d,r)}$. If $\boldsymbol{\lambda}_0 \in \mathbf{\Lambda}^{\bullet}_{(d,r)} \cap \mathbf{\Lambda}^{\circ}_{(d,r)}$, there is nothing to prove, so suppose that $\boldsymbol{\lambda}_0 \notin \mathbf{\Lambda}^{\bullet}_{(d,r)} \cap \mathbf{\Lambda}^{\circ}_{(d,r)}$. Since $\boldsymbol{\lambda}_0 \in \mathbf{\Lambda}^{\circ}_{(d,r)}$, we must have $\boldsymbol{\lambda}_0 \notin \mathbf{\Lambda}^{\bullet}_{(d,r)}$. Fix any open subset $\mathfrak{O}$ of $\mathbf{\Lambda}^{*}_{(d,r)}$ containing $\boldsymbol{\lambda}_0$.

Let $\mathcal{H}_1$ be the set of all histories of announcements h in $G^{*}_{(d,r)}(\boldsymbol{\lambda}_0)$ satisfying the following: (i) the last pair of announcements (g^A, g^B) in h has $g^A \neq g^B$; (ii) there is a player i such that for any announcement f^i that is feasible for i in the round that follows h (i.e., any member of $\mathfrak{F}^{*}_{(d,r)}(g^i)$), there exists an announcement f^j that the opponent j may choose (in the same round, in $\mathfrak{F}^{*}_{(d,r)}(g^j)$) such that $v^{*\left((h,f^A,f^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_0) < 0$; (iii) $v^{*h}_{(d,r)}(\boldsymbol{\lambda}_0) = 0$; and (iv) there is no super-history h' of h different from h such that (1) the last pair of announcements (t^A, t^B) in h' has $t^A \neq t^B$, (2) there is a player i such that for any announcement $\tau^i \in \mathfrak{F}^{*}_{(d,r)}(t^i)$, there exists an announcement τ^j that the opponent j may choose (in the same round, in $\mathfrak{F}^{*}_{(d,r)}(t^j)$) such that $v^{*\left((h',\tau^A,\tau^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_0) < 0$, and (3) $v^{*h'}_{(d,r)}(\boldsymbol{\lambda}_0) = 0$. Because $\boldsymbol{\lambda}_0 \notin \mathbf{\Lambda}^{\bullet}_{(d,r)}$, since $\boldsymbol{\lambda}_0$

is symmetric, we have $\mathcal{H}_1 \neq \emptyset$. Let $\boldsymbol{\lambda}_{1\epsilon} = (\boldsymbol{\chi}_{1\epsilon}, \boldsymbol{\vartheta}_{1\epsilon})$ be an element of $\boldsymbol{\Lambda}^*_{(d,r)}$ exactly like $\boldsymbol{\lambda}_0$ except for the following:

- Take $h_1 \in \mathcal{H}_1$. For some $1 > \epsilon > 0$,

$$\chi_{1\epsilon}^{(h_1,A)}(m) = \begin{cases} 1-\epsilon & \text{if } m = (0,1), \\ \epsilon & \text{if } m = (0,0), \\ 0 & \text{elsewhere,} \end{cases}$$

  and

$$\chi_{1\epsilon}^{(h_1,B)}(m) = \begin{cases} 1 & \text{if } m = (1,0), \\ 0 & \text{elsewhere.} \end{cases}$$

  Moreover, for all $h \in \mathfrak{H}_{(d,r)}(h_1)$, $\chi_{1\epsilon}^{(h_1,A)} = \chi_{1\epsilon}^{(h,B)}$ and $\chi_{1\epsilon}^{(h_1,B)} = \chi_{1\epsilon}^{(h,A)}$.

- Let $h_2 \in \mathcal{H}_1 \setminus \left(\{h_1\} \cup \mathfrak{H}_{(d,r)}(h_1)\right)$. Then, $\chi_{1\epsilon}^{(h_2,A)} = \chi_{1\epsilon}^{(h_1,A)}$ and $\chi_{1\epsilon}^{(h_2,B)} = \chi_{1\epsilon}^{(h_1,B)}$. Moreover, for all $h \in \mathfrak{H}_{(d,r)}(h_2)$, $\chi_{1\epsilon}^{(h_2,A)} = \chi_{1\epsilon}^{(h,B)}$ and $\chi_{1\epsilon}^{(h_2,B)} = \chi_{1\epsilon}^{(h,A)}$.

- Let $h_3 \in \mathcal{H}_1 \setminus \left(\cup_{l=1}^2 \left(\{h_l\} \cup \mathfrak{H}_{(d,r)}(h_l)\right)\right)$ and proceed as in the previous steps until $\chi_{1\epsilon}^{(h,i)}$ has been defined for each $h \in \mathcal{H}_1$ and each i.

Claim 1. *There exists a symmetric* $\boldsymbol{\lambda}_1 = (\boldsymbol{\chi}_1, \boldsymbol{\vartheta}_1) \in \boldsymbol{\Lambda}^\circ_{(d,r)}$ *such that* $v^{*h}_{(d,r)}(\boldsymbol{\lambda}_1) \neq 0$ *for each* $h \in \mathcal{H}_1$*, and* $\boldsymbol{\lambda}_1 \in \mathfrak{D}$.

Proof. Let $\mathcal{H}_1^*$ be the set of all $h \in \mathcal{H}_1$ for which, given ϵ, there exists $\epsilon' < \epsilon$ such that $v^{*h}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon'}) \neq 0$. Define $\mathcal{H}_1' := \mathcal{H}_1 \setminus \mathcal{H}_1^*$. Suppose that $\mathcal{H}_1' \neq \emptyset$ (if this set is empty, ignore the ensuing argument and go straight to the next paragraph). Let $\boldsymbol{\lambda}'_{1\epsilon} = (\boldsymbol{\chi}'_{1\epsilon}, \boldsymbol{\vartheta}'_{1\epsilon})$ be an element of $\boldsymbol{\Lambda}^*_{(d,r)}$ that is exactly like $\boldsymbol{\lambda}_{1\epsilon}$ except for the following:

- Take $h_1 \in \mathcal{H}_1'$. Then

$${\chi'}_{1\epsilon}^{(h_1,B)}(m) = \begin{cases} 1-\epsilon & \text{if } m = (1,0), \\ \epsilon & \text{if } m = (0,0), \\ 0 & \text{elsewhere,} \end{cases}$$

  and

$${\chi'}_{1\epsilon}^{(h_1,A)}(m) = \begin{cases} 1 & \text{if } m = (0,1), \\ 0 & \text{elsewhere.} \end{cases}$$

  Moreover, for all $h \in \mathfrak{H}_{(d,r)}(h_1)$, ${\chi'}_{1\epsilon}^{(h_1,A)} = {\chi'}_{1\epsilon}^{(h,B)}$ and ${\chi'}_{1\epsilon}^{(h_1,B)} = {\chi'}_{1\epsilon}^{(h,A)}$.

- Let $h_2 \in \mathcal{H}'_1 \setminus \left(\{h_1\} \cup \mathfrak{H}_{(d,r)}(h_1)\right)$. Then, ${\chi'}_{1\epsilon}^{(h_2,A)} = {\chi'}_{1\epsilon}^{(h_1,A)}$ and ${\chi'}_{1\epsilon}^{(h_2,B)} = {\chi'}_{1\epsilon}^{(h_1,B)}$. Moreover, for all $h \in \mathfrak{H}_{(d,r)}(h_2)$, ${\chi'}_{1\epsilon}^{(h_2,A)} = {\chi'}_{1\epsilon}^{(h,B)}$ and ${\chi'}_{1\epsilon}^{(h_2,B)} = {\chi'}_{1\epsilon}^{(h,A)}$.
- Let $h_3 \in \mathcal{H}'_1 \setminus \left(\cup_{l=1}^{2} \left(\{h_l\} \cup \mathfrak{H}_{(d,r)}(h_l)\right)\right)$ and proceed as in the previous steps until ${\chi'}_{1\epsilon}^{(h,i)}$ has been defined for each $h \in \mathcal{H}'_1$ and each i.

Clearly, each $\boldsymbol{\lambda}'_{1\epsilon}$ belongs to $\boldsymbol{\Lambda}^{\circ}_{(d,r)}$. We claim that, for ϵ sufficiently small, ${v^*}^{h}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) \neq 0$ for each $h \in \mathcal{H}'_1$. We momentarily take this statement for granted and relegate its proof to the next paragraph. If, for ϵ sufficiently small, ${v^*}^{h}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) \neq 0$ for each $h \in \mathcal{H}'_1$, then, for ϵ sufficiently small, ${v^*}^{h}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) \neq 0$ for each $h \in \mathcal{H}_1$, and $\boldsymbol{\lambda}'_{1\epsilon} \in \mathfrak{D}$. Thus, since $\boldsymbol{\lambda}'_{1\epsilon}$ is symmetric by construction, the proof of Claim 1 is completed by taking $\boldsymbol{\lambda}_1 := \boldsymbol{\lambda}'_{1\epsilon}$ (ϵ sufficiently small).

We now turn to showing that ${v^*}^{h}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) \neq 0$ for each $h \in \mathcal{H}'_1$, for ϵ sufficiently small. Fix $h \in \mathcal{H}'_1$. For ϵ sufficiently small, $\boldsymbol{\lambda}_{1\epsilon}$ is accurate enough that the players' equilibrium behavior along the equilibrium path in $G^*_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon})$, in the round that follows h, is as follows:

(a) For each i, if the content of i's first message is 0, then i's equilibrium behavior strategy α_i has support within

$$\arg\min_{f^i} \left(\max_{f^{-i}} {v^*}^{\left((h,f^A,f^B),-i\right)}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon}) \right), \quad -i \neq i,$$

at the corresponding information set of i.

(b) For each i, if the content of i's first message is 1 and the second message says 'f^{-i}' ($-i \neq i$), i's equilibrium behavior strategy has support within

$$\arg\max_{f^i} {v^*}^{\left((h,f^A,f^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon})$$

at the corresponding information set of i.

Analogous statements hold for $G^*_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon})$.

Since $h \in \mathcal{H}'_1$, $h \in \mathcal{H}_1$, and so ${v^*}^{h}_{(d,r)}(\boldsymbol{\lambda}_0) = 0$. Hence, in view of (a) and (b),

$$\begin{aligned} &{v^*}^{(h,A)}_{(d,r)}(\boldsymbol{\lambda}_0) \\ &= \tfrac{1}{2} \min_{f^B} \max_{f^A} {v^*}^{\left((h,f^A,f^B),A\right)}_{(d,r)}(\boldsymbol{\lambda}_0) + \tfrac{1}{2} \max_{f^A} \min_{f^B} {v^*}^{\left((h,f^A,f^B),A\right)}_{(d,r)}(\boldsymbol{\lambda}_0) \qquad (3) \\ &= 0. \end{aligned}$$

Moreover, since $h \in \mathcal{H}_1'$, $v^{*h}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon}) = 0$, and therefore

$$\begin{aligned} 0 = v^{*(h,A)}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon}) = \tfrac{1}{2} \min_{f^B} \max_{f^A} v^{*\left((h,f^A,f^B),A\right)}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon}) \\ + \tfrac{1}{2}\Bigg(\epsilon \sum_{(f^A,f^B)} \alpha_A(f^A)\alpha_B(f^B) v^{*\left((h,f^A,f^B),A\right)}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon}) \\ + (1-\epsilon) \max_{f^A} \min_{f^B} v^{*\left((h,f^A,f^B),A\right)}_{(d,r)}(\boldsymbol{\lambda}_{1\epsilon}) \Bigg). \end{aligned}$$

We can replace $\boldsymbol{\lambda}_{1\epsilon}$ by $\boldsymbol{\lambda}_0$ in this equation, for $\boldsymbol{\lambda}_{1\epsilon}$ and $\boldsymbol{\lambda}_0$ coincide at the beginning of the round that follows any history of announcements of the form (h, f^A, f^B) . Hence, using (3), we obtain

$$\sum_{(f^A,f^B)} \alpha_A(f^A)\alpha_B(f^B) v^{*\left((h,f^A,f^B),A\right)}_{(d,r)}(\boldsymbol{\lambda}_0) = \max_{f^A} \min_{f^B} v^{*\left((h,f^A,f^B),A\right)}_{(d,r)}(\boldsymbol{\lambda}_0).$$

Since we know that the right-hand side of this equation is negative, so is the left-hand side, and therefore

$$\sum_{(f^A,f^B)} \alpha_A(f^A)\alpha_B(f^B) v^{*\left((h,f^A,f^B),B\right)}_{(d,r)}(\boldsymbol{\lambda}_0) > 0. \tag{4}$$

Moreover,

$$\begin{aligned} v^{*(h,B)}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) = \tfrac{1}{2} \min_{f^A} \max_{f^B} v^{*\left((h,f^A,f^B),B\right)}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) \\ + \tfrac{1}{2}\Bigg(\epsilon \sum_{(f^A,f^B)} \alpha_A(f^A)\alpha_B(f^B) v^{*\left((h,f^A,f^B),B\right)}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) \\ + (1-\epsilon) \max_{f^B} \min_{f^A} v^{*\left((h,f^A,f^B),B\right)}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) \Bigg). \end{aligned}$$

As before, we can replace $\boldsymbol{\lambda}_{1\epsilon}$ by $\boldsymbol{\lambda}_0$ in this equation. Therefore, since equation (3) holds true and $\min_{f^A} \max_{f^B} v^{*\left((h,f^A,f^B),B\right)}_{(d,r)}(\boldsymbol{\lambda}_0) > 0$, it follows (in view of (4)) that $v^{*(h,B)}_{(d,r)}(\boldsymbol{\lambda}'_{1\epsilon}) > 0$, as we sought. □

Claim 2. *For each $h \in \mathcal{H}_1$, there is no super-history h' (including h) of h such that (1) the last pair of announcements (t^A, t^B) in h' has $t^A \neq t^B$, (2) there is a player i such that for any announcement $\tau^i \in \mathfrak{F}^*_{(d,r)}(t^i)$, there exists an announcement τ^j that the opponent j may choose (in the same round, in $\mathfrak{F}^*_{(d,r)}(t^j)$) such that $v^{*\left((h',\tau^A,\tau^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_1) < 0$, and (3) $v^{*h'}_{(d,r)}(\boldsymbol{\lambda}_1) = 0$.*

Proof. Take $h \in \mathcal{H}_1$ and a super-history h' of h. If $h' = h$, the statement follows immediately from Claim 1. If $h \neq h'$, the statement follows from (iv) and the fact that when $h \neq h'$ the choice of $\boldsymbol{\lambda}_1$ (in the proof of Claim 1) entails $v^{*\left((h',\tau^A,\tau^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_0) = v^{*\left((h',\tau^A,\tau^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_1)$ and $v^{*h'}_{(d,r)}(\boldsymbol{\lambda}_0) = v^{*h'}_{(d,r)}(\boldsymbol{\lambda}_1)$. □

Since $\boldsymbol{\lambda}_1 \in \mathfrak{D} \cap \boldsymbol{\Lambda}^{\circ}_{(d,r)}$ (Claim 1), if $\boldsymbol{\lambda}_1 \in \boldsymbol{\Lambda}^{\bullet}_{(d,r)}$, the proof is complete. Suppose that $\boldsymbol{\lambda}_1 \notin \boldsymbol{\Lambda}^{\bullet}_{(d,r)}$. Let $\mathcal{H}_2$ be the set of all histories of announcements h in $G^*_{(d,r)}(\boldsymbol{\lambda}_1)$ satisfying the following: (1) the last pair of announcements (g^A, g^B) in h has $g^A \neq g^B$; (2) there is a player i such that for any announcement $f^i \in \mathfrak{F}^*_{(d,r)}(g^i)$, there exists an announcement f^j that the opponent j may choose (in the same round, in $\mathfrak{F}^*_{(d,r)}(g^j)$) such that $v^{*\left((h,f^A,f^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_1) < 0$; (3) $v^{*h}_{(d,r)}(\boldsymbol{\lambda}_1) = 0$; and (4) there is no super-history h' of h different from h such that (a) the last pair of announcements (t^A, t^B) in h' has $t^A \neq t^B$, (b) there is a player i such that for any announcement $\tau^i \in \mathfrak{F}^*_{(d,r)}(t^i)$, there exists an announcement τ^j that the opponent j may choose (in the same round, in $\mathfrak{F}^*_{(d,r)}(t^j)$) such that $v^{*\left((h',\tau^A,\tau^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}_1) < 0$, and (c) $v^{*h'}_{(d,r)}(\boldsymbol{\lambda}_1) = 0$. Because $\boldsymbol{\lambda}_1 \notin \boldsymbol{\Lambda}^{\bullet}_{(d,r)}$, since $\boldsymbol{\lambda}_1$ is symmetric (Claim 1), we have $\mathcal{H}_2 \neq \emptyset$. Moreover, by Claim 2, and letting $\mathcal{H}$ and $\mathcal{H}^{\circ}_1$ represent, respectively, the set of all histories of announcements in $G^*_{(d,r)}$ and the set of all super-histories (in $G^*_{(d,r)}$) of the members of $\mathcal{H}_1$ (including $\mathcal{H}_1$), we must have $\mathcal{H}_2 \subseteq \mathcal{H} \setminus \mathcal{H}^{\circ}_1$. Reasoning as before, one can obtain the analogue of $\boldsymbol{\lambda}_1$, $\boldsymbol{\lambda}_2$, in $\mathfrak{D} \cap \boldsymbol{\Lambda}^{\circ}_{(d,r)}$. If $\boldsymbol{\lambda}_2 \in \boldsymbol{\Lambda}^{\bullet}_{(d,r)}$, the proof is complete. Otherwise, one can define the analogue of $\mathcal{H}_2$, $\mathcal{H}_3$, and show that $\mathcal{H}_3 \subseteq \mathcal{H} \setminus \mathcal{H}^{\circ}_2$, where $\mathcal{H}^{\circ}_2$ is the analogue of $\mathcal{H}^{\circ}_1$. Eventually, there is an element of the sequence $\mathcal{H} \setminus \mathcal{H}^{\circ}_1, \mathcal{H} \setminus \mathcal{H}^{\circ}_2, \ldots$ that becomes $\emptyset$, and some $\boldsymbol{\lambda}_k$ is obtained in $\mathfrak{D} \cap \boldsymbol{\Lambda}^{\circ}_{(d,r)}$ with $\boldsymbol{\lambda}_k \in \boldsymbol{\Lambda}^{\bullet}_{(d,r)}$. ■

Lemma 2. *Suppose that $+\infty > \eta > 0$ and $1 > \theta > 0$. Then $\mathcal{M}\left(\varepsilon_{(\eta,\theta)}, \theta\right) \subseteq \mathcal{M}^*_\eta$ for some $+\infty > \varepsilon_{(\eta,\theta)} > 0$.*

Proof. Suppose that $+\infty > \eta > 0$. It is easy to see that there exists $+\infty > \varepsilon_\eta > 0$ such that any member of $\{(d,r) \in \mathcal{M} : \varepsilon(d) \leq \varepsilon_\eta\}$ satisfies the second condition in the definition of $\mathcal{M}^*_\eta$ (see (2)).[41]

We show that there exists $+\infty > \varepsilon_\theta > 0$ such that any member of $\left\{(d,r) \in \mathcal{M} : \varepsilon(d) \leq \varepsilon_\theta \text{ and } \frac{r}{\sum_{(x,a)} d(x,a)x} \leq \theta\right\}$ satisfies the first condition in

[41] As $\varepsilon(d)$ decreases, given $f \in \mathcal{T}^{\circ}_{(d,r)}$, (x,a) with $f(x,a) > 0$, and (z,c) with $d(z,c) < d(x,a)$, $z - f(z,c)$ decreases, and one can take $z - f(z,c)$ below any positive number if $\varepsilon(d)$ is sufficiently small.

the definition of $\mathcal{M}_\eta^*$. Given $(d,r) \in \mathcal{M}$, define the sets $X^0_{(d,r)}, X^1_{(d,r)}, \ldots$ inductively as follows: $X^0_{(d,r)} := \arg\min_{(x,a)\in\{d>0\}} d(x,a)$, and, for $k = 1, 2, \ldots$,

$$X^k_{(d,r)} := \arg \min_{(x,a)\in\{d>0\}\setminus\left(\bigcup_{l=0}^{k-1} X^l_{(d,r)}\right)} d(x,a).$$

Define $\tau_{(d,r)}$ inductively as follows:

- $\tau_{(d,r)}(x,a) := \max_{\substack{\mathfrak{f}\in\mathcal{P}_{(d,r)}: \\ (x-\mathfrak{f}(x,a)-\varepsilon(d))d(x,a) \\ \geq \max_{(y,b)} d(y,b)\varepsilon(d)}} \mathfrak{f}(x,a)$ for all $(x,a) \in X^0_{(d,r)}$;
- for $k = 1, 2, \ldots$, let

$$\tau_{(d,r)}(x,a) := \max_{\substack{\mathfrak{f}\in\mathcal{P}_{(d,r)}: \\ (x-\mathfrak{f}(x,a)-\varepsilon(d))d(x,a)\geq \max_{(y,b)} d(y,b)\varepsilon(d) \\ +\sum_{(y,b):d(y,b)<d(x,a)}(y-\mathfrak{f}(y,b))d(y,b)}} \mathfrak{f}(x,a)$$

  for all $(x,a) \in X^k_{(d,r)}$.

Note that the choice of $\tau_{(d,r)}$ entails that, for any (d,r) with $\varepsilon(d)$ sufficiently small, say $\varepsilon(d) \leq \varepsilon_\theta$,

$$\frac{\sum_{(x,a)} \tau_{(d,r)}(x,a)d(x,a)}{\sum_{(x,a)} xd(x,a)} > \theta. \tag{5}$$

Now fix any $(d,r) \in \mathcal{M}$ with $\varepsilon(d) \leq \varepsilon_\theta$ and $\frac{r}{\sum_{(x,a)} d(x,a)x} \leq \theta$. The first condition in the definition of $\mathcal{M}_\eta^*$ says that for any non-final $f \in \mathcal{P}_{(d,r)}$ for which $f \leq g$, some $g \in \mathcal{O}_{(d,r)}$, $Z_{(d,r)}(f) \neq \emptyset$. Fix a non-final $f \in \mathcal{P}_{(d,r)}$ for which $f \leq g$, some $g \in \mathcal{O}_{(d,r)}$. Because $\varepsilon(d) \leq \varepsilon_\theta$ (so (5) holds) and $\frac{r}{\sum_{(x,a)} d(x,a)x} \leq \theta$, there exists $(\mathfrak{x},\mathfrak{a})$ such that $f(\mathfrak{x},\mathfrak{a}) \leq \tau_{(d,r)}(\mathfrak{x},\mathfrak{a})$. Clearly, the choice of $\tau_{(d,r)}$ entails $Z_{(d,r)}\left(\tau_{(d,r)}\right) \ni (x,a)$ for all $(x,a) \in \{d > 0\}$. In particular, $Z_{(d,r)}\left(\tau_{(d,r)}\right) \ni (\mathfrak{x},\mathfrak{a})$. Therefore, since $f(\mathfrak{x},\mathfrak{a}) \leq \tau_{(d,r)}(\mathfrak{x},\mathfrak{a})$, it follows that $Z_{(d,r)}(f) \neq \emptyset$.

We conclude that any member of

$$\left\{(d,r) \in \mathcal{M} : \varepsilon(d) \leq \varepsilon_\theta \text{ and } \frac{r}{\sum_{(x,a)} d(x,a)x} \leq \theta\right\}$$

satisfies the first condition in the definition of $\mathcal{M}_\eta^*$. It only remains to observe that $\mathcal{M}\left(\varepsilon_{(\eta,\theta)},\theta\right) \subseteq \mathcal{M}_\eta^*$ holds for any $0 < \varepsilon_{(\eta,\theta)} < \min\{\varepsilon_\eta, \varepsilon_\theta\}$. ■

Lemma 3. *Suppose that* $(d,r) \in \mathcal{M}^\circ$ *and* $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)}$. *Suppose that* $h = (f_1^A, f_1^B, ..., f_k^A, f_k^B)$ *is a history of announcements in* $G^*_{(d,r)}(\boldsymbol{\lambda})$ *such that* $f_k^A \neq f_k^B$. *If the value of* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *is zero, there exists a Nash equilibrium in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *that generates a unique history of announcements in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *whose last two elements are identical.*

Proof. Suppose that $(d,r) \in \mathcal{M}^\circ$ and $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)}$. Suppose that $h = (f_1^A, f_1^B, ..., f_k^A, f_k^B)$ is a history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ such that $f_k^A \neq f_k^B$, and let the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ be zero. We proceed by induction on the number of rounds left until the end of the game.

If the set $\mathfrak{F}^*_{(d,r)}(f_k^A) \times \mathfrak{F}^*_{(d,r)}(f_k^B)$ contains only pairs of final announcements, then, since (1) $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)}$, (2) $h = (f_1^A, f_1^B, ..., f_k^A, f_k^B)$ is a history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ such that $f_k^A \neq f_k^B$, and (3) the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ is zero, for any player i, there exists $f^i \in \mathfrak{F}^*_{(d,r)}(f_k^i)$ such that for all $\tau^j \in \mathfrak{F}^*_{(d,r)}(f_k^j)$ $(j \neq i)$, $v^{*((h,f^i,\tau^j),i)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$.[42] It follows that there is an equilibrium in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ that prescribes play of (f^A, f^B) (with probability one) in the first round of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$. Moreover, since $v^{*((h,f^A,f^B),i)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$ for each i, we have $v^{*(h,f^A,f^B)}_{(d,r)}(\boldsymbol{\lambda}) = 0$. Therefore, since $(d,r) \in \mathcal{M}^\circ$ and, by assumption, each f^i is final, $f^A = f^B$ must hold. We have thus obtained a Nash equilibrium in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical.

Now suppose that the following has been proven: Suppose that $h = (f_1^A, f_1^B, ..., f_k^A, f_k^B)$ is a history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ such that $f_k^A \neq f_k^B$, and let the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ be zero. Suppose that $(\tau_\kappa^A, \tau_\kappa^B)$ is a pair of final announcements for any history $(\tau_1^A, \tau_1^B, ..., \tau_\kappa^A, \tau_\kappa^B)$ of announcements in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ having length κ. Then there exists a Nash equilibrium in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical.

Suppose that $h = (f_1^A, f_1^B, ..., f_k^A, f_k^B)$ is a history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ such that $f_k^A \neq f_k^B$, and let the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ be zero. Suppose that $(\tau_{\kappa+1}^A, \tau_{\kappa+1}^B)$ is a pair of final announcements for any history $(\tau_1^A, \tau_1^B, ..., \tau_{\kappa+1}^A, \tau_{\kappa+1}^B)$ of announcements in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ having length $\kappa+1$. We show that there exists a Nash equilibrium in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical.

As before, there exists an equilibrium in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ that prescribes play of

[42]By $v^{*((h,f^i,\tau^j),i)}_{(d,r)}(\boldsymbol{\lambda})$ we mean $v^{*((h,f^A,\tau^B),i)}_{(d,r)}(\boldsymbol{\lambda})$ if $i = A$ and $v^{*((h,\tau^A,f^B),i)}_{(d,r)}(\boldsymbol{\lambda})$ if $i = B$.

(f^A, f^B) in the first round of the game, and $v^{*(h,f^A,f^B)}_{(d,r)}(\boldsymbol{\lambda}) = 0$. If $f^A = f^B$, the proof is complete. Otherwise, we have a history (h, f^A, f^B) with $f^A \neq f^B$ and $v^{*(h,f^A,f^B)}_{(d,r)}(\boldsymbol{\lambda}) = 0$, and the induction hypothesis gives the desired conclusion. ■

Lemma 4. *Suppose that* $+\infty > \eta > 0$ *and* $(d, r) \in \mathcal{M}^*_\eta$. *Let* h *be a history of announcements in* $G^*_{(d,r)}(\boldsymbol{\lambda})$, *where* $\boldsymbol{\lambda}$ *is symmetric in* $\boldsymbol{\Lambda}^*_{(d,r)}$. *Suppose that the last two announcements in* h *are identical and equal to* f, *and let* $\mathbf{T}_{(d,r)}(f) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$. *Then either* $\mathbf{T}_{(d,r)}(f) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ *or there exists a Nash equilibrium in* ${\Gamma^*}^h_{(d,r)}(\boldsymbol{\lambda})$ *that generates a unique history of announcements in* ${\Gamma^*}^h_{(d,r)}(\boldsymbol{\lambda})$ *whose last two elements are identical.*

Proof. While the general proof is long, the essence of the argument is relatively straightforward, and can be illustrated within the example of Section 3.1 (the details are available upon request). In this example, assuming the antecedent of the lemma, it is easy to see that, if $\mathbf{T}_{(d,r)}(f) \nsubseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$, then, for some $(y, b) \in Z_{(d,r)}(f)$, the following map is not final:

$$(x, a) \mapsto \begin{cases} x & \text{if } d(x, a) < d(y, b), \\ f(x, a) & \text{elsewhere.} \end{cases} \tag{6}$$

(This is also true in general.) Any player i's choice of $\varphi^{(y,b)}_{(d,r)}(f)$ in the round that follows h is optimal. In fact, if the opponent chooses $\varphi^{(z,c)}_{(d,r)}(f)$ with $d(z, c) = d(y, b)$, then $(z, c) = (y, b)$;[43] since $\boldsymbol{\lambda}$ is symmetric, choosing the same action gives both players a payoff of zero, and zero is the value of ${\Gamma^*}^h_{(d,r)}(\boldsymbol{\lambda})$ (since $\boldsymbol{\lambda}$ is symmetric and the last two announcements in h equal f). If the opponent chooses $\varphi^{(z,c)}_{(d,r)}(f)$ with $d(z, c) < d(y, b)$, then, since the map (6) is not final, she cannot win the election outright. One can show that, by choosing $\varphi^{(z,c)}_{(d,r)}(f)$ with $d(z, c) < d(y, b)$, the opponent can secure at most a payoff of zero. If the opponent chooses $\varphi^{(z,c)}_{(d,r)}(f)$ with $d(z, c) > d(y, b)$, then two cases are possible: (1) i can keep taxing group (y, b) until the revenue constraint is met, thereby defeating her opponent; (2) otherwise, the map $(x, a) \mapsto x$ if $d(x, a) < d(y, b)$, $(x, a) \mapsto \varphi^{(z,c)}_{(d,r)}(f)(x, a)$ elsewhere, is not final, and i can choose $\varphi^{(z,c)}_{(d,r)}\left(\varphi^{(y,b)}_{(d,r)}(f)\right)$ after $\varphi^{(y,b)}_{(d,r)}(f)$. In this case, the argument can be repeated (with $\varphi^{(y,b)}_{(d,r)}(f)$ replacing f). A finite number of iterations will give the desired result. ■

[43] This is true in the example and, in general, if $(d, r) \in \mathcal{M}^\circ$. While the statement of Lemma 4 covers in principle cases where $(d, r) \notin \mathcal{M}^\circ$, proving it only for the cases where $(d, r) \in \mathcal{M}^\circ$ suffices for the proof of Theorem 2.

Lemma 5. *Suppose that* $+\infty > \eta > 0$ *and* $(d,r) \in \mathcal{M}^*_\eta \cap \mathcal{M}^\circ$. *Let* $\boldsymbol{\lambda}$ *be symmetric in* $\boldsymbol{\Lambda}^\circ_{(d,r)}$. *Let* $h = (f^A_1, f^B_1, ..., f^A_k, f^B_k)$ *be a history of announcements in* $G^*_{(d,r)}(\boldsymbol{\lambda})$ *satisfying the following: (1) there exist* (x,a) *and* (y,b) *and* i *and* j *such that*

$$f^j_k(x,a) = f^i_k(x,a) + \varepsilon(d), f^i_k(y,b) = f^j_k(y,b) + \varepsilon(d), f^i_k = f^j_k \text{ elsewhere,} \quad (7)$$

and $f^i_k(y,b) > t(y,b)$ *for all* $t \in \mathcal{E}^*_{(d,r)}$, *and (2)* f^A_k *and* f^B_k *are not final. Let* σ_i *be a strategy profile in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *that prescribes* $\varphi^{(x,a)}_{(d,r)}(f^i_k)$ *in the first round of* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$. *Then, there is a strategy* σ_j *in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *that secures a positive payoff in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *against* σ_i.

Proof. The details are cumbersome (and available upon request), but the general idea of the proof is simple. Since (7) holds and σ_i prescribes $\varphi^{(x,a)}_{(d,r)}(f^i_k)$ in the first round of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$, it suffices to show that there is a strategy ν_j in $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$ that secures a positive payoff against ν_i, where the last two elements of h' are identical and equal to f^j_k, and ν_i prescribes $\varphi^{(y,b)}_{(d,r)}(f^j_k)$ in the first round of $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$. Observe that, because (1) f^j_k is not final, (2) $f^j_k \leq g$ for some $g \in \mathcal{O}_{(d,r)}$ (for h is a history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$), and (3) $(d,r) \in \mathcal{M}^*_\eta$, we have $Z_{(d,r)}(f^j_k) \neq \emptyset$. Therefore, the argument in the proof of Lemma 4 can be used to see that there is a strategy ν_j in $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$ that secures a nonnegative payoff against ν_i. Imagine for a moment that $\boldsymbol{\lambda}$ is perfect. Consider a strategy μ_j in $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$ that mimics ν_j except that, each time it is j's time to move second, j increases her tax policy by $\varepsilon(d)$ at the point where i increased her tax policy as a first mover. Because $f^i_k(y,b) > t(y,b)$ for all $t \in \mathcal{E}^*_{(d,r)}$, for each non-final $\mathfrak{f} \in \mathbf{T}_{(d,r)}(f^j_k)$ and every $(z,c) \in Z_{(d,r)}(\mathfrak{f})$, we have $d(z,c) < d(y,b)$. Thus, μ_j ensures, against ν_i, a zero payoff when it mimics ν_j and a positive payoff in the event in which j moves second in each round of $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$. The same argument is valid if $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\circ_{(d,r)}$. ■

Lemma 6. *Suppose that* $(d,r) \in \mathcal{M}$. *Suppose that* $\boldsymbol{\lambda}$ *is symmetric in* $\boldsymbol{\Lambda}^*_{(d,r)}$. *Let* h *be a history in* $G^*_{(d,r)}(\boldsymbol{\lambda})$ *whose last pair of announcements is* (f^A, f^B), *where* $f^A \neq f^B$. *Let* $\sigma = (\sigma_A, \sigma_B)$ *be a strategy profile in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *with following property:* σ *is a Nash equilibrium in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$, *and play of* σ *generates a unique history of announcements in* $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ *whose last two elements are identical. Let* $(f^A_1, f^B_1, ..., f^A_k, f^B_k)$ *be a history generated under play of* σ. *Suppose that* $k > 1$, $f^A_\ell \neq f^B_\ell$ *for all* $\ell = 1, ..., k-1$, *and* $f^A_k = f^B_k$. *Suppose that* $f^A_1 = \varphi^{(x,a)}_{(d,r)}(f^A)$ *and* $f^A_2 = \varphi^{(y,b)}_{(d,r)}(f^A_1)$, *where* $(x,a) \neq (y,b)$. *Then the value of* $\Gamma^{*\left(h, \varphi^{(y,b)}_{(d,r)}(f^A), f^B_1\right)}_{(d,r)}(\boldsymbol{\lambda})$ *is zero. Analogous statements are true for player* B.

Proof. Assume the antecedent. We prove the statement for player A (the argument for B is similar). By the property of σ and the fact that $\boldsymbol{\lambda}$ is symmetric, the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ is zero. Let $h' := \left(h, \varphi^{(y,b)}_{(d,r)}(f^A), f^B_1\right)$. We show that $v^{*(h',A)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$. To see this, it suffices to show that

$$v^{*\left(\left(h',\varphi^{(x,a)}_{(d,r)}\left(\varphi^{(y,b)}_{(d,r)}(f^A)\right),\tau^B\right),A\right)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$$

for any announcement τ^B that is feasible for B in the first round of $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$. Observe that $\left(\varphi^{(x,a)}_{(d,r)}\left(\varphi^{(y,b)}_{(d,r)}(f^A)\right), \tau^B\right) = (f^A_2, \tau^B)$. Further,

$$v^{*\left(\left(h,f^A_1,f^B_1,f^A_2,\tau^B\right),A\right)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0,$$

since the choice of f^A_2 is optimal in the round that follows (h, f^A_1, f^B_1) (and σ is a Nash equilibrium in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements $(f^A_1, f^B_1, ..., f^A_k, f^B_k)$ in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical). Therefore, $v^{*\left(\left(h',\varphi^{(x,a)}_{(d,r)}\left(\varphi^{(y,b)}_{(d,r)}(f^A)\right),\tau^B\right),A\right)}_{(d,r)}(\boldsymbol{\lambda}) = v^{*\left(\left(h,f^A_1,f^B_1,f^A_2,\tau^B\right),A\right)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$, as desired.

Because the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ is zero and there exists an optimal play of B that prescribes f^B_1 in the first round of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$, regardless of nature's choices (of the order of moves and the signals), we must have $v^{*(h',B)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$. Hence, since we know that $v^{*(h',A)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$, we obtain $v^{*h'}_{(d,r)}(\boldsymbol{\lambda}) = 0$, as we sought. ■

Lemma 7. *Suppose that $(d, r) \in \mathcal{M}^\circ$ and $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)} \cap \boldsymbol{\Lambda}^\circ_{(d,r)}$. Let $h = (f^A_1, f^B_1, ..., f^A_k, f^B_k)$ be a history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ such that $\mathbf{T}_{(d,r)}(f^i_k) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ for each i. Let $\sigma = (\sigma_A, \sigma_B)$ be a Nash equilibrium in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ whose last two components are identical. Given a player i, suppose that σ_i prescribes g^i in the round that follows h. Then $\mathbf{T}_{(d,r)}(g^i) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$.*

Proof. Assume the antecedent. We fix a player i, assume that $\mathbf{T}_{(d,r)}(g^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$, and derive a contradiction. Because $\mathbf{T}_{(d,r)}(f^i_k) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ and $\mathbf{T}_{(d,r)}(g^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$, $g^i = \varphi^{(x,a)}_{(d,r)}(f^i_k)$ for some (x, a). The choice of (x, a), along with the fact that $\mathbf{T}_{(d,r)}(f^i_k) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ and $\mathbf{T}_{(d,r)}(g^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$, entails $g^i(x, a) > g_\tau(x, a)$ for all $\tau \in \mathcal{T}^\circ_{(d,r)}$ (g_τ was defined in Section 5.1). By assumption, play of σ in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ generates a unique history of announcements $(\tau^A_1, \tau^B_1, ..., \tau^A_l, \tau^B_l)$ in $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical ($\tau^A_l =$

τ_l^B). Either $l = 1$ or $l > 1$. Both cases can be dealt with using a similar argument, so we only consider the case when $l > 1$.

Without loss of generality, assume $\tau_\ell^A \neq \tau_\ell^B$ for all $\ell = 1, ..., l-1$. Since $\tau_1^i = g^i$ and $g^i(x,a) > g_\tau(x,a)$ for all $\tau \in \mathcal{T}^\circ_{(d,r)}$, we may pick the longest sub-history $(\tau_1^A, \tau_1^B, ..., \tau_m^A, \tau_m^B)$ of $(\tau_1^A, \tau_1^B, ..., \tau_{l-1}^A, \tau_{l-1}^B)$ such that the following is true:

($\star$) There exist a player ι and (y,b) with $\tau_{m+1}^\iota = \varphi_{(d,r)}^{(y,b)}(\tau_m^\iota)$ and $\tau_{m+1}^\iota(y,b) > g_\tau(y,b)$ for all $\tau \in \mathcal{T}^\circ_{(d,r)}$.

Let $h' := (h, \tau_1^A, \tau_1^B, ..., \tau_m^A, \tau_m^B)$. Two cases are possible: $\tau_m^A(y,b) = \tau_m^B(y,b)$ and $\tau_m^\iota(y,b) < \tau_m^j(y,b)$, where $j \neq \iota$ (the case when $\tau_m^\iota(y,b) > \tau_m^j(y,b)$ is not possible, for in this case $(\tau_1^A, \tau_1^B, ..., \tau_m^A, \tau_m^B)$ would not be the longest sub-history of $(\tau_1^A, \tau_1^B, ..., \tau_{l-1}^A, \tau_{l-1}^B)$ satisfying ($\star$) (recall that $\tau_l^A = \tau_l^B$)). We consider each case in turn.

Case 1. $\tau_m^A(y,b) = \tau_m^B(y,b)$. Since $\tau_l^A = \tau_l^B$ and $(\tau_1^A, \tau_1^B, ..., \tau_m^A, \tau_m^B)$ is the longest sub-history of $(\tau_1^A, \tau_1^B, ..., \tau_{l-1}^A, \tau_{l-1}^B)$ such that ($\star$) holds, in this case we must have $\tau_{m+1}^A(y,b) = \tau_{m+1}^B(y,b)$. On the other hand, we know that $\tau_m^A \neq \tau_m^B$ (for $m < l$). Since (1) $(\tau_1^A, \tau_1^B, ..., \tau_m^A, \tau_m^B)$ is the longest sub-history of $(\tau_1^A, \tau_1^B, ..., \tau_{l-1}^A, \tau_{l-1}^B)$ such that ($\star$) holds, (2) $\tau_m^A \neq \tau_m^B$, (3) $\tau_m^A(y,b) = \tau_m^B(y,b)$, and (4) $\tau_l^A = \tau_l^B$, letting j be ι's opponent, τ_{m+1}^j must be not final, so we may write $\tau_{m+2}^j = \varphi_{(d,r)}^{(z,c)}(\tau_{m+1}^j)$ for some (z,c). Note that the fact that $(\tau_1^A, \tau_1^B, ..., \tau_m^A, \tau_m^B)$ is the longest sub-history of $(\tau_1^A, \tau_1^B, ..., \tau_{l-1}^A, \tau_{l-1}^B)$ such that ($\star$) holds implies that $(z,c) \neq (y,b)$. Thus, we have $\tau_{m+1}^j = \varphi_{(d,r)}^{(y,b)}(\tau_m^j)$, $\tau_{m+2}^j = \varphi_{(d,r)}^{(z,c)}(\tau_{m+1}^j)$, and $(z,c) \neq (y,b)$. Moreover, $m+2 \leq l$ (indeed, we have $\tau_m^A \neq \tau_m^B$ and $\tau_{m+1}^A \neq \tau_{m+1}^B$). Therefore, by Lemma 6, the value of $\Gamma^{*\left(h', \varphi_{(d,r)}^{(z,c)}(\tau_m^j), \tau_{m+1}^\iota\right)}_{(d,r)}(\boldsymbol{\lambda})$ is zero.[44] Since (1) $(d,r) \in \mathcal{M}^\circ$, (2) $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)}$, (3) $\varphi_{(d,r)}^{(z,c)}(\tau_m^j) \neq \tau_{m+1}^\iota$, and (4) the value of $\Gamma^{*\left(h', \varphi_{(d,r)}^{(z,c)}(\tau_m^j), \tau_{m+1}^\iota\right)}_{(d,r)}(\boldsymbol{\lambda})$ is zero,[45] Lemma 3 gives a Nash equilibrium $\nu = (\nu_A, \nu_B)$ in $\Gamma^{*\left(h', \varphi_{(d,r)}^{(z,c)}(\tau_m^j), \tau_{m+1}^\iota\right)}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements $(t_1^A, t_1^B, ..., t_\kappa^A, t_\kappa^B)$ whose last two elements are identical ($t_\kappa^A = t_\kappa^B$). We consider two sub-cases:

[44]By $\Gamma^{*\left(h', \varphi_{(d,r)}^{(z,c)}(\tau_m^j), \tau_{m+1}^\iota\right)}_{(d,r)}(\boldsymbol{\lambda})$ we mean $\Gamma^{*\left(h', \varphi_{(d,r)}^{(z,c)}(\tau_m^j), \tau_{m+1}^\iota\right)}_{(d,r)}(\boldsymbol{\lambda})$ if $j = A$ and $\Gamma^{*\left(h', \tau_{m+1}^\iota, \varphi_{(d,r)}^{(z,c)}(\tau_m^j)\right)}_{(d,r)}(\boldsymbol{\lambda})$ if $j = B$.

[45]See footnote 44.

Case 1.1. $\kappa = 1$. To ease notation, let $t_0^j := \varphi_{(d,r)}^{(z,c)}(\tau_m^j)$ and $t_0^\iota := \tau_{m+1}^\iota$. We have $t_0^A \neq t_0^B$, t_0^A and t_0^B non-final, and (if $\kappa = 1$) $t_1^A = t_1^B$. Therefore, there exist $(\mathfrak{x}, \mathfrak{a})$ and $(\mathfrak{y}, \mathfrak{b})$ and i_1 and i_2 such that

$$t_0^{i_1}(\mathfrak{x}, \mathfrak{a}) = t_0^{i_2}(\mathfrak{x}, \mathfrak{a}) + \varepsilon(d), t_0^{i_2}(\mathfrak{y}, \mathfrak{b}) = t_0^{i_1}(\mathfrak{y}, \mathfrak{b}) + \varepsilon(d), t_0^A = t_0^B \text{ elsewhere.} \quad (8)$$

Moreover, we have the following: (1) t_0^A and t_0^B are non-final; (2) $t_0^\iota(y, b) = \tau_{m+1}^\iota(y, b) > t(y, b)$ for all $t \in \mathcal{E}_{(d,r)}^*$ (since $\tau_{m+1}^\iota(y, b) > g_\tau(y, b)$ for all $\tau \in \mathcal{T}_{(d,r)}^\circ$); (3) since $\tau_m^A(y, b) = \tau_m^B(y, b)$ and $(z, c) \neq (y, b)$,

$$t_0^j(y, b) = \varphi_{(d,r)}^{(z,c)}(\tau_m^j)(y, b) < \varphi_{(d,r)}^{(y,b)}(\tau_m^j)(y, b) = \tau_{m+1}^\iota(y, b) = t_0^\iota(y, b),$$

and so, in view of (8), we have either $(y, b) = (\mathfrak{x}, \mathfrak{a})$ or $(y, b) = (\mathfrak{y}, \mathfrak{b})$. Without loss of generality, say $(y, b) = (\mathfrak{x}, \mathfrak{a})$, so that $j = i_2$ (and $\iota = i_1$); (4) and, since $t_1^A = t_1^B$, ν_ι is a strategy in $\Gamma^{*\left(h', t_0^A, t_0^B\right)}_{(d,r)}(\boldsymbol{\lambda})$ that prescribes $\varphi_{(d,r)}^{(\mathfrak{y},\mathfrak{b})}(t_0^\iota)$ in the first round of $\Gamma^{*\left(h', t_0^A, t_0^B\right)}_{(d,r)}(\boldsymbol{\lambda})$. Consequently, by Lemma 5, ν_ι is not optimal in $\Gamma^{*\left(h', t_0^A, t_0^B\right)}_{(d,r)}(\boldsymbol{\lambda})$, thereby contradicting the fact that ν is a Nash equilibrium in $\Gamma^{*\left(h', t_0^A, t_0^B\right)}_{(d,r)}(\boldsymbol{\lambda})$.

Case 1.2. $\kappa > 1$. In this case, one may proceed as before, i.e., one may pick the longest sub-history $(t_1^A, t_1^B, ..., t_n^A, t_n^B)$ of $(t_1^A, t_1^B, ..., t_{\kappa-1}^A, t_{\kappa-1}^B)$ such that the analogue of $(\star)$ is satisfied. One may then consider the analogues of the current Case 1 and Case 2 (below) and either obtain a contradiction or repeat the argument once more, until a point is reached in which a contradiction arises.

Case 2. $\tau_m^\iota(y, b) < \tau_m^j(y, b)$, where $j \neq \iota$. We consider three sub-cases:

Case 2.1. $m + 1 = l$ and τ_m^j is final. In this case, $\tau_{m+1}^A = \tau_{m+1}^B$, and ι has an optimal strategy in $\Gamma^{*h'}_{(d,r)}$, σ_ι, which chooses $\tau_{m+1}^\iota = \varphi_{(d,r)}^{(y,b)}(\tau_m^\iota)$ in the first round of this game (regardless of nature's choices). Therefore, since τ_m^j is final by assumption, we must have $\tau_m^A = \tau_m^B$ everywhere except at (y, b), where $\tau_m^\iota(y, b) < \tau_m^j(y, b)$. Because $\tau_{m+1}^\iota(y, b) > g_\tau(y, b)$ for all $\tau \in \mathcal{T}_{(d,r)}^\circ$, $\tau_m^\iota(y, b) \geq g_\tau(y, b)$ for all $\tau \in \mathcal{T}_{(d,r)}^\circ$. We claim that this implies that there exists (y, b) such that $d(\mathsf{y}, \mathsf{b}) < d(y, b)$ and, for some positive integer l, the following map is final:

$$(\mathsf{x}, \mathsf{a}) \mapsto \begin{cases} \tau_m^\iota(\mathsf{x}, \mathsf{a}) & \text{if } (\mathsf{x}, \mathsf{a}) \neq (y, b), \\ \tau_m^\iota(\mathsf{x}, \mathsf{a}) + l\varepsilon(d) & \text{if } (\mathsf{x}, \mathsf{a}) = (y, b). \end{cases} \quad (9)$$

To see this, observe that we must have $(\mathsf{y}, \mathsf{b}) \in Z_{(d,r)}(\tau_m^\iota)$ for some (y, b) with $d(\mathsf{y}, \mathsf{b}) < d(y, b)$ (otherwise, since $\tau_m^\iota(y, b) \geq g_\tau(y, b)$ for all $\tau \in \mathcal{T}_{(d,r)}^\circ$ and

$\arg\max_{(\mathsf{x},\mathsf{a}):\tau(\mathsf{x},\mathsf{a})>0} d(\mathsf{x},\mathsf{a})$ is a singleton for all $\tau \in \mathcal{T}^{\circ}_{(d,r)}$ (because $(d,r) \in \mathcal{M}^{\circ}$), τ^{ι}_{m} would be final, thereby contradicting that $\tau^{\iota}_{m+1} = \varphi^{(y,b)}_{(d,r)}(\tau^{\iota}_{m})$). But since $(\mathsf{y},\mathsf{b}) \in Z_{(d,r)}(\tau^{\iota}_{m})$, we have

$$(\mathsf{y} - \tau^{\iota}_{m}(\mathsf{y},\mathsf{b}) - \varepsilon(d))\, d(\mathsf{y},\mathsf{b}) \geq \max_{(\mathsf{x},\mathsf{a})} d(\mathsf{x},\mathsf{a})\varepsilon(d) \geq d(y,b)\varepsilon(d).$$

Therefore, since τ^{j}_{m} is final and $\tau^{A}_{m} = \tau^{B}_{m}$ everywhere except at (y,b), where $\tau^{\iota}_{m}(y,b) < \tau^{j}_{m}(y,b)$, it follows that the map in (9) is final for some positive integer l and some (y,b) with $d(\mathsf{y},\mathsf{b}) < d(y,b)$. But then, σ_{ι}, which chooses $\tau^{\iota}_{m+1} = \varphi^{(y,b)}_{(d,r)}(\tau^{\iota}_{m})$ in the first round of $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$, cannot be optimal (in fact, choosing the map in (9) in $\Gamma^{*h'}_{(d,r)}(\boldsymbol{\lambda})$ gives ι a higher payoff).

Case 2.2. $m+1 = l$ and τ^{j}_{m} is not final. In this case, $\tau^{A}_{m+1} = \tau^{B}_{m+1}$ (as in Case 2.1), and there exists (z,c) such that

$$\tau^{j}_{m}(y,b) = \tau^{\iota}_{m}(y,b) + \varepsilon(d), \tau^{\iota}_{m}(\mathsf{z},\mathsf{c}) = \tau^{j}_{m}(\mathsf{z},\mathsf{c}) + \varepsilon(d), \text{ and } \tau^{A}_{m} = \tau^{B}_{m} \text{ elsewhere.}$$

We omit the rest of the argument, which is similar to that for Case 1.1.

Case 2.3. $m+1 < l$. In this case, since $(\tau^{A}_{1}, \tau^{B}_{1}, ..., \tau^{A}_{m}, \tau^{B}_{m})$ is the longest subhistory of $(\tau^{A}_{1}, \tau^{B}_{1}, ..., \tau^{A}_{l-1}, \tau^{B}_{l-1})$ such that $(\star)$ holds, there must exist $(\mathfrak{z},\mathfrak{c})$ such that $\tau^{\iota}_{m+2} = \varphi^{(\mathfrak{z},\mathfrak{c})}_{(d,r)}(\tau^{\iota}_{m+1})$ and $(\mathfrak{z},\mathfrak{c}) \neq (y,b)$. Thus, we have $\tau^{\iota}_{m+1} = \varphi^{(y,b)}_{(d,r)}(\tau^{\iota}_{m})$, $\tau^{\iota}_{m+2} = \varphi^{(\mathfrak{z},\mathfrak{c})}_{(d,r)}(\tau^{\iota}_{m+1})$, and $(\mathfrak{z},\mathfrak{c}) \neq (y,b)$. Now, by Lemma 6, the value of $\Gamma^{*\left(h', \varphi^{(\mathfrak{z},\mathfrak{c})}_{(d,r)}(\tau^{\iota}_{m}), \tau^{j}_{m+1}\right)}_{(d,r)}(\boldsymbol{\lambda})$ is zero. Since $(d,r) \in \mathcal{M}^{\bullet} \cap \mathcal{M}^{\circ}$, $\varphi^{(\mathfrak{z},\mathfrak{c})}_{(d,r)}(\tau^{\iota}_{m}) \neq \tau^{j}_{m+1}$, and the value of $\Gamma^{*\left(h', \varphi^{(\mathfrak{z},\mathfrak{c})}_{(d,r)}(\tau^{\iota}_{m}), \tau^{j}_{m+1}\right)}_{(d,r)}(\boldsymbol{\lambda})$ is zero, Lemma 3 gives a Nash equilibrium in $\Gamma^{*\left(h', \varphi^{(\mathfrak{z},\mathfrak{c})}_{(d,r)}(\tau^{\iota}_{m}), \tau^{j}_{m+1}\right)}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements whose last two elements are identical. One may now formulate the analogues of Case 1 and Case 2 for the new history, and repeat the argument as needed, until a contradiction is obtained. ■

We omit the proof of the following lemma, which is similar to that of Lemma 7.

Lemma 8. *Suppose that $(d,r) \in \mathcal{M}^{\circ}$ and $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^{\bullet}_{(d,r)}$. Let $h = (f^{A}_{1}, f^{B}_{1}, ..., f^{A}_{k+1}, f^{B}_{k+1})$ be a history of announcements in $G^{*}_{(d,r)}(\boldsymbol{\lambda})$ such that $\mathbf{T}_{(d,r)}(f^{\iota}_{k}) \cap \mathcal{E}^{*}_{(d,r)} \neq \emptyset$ for each ι, $\mathbf{T}_{(d,r)}(f^{i}_{k+1}) \cap \mathcal{E}^{*}_{(d,r)} \neq \emptyset$ and $\mathbf{T}_{(d,r)}(f^{j}_{k+1}) \cap \mathcal{E}^{*}_{(d,r)} = \emptyset$ for some i, j. Then the value of $\Gamma^{*h}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero.*

The proof of the following lemma is tedious but otherwise relatively straightforward. It is therefore omitted (yet available from the author upon request).

Lemma 9. *Suppose that* $+\infty > \eta > 0$ *and* $(d,r) \in \mathcal{M}_\eta^* \cap \mathcal{M}^\circ$. *Then* $\mathcal{E}^*_{(d,r)} \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$.

Lemma 10. *Suppose that* $+\infty > \eta > 0$, $(d,r) \in \mathcal{M}_\eta^* \cap \mathcal{M}^\circ$, *and* $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)} \cap \boldsymbol{\Lambda}^\circ_{(d,r)}$. *Then the tax policy implemented at any Nash equilibrium of* $G^*_{(d,r)}(\boldsymbol{\lambda})$ *lies in* $N_\eta\left(\mathcal{E}_{(d,r)}\right)$.

Proof. Suppose that $+\infty > \eta > 0$, $(d,r) \in \mathcal{M}_\eta^* \cap \mathcal{M}^\circ$, and $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)} \cap \boldsymbol{\Lambda}^\circ_{(d,r)}$. Suppose that $\sigma = (\sigma_A, \sigma_B)$ is a strategy profile in $G^*_{(d,r)}(\boldsymbol{\lambda})$. Let $h = (f_1^A, f_1^B, ..., f_k^A, f_k^B)$ be any history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ generated under play of σ. We know that $f_1^i = 0$ for each i, and therefore $\mathbf{T}_{(d,r)}(f_1^i) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ for each i. If $\mathbf{T}_{(d,r)}(f_k^\iota) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ for each ι, there is nothing to prove, so suppose that it is not true that $\mathbf{T}_{(d,r)}(f_k^\iota) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ for each ι. Let $(h, f_{k+1}^A, f_{k+1}^B)$ be a super-history of h in $G^*_{(d,r)}(\boldsymbol{\lambda})$ generated under play of σ, and suppose that $\mathbf{T}_{(d,r)}(f_{k+1}^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$ for some i. We first show that σ is not a Nash equilibrium of $G^*_{(d,r)}(\boldsymbol{\lambda})$.

Observe that, because $\mathbf{T}_{(d,r)}(f_{k+1}^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$ and $\mathbf{T}_{(d,r)}(f_1^i) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$, there must exist some $\kappa = 1, ..., k$ with $\mathbf{T}_{(d,r)}(f_\kappa^i) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ and $\mathbf{T}_{(d,r)}(f_{\kappa+1}^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$. There is no loss of generality in assuming that i's opponent j satisfies $\mathbf{T}_{(d,r)}(f_\kappa^j) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$. Let $h_l := (f_1^A, f_1^B, ..., f_l^A, f_l^B)$ for each $l = 1, ..., \kappa$. We consider two cases:

Case 1. Either $f_\kappa^A \neq f_\kappa^B$ *and the value of* $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ *is zero or* $f_\kappa^A = f_\kappa^B$. If $f_\kappa^A = f_\kappa^B$, since $\mathbf{T}_{(d,r)}(f_\kappa^\iota) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ for each ι, Lemma 4 implies that either $\mathbf{T}_{(d,r)}(f_\kappa^\iota) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ for each ι or there exists a Nash equilibrium in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical. Thus, if $f_\kappa^A = f_\kappa^B$, it suffices to consider the case where $\mathbf{T}_{(d,r)}(f_\kappa^\iota) \nsubseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ (recall that we are assuming that it is not true that $\mathbf{T}_{(d,r)}(f_k^\iota) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ for each ι).

There exists a Nash equilibrium μ in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical (this follows from Lemma 3 if $f_\kappa^A \neq f_\kappa^B$ and the value of $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ is zero, and from Lemma 4 if $f_\kappa^A = f_\kappa^B$ and $\mathbf{T}_{(d,r)}(f_\kappa^\iota) \nsubseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$). Therefore, since $\boldsymbol{\lambda}$ is symmetric and $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ is zero-sum, the value of $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ is zero. Since the value of $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ is zero, letting (t^A, t^B) be a history in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ that is generated with probability one under play of μ, and letting j be i's opponent, we must have $v^{*((h_\kappa,t^j,t),j)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$ for any $t \in \mathfrak{F}^*_{(d,r)}(f^i_\kappa)$.[46] Since $\Gamma^{*h_k}_{(d,r)}(\boldsymbol{\lambda})$ is zero-sum, this implies $v^{*((h_\kappa,t^j,t),i)}_{(d,r)}(\boldsymbol{\lambda}) \leq 0$ for any such t. Observe that we cannot have $\mathbf{T}_{(d,r)}(t^j) \cap \mathcal{E}^*_{(d,r)} = \emptyset$, for, if this equality held, since $\mathbf{T}_{(d,r)}(f^\iota_\kappa) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ for each ι, μ would not be a Nash equilibrium in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ by Lemma 7. So we must have $\mathbf{T}_{(d,r)}(t^j) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$. Because $\mathbf{T}_{(d,r)}(f^\iota_\kappa) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ for each ι, $\mathbf{T}_{(d,r)}(t^j) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$, and $\mathbf{T}_{(d,r)}(f^i_{\kappa+1}) \cap \mathcal{E}^*_{(d,r)} = \emptyset$, Lemma 8 implies that the value of $\Gamma^{*(h_\kappa,t^j,f^i_{\kappa+1})}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero.[47] Since $v^{*((h_\kappa,t^j,t),i)}_{(d,r)}(\boldsymbol{\lambda}) \leq 0$ for any $t \in \mathfrak{F}^*_{(d,r)}(f^i_\kappa)$, we must have, in particular, $v^{*\left((h_\kappa,t^j,f^i_{\kappa+1}),i\right)}_{(d,r)}(\boldsymbol{\lambda}) \leq 0$. This inequality, together with the fact that the value of $\Gamma^{*(h_\kappa,t^j,f^i_{\kappa+1})}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero, gives $v^{*\left((h_\kappa,t^j,f^i_{\kappa+1}),i\right)}_{(d,r)}(\boldsymbol{\lambda}) < 0$. Therefore, since $v^{*\left((h_\kappa,t^A,t^B),i\right)}_{(d,r)}(\boldsymbol{\lambda}) = 0$ and h_κ is reached with positive probability under play of σ in $G^*_{(d,r)}(\boldsymbol{\lambda})$, if player i chooses $f^i_{\kappa+1}$ under play of σ as a first mover in the round that follows h_κ, then σ is not a Nash equilibrium in $G^*_{(d,r)}(\boldsymbol{\lambda})$.

Suppose that player i chooses $f^i_{\kappa+1}$ under play of σ as a second mover in the round that follows h_κ. Since $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\circ_{(d,r)}$ (so $\boldsymbol{\lambda}$ is symmetric) and there exists a Nash equilibrium in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical, at any equilibrium of $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$, in the first round of the game, the first mover, say player ι, must choose some $\tau^\iota \in \mathfrak{F}^*_{(d,r)}(f^\iota_\kappa)$ with $v^{*((h_\kappa,\tau^\iota,\tau),\iota)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$ for all $\tau \in \mathfrak{F}^*_{(d,r)}(f^{-\iota}_\kappa)$, $-\iota \neq \iota$.[48] Since we are assuming that player i chooses $f^i_{\kappa+1}$ under play of σ as a second mover in the round that follows h_κ, $f^j_{\kappa+1}$ must have the properties of τ^ι. We claim that $\mathbf{T}_{(d,r)}(f^j_{\kappa+1}) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$. To show this, we assume $\mathbf{T}_{(d,r)}(f^j_{\kappa+1}) \cap \mathcal{E}^*_{(d,r)} = \emptyset$ and derive a contradiction. Suppose that $\mathbf{T}_{(d,r)}(f^j_{\kappa+1}) \cap \mathcal{E}^*_{(d,r)} = \emptyset$. Recall that (t^A, t^B) represents the history in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ that is generated with probability one under play of μ. Observe that we cannot have $\mathbf{T}_{(d,r)}(t^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$, for, if this equality held, since $\mathbf{T}_{(d,r)}(f^\iota_\kappa) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ for each ι, μ would not be a Nash equilibrium in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ by Lemma 7. So we must have $\mathbf{T}_{(d,r)}(t^i) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$. But then Lemma 8 implies that the value of $\Gamma^{*\left(h_\kappa,f^j_{\kappa+1},t^i\right)}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero, thereby contradicting the

[46] By $v^{*((h_\kappa,t^j,t),j)}_{(d,r)}(\boldsymbol{\lambda})$ we mean $v^{*((h_\kappa,t^j,t),j)}_{(d,r)}(\boldsymbol{\lambda})$ if $j = A$ and $v^{*((h_\kappa,t,t^j),j)}_{(d,r)}(\boldsymbol{\lambda})$ if $j = B$.

[47] By $\Gamma^{*(h_\kappa,t^j,f^i_{\kappa+1})}_{(d,r)}(\boldsymbol{\lambda})$ we mean $\Gamma^{*(h_\kappa,f^i_{\kappa+1},t^j)}_{(d,r)}(\boldsymbol{\lambda})$ if $i = A$ and $\Gamma^{*(h_\kappa,t^j,f^i_{\kappa+1})}_{(d,r)}(\boldsymbol{\lambda})$ if $i = B$.

[48] By $v^{*((h_\kappa,\tau^\iota,\tau),\iota)}_{(d,r)}(\boldsymbol{\lambda})$ we mean $v^{*((h_\kappa,\tau^\iota,\tau),\iota)}_{(d,r)}(\boldsymbol{\lambda})$ if $\iota = A$ and $v^{*((h_\kappa,\tau,\tau^\iota),\iota)}_{(d,r)}(\boldsymbol{\lambda})$ if $\iota = B$.

fact that $f^j_{\kappa+1}$ has the properties of τ^ι. Thus, $\mathbf{T}_{(d,r)}(f^j_{\kappa+1}) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$. In this case, Lemma 8 implies that the value of $\Gamma^{*\left(h_\kappa, f^A_{\kappa+1}, f^B_{\kappa+1}\right)}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero. Since there exists a Nash equilibrium in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical, there exists $(\tau^A, \tau^B) \in \mathfrak{F}^*_{(d,r)}(f^A_\kappa) \times \mathfrak{F}^*_{(d,r)}(f^B_\kappa)$ such that for each ι, $v^{*((h_\kappa, \tau^\iota, \tau), \iota)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$ for all $\tau \in \mathfrak{F}^*_{(d,r)}(f^{-\iota}_\kappa)$, $-\iota \neq \iota$.[49] This, together with the fact that the value of $\Gamma^{*\left(h_\kappa, f^A_{\kappa+1}, f^B_{\kappa+1}\right)}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero, implies that the restriction of σ to $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ is not a Nash equilibrium in $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$. Hence, because h_κ is reached with positive probability under play of σ in $G^*_{(d,r)}(\boldsymbol{\lambda})$, σ cannot be a Nash equilibrium in $G^*_{(d,r)}(\boldsymbol{\lambda})$.

Case 2. $f^A_\kappa \neq f^B_\kappa$ and the value of $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero. Since the value of $G^*_{(d,r)}(\boldsymbol{\lambda})$ is zero and the value of $\Gamma^{*h_\kappa}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero, there must exist some $\ell = 1, ..., \kappa - 1$ (note that $\kappa > 1$) such that the value of $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ is zero and the value of $\Gamma^{*h_{\ell+1}}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero. If the last two announcements in h_ℓ are identical and equal to f, Lemma 4 says that either $\mathbf{T}_{(d,r)}(f) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ or there exists a Nash equilibrium in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical. Since we are assuming that it is not true that $\mathbf{T}_{(d,r)}(f^\iota_k) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ for each ι, in this case there must exist a Nash equilibrium in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical. The same is true if the last two announcements in h_ℓ are not identical, for, in this case, since $\Gamma^{*h_\ell}_{(d,r)}$ has value zero, one can use Lemma 3 to reach that conclusion.

Therefore, in all cases, there is a Nash equilibrium in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ that generates a unique history of announcements in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ whose last two elements are identical. This implies that each player ι can choose an announcement t^ι in $\mathfrak{F}^*_{(d,r)}(f^\iota_\ell)$ such that $v^{*((h_\ell, t^\iota, t), \iota)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$ for all $t \in \mathfrak{F}^*_{(d,r)}(f^{-\iota}_\ell)$, $-\iota \neq \iota$.[50] At any equilibrium of $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$, in the first round of the game, the first mover, say player ι, must choose one such t^ι. And the second mover, say player $-\iota$, must choose, in equilibrium, a strategy that prescribes, for any announcement t of the first mover in the first round of $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$, some τ_t such that $v^{*((h_\ell, \tau_t, t), -\iota)}_{(d,r)}(\boldsymbol{\lambda}) \geq 0$. Hence, for any history of announce-

[49]Here, the analogue of footnote 48 applies.

[50]See footnote 49.

ments $(t_1^A, t_1^B, ..., t_l^A, t_l^B)$ in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$ generated under play of an equilibrium in $\Gamma^{*h_\ell}_{(d,r)}(\boldsymbol{\lambda})$, $\Gamma^{*(h_\ell, t_1^A, t_1^B)}_{(d,r)}(\boldsymbol{\lambda})$ must have value zero. But then, since the value of $\Gamma^{*h_{\ell+1}}_{(d,r)}(\boldsymbol{\lambda})$ is nonzero and h_κ is reached with positive probability under play of σ in $G^*_{(d,r)}(\boldsymbol{\lambda})$, σ cannot be a Nash equilibrium of $G^*_{(d,r)}(\boldsymbol{\lambda})$.

We have shown that if (1) $h = (f_1^A, f_1^B, ..., f_k^A, f_k^B)$ is a history of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ generated under play of σ, (2) it is not true that $\mathbf{T}_{(d,r)}(f_k^\iota) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ for each ι, (3) $(h, f_{k+1}^A, f_{k+1}^B)$ is a super-history of h in $G^*_{(d,r)}(\boldsymbol{\lambda})$ generated under play of σ, and (4) $\mathbf{T}_{(d,r)}(f_{k+1}^i) \cap \mathcal{E}^*_{(d,r)} = \emptyset$ for some i, then σ is not a Nash equilibrium in $G^*_{(d,r)}(\boldsymbol{\lambda})$. Since σ was an arbitrary profile in $G^*_{(d,r)}(\boldsymbol{\lambda})$, we conclude that any history $(t_1^A, t_1^B, ..., t_l^A, t_l^B)$ of announcements in $G^*_{(d,r)}(\boldsymbol{\lambda})$ generated under play of a Nash equilibrium in $G^*_{(d,r)}(\boldsymbol{\lambda})$ must satisfy either $\mathbf{T}_{(d,r)}(t_l^\iota) \subseteq N_\eta\left(\mathcal{E}_{(d,r)}\right)$ for each ι or $\mathbf{T}_{(d,r)}(t_l^\iota) \cap \mathcal{E}^*_{(d,r)} \neq \emptyset$ for each ι. Now Lemma 10 is obtained via Lemma 9. ■

Proof of Theorem 2'. Suppose that $+\infty > \eta > 0$ and $1 > \theta > 0$. Take any $(d,r) \in \mathcal{M}^*_\eta \cap \mathcal{M}^\circ$. By Lemma 1, there exists $G_{(d,r)}(\boldsymbol{\lambda})$ arbitrarily close to $G_{(d,r)}$ such that $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^\bullet_{(d,r)} \cap \boldsymbol{\Lambda}^\circ_{(d,r)}$. By Lemma 10, The tax policy implemented at any Nash equilibrium of $G^*_{(d,r)}(\boldsymbol{\lambda})$ lies in $N_\eta\left(\mathcal{E}_{(d,r)}\right)$. By Lemma 2, $\mathcal{M}\left(\varepsilon_{(\eta,\theta)}, \theta\right) \subseteq \mathcal{M}^*_\eta$ for some $+\infty > \varepsilon_{(\eta,\theta)} > 0$. Moreover, $\mathcal{M}\left(\varepsilon_{(\eta,\theta)}, \theta\right) \cap \mathcal{M}^\circ$ is clearly open and dense in $\mathcal{M}\left(\varepsilon_{(\eta,\theta)}, \theta\right)$. Therefore, because (d,r) was arbitrary in $\mathcal{M}^*_\eta \cap \mathcal{M}^\circ$ and $G_{(d,r)}(\boldsymbol{\lambda})$ was arbitrarily close to $G_{(d,r)}$, the desired result follows. ■

References

[1] Aumann, R. J., and M. Kurz (1977), "Power and taxes," *Econometrica*, 45, 1137-1161.

[2] Austen-Smith, D. (2000), "Redistributing income under proportional representation," *Journal of Political Economy*, 108, 1235-1269.

[3] Austen-Smith, D., and J. S. Banks (2000), *Positive Political Theory I: Collective Preference*, University of Michigan Press.

[4] Austen-Smith, D., and J. S. Banks (2005), *Positive Political Theory II: Strategy and Structure*, University of Michigan Press.

[5] Barry, B. (1970), *Sociologists, Economists, and Democracy*, Chicago: University of Chicago Press.

[6] Benabou, R. (2000), "Unequal societies: income distribution and the social contract," *American Economic Review*, 90, 96-129.

[7] Berliant, M. C., and M. Gouveia (1999), "On the political economy of income taxation," mimeo.

[8] Berliant, M. C., and R. P. Strauss (1996), "On recent expositions of horizontal and vertical equity," *Public Economics Review*, 1, 129-150.

[9] Burman, L. E., and M. A. Saleem (2004), "Income tax statistics for sample families, 2003," *Tax Notes*, January 19, 2004.

[10] Carbonell-Nicolau, O., and E. F. Klor (2003), "Representative democracy and marginal rate progressive income taxation," *Journal of Public Economics*, 87, 2339-2366.

[11] Carbonell-Nicolau, O., and E. A. Ok (2007), "Voting over income taxation," *Journal of Economic Theory*, 134, 249-286.

[12] Carbonell-Nicolau, O. (2008), "A positive theory of income taxation," Rutgers University.

[13] Chen, Y. (2000), "Electoral systems, legislative process, and income taxation," *Journal of Public Economic Theory*, 2, 71-100.

[14] Cukierman, A., and A. H. Meltzer (1991), "A political pheory of progressive taxation." In: Meltzer, A., Cukierman, A., and S. F. Richard (Eds.), *Political Economy*. Oxford University Press, Oxford, pp. 76-108.

[15] De Donder, P., and J. Hindriks (2003), "The politics of progressive income taxation with incentive effects," *Journal of Public Economics*, 87, 2491-2505.

[16] Dekel, E., Jackson, M. O., and A. Wolinsky (2008), "Vote buying: general elections," *Journal of Political Economy*, 116, 351-380.

[17] Dixit, A., and J. Londregan (1998), "Ideology, tactics, and efficiency in redistributive politics," *Quarterly Journal of Economics*, 113, 497-529.

[18] Downs, A. (1957), *An Economic Theory of Democracy*, New York, Harper and Row.

[19] Gouveia, M., and D. Oliver (1996), "Voting over flat taxes in an endowment economy," *Economics Letters*, 50, 251-258.

[20] Hettich, W., and S. L. Winer (1988), "Economic and political foundations of tax structure," *American Economic Review*, 78, 701-712.

[21] Hindriks, J. (2001), "Is there a demand for income tax progressivity?" *Economics Letters*, 73, 43-50.

[22] Kranich, L. (2001), "Altruism and the political economy of income taxation," *Journal of Public Economic Theory*, 3, 455-469.

[23] Laslier, J.-F., and N. Picard (2002), "Distributive politics and electoral competition," *Journal of Economic Theory*, 103, 106-130.

[24] Ledyard, J. O. (2005), "Voting and efficient public good mechanisms," mimeo.

[25] Lindbeck, A., and J. W. Weibull (1987), "Balanced-budget redistribution as the outcome of political competition," *Public Choice*, 52, 273-297.

[26] Lizzeri, A., and N. Persico (2001), "The provision of public goods under alternative electoral systems," *American Economic Review*, 91, 225-239.

[27] Mamer, J. W., and K. E. Schilling (1986), "A zero-sum game with incomplete information and compact action spaces," *Mathematics of Operations Research*, 11, 627-631.

[28] Marhuenda, F., and I. Ortuño-Ortín (1995), "Popular support for progressive income taxation," *Economics Letters*, 48, 319-324.

[29] Marhuenda, F., and I. Ortuño-Ortín (1998), "Income taxation, uncertainty, and stability," *Journal of Public Economics*, 67, 285-300.

[30] Meltzer, A. H., and S. F. Richard (1981), "A rational theory of the size of the government," *Journal of Political Economy*, 89, 914-927.

[31] Miller, A. C., and T. R. Rice (1983), "Discrete approximations of probability distributions," *Management Science*, 29, 352-362.

[32] Mirrlees, J. A. (1971), "An exploration in the theory of optimal income taxation," *Review of Economic Studies*, 38, 175-208.

[33] Myerson, R. B. (1993), "Incentives to cultivate favored minorities under alternative electoral systems," *American Political Science Review*, 87, 856-869.

[34] Roberts, K. W. S. (1977), "Voting over income tax schedules," *Journal of Public Economics*, 8, 329-340.

[35] Roemer, J. E. (1999), "The democratic political economy of progressive income taxation," *Econometrica*, 67, 1-19.

[36] Romer, T. (1975), "Individual welfare, majority voting, and the properties of a linear income tax," *Journal of Public Economics*, 4, 163-185.

[37] Sion, M. (1957), "On general minimax theorems," *Pacific Journal of Mathematics*, 8, 171-176.

[38] Snyder, J. M., and G. H. Kramer (1988), "Fairness, self-interest, and the politics of the progressive income tax," *Journal of Public Economics*, 36, 197-230.

[39] Stigler, G. J. (1970), "Director's law of public income redistribution," *Journal of Law and Economics*, 13, 1-10.