Stockholm School of Economics
Department of Economics
5350 Master's thesis in economics
Spring 2026

All The Love That I Can Get: Discriminatory Service Provision Under Voice and Exit

Maria Bohush (42845) and Kiryl Platonau (42860)

Abstract
Municipal governments routinely provide public services unevenly across demographic groups. While such patterns are often attributed to prejudice, they may also arise from the electoral incentives of office-motivated incumbents. Political bargaining can mediate the discriminatory distribution of a government budget. Following Kiss (2012), we build a game-theoretic model of municipal elections in a polity composed of two demographic groups of differing political preferences, and show how discriminatory provision can be sustained in an equilibrium of a simple retrospective voting game. We then extend our model to allow discriminated voters to exit, and show that the possibility of exit strengthens incumbent's incentives for discrimination.

Keywords: political accountability; voice and exit; residential mobility; municipal service provision; targeted provision; discriminatory provision; retrospective voting; Tiebout sorting; Curley effect; noncooperative game

JEL: C72, D72, H42, H73, R23

Supervisor:	Karl Wärneryd
Date submitted:	18 May, 2026
Date examined:	28 May, 2026
Discussants:	Dávid Miženko (42847) and Petter Hallqvist (42843)
Examiner:	Elena Paltseva

Acknowledgements

We are deeply grateful to Jörgen Weibull for his advice, guidance, and thoughtful comments on this paper. His support helped us clarify both the economic argument and the structure of the thesis.

AI Disclosure

We used AI tools, including ChatGPT and Codex by OpenAI, as well as Claude AI, during the preparation of this thesis. These tools were used to search for thematic literature, with the accuracy of the suggestions verified by reading the relevant literature directly. They were also used to search for mathematical properties of some relevant objects, such as probability distributions; aid in creating LaTeX figures and charts based on verbal descriptions; to improve sentence structure and grammar in selected passages; to perform grammar checks; and to translate handwritten formulas into LaTeX.

AI tools do not serve as a primary source for any information in this paper.

1 Introduction

Municipal governments make decisions that affect citizens in very direct and visible ways. They may decide how local services are distributed, which neighborhoods receive investment, which groups benefit from public programs, and they may also contribute to communities being neglected. In principle, democratic elections on the municipal level should discipline these choices. If an incumbent politician provides poor services or treats citizens unfairly, voters can punish him at the next election. In local politics, citizens may also have another option: they may leave the municipality and move somewhere else. At first glance, both of these mechanisms look like ways in which citizens can hold local governments accountable: an incumbent who governs badly can be voted out, or abandoned. To describe these two responses, we borrow the convenient terminology of Hirschman (1970): voice and exit. In Hirschman's original use, exit refers broadly to leaving or switching away from a declining firm, organization, party, municipality, or state. Voice refers broadly to trying to improve the situation from within, for example by complaining, protesting, voting, organizing, or otherwise expressing dissatisfaction. In this thesis, we use the terms more narrowly. We restrict "voice" to mean electoral discipline through voting, while using "exit" to describe residential mobility out of the municipality. However, in our paper, these two responses are less alike than this parallel suggests, and neither necessarily forces politicians to serve all citizens equally. Voting and leaving bear on an incumbent in very different ways — indeed, Hirschman (1970) already observed that the availability of exit can crowd out the exercise of voice rather than reinforce it. We make this asymmetry precise once the incumbent's incentives have been laid out; for now it is enough to note that residential mobility need not discipline local government in the direction one might expect. In a divided municipality, where citizens belong to different social, ethnic, economic, or political groups, an incumbent may have stronger electoral ties to one group than to another. In that setting, the incumbent may not need to satisfy the whole population in order to remain in office. It may be enough to satisfy a loyal or electorally important group, while neglecting another one. More importantly, biased service provision may itself become a political strategy. By choosing who receives services and who does not, the incumbent may change the future electoral environment in his own favor.

Previous research has already studied several parts of this problem. One line of work studies elections as a mechanism of accountability. It describes how voters may discipline incumbents when politicians care about holding office and when campaign promises are not fully credible. One of the key ideas in this literature branch is that voters can use backward-looking (retrospective) voting rules: they reward politicians who perform well when holding the office and remove those who fail to produce satisfactory results. This helps explain how democratic elections may create incentives for politicians to provide effort, public goods, or good governance.

Another body of literature studies local mobility and residential choice. In local politics, citizens are not only voters; they are also residents who may decide whether to stay or leave. This literature shows that local policy and population composition are connected. Taxes, services, housing conditions, public goods, and redistribution can affect who lives in a municipality. As a result, local political decisions may change not only current welfare, but also the future composition of the electorate.

Finally, there is work that studies divided societies, polarization, and discriminatory politics. This work shows that politicians may benefit from group divisions.

When voters are attached to groups or parties, the incumbent may have an incentive to appeal to one side rather than govern in the common interest. In such settings, societal division can weaken accountability, since citizens of the incumbent's preferred identity may be unwilling to punish the politician for neglecting another group.

These literature blocs, together with their main theoretical mechanisms and contributions, are discussed in greater detail in the literature review section (2). The accountability literature (Ferejohn, 1986; Kiss, 2012) models voice over a fixed electorate, with explicit strategic interaction over how much the incumbent must provide. The mobility literature (Hindriks, 2001; Glaeser and Shleifer, 2005) lets policy reshape the electorate, but represents the electoral response in reduced form, without bargaining between the groups over which provision is contested. This thesis embeds retrospective-threshold bargaining into a single model in which the same incumbent can also reshape his electorate through exit, which lets us measure the marginal effect of exit against an otherwise identical voice-only benchmark. The contribution is one of synthesis and formalization in terms of a game-theoretical model. *The research question of the paper is: when do reelection incentives lead an office-motivated municipal incumbent to provide services to one identity group rather than universally — and does the channel through which citizens respond, voice or exit, make such discrimination more or less likely?*

The thesis takes a game-theoretic approach to answer this question. The problem is modeled as a strategic interaction between an incumbent, and two citizens groups, A and B, which may represent ethnic, linguistic, religious, class-based, or other politically relevant identities. The incumbent is associated with group A, while the challenger is associated with group B. At the same time, citizens also differ in political allegiance: some are partisan supporters of the incumbent or the challenger, while others are independent voters. This means that group identity and political preference are related but not identical. The incumbent therefore does not face an undifferentiated electorate; instead, he deals with two groups with potentially different political compositions and different electoral importance.

The incumbent values reelection and chooses between three types of service provision. No provision means that the incumbent does not provide any services to either group. Universal provision means that services are provided to both groups uniformly. Targeted provision means that services are provided only to group A, the group electorally associated with the incumbent. The incumbent does not necessarily discriminate because he dislikes group B. Still, we assume that the incumbent cannot outwardly discriminate against his own group. As a result, he chooses between "no discrimination" and "discrimination in favor of A". While that is an assymetrical assumption, two reasons compel us to make it. The first is that we are seeking to model ingroup favoritism, and view the choice between universal and ingroup provision as more realistic in that context. The other is that we are following the literature, in particular the model in Glaeser and Shleifer (2005), where a similar assumption is made. Citizens respond either by staying or leaving, or by voting according to whether their group receives enough services. The election then determines whether the incumbent remains in office. [1] Each actor is assumed to act strategically, taking

[1]We abstract from turnout incentives. This is a standard simplification, but it is not fully realistic: the rational-choice literature on turnout emphasizes the paradox of voting, namely that in a large electorate an individual voter has little instrumental probability of being pivotal. See Riker and Ordeshook (1968), Ferejohn and Fiorina (1974), and Feddersen (2004). We therefore assume full turnout throughout and leave endogenous turnout outside the model.

into account how the others will respond. We will also treat each citizen group as a player, that is, as if it could make decisions on behalf of its members. The analysis focuses on equilibrium outcomes: under what conditions does the incumbent provide no services, provide services to both groups, or provide services only to his own group? In line with the literature, the opposition or challenger will not be treated as a player.

The general mechanisms described above are formalized through two mechanisms. The first one is a voice model. In this setup, citizens cannot leave the municipality, so voting is their only instrument to discipline the incumbent. To do that, independent (non-partisan) voters in each group set group-specific retrospective thresholds, t_A and t_B. They support the incumbent only if their own group receives enough service provision. The incumbent knows these thresholds and chooses a service mode and exact quantity of provision accordingly. In this setting, biased provision can arise because the incumbent compares the electoral value of satisfying group A with the electoral value of satisfying group B, taking into account the cost of service provision and the political composition of each group. Then, the incumbent may provide services to the group whose support is pivotal for reelection, while neglecting a group whose votes are unattainable, or unnecessary. An important thing to emphasize is that voice operates as a genuine sanction: the only thing that moves the office-motivated incumbent is the prospect of losing office, so conditional voters discipline him precisely by threatening to withhold the support he needs to win.

The second model is an exit model. In this model, some citizens are mobile. If they receive sufficiently low utility from remaining in the municipality, they may leave before the election. The incumbent understands this and chooses the policy strategically. Universal provision keeps both groups in the municipality, while targeted provision keeps the incumbent's group and may allow the mobile segment of the opposing group to leave. As a result, targeted provision may be attractive because it changes the composition of the electorate (Curley effect, see below). The incumbent may prefer to provide services only to his own group if doing so increases his chance of reelection enough to justify the cost. This is where the voice and exit mechanisms come apart. Exit in the paper does not act exactly as a second disciplining device running in parallel to voting. In the voice model, citizens discipline the incumbent through their electoral response; in the exit model, a resident leaves only because the incumbent's policy makes her utility from remaining in the municipality fall below the outside option. If the residents who leave belong to a group whose support the incumbent was not relying on (which is precisely the region of the cost function we are interested in), their exit does not punish him. Instead, it changes the composition of the electorate in his favor and may increase his probability of reelection. In a divided municipality, exit can therefore become part of the incumbent's strategic problem: by making staying unattractive for the opposing group, targeted provision may end up doing part of the incumbent's work for him.

The main contribution of the thesis is to isolate how exit changes the discipline that voice would otherwise impose. We first build a voice-only benchmark in which conditional voters in each group use retrospective thresholds to bargain over provision, and characterize when this bargaining sustains universal provision and when it tips into targeted provision to the incumbent's group. We then let the same incumbent reshape his electorate by allowing dissatisfied members of the opposing group to exit, holding everything else fixed; comparing the two settings isolates the marginal effect of exit. Our central finding is that this effect is one-directional: exit, as designed in our model, does not add a second channel of accountability but weakens the first, enlarging the region of the parameter space in which discriminatory provision is sustained. Voice can hold the office-motivated incumbent to account; exit actually aids the incumbent by removing the voters he would otherwise have had to satisfy. We then relate this shift to the assumptions and parameters of the model and to the real-world determinants they map onto.

The rest of the paper is structured as follows. We first discuss the literature foundations of this thesis in Section 2. We then move to Section 3, where we begin constructing our model (the curious reader may choose to start there). In that section (3.2), we begin by restating one of our main foundational reference models, developed by Kiss (2012), and reformulate it in a more formal game-theoretic fashion. We then proceed to introduce and solve our own voice model in Section 3.3. There, we first restrict the incumbent to inducing only universal provision that satisfies voters from both groups. In Section 3.4, we lift this restriction, thereby allowing reverse clientelism toward group B through universal provision: the incumbent is now able to induce universal provision at a level that satisfies only group B, but not group A. Since incumbents typically do not bear the personal costs of local service provision, we introduce taxation and budgeting in Section 3.5. Finally, in Section 3.6, we explore the exit mechanism by allowing citizens to leave the polity before the election if their utility becomes too low. We conclude by discussing the implications of our findings, their connection to the broader literature, and potential limitations.

2 Related Literature

We will begin with presenting and discussing the existing literature body related to this thesis. The question of how reelection incentives shape biased municipal service provision lies at the intersection of residential mobility, electoral accountability, and strategic political behavior. The central idea that motivates this thesis is that municipal services are not only welfare instruments; they can also be electoral instruments. In a divided municipality, an incumbent may use service provision to influence either who remains in the municipality or which voters are electorally decisive. The literature can therefore be organized around two mechanisms that we have defined: voice, where citizens respond to incumbent's policy through voting, and exit, where citizens respond through mobility.

We start with exploring the literature related to the voice mechanism. This set-up poses a following problem: how voters can discipline incumbents when they cannot or do not leave the polity. The relevant starting point here is the literature on electoral accountability. Alesina (1988) challenges the implication that electoral competition necessarily produces policy convergence. In his model, parties care not only about winning office but also about the policy implemented after the election, and campaign platforms are not always fully credible. Rational voters therefore anticipate possible post-election deviations from announced platforms. This is relevant for our paper motivationally, as it highlights why the analysis should focus not only on electoral promises or formal platforms, but also on the incentives politicians face once in office.

Ferejohn (1986) provides the core model of electoral accountability. Author studies how voters can control incumbents when campaign promises are not credible. Standard electoral competition assumes that candidates announce platforms and then implement them, but Ferejohn (1986) argues that this assumption is weak because politicians may deviate once in office. The paper therefore asks whether repeated elections can create an incentive for incumbents to act in voters' interests. The paper builds a dynamic accountability model. Voters do not observe the incumbent's action directly; instead, they observe performance. The voter uses a retrospective rule: reelect the incumbent if performance is above a threshold K, and remove him otherwise. The incumbent then chooses effort by comparing the cost of satisfying the threshold with the future value of staying in office. The optimal incumbent strategy is to exert just enough effort to meet the voter's threshold only when the realized state is favorable enough. The optimal voter threshold depends positively on the value of office to the incumbent. The more valuable reelection is, the more voters can demand from the incumbent. In the uniform special case, the retrospective threshold becomes stationary, and voter welfare rises with the incumbent's value of remaining in office up to a point. For this thesis, Ferejohn (1986) provides the foundation for the voice mechanism, where citizens use voting to reward or punish the incumbent. Importantly, Ferejohn (1986) also proposes the idea that his disciplining result is clearest in the homogeneous-electorate case; when voters are heterogeneous, he shows that individualistic or group-based retrospective criteria can weaken electoral control unless voters use an aggregate performance criterion. Banks and Sundaram (1998) reinforce this accountability logic in a broader agency framework. They study retention as a disciplining device when a long-lived principal decides whether to keep or replace short-lived agents. Their model shows that a threshold-based retention rule can motivate better behavior even without full precommitment.

In the analysis section, we rely heavily on Kiss (2012), who develops a model with partisan and independent voters and uses an accountability mechanism similar to that in Ferejohn (1986). The main interest of the model is politicians engaging in divisive politics even when it does not clearly increase their direct vote share, and even when it may also benefit the opposition. The model has an incumbent, an opponent, partisan voters, and independent voters. Independent voters are important because they can discipline the incumbent by setting a retrospective effort threshold. The incumbent first chooses whether to engage in divisive politics. Divisive politics converts some independent voters into partisans of either side. After that, independent voters set an effort threshold, the incumbent chooses effort, and the election occurs. Independent voters set the reelection threshold so that the incumbent is just indifferent between exerting effort and choosing zero effort. The incumbent then chooses exactly that effort. Divisive politics reduces the share of independent voters and increases the probability for the incumbent to win with zero effort, and the incumbent chooses divisive politics for all admissible values of the model parameters. Divisive politics weakens accountability because it reduces the number of independent voters able to discipline the incumbent. Kiss (2012) also shows that the opposition can benefit from divisive politics because the weakening of accountability can increase the opposition's probability of winning. Thus, both sides may gain from polarization, while voters lose because effort and accountability decline. This paper is relevant to the thesis mainly because it gives a useful voice setup where partisan voters are not the main source of accountability, while neutral or independent voters can condition support on performance.

The natural starting point for exit mechanism is Tiebout (1956), who provides the foundation for thinking about local public goods and residential mobility. Tiebout (1956) argues that the problem of preference revelation for public goods is different at the local level than at the national level. At the national level, citizens cannot easily choose among alternative governments offering different tax-service bundles. At the local level, however, citizens may choose among municipalities. A resident can move to the community whose bundle of taxes, local goods and services best matches his preferences. In this sense, residential mobility becomes a way for citizens to reveal demand for local public goods.

Tiebout (1956) is therefore the conceptual base block for the exit mechanism. In his model, mobility disciplines local governments because municipalities compete for residents, and citizens "vote with their feet" by selecting the community of residence. Under strong assumptions—full mobility, full information, many communities, no employment restrictions, and no interjurisdictional externalities—local public-good provision can approximate a market-like allocation. However, Tiebout (1956) also recognizes that these assumptions are demanding. Mobility is costly, citizens do not have perfect information, and local policies can generate external effects across communities.

Local public finance literature develops this connection between local policy and community composition. Epple et al. (1984) integrate voting and residential choice in a model of local jurisdictions. One of their contributions is to analyze equilibrium when households simultaneously choose where to live in response to local tax-public-good bundles and, once resident, participate in determining those bundles through voting. The paper provides a theoretical mechanism through which municipal policy and the composition of the local electorate are jointly determined.

Fernandez and Rogerson (1996) make a related point in the context of public education. They study communities where individuals sort by income, local education is financed by local taxation, and tax rates are chosen politically. Their conclusion is that policies affecting the attractiveness of a community can change its income composition, tax base, public-service quality, and welfare. The relevance of Fernandez and Rogerson (1996) for our paper is motivational: it shows that local public-service policy can affect sorting across communities and that changes in community composition can have important welfare effects.

Hindriks (2001) is more directly useful for the exit side because it shows how mobility can create political inefficiency. The paper shows that when rich and poor residents are mobile to different degrees and policy is chosen by majority rule, a jurisdiction's policy affects who enters or exits, and this changed population then affects future policy. One interesting result is that a majority may reject a Pareto-improving tax reduction if it fears that the policy would attract new residents and shift the political majority. This is highly relevant for the thesis because it shows that mobility can generate political failures: it is possible to rationally choose inefficient policies because one anticipates how policy affects the future electorate.

The exit logic is also directly motivated by Glaeser and Shleifer (2005). Their Curley-effect model studies an incumbent who uses inefficient redistribution, hostile rhetoric, or biased public-service provision to encourage an opposing group to leave the jurisdiction. They argue that the Curley effect challenges two standard assumptions about government: that forward-looking leaders are less likely to adopt socially harmful policies, and that voter or resource mobility disciplines bad local policy. In their account, longer political time horizons and stronger mobility responses can instead make harmful incumbent strategies more attractive. Specifically, the authors explore the situation where incumbent would adopt inefficient, wealth-reducing policies that drive some residents away. They say that these policies can be electorally rational if they reshape the electorate in the incumbent's favor. The paper studies a two-group jurisdiction where the incumbent belongs to, or has an electoral advantage with, one group. He can impose redistribution or service bias against the other group. The policy is inefficient because redistribution wastes resources, but it may still increase the incumbent's expected vote share if members of the opposing group leave. The authors explicitly note that the tax or redistribution variable that they use can also be interpreted as biased public-service provision toward the favored group. The incumbent is more likely to choose discriminatory redistribution when group-based voting is strong, when redistribution is less wasteful, when policy has a large effect on migration, and when the challenger is likely to come from the opposing group. In other words, the policy is chosen when the electoral benefit of changing the electorate exceeds the social and electoral costs of redistribution. Glaeser and Shleifer (2005) may therefore be interpreted as a reverse of the optimistic Tiebout logic. In Tiebout (1956), mobility helps citizens find better-fitting municipalities and can discipline local governments. In Glaeser and Shleifer (2005), mobility can make discriminatory policy more attractive because the incumbent benefits when political opponents exit. It is also worth noting that we treat this paper as a motivating contribution rather than as direct evidence for our proposed mechanism.

Finally, we should discuss a strand of literature that approaches a similar research question in a way different from our paper. Firstly, Lindbeck and Weibull (1987) shows that in an equilibrium where candidates maximize the probability of a plurality, the expected value of the party bias within a voting group can drive down the consumption of that group. This answers one of the questions that would become extremely salient as we develop our model. Dixit and Londregan (1995) state a similar result explicitly, defining a "greed coefficient" for agents in the model and showing that being relatively indifferent with regard to ideological closeness while having strong preferences over material provisions helps a group get more material provisions sent their way.

Taken together, these papers provide the building blocks for the thesis. Tiebout (1956) provides the benchmark idea of exit through local mobility. Epple et al. (1984), Fernandez and Rogerson (1996), and Hindriks (2001) show that mobility and local political outcomes are jointly determined, with Hindriks (2001) especially showing that endogenous electorates can generate political inefficiency. Glaeser and Shleifer (2005) then provide the direct exit mechanism: an incumbent may use biased policy to encourage an opposing group to leave. On the voice side, Ferejohn (1986) provides the retrospective accountability logic, while Kiss (2012) provides a partisan/independent voter structure that can be adapted to service thresholds.

3 The model

3.1 Overview of the Game Variants

We are now ready to proceed to the analytical part of the paper. In this section, we present the game-theoretic framework used throughout the thesis. We first introduce the setup, players, timing, strategies, and payoff structures of the games under consideration. We then define the relevant solution concepts and derive a number of propositions characterizing the equilibrium outcomes.

We begin with Kiss (2012), because it provides the foundation for our framework.

The main idea there is that voters are allowed to choose strategies of the form "vote for the incumbent unless service provision is less than some threshold T". This allows us to model citizens having "voice" in a very parsimonious way. We restate the setup in Kiss (2012) in more formal game-theoretic terms, making explicit its structure as a finite sequential game of political accountability with endogenous divisive politics.

On this basis, we then introduce our voice case, where citizens are not allowed to leave the municipality and therefore discipline the incumbent only through voting. The model follows the accountability logic of Ferejohn (1986): voters do not rely on campaign promises, but instead use retrospective thresholds, meaning that they support the incumbent only if his observed performance is high enough.

We then introduce the restricted voice game, G_{vr}, which is the version closest to Kiss (2012). The restricted voice game is deliberately simple. Citizens cannot exit, so voting is their only instrument of discipline. Universal provision is interpreted restrictively: if the incumbent chooses to provide service to everyone, he must choose enough quantity to satisfy both groups whose thresholds are relevant.

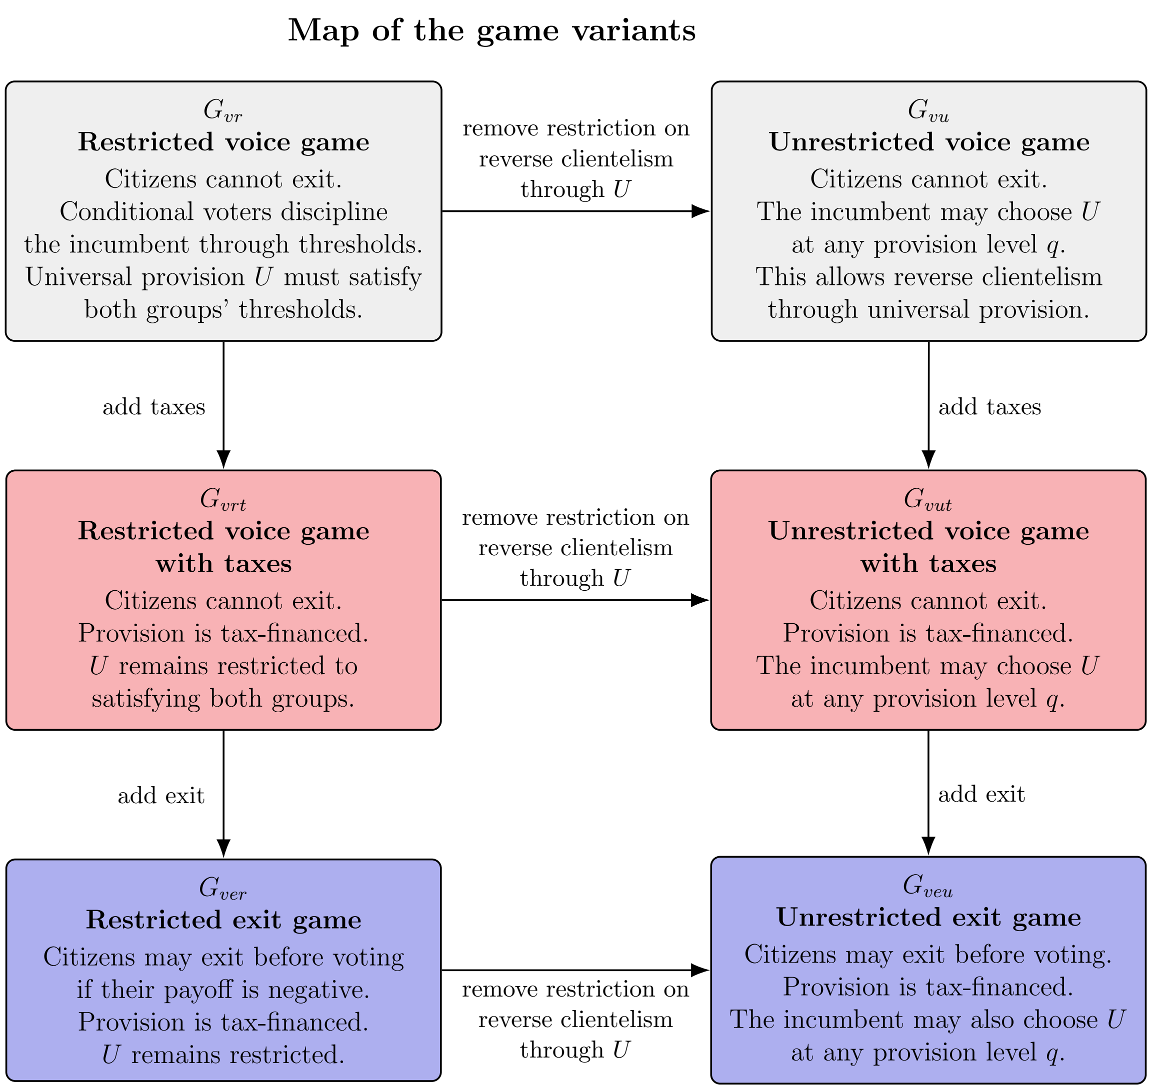

Figure 1: Relationship between the six game variants. We begin with the voice games without taxes, then introduce tax financing, and finally allow for exit. Horizontal extensions remove the restriction on reverse clientelism through universal provision U.

We then relax this restriction and study the problem of reverse clientelism. In the unrestricted voice game, G_{vu}, the incumbent may choose universal provision at a lower level that satisfies only one group. This matters because group B may then be able to attract provision by setting a sufficiently low threshold, even though the incumbent is politically associated with group A.

After we do that, we introduce taxation and budgeting in games G_{vrt} and G_{vut}. This step does not change the logic of the model, but it is necessary to make a clean transit to the exit model: now that services are tax-financed, citizens may be harmed not only by being excluded from provision, but also by having to pay for provision received by others.

Finally, we add exit. In the exit game, G_{ver}, citizens who receive negative utility may leave the municipality before the election. This changes the incumbent's problem, since now he is able to affect the composition of his future electorate, as in Glaeser and Shleifer (2005). The most general version, G_{veu}, combines exit with the unrestricted use of universal provision. Thus, the model moves from the cleanest voice benchmark, closest to Kiss (2012), toward richer versions that allow reverse clientelism, taxation, and finally exit. The figure 1 shows the relations between the game modifications that we propose.

Now we are ready to start with exploring the setup proposed by Kiss (2012).

3.2 The Model in Kiss (2012)

To stay close to the original paper, this subsection uses Kiss's notation: A denotes the incumbent politician and B denotes the opponent. The economy contains a large number of voters and two politicians. The incumbent A can choose costly effort $e \in \mathbb{R}_+$, which should be interpreted as policy effort in the common interest. Voters observe this effort before the election. Incumbent is office-motivated: the winner of the election receives office rent $R > 0$, while the incumbent also bears the cost of effort.

There are three types of voters $\theta \in \{A, 0, B\}$. Some voters are partisan of A, some are partisan of B, others are independent voters 0. Partisan voters always vote for their preferred politician. By contrast, independent voters do not care directly which politician wins. They care about the incumbent's effort, and therefore are able to discipline the politician by setting a retrospective voting rule based on the provision that they receive from the incumbent.

At the beginning of the game, the incumbent chooses whether to engage in divisive politics. Let

$$D \in \{0, 1\},$$

where $D = 1$ means divisive politics and $D = 0$ means no divisive politics. When incumbent chooses $D = 0$, the expected partisan shares remain at their initial levels, namely $\bar{s}_A = \bar{s}_B = b$. If he does use divisive politics, a part Δ of independent voters is converted into partisan voters. Of these newly partisan voters, a fraction $\lambda\Delta$ becomes partisan for A, while the remaining fraction $(1 - \lambda)\Delta$ becomes partisan for B. Hence, after divisive politics, the expected partisan shares become

$$\bar{s}_A = b + \lambda\Delta, \qquad \bar{s}_B = b + (1 - \lambda)\Delta.$$

Divisive politics therefore reduces the share of independent voters and increases the share of partisan voters on both sides. At the election stage, however, the realized partisan shares are subject to uncertainty. In particular, a zero-mean random shock $\varepsilon \in [-1, 1]$ shifts voters between the two partisan camps. As a result, the realized partisan shares are given by $s_A = \bar{s}_A - \varepsilon$ and $s_B = \bar{s}_B + \varepsilon$. The figure 2 presents the population split in the given model. Therefore, the positive realization of ε is favorable for B, and negative - for A.

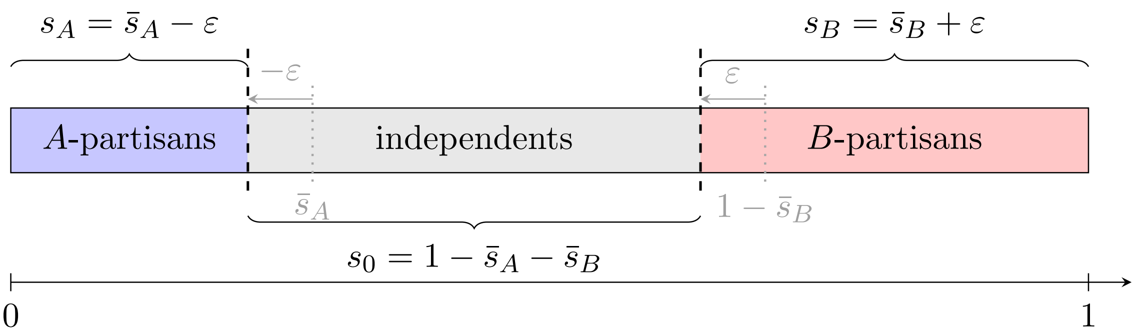

Figure 2: Population distribution before and after the shock in the Kiss (2012) model.

After observing D, the independent voters choose a simple retrospective voting rule, which is described by the effort threshold $\bar{e}$. They announce that they will vote for the incumbent if and only if his effort is at least as large as their decided threshold. If the incumbent chooses effort level below the announced threshold, the independent type votes for B. The incumbent observes the announced threshold and then chooses effort, and this choice is publicly observed. He also cannot deviate from the chosen effort later.
An election is then held. Each voter casts exactly one vote, either for A or for B, and the winner is determined by majority. After the election and payoffs are realized. The timeline is visualized on the figure 3.

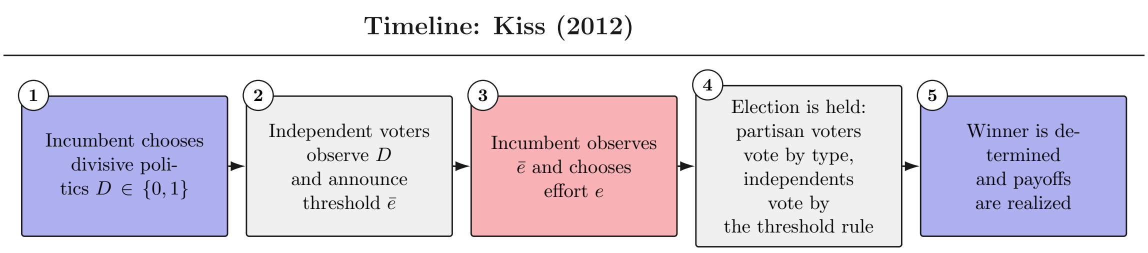

Figure 3: The timeline of the Kiss (2012) game. Each cell corresponds to a stage.

Given the setup above, this game can be described more formally.

Players. The payoff-maximizing strategic players are the incumbent A and the bloc of independent voters $\mathcal{I}$. The extensive-form game also includes moves by nature, which draws the shock ε. The opponent B and the partisan voters have no strategic choices in the game.

Strategies. Incumbent A moves twice. First, at the initial node, incumbent chooses D (divisive politics or not), and then, after observing the independents' announcement chooses effort $e \in \mathbb{R}_+$. A pure strategy of the incumbent is

$$\sigma_A = (D, e_0, e_1),$$

where $D \in \{0, 1\}$ is the divisive-politics choice and, for each $d \in \{0, 1\}$, $e_d : \mathbb{R}_+ \to \mathbb{R}_+$ specifies the effort chosen after any announced threshold $\bar{e}$. As for the independents, after observing D, they collectively choose and announce a retrospective voting rule $\bar{e} \in \mathbb{R}_+$. They vote for A iff $\bar{e} \leq e$, and vote for B otherwise. A pure strategy of the independent voters is then

$$\sigma_{\mathcal{I}} = (\bar{e}_0, \bar{e}_1),$$

where $\bar{e}_D \in \mathbb{R}_+$ is the announced retrospective threshold.

Payoffs. If A is re-elected, he gains the value R from holding the office, and gains nothing in case of a loss. Additionally, he incurs cost of the effort e he decided to exert. The incumbent's payoff is

$$U_A = R\mathbf{1}_{\{A \text{ wins}\}} - e.$$

The independent bloc's payoff is

$$U_{\mathcal{I}} = e.$$

This does not mean that their strategy is irrelevant to the payoffs. Their ex-post payoff depends only on effort, but their announced threshold affects how much effort the incumbent is willing to choose in order to obtain their votes.

Now that we have explored the original setup proposed by Kiss (2012), we are ready to introduce our first model variation: the restricted case of the voice game.

3.3 The Restricted Voice Game

As in Kiss (2012), our polity is inhabited by a very large, finite, number N of citizen-voters. We assume that the entire voting population always votes and that the actual game that plays out merely determines who gets which votes. As far as voting, the voters have two options, that is, voting for the incumbent and voting for the abstract "Challenger", whom we do not model beyond stating that there is, in fact, such an option. The incumbent needs to obtain more than a half of all votes in order to get reelected, although our specification doesn't significantly change when the threshold does.

Since we are looking to deal with group-based discrimination, it is needed to properly present the demographics of our fictional polity. Citizens differ along two dimensions. First, they have a political identity $p \in P$: a citizen can be partisan for the incumbent, partisan for the challenger, or independent. In a significant departure from Kiss (2012), any voter-citizen now also has a group identity $g \in \{A, B\} = G$. These can be viewed as ethnolinguistic, religious, or class-based groups. The group identity is related to, but explicitly made distinct from, political preferences.

An important fact about this setup is that the incumbent knows the group demographics exactly. This is realistic, because demographic data about fixed and observable characteristics (race, ethnicity, religion) are usually extremely precise. Furthermore, we follow Kiss (2012) to introduce uncertainty in the composition of the electorate: in the original model, there is a shock that affects the number of voters who belong to each political camp. We extend this structure to fit the model with two distinct identity groups, each with its own political split. For mathematical reasons, we use two shocks instead of one to capture the fact that the incumbent, along with everyone else, is unaware of the exact political preferences of people in the groups. Up to $\lambda_A N_A$ voters in group A can actually end up having been supporters of the challenger all along, and up to $\lambda_B N_B$ voters in group B can end up having been for "our guy". Equivalently, and perhaps more intuitively, the shock can be interpreted as changes in partisan alignment before the elections: some supporters of party A become independents, while a similarly sized share of previously neutral voters aligns with party B. The same happens for the share of B-partisans.

We can now describe the composition of the electorate before and after the shocks are realized. Initially, the population is divided into four group-political subdemographics. In group A, voters are either independents or partisans of the incumbent. In group B, voters are either independents or partisans of the challenger. After the shocks are realized, this need no longer be true. Some voters in group A may switch to the challenger-partisan category, and some voters in group B may switch to the incumbent-partisan category. As such, the realized voting population may contain six subdemographics, corresponding to the elements of the product set $G \times P$. The figure 4 showcases the population composition in our model.

We also believe that shocks of this kind are empirically plausible and provide a useful way to capture real-world changes in electoral alignment. For example, Fieldhouse et al. (2020) argue that electoral shocks have become more consequential as partisan attachments have weakened over time. In particular, increasing partisan dealignment has reduced strong party identities, making voters more volatile and more willing to switch political alignment. The reason why we choose the stated supports for the shocks in our model is to simplify the mathematics by ruling out cases in which the incumbent loses more votes than he initially has.

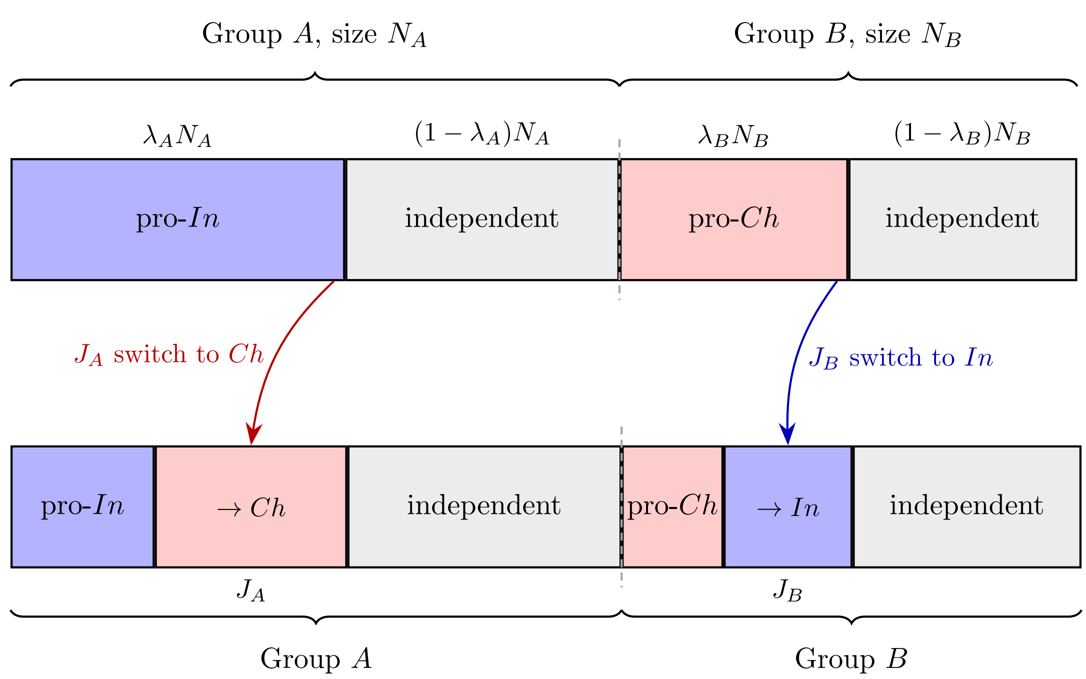

Figure 4: Assumed and realized population distribution in the voice model.

3.3.1 Timeline

We can now turn to the game that we use for the voice model. The logic is close to Kiss (2012): conditional voters announce retrospective thresholds, the incumbent observes these thresholds, chooses policy, and an election is then held. However, the order of moves is not identical to the original model.

In Kiss (2012), the incumbent first chooses whether to engage in divisive politics, independent voters then announce a retrospective effort threshold, and the incumbent finally chooses effort. In our model, instead, the first movers are the conditional voters in groups A and B, who announce group-specific thresholds t_A and t_B. These thresholds describe how much service provision each group requires in order to support the incumbent.

After observing (t_A, t_B), instead of divisive politics choice, the incumbent chooses a service-provision policy. This choice has two components, chosen simultaneously. First, the incumbent chooses a service mode

$$M \in \{U, T, N\},$$

where N is the "do nothing" mode, U is universal provision of services to both groups, and T is targeted provision to group A only. Second, the incumbent chooses the aggregate quantity of provision Q. The quantity Q of service is distributed equally between citizens who are chosen to receive it. We will further discuss the features of the provision in our model in the 3.3.3 subsection.

Once the incumbent has chosen (M, Q), the election is held. Partisan voters vote according to their political allegiance. Conditional voters support the incumbent if and only if the provision received by their group meets their announced threshold. The winner is then determined by majority rule, and payoffs are realized.

The timeline of the voice game is shown in Figure 5. One technical point concerns the first stage. Both groups announce thresholds before the incumbent moves, but we still need to specify the exact order in which t_A and t_B are chosen. In the 3.3.6 section, we showcase the impact of making different assumptions on the order of the moves on the game outcomes.

It is equally important to be explicit about when the identity shock is realized and what each player knows when it moves. Nature draws the pair (J_A, J_B) only at the election stage, after both groups have announced their thresholds and after the incumbent has chosen (M, Q); this is the order shown in the extensive form of Figure 6, where Nature moves last. No player observes the realisation of (J_A, J_B) before acting, so the two blocs and the incumbent all move under the same prior over the shock, summarized by the distribution F_Δ of the composite term $\Delta = J_B - J_A$. In particular, the conditional voters do not know, when they announce t_A and t_B, how the shock will reclassify partisan labels in their group; they choose the threshold that maximises their group's expected provision, just as the incumbent chooses (M, Q) to maximise an expected payoff in which the shock enters only through (F_Δ). The game is thus one of perfect information with respect to the strategic moves and of symmetric uncertainty with respect to Nature.

We are now ready to define the game in its normal and extensive forms. Like any normal-form game, ours is defined by a triplet:

$$G_v = (N_v,\ S_v,\ U_v).$$

We now go on to showcase each of its elements.

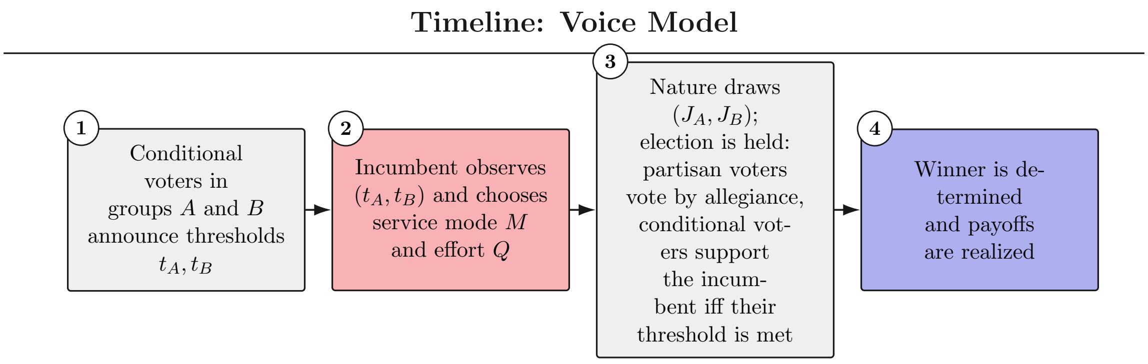

Figure 5: The timeline of the "voice" game. Each cell corresponds to a stage.

3.3.2 Players

Our game now has three players. We continue to assume that voters within each group make collective decisions; however, these decisions are now made separately by independents in groups A and B. We denote the independents in identity group A by $\mathcal{I}_A$, and those in group B by $\mathcal{I}_B$. Our third player is the incumbent, denoted by In. This differs from the original notation in Kiss (2012) and is introduced here to avoid confusion with identity group A. We write:

$$N_v = \{In, \mathcal{I}_A, \mathcal{I}_B\}.$$

It is worth noting that the identity shocks drawn by Nature leave the independent voters of both groups intact, reshuffling only some partisan alignments. For the model, this means that the composition of each "player" bloc is fixed from the outset, which lends additional plausibility to the assumption that its members coordinate on a single decision.

3.3.3 Strategies and actions

Since this is a sequential game with perfect information, there is a difference between action sets and strategy sets. We first showcase the former. Both of the independent blocs have the positive part of the real line $\mathbb{R}^+$ as their action sets. The incumbent's action set is the union-product set:

$$\Gamma = (\mathbb{R}^+ \times \{\ U,\ T\})\ \cup (0, N),$$

where $N, T, U\ \in\ M$. The service discussed in this game is modeled as a private good rather than a pure public good: one voter's utility depends on the per-capita amount of service she receives, q, rather than on a municipality-wide service level that is consumed by all residents. Thus, if the same aggregate effort Q is distributed among fewer voters, each voter receives a higher per-capita level of service.

We now look at the strategies. In a sequential game with perfect information, each player has to choose a function from the set of possible histories where they get the move, to the action set. The first player is whichever voting bloc gets to state their threshold first. All of the histories where they get the move are the same, and so they just need to choose a number on the positive real line $\mathbb{R}^+$.

The second player is the second voting bloc, whether A or B. They always get the move, and any history up to that point is uniquely characterized by just one number, that is the first group's threshold t_1. The strategy set of the second mover is the set of all functions :

$$f_2 : \mathbb{R}^+ \to \mathbb{R}^+.$$

Finally, the incumbent observes both thresholds. As usual, his strategy set needs to correspond to the set of all functions from the possible histories where the move is his, to the action set. It is clear that any possible history is uniquely characterized by an ordered pair (t_1, t_2), and so the strategy set of the incumbent corresponds to the set of all functions of form:

$$f_{In} : \mathbb{R}^2 \to [(\mathbb{R}^+ \times \{\ U,\ T\})\ \cup (0, N)].$$

3.3.4 Payoffs and utilities

We are going to define the payoffs and utilities of our game. These are distinct since we have a sequential game with a random element; a reasonable solution concept would include preferences on the distribution of outcomes. For now, we assume that both the incumbent and the voters are risk-neutral and maximize the expected payoff. The payoffs of the independent voters do not depend on the outcome of the election, so their payoff is a direct consequence of the incumbent choice.

Let us look at the payoffs of the independent voters first. In line with the original model in Kiss (2012), we view independent voters for having preferences that monotonically increase with the provision of the good. The exact shape doesn't matter under the typical utility assumptions. Denote the strategy of the incumbent as s_{In}. The elements of the tuple s_{In} are s^1_{In} (strategy on the choice of provision mode) and s^2_{In} (strategy on the choice of provision quantity). The utility V of A's independent voters are:

$$V_{\mathcal{I}A} = u(q_A) = q_A,$$

where $q_A = q$ if $s^1_{In} \in \{T, U\}$ and 0 otherwise. In the same way, the utility-payoff of the playing voters in B are:

$$V_{\mathcal{I}B} = u(q_B) = q_B,$$

where $q_B = q$ if $s^1_{In} \in \{U\}$ and 0 otherwise. Finally, let us look at the incumbent's payoff. The incumbent is still eager to get the office, which he values at R. The aggregate quantity of the service Q costs him according to a cost function $c(Q)$ such that $c_1(Q) > 0$ and $c_2(Q) \geq 0\ \forall\ Q \in \mathbb{R}^+$. As such, the payoff function is

$$B_{In}(Q) = \begin{cases} R - c(Q) \text{ if In wins;} \\ -c(Q) \text{ otherwise.} \end{cases}$$

Denote the probability of an incumbent victory as $\mathbb{P}_{In}$ Under incumbent risk-neutrality, the expected utility is then plainly:

$$V_{In}(Q) = \mathbb{E}\ [B_{In}(Q)] = \mathbb{P}_{In}\ [R] - c(Q).$$

3.3.5 Extensive form diagram

We are now ready to draw the extensive form tree of the game - it is presented in the figure 6. $B(Q)$ denotes the realized electoral payoff after the incumbent decides on a Q and Nature draws a pair of random shocks (J_A, J_B). Nature is placed last on purpose: (J_A, J_B) is realized only after every strategic choice has been made, so neither bloc nor the incumbent conditions its move on the realized shock.

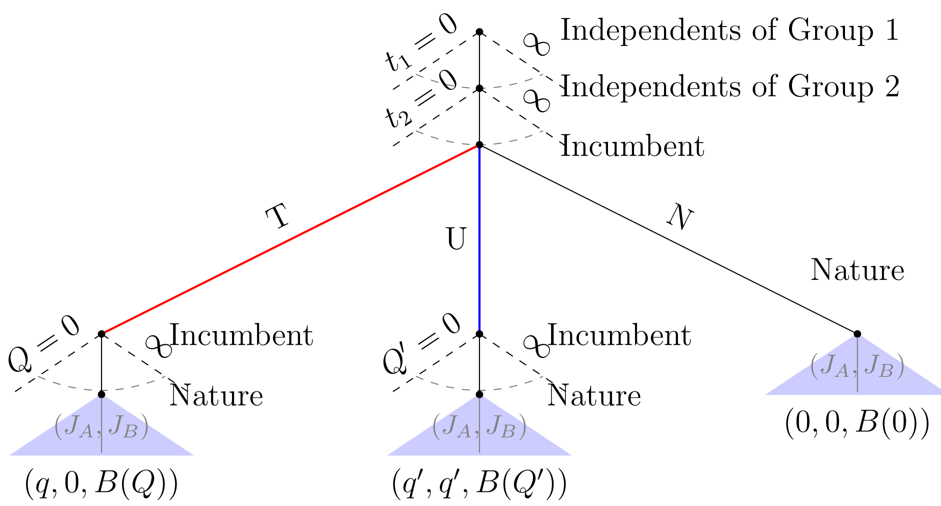

Figure 6: The extensive form tree of the voice game.

3.3.6 Solving the Restricted Voice Game

We can now proceed to solving the restricted voice game by backward induction. The strategy profiles constructed in this way will be subgame-perfect, and therefore also Nash equilibria. Let us once again make explicit what is being restricted. Once the thresholds t_A and t_B have been announced, in whichever order, the incumbent can choose among three relevant modes: no provision N, targeted provision to group A, denoted T_A, and universal provision U. In the restricted game, universal provision means provision at a level high enough to satisfy both independent blocs. Thus, if the incumbent chooses U, he must provide

$$q_A = q_B = \max\{t_A, t_B\},$$

so that aggregate provision is

$$Q_U^* = (N_A + N_B)\max\{t_A, t_B\}.$$

If he chooses targeted provision to group A, he provides exactly

$$q_A = t_A, \qquad q_B = 0,$$

so that aggregate provision is

$$Q_T^* = N_A t_A.$$

Finally, if he chooses no provision,

$$Q_N^* = 0.$$

Two points are worth spelling out because they define the restricted model. First, it does not make sense for the incumbent to choose a provision level above the relevant threshold: doing so does not win any additional votes, but it does increase costs. Second, the incumbent is not allowed here to choose universal provision at the lower of the two thresholds when doing so would satisfy only one independent bloc. That is precisely the restriction that is relaxed in the unrestricted voice game.

This observation lets us get rid of the continuum branches in the incumbent's last move. For each mode, the relevant aggregate quantity is pinned down by the thresholds. The incumbent's continuation problem is therefore a comparison among three values. For a given strategy s, write

$$V_{\text{In}}(s) = E[W(s)] - c(Q_s^*),$$

where $W(s)$ is the realized office payoff and Q_s^* is the aggregate provision associated with the chosen mode.

Let us now turn to the random nature of the game. Unwrapping the expectation operator is necessary if we want to understand what the solutions look like. Remember that we defined a pair of voter-composition shocks J_A and J_B. The only distributional object we need for the solution is their difference.

Lemma 1. *Let J_A and J_B be a pair of independent random variables with supports $[0, K_A]$ and $[0, K_B]$, respectively. Let f_A and f_B denote their density functions. The difference $\Delta = J_B - J_A$ is a random variable with support $[-K_A, K_B]$, and its density is given by the convolution $f_\Delta = f_B * f_{-A}$, where f_{-A} is the density of $-J_A$, so that*

$$f_{-A}(-x) = f_A(x) \qquad \forall x \in [0, K_A].$$

Proof. Since $\Delta = J_B + (-J_A)$, and J_B and $-J_A$ are independent, the density of Δ is the convolution of the two corresponding densities. The support of J_B is $[0, K_B]$, while the support of $-J_A$ is $[-K_A, 0]$. Hence the support of their sum is $[-K_A, K_B]$. □

Lemma 1 allows us to work with a single composite shock, $\Delta = J_B - J_A$, rather than with the two shocks separately. Let f_Δ denote its density and let F_Δ denote its cumulative distribution function. If, for example, f_A and f_B are continuous, then the convolution density f_Δ is well behaved, and so is the associated CDF F_Δ. This is enough for our purposes, where the incumbent's winning probability will be written directly in terms of F_Δ.

Let $\mathcal{M}(s)$ denote the number of votes the incumbent receives before the composite shock is realized. Then

$$\mathcal{M}(s) = \lambda_A N_A + (1-\lambda_A) N_A \mathbf{1}\{q_A(s) \geq t_A\} + (1-\lambda_B) N_B \mathbf{1}\{q_B(s) \geq t_B\}.$$

The realized incumbent vote total is $\mathcal{M}(s) + \Delta$. Hence the probability of incumbent victory is

$$\mathbb{P}[\textit{In wins} \mid s] = \Pr\left[\mathcal{M}(s) + \Delta > \frac{N_A + N_B}{2}\right],$$

or, equivalently,

$$\mathbb{P}[\textit{In wins} \mid s] = 1 - F_\Delta\left(\frac{N_A + N_B}{2} - \mathcal{M}(s)\right).$$

Since the incumbent's relevant choice set has collapsed to N, T, and U, we can define three gross electoral values. Let

$$G_N = R\left[1 - F_\Delta\left(\frac{N_A + N_B}{2} - \lambda_A N_A\right)\right],$$

$$G_T = R\left[1 - F_\Delta\left(\frac{N_A + N_B}{2} - N_A\right)\right],$$

and

$$G_U = R\left[1 - F_\Delta\left(\frac{N_A + N_B}{2} - N_A - (1 - \lambda_B) N_B\right)\right].$$

These are the incumbent's gross gains from no provision, targeted provision to group A, and universal provision, respectively. The corresponding net payoffs (that we will later define as "bid functions") are

$$V_{\text{In}}(N) = G_N,$$

$$V_{\text{In}}(T_A \mid t_A) = G_T - c(N_A t_A),$$

and

$$V_{\text{In}}(U \mid t_A, t_B) = G_U - c\left((N_A + N_B)\max\{t_A, t_B\}\right).$$

The incumbent's best-response correspondence is therefore

$$BR_{\text{In}}(t_A, t_B) = \arg\max_{m \in \{N,T,U\}} V_{\text{In}}(m \mid t_A, t_B).$$

We assume throughout the thesis that, when indifferent, the incumbent chooses the option with provision to larger number of citizens. Thus, he chooses U over both T and N, and T over N.

Before we start with subgame perfection, let us show why we need it.

Proposition 1 (Degenerate Nash equilibrium)**.** *Consider an instance G^*_{vr} of the restricted voice game. For any $R > 0$, there exists a number $\bar{t} > 0$ such that any strategy profile in which both groups set thresholds at least $\bar{t}$ at all relevant histories and the incumbent always chooses N is a Nash equilibrium.*

Proof. Choose $\bar{t}$ such that

$$c(N_A \bar{t}) > R.$$

Then targeted provision to group A at any threshold $t_A \geq \bar{t}$ costs more than the value of office. Universal provision is at least as costly, since it requires

$$(N_A + N_B)\max\{t_A, t_B\} \geq N_A \bar{t}.$$

Thus, even if provision guaranteed victory with probability one, the incumbent would not gain by deviating from N.

Given that the incumbent always chooses N, neither group receives provision, regardless of the threshold it sets. Hence neither group can profitably deviate either. The strategy profile is therefore a Nash equilibrium. It is not generally subgame perfect, since the incumbent's promise to choose N after lower off-path thresholds need not be sequentially rational. □

Proposition 1 motivates our need for subgame perfection. We want every player to make time-consistent moves, rather than relying on non-credible off-path behavior.

For finding subgame-perfect equilibria, we use a particular device. The voters in this game can be thought of as bidding for the incumbent's choice. Group A's bid depends on whether it can make targeted provision valuable enough for the incumbent. Group B, in the restricted game, does not have its own targeted provision option. It can only work through universal provision. Finally, neither group wants to set its threshold so high that the incumbent gives up on provision altogether and reverts to his outside option, namely playing N and banking on his partisans and a favorable shock.

Consider the bid functions:

$$B_A(t) = G_T - c(N_A t),$$

and

$$B_U(t) = G_U - c((N_A + N_B)t).$$

The first bid function is the incumbent's net payoff from targeted provision to group A at per-capita threshold t. The second is the incumbent's net payoff from universal provision at common per-capita threshold t. The no-provision payoff is the horizontal line

$$B_N = G_N$$

.

Notice that t appears only in the cost term. Hence both bid functions are downward sloping. Their derivatives are

$$B'_A(t) = -N_A c'(N_A t),$$

and

$$B'_U(t) = -(N_A + N_B)c'((N_A + N_B)t).$$

Universal provision starts from the higher gross electoral value, since $G_U > G_T$, but it also falls faster in t, because the same per-capita threshold must be provided to the whole population. The relevant question is therefore not simply which curve starts higher, but which curve can stay above the no-provision payoff for longer.

Define k_A as the point where the targeted bid function reaches the no-provision payoff:

$$B_A(k_A) = G_N.$$

Similarly, define k_U as the point where the universal bid function reaches the no-provision payoff:

$$B_U(k_U) = G_N.$$

Proposition 2 (Thresholds on thresholds)**.** *Consider an instance G^*_{vr} of the restricted voice game. Assume that F_Δ is strictly increasing on the interval containing*

$$\frac{N_A + N_B}{2} - N_A - (1 - \lambda_B)N_B, \qquad \frac{N_A + N_B}{2} - N_A, \qquad \frac{N_A + N_B}{2} - \lambda_A N_A,$$

and that the cost function $c(Q)$ is continuous, strictly increasing, and weakly convex. For a given $R > 0$, group A can find a threshold value t_A that ensures

$$V_{\text{In}}(T_A \mid t_A) \geq V_{\text{In}}(N).$$

The largest such threshold is

$$k_A = \frac{1}{N_A} c^{-1}(G_T - G_N).$$

Proof. Since F_Δ is strictly increasing on the relevant interval, and since

$$\frac{N_A + N_B}{2} - N_A - (1 - \lambda_B)N_B < \frac{N_A + N_B}{2} - N_A < \frac{N_A + N_B}{2} - \lambda_A N_A,$$

we have

$$G_U > G_T > G_N.$$

In particular,

$$G_T - G_N > 0.$$

Group A can therefore choose a threshold satisfying

$$c(N_A t_A) \leq G_T - G_N.$$

For any such threshold,

$$V_{\text{In}}(T_A \mid t_A) = G_T - c(N_A t_A) \geq G_N = V_{\text{In}}(N).$$

The largest threshold that still makes targeted provision weakly optimal is obtained by imposing equality:

$$G_T - c(N_A k_A) = G_N.$$

Equivalently,

$$c(N_A k_A) = G_T - G_N.$$

Applying c^{-1} gives

$$N_A k_A = c^{-1}(G_T - G_N),$$

and therefore

$$k_A = \frac{1}{N_A} c^{-1}(G_T - G_N).$$

□

This tells us a couple of things. The first is that group A has a genuine targeted outside option in the restricted voice game. It can set a threshold that makes targeted provision weakly attractive to the incumbent. Proposition 2 also identifies the maximal threshold that group A can impose while still making targeted provision weakly preferable to no provision. We now compare this outside option with the maximal threshold that can be sustained under universal provision. We now define the analogous cutoff for universal provision. Similarly to k_A, derive

$$k_U = \frac{1}{N_A + N_B} c^{-1}(G_U - G_N),$$

so that

$$B_U(k_U) = G_N.$$

The comparison between k_A and k_U determines the equilibrium outcome. Consider first the case in which group A's targeted cutoff is higher.

Proposition 3 (Targeted provision to group A)**.** *Consider an instance G^*_{vr} of the restricted voice game. Let $\lambda_A, \lambda_B, N_A, N_B$ be parameters, let F_Δ be a strictly increasing cumulative distribution function of the random shock, and let $c(Q)$ be the cost function. Suppose $k_U < k_A$. Then, for either order of threshold moves, there exists a subgame-perfect equilibrium with $t_A = k_A$, and with group B setting a sufficiently small positive threshold $t_B = \varepsilon \in (0, k_U)$, in which the incumbent chooses targeted provision to group A. The equilibrium service levels are*

$$q_A = k_A, \qquad q_B = 0.$$

Proof. At $t_A = k_A$, targeted provision gives

$$B_A(k_A) = G_T - c(N_A k_A) = G_N.$$

No provision also gives $V_{\text{In}}(N) = G_N$.

Universal provision must satisfy both groups in the restricted game. Hence, no matter what threshold group B sets, universal provision requires at least $q \geq k_A$. Since $k_A > k_U$ and B_U is decreasing, we have

$$B_U(k_A) < B_U(k_U) = G_N.$$

Therefore universal provision is strictly worse than no provision whenever group A sets k_A. Targeted provision is weakly optimal relative to no provision and strictly better than universal provision. By the tie-breaking rule, the incumbent chooses T over N.

Let group A have the second move. After any threshold chosen by group B, group A can set $t_A = k_A$. If group B chose $t_B \leq k_U$, group A's choice of k_A prevents profitable universal provision, because universal provision would have to be made at least at level k_A. If group B chose $t_B > k_U$, universal provision is already below the no-provision payoff. In both cases, the incumbent chooses targeted provision to group A. Group A cannot do better by raising its threshold, because then $B_A(t_A) < G_N$ and targeted provision is no longer weakly optimal. It cannot do better by lowering its threshold, because this either gives it targeted provision at a lower level or allows universal provision at a level no higher than $k_U < k_A$.

Now let group A have the first move. If it sets $t_A = k_A$, group B cannot induce universal provision by lowering its own threshold. Universal provision must still satisfy group A, and therefore must still provide at least k_A per capita, which is too costly for the incumbent. Group B is consequently indifferent among its thresholds; let it set a small positive threshold $\varepsilon \in (0, k_U)$. Group A cannot profitably raise its threshold, because targeted provision would no longer be weakly optimal. It cannot profitably lower its threshold either: any universal provision that group B could induce after such a deviation would be at a level weakly below k_U, while targeted provision would give group A only its lower threshold. Both are below k_A. Thus the described strategies are sequentially rational under either order of moves. □

We see that the restricted case is indeed quite restrictive. If the parameters of the model are not in group B's favor, there is little that the group can do to counteract targeted provision.

It is more or less clear what happens in the opposite case: universal provision can be sustained. We formalize this as well.

Proposition 4 (Universal-provision equilibrium)**.** *Consider an instance* G^*_{vr} *of the restricted voice game. Let* $\lambda_A, \lambda_B, N_A, N_B$ *be parameters, let* F_Δ *be a strictly increasing cumulative distribution function of the random shock, and let* $c(Q)$ *be the cost function. Suppose* $k_U \geq k_A$*. Then, for either order of threshold moves, there exists a subgame-perfect equilibrium with* $t_A = t_B = k_U$*, in which the incumbent chooses universal provision. The equilibrium service levels are* $q_A = q_B = k_U$*.*

Proof. At $t_A = t_B = k_U$, universal provision gives $B_U(k_U) = G_N$. Targeted provision gives $B_A(k_U)$. Since $k_U \geq k_A$ and B_A is strictly decreasing,

$$B_A(k_U) \leq B_A(k_A) = G_N.$$

Thus

$$B_U(k_U) = G_N \geq B_A(k_U).$$

Universal provision is weakly optimal for the incumbent, and the tie-breaking rule selects U over both T and N.

Now let group B move first. If B sets any threshold $t_B \leq k_U$, group A can set $t_A = k_U$. Universal provision at k_U is then chosen, and group B receives k_U. If group B instead sets $t_B > k_U$, universal provision becomes strictly worse than no provision. Group A can then fall back on its targeted cutoff k_A, while group B receives zero. Hence group B does not gain by raising its threshold above k_U, and lowering it below k_U does not increase its provision.

Now let group A move first. If A sets $t_A = k_U$, then group B can set $t_B = k_U$, after which universal provision is chosen. Group B cannot gain by lowering its threshold: since group A's threshold remains k_U, universal provision must still be provided at level k_U if it is chosen. Thus lowering t_B does not increase group B's provision. Group B also cannot gain by raising its threshold above k_U, because then universal provision becomes strictly worse than no provision, and any switch to targeted provision gives group B zero.

Group A also has no profitable deviation. If it raises its threshold above k_U, universal provision becomes too costly, and targeted provision is also worse than no provision because $k_U \geq k_A$. Hence group A receives zero. If it lowers its threshold, it cannot obtain more than k_U. If universal provision is still chosen, the highest relevant universal level is already k_U. If lowering the threshold instead induces targeted provision, group A receives only its lowered threshold. Thus group A cannot improve by deviating from k_U.

Therefore the proposed strategies are sequentially rational under either order of threshold moves, and the profile is a subgame-perfect equilibrium. □

Among other things, we have managed to show that $k_A > k_U$ is a **necessary and sufficient condition** for the existence of a targeted-provision SPE in this game. First of all, if $k_U \geq k_A$ then no self-interested player would play anything but k_U, and the incumbent would oblige. Secondly, if $k_A > k_U$, there needs to be a targeted-provision equilibrium: one of the group A and the incumbent would have to play a completely suboptimal move for the game to end up otherwise. In the future, we are going to use that fact to show that certain modifications of the game make this condition easier to fulfill and thus hurt group B.

Figure 7 provides useful visual intuition for the two cases described above. In panel a), we can clearly see that when $k_A > k_U$, group A has a segment on the bidding curve where it can receive a higher level of provision than under universal provision while still providing the incumbent with a higher payoff. This gives the incumbent an incentive to choose targeted provision. Hence, such cases can only sustain targeted equilibria.

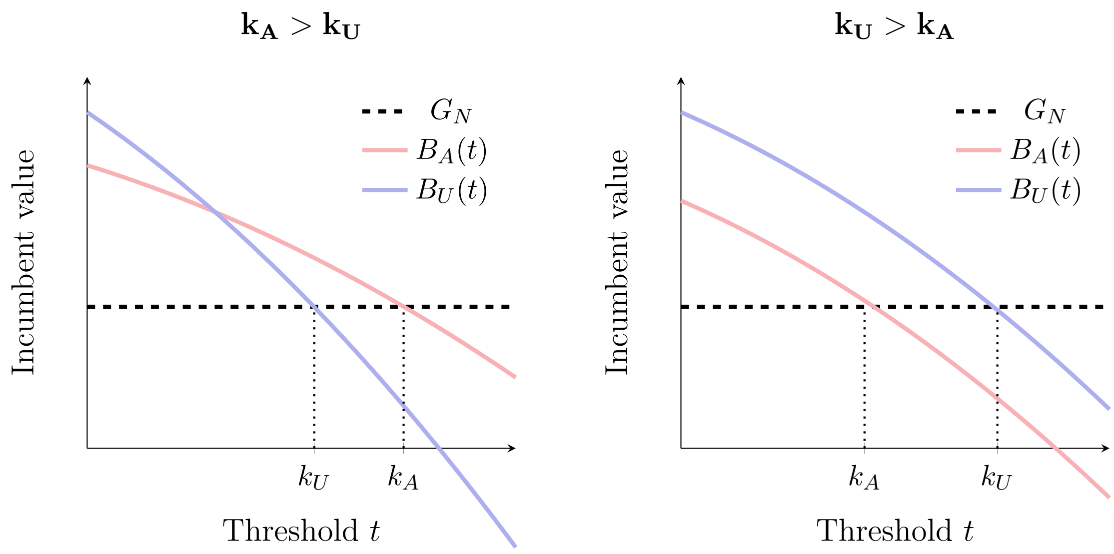

(a) Targeted provision case

(b) Universal provision case

Figure 7: Bidding curves in the restricted voice game. The dashed line represents the incumbent's payoff from no provision, G_N. The cutoffs k_A and k_U are the threshold levels at which targeted and universal provision, respectively, leave the incumbent indifferent between providing and doing nothing.

3.3.7 Restricted model: discussion

The restricted voice game comes down to the comparison between the two threshold values k_A and k_U. If $k_U \geq k_A$, then universal provision can be sustained. If $k_A > k_U$, then group A can use its unconditional option to induce large, targeted provision of benefits. Let us see when that happens.

Substituting the definitions, $k_A > k_U$, means

$$\frac{1}{N_A + N_B} c^{-1}(G_U - G_N) < \frac{1}{N_A} c^{-1}(G_T - G_N).$$

Equivalently,

$$k\left(\frac{N_A + N_B}{N_A} k^{-1}(G_T - G_N)\right) > G_U - G_N.$$

So targeted provision is more likely when N_A is smaller, because it is cheaper to serve group A alone; when $B_A - G_N$ is larger, because targeted provision becomes more electorally valuable; and when $B_U - G_N$ is smaller, because group B does not add much electoral value through universal provision.

Written out in terms of the composite shock, these electoral gains are

$$G_T - G_N = R\left[F_\Delta\left(\frac{N_A + N_B}{2} - \lambda_A N_A\right) - F_\Delta\left(\frac{N_A + N_B}{2} - N_A\right)\right],$$

and

$$G_U - G_N = R\left[F_\Delta\left(\frac{N_A + N_B}{2} - \lambda_A N_A\right) - F_\Delta\left(\frac{N_A + N_B}{2} - N_A - (1 - \lambda_B) N_B\right)\right].$$

Thus, what matters is not only the size of the groups, but also where the relevant vote totals fall in the distribution of the composite shock. Group A is advantaged when satisfying it alone moves the incumbent substantially closer to reelection, while group B is advantaged when adding its independent voters through universal provision substantially increases the incumbent's winning probability. In general, it is important to note that being too partisan reduces the bargaining power of a group. Since the independents are the ones doing the bargaining, the group at large benefits from there being more independents and fewer partisans.

We can also notice some additional comparative statics once we impose more structure on the cost function. The discussion above used only the fact that k is increasing and weakly convex. This is enough to derive the general condition for targeted provision. However, if we assume a more specific functional form, or at least a more specific scaling property of the cost function, the condition becomes easier to interpret. In particular, suppose that the cost function is homogeneous of degree $\eta \geq 1$, meaning that for any $a > 0$ and any $Q \geq 0$,

$$c(aQ) = a^{\eta} c(Q).$$

This assumption captures how fast the cost of provision increases when the same per-capita service level is extended to a larger population. A larger η means that scaling up provision is more costly, and therefore that universal provision becomes relatively less attractive compared to targeted provision.

Proposition 5. *(Homogeneous costs and targeted provision) Suppose that the cost function c is strictly increasing and homogeneous of degree $\eta \geq 1$ and population $N_A + N_B$ is normalized to 1. Then the condition*

$$c\left(\frac{1}{N_A} c^{-1}(B_A - G_N)\right) > B_U - G_N.$$

is equivalent to

$$\frac{1}{N_A^{\eta}}(B_A - G_N) > B_U - G_N.$$

Proof. Starting from the targeted-provision condition,

$$B_U - G_N < c\left(\frac{1}{N_A} c^{-1}(B_A - G_N)\right).$$

Using homogeneity of degree η, we obtain

$$c\left(\frac{1}{N_A} c^{-1}(B_A - G_N)\right) = \left(\frac{1}{N_A}\right)^{\eta} c\big(c^{-1}(B_A - G_N)\big).$$

Now we can see that

$$\left(\frac{1}{N_A}\right)^{\eta} (B_A - G_N) > B_U - G_N.$$

□

The term $\left(\frac{1}{N_A}\right)^{\eta}$ is the cost-scaling penalty from extending a service level that could be targeted to group A to the whole municipality. Since

$$\frac{1}{N_A} > 1,$$

this penalty is increasing in η. Therefore, when the cost function is more convex in this homogeneous sense, targeted provision becomes easier to sustain. Conversely, when η is close to one, scaling provision to the whole population is less costly, and universal provision becomes more attractive.

3.4 The Unrestricted Voice Game

Let us now consider the unrestricted model. The restriction being removed is the prohibition on the incumbent choosing the lesser, rather than the larger, of the two thresholds when playing U. Removing it enables independent bloc of B to make threats. Now, if $t_B < t_A$, the incumbent can choose universal provision at $q = t_B$. Then $\mathcal{I}_B$'s threshold is met, while $\mathcal{I}_A$'s threshold is not. This is not targeted provision to group B in the technological sense, but electorally it behaves like reverse clientelism. There is no additional targeted-provision action T_B in the game.

The new relevant gross electoral value is

$$G_B = R\left[1 - F_{\Delta}\left(\frac{N_A + N_B}{2} - \lambda_A N_A - (1-\lambda_B)N_B\right)\right].$$

The associated reverse-through-U bidding function is

$$B_B(t) = G_B - c((N_A + N_B)t).$$

This is the incumbent's payoff from playing U at a level that satisfies $\mathcal{I}_B$, but not $\mathcal{I}_A$.

The case $q = t_A < t_B$ can be ignored, since it captures the same independent voters as targeted provision to group A, but at higher cost. For every $t_A > 0$,

$$G_T - c((N_A + N_B)t_A) < G_T - c(N_A t_A).$$

Thus universal provision that satisfies only $\mathcal{I}_A$ is strictly dominated by T_A.

3.4.1 Condition for reverse clientelism and compatibility with full-extraction targeted equilibrium

A sufficient condition for reverse clientelism to be irrelevant would be $G_B \leq G_N$. Intuitively, this would mean that satisfying $\mathcal{I}_B$ alone is not electorally valuable enough to beat the incumbent's outside option of no provision. In that case, even an arbitrarily cheap reverse-through-U offer would not be attractive to the incumbent. Formally, the reverse-through-U payoff at threshold t_B is

$$B_B(t_B) = G_B - c((N_A + N_B)t_B),$$

while no provision gives

$$V(N) = G_N.$$

Hence, if $G_B \leq G_N$, then for every $t_B > 0$,

$$B_B(t_B) = G_B - c((N_A + N_B)t_B) < G_N.$$

Thus reverse-through-U provision is strictly dominated by no provision. If $t_B = 0$ is allowed, the conclusion is weak when $G_B = G_N$; strict domination then requires either $G_B < G_N$ or a positive lower bound on admissible thresholds.

However, under the maintained strict monotonicity of F_Δ, the condition $G_B \leq G_N$ cannot hold whenever $(1-\lambda_B)N_B > 0$. Recall that

$$G_B = R\left[1 - F_\Delta\left(\frac{N_A+N_B}{2} - \lambda_A N_A - (1-\lambda_B)N_B\right)\right]$$

and

$$G_N = R\left[1 - F_\Delta\left(\frac{N_A+N_B}{2} - \lambda_A N_A\right)\right].$$

Since

$$\frac{N_A+N_B}{2} - \lambda_A N_A - (1-\lambda_B)N_B < \frac{N_A+N_B}{2} - \lambda_A N_A,$$

strict monotonicity of F_Δ implies

$$F_\Delta\left(\frac{N_A+N_B}{2} - \lambda_A N_A - (1-\lambda_B)N_B\right) < F_\Delta\left(\frac{N_A+N_B}{2} - \lambda_A N_A\right).$$

Therefore

$$G_B > G_N.$$

Equivalently, the condition $G_B \leq G_N$ could hold only if F_Δ were flat on the interval

$$\left[\frac{N_A+N_B}{2} - \lambda_A N_A - (1-\lambda_B)N_B, \frac{N_A+N_B}{2} - \lambda_A N_A\right].$$

This is ruled out by strict monotonicity whenever $\mathcal{I}_B$ has positive mass.

This matters for the targeted-A equilibrium from the restricted game. In that equilibrium, group A extracts the full amount that is compatible with making targeted provision weakly attractive relative to no provision. The corresponding threshold is k_A, defined by

$$B_A(k_A) = G_N.$$

But if $G_B > G_N$, then at $t_A = k_A$ the incumbent can be offered a strictly better continuation payoff by group B. Indeed, group B can set a very small positive threshold $t_B = \varepsilon$, with $\varepsilon < k_A$. The incumbent can then play U at $q = \varepsilon$, satisfying $\mathcal{I}_B$ but not $\mathcal{I}_A$, and obtain

$$B_B(\varepsilon) = G_B - c((N_A+N_B)\varepsilon).$$

As $\varepsilon \downarrow 0$, this payoff converges to

$$B_B(0) = G_B.$$

Since $G_B > G_N = B_A(k_A)$, for sufficiently small $\varepsilon > 0$ we have

$$B_B(\varepsilon) > B_A(k_A).$$

The incumbent would then choose reverse-through-U provision rather than targeted provision to group A. Hence full rent extraction by $\mathcal{I}_A$ is incompatible with no reverse clientelism whenever $G_B > G_N$.

The unrestricted model therefore changes the meaning of the targeted-A outcome. It is no longer enough that targeted provision beats no provision. It must also beat the best reverse-through-U offer available to group B. In particular, a targeted-A equilibrium with no reverse clientelism requires a threshold $t_A > 0$ such that

$$G_T - c(N_A t_A) \geq G_N,$$
$$G_T - c(N_A t_A) \geq G_B,$$

and

$$G_T - c(N_A t_A) \geq G_U - c((N_A + N_B)t_A).$$

The first inequality says that targeted provision to group A beats no provision. The second says that it beats reverse-through-U provision to $\mathcal{I}_B$. The third says that it beats full universal provision at the cheapest level that satisfies $\mathcal{I}_A$, namely $q = t_A$.

Equivalently, the first two inequalities imply

$$c(N_A t_A) \leq G_T - \max\{G_N, G_B\}.$$

The third inequality can be written as

$$c((N_A + N_B)t_A) - c(N_A t_A) \geq G_U - G_T.$$

Thus a no-reverse targeted-A equilibrium requires a threshold that is low enough to keep T_A more attractive than both no provision and reverse-through-U, but high enough to make full universal provision too expensive. Since the usual case is $G_B > G_N$, the binding upper bound is

$$c(N_A t_A) \leq G_T - G_B.$$

In words, group A must leave the incumbent at least $G_B - G_N$ of surplus relative to the restricted full-extraction threshold. The full-extraction threshold k_A is therefore replaced by the reverse-proof cutoff k_A^R, defined by

$$B_A(k_A^R) = G_B.$$

3.4.2 Solving the Unrestricted Voice Game

For a more intuitive understanding of this section, we use Figure 8, which explicitly maps out the relevant bidding curves for each discussed case.

Proposition 6. *Consider an instance G^*_{vu} of the unrestricted voice game. Let F_Δ be strictly increasing on the relevant support, and let c be continuous, strictly increasing, and satisfy $c(0) = 0$. If $k_U \geq k_A$, then universal provision at level k_U will be chosen in a subgame-perfect equilibrium.*

Proof. Let $t_A = t_B = k_U$. Then $B_U(k_U) = G_N$, while $B_A(k_U) \leq B_A(k_A) = G_N$. By the tie-breaking rule, the incumbent chooses U.

The reverse-through-U option does not upset this outcome. Since satisfying both independent blocs is weakly more valuable than satisfying only $\mathcal{I}_B$, $G_U \geq G_B$, with strict inequality whenever $(1 - \lambda_A)N_A > 0$. Therefore

$$B_U(t) - B_B(t) = G_U - G_B \geq 0.$$

Thus B_B is never above the full-universal bid. If group B lowers its threshold, it can only induce a lower level of universal provision. If it raises its threshold above k_U, it cannot obtain more than k_U. Group A cannot gain by raising its threshold above k_U, and lowering it cannot give it more than k_U. Hence the proposed strategy profile is sequentially rational. □

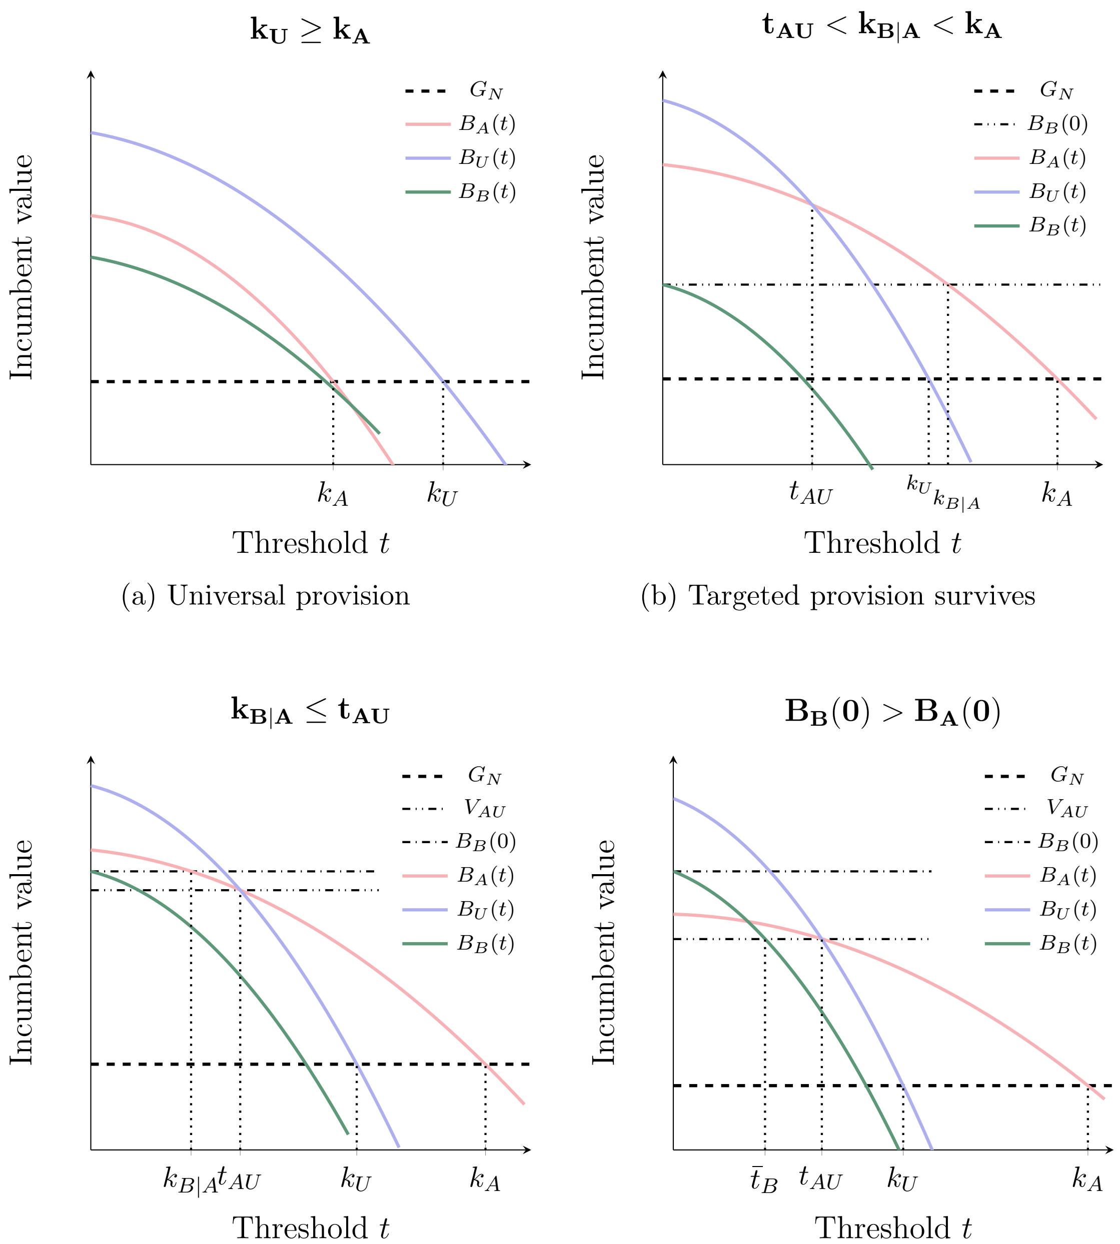

(a) Universal provision

(b) Targeted provision survives

(c) Reverse-through-U threat binds, projection exists

(d) Reverse-through-U threat binds, no projection

Figure 8: Bidding curves in the unrestricted voice game. The red curve $B_A(t)$ is the incumbent payoff from targeted provision to group A, the blue curve $B_U(t)$ is the payoff from universal provision satisfying both groups, and the green curve $B_B(t)$ is the reverse-through-universal bidding curve, where the incumbent provides a low universal level that satisfies group B but not group A. The black dashed line is the payoff from no provision, G_N. In panels (b)–(d), the additional horizontal lines mark the relevant comparison levels $B_B(0)$ and V_{AU}. $\bar{t}_B$ is the largest threshold at which group B's reverse-through-U bid gives the incumbent at least V_{AU}

We now turn to the case $k_A > k_U$. In this case, full universal provision at k_U is not protected against targeted provision to A.

Proposition 7. *Consider an instance G_{vu} of the unrestricted voice game. Suppose $k_A > k_U$, and suppose $\mathcal{I}_B$ chooses its threshold before $\mathcal{I}_A$. Let t_{AU} be defined by*

$$B_A(t_{AU}) = B_U(t_{AU}),$$

and write

$$V_{AU} = B_A(t_{AU}) = B_U(t_{AU}).$$

Define $k_{B|A}$, whenever it exists, by

$$B_A(k_{B|A}) = B_B(0).$$

Then the following cases obtain.

1. *$k_A > k_{B|A} > k_U$: A can sustain targeted provision at any threshold strictly below $k_{B|A}$. In the continuous game the supremum targeted threshold is $k_{B|A}$, approached from the left. If thresholds are discretized and*

$$\hat{k} = \max\{t < k_{B|A}\}$$

exists with $\hat{k} > t_{AU}$, then targeted provision to A at $t_A = \hat{k}$ is sustained in a subgame-perfect equilibrium.

2. *if $k_A > k_U$ and $B_B(0) > B_A(0)$, then the equilibrium depends on the order of moves. When moving first, A can ensure universal provision at k_U. When B moves first, the game depends on whether $B_B(0) > V_{AU}$. If so, there exists a universal provision equilibrium, otherwise any move by B results in targeted provision to A.*

Proof. Let us first tackle the simple first case. When the "projection" of $B_B(0)$ onto B_A lies further out than k_U, the group A has both the leverage and the incentive to fight for targeted provision. When going first, A can just play that threshold. It is the largest value they can get (since $k_{B|A} > k_U$) and B cannot fight (since any of their bids would be beat). Same goes for when they go second, except here B does not have a best strategy. Everything is guaranteed to bring them 0, so for instance k_B satisfying $B_B(k_B) = V_N$ and k_A can be sustained.

Let us move to the second case, first showing that k_U units of the good is provided universally when A moves first. For a certain provision quantity to get realized, it needs to be played by someone as a threshold. If A plays anything above k_U, B would have to bid on its own bidding curve, as proper universal provision can no longer be realized. If A plays a threshold up to k_U, B can play k_U, because k_U is the most value the group B can ever get. As such, k_U is realized when A moves first.

The second subcase of this case is more difficult. When the player B moves first, it expends its leverage prematurely. To see that, consider what happens when k_U is offered. Since $k_A > k_U$, the player A would just play k_A. Since $B_B(k_U)$ is smaller than G_N, this gets us targeted provision. Clearly, it cannot be optimal for B to open with k_U. Instead, something on its own bidding curve should be played. What if $B_B(0)$ is larger than V_{AU}? Then, there exists a set of bids for B that, when overbid by A, still land us on the universal provision curve. If $B_B(0)$ is smaller than V_{AU}, then any move B plays lets A get targeted provision.

□

3.5 Taxation and budgeting: Restricted and Unrestricted cases

Another problem that has to be addressed before we can be done with this model is the problem of budgeting. The incumbent is not spending his own money, and the budget balance needs to be respected. We show that under a flat tax-and-transfer scheme the solutions of the original game do not change meaningfully. This enables us to make the Exit game the way that it is, thus partly transferring the Glaeser and Shleifer (2005) logic of taxation into our model. Moreover, it allows for a cleaner reformulation of incumbent's utility, now in terms of "managerial costs": he is not spending his own money, but his effort costs are increasing in the size of required budget expansion.

Formally, we set the pre-existing target budget to zero and require the service expenditure $c(Q)$ to be financed by a flat per-capita tax. Thus, after the incumbent chooses aggregate provision Q, each citizen pays

$$\tau(Q) = \frac{c(Q)}{N_A + N_B}.$$

If $N_A + N_B = 1$, this reduces to $\tau(Q) = c(Q)$.

The tax-financed game keeps the same timing, strategies, voting rule, and incumbent payoff as the corresponding voice game. The only change is the payoff of independent voters. In G_{vrt} and G_{vtu}, terminal payoffs become

$$B_{0A} = V_{0A} = u(q_A) - \tau(Q),$$

and

$$B_{0B} = V_{0B} = u(q_B) - \tau(Q).$$

Hence, under targeted provision to group A,

$$B_{0A} = V_{0A} = u(q_A) - \frac{c(Q)}{N_A + N_B}, \qquad B_{0B} = V_{0B} = u(0) - \frac{c(Q)}{N_A + N_B},$$

while under universal provision both groups receive

$$V_{\mathcal{I}A} = V_{\mathcal{I}B} = u(q) - \frac{c(Q)}{N_A + N_B}.$$

Threshold voting is still based on gross service provision: conditional voters in group g support the incumbent if and only if $q_g \geq t_g$. Taxes therefore do not change which voters are satisfied at a given terminal history; they only change the payoff from that terminal history. We denote the tax-financed versions of G_{vr} and G_{vu} by G_{vrt} and G_{vut}, respectively. A good thing about this game variation is that it has the same targeted-provision subgame-perfect equilibria, both in restricted and unrestricted variants. Let us state and prove that formally.

Proposition 8. *For a given restricted voice game G_{vr}, define a tax-and-transfer voice game G_{vrt}. Let the cost function c grow slower than linear ($c'(Q) < 1$) on the interval $[0, \frac{1}{N_A}c^{-1}(R)]$. If targeted provision is attainable in SPE of G_v ($k_A > k_U$), it is attainable in SPE of G_{vrt}.*

Proof. Observe that if k_A is offered to the incumbent in this new case, he is going to accept. The derivative condition means that the players have leeway to increase the amount that gets paid out net of tax at least until $\frac{1}{N_A}c^{-1}(R)$. Since all thresholds k_A, k_U are smaller than $\frac{1}{N_A}c^{-1}(R)$ (they are all monotonically increasing inverse cost functions of various fractions of R), we know that $k_A - c(k_A * N_A) > 0$. Moreover, it is larger than $k_U - c(k_U)$, because $k_A > k_U$ by assumption, the gap $Q - c(Q)$ increases with Q, and we are additionally scaling the costs down by N_A^{η}. As such, k_A is still the most attractive option of the group A. As established before, group B cannot do anything but pay the bill.

□

Proposition 8 tells us that if we have a suitable subgame-perfect equilibrium in the restricted voice game without taxation, we will have one in the restricted voice game with taxation. This is important because we are going to build our exit model on taxation. Let us now show that the same is true for the unrestricted game.

Proposition 9. *For a given unrestricted voice game G_{vu}, define a tax-and-transfer voice game G_{vut}. Let the cost function c grow slower than linear ($c'(Q) < 1$) on the interval $[0, c^{-1}(R)]$. If targeted provision is attainable in SPE of G_{vu} ($k_A > k_{B|A} > k_U$), it is attainable in SPE of G_{vut}.*

Proof. Observe that if k_A is offered to the incumbent in this new case, he is going to accept. The group A cannot get anything better, as it cannot increase the threshold (would have done it in the original game) and it cannot play k_U for the same reason. Moreover, the payoff gap between k_U and A's pick (the projection $k_{B|A}$) is now larger in A's perception, because when A offers k_U to the incumbent, it has to pay taxes on the common provision costs $c(k_U)$, and when offering $k_{B|A}$, it merely has to pay $c(N_A * k_{B|A})$. □

Propositions 8 and 9 tell us that the exit game is going to have some of the same equilibria with targeted provision. In general, our aim is to show that the exit model does not have *fewer* subgame-perfect equilibria with targeted provision than the voice model.

3.5.1 Taxation: discussion

The tax modification to the model is conservative with respect to results on targeted provision. In both the restricted and unrestricted variants, targeted-provision equilibria that were attainable in the corresponding voice game remain attainable after taxation is introduced. Proposition 8 gives this result for the restricted case, while Proposition 9 gives the analogous result for the unrestricted case under the relevant targeted-provision conditions.

The reason is that the flat tax changes the groups' net payoffs, but does not remove the incumbent's incentive to accept the relevant targeted offer. In the restricted case, if group A can sustain k_A in the original game, it can still do so in the existence of taxes, provided that the cost function grows slower than linearly on the relevant interval. Group B, by contrast, still has no targeted instrument of its own and can do little except pay the bill. In the unrestricted case, the same preservation result holds for targeted provision that already survives the reverse-through-U threat.

Note, however, that the introduction of budgeting to the model has adverse effect on universal provision equilibria: some of them may dissolve once the taxation is added.

3.6 The Exit Game

We now introduce the exit mechanism. In the voice model, citizens could discipline the incumbent only through voting. In the exit model, citizens have an additional response: if the incumbent's policy makes them sufficiently worse off, they can leave the municipality before the election.

This is relevant, since taxation was introduced. A resident's payoff now depends not only on whether her group receives public services, but also on the tax cost she has to bear. As a result, some residents may receive negative utility from remaining in the municipality. For example, if the incumbent provides services only to group A, members of group B receive no service benefit but still bear the tax burden. For sufficiently high taxes, this may make staying unattractive for them. We assume that exit happens after taxes are paid. In making this assumption, we have followed Glaeser and Shleifer (2005). The problem that this introduces is that it becomes difficult to imagine why voters who observe all parameters would not just leave before any taxes are levied. An explanation that can be incorporated into the future version of the exit model without distorting the results is that the exit is associated with a cost, the voters are impatient, and they get fully convinced that taxes will continue to get levied in the next period if the incumbent wins.

We capture this idea by giving citizens an outside option with reservation utility normalized to zero. If a citizen's utility from staying in the municipality is below zero, she would prefer to exit. However, exit is not automatic. To capture the idea that people are attached to where they live, and that moving is often difficult, we assume that a citizen with negative utility V receives an opportunity to exit only with probability

$$\chi(Q) \in [0, 1] \; \forall \; Q.$$

If she receives this opportunity, she leaves the municipality before the election. Otherwise, she remains and votes. Since the model treats each group as a collective, the probability function $\chi(Q)$ also has a population-share interpretation. If a group receives negative utility, then a share $\chi(Q)$ of that group exits the municipality, while the remaining share $1 - \chi(Q)$ stays. We do not otherwise model moving costs explicitly. The function χ is a reduced-form way of capturing mobility frictions, local attachment, family and social ties, housing constraints, and other reasons why citizens may remain in a municipality even when they are dissatisfied with incumbent policy. Note that most of our constructions in this section are compatible with χ being a constant as well as an increasing function of Q. New timeline of the game is presented on the figure 9.

This modification changes the incumbent's strategic problem as well. Now, by denying services to one group, while imposing taxes on everyone, incumbent can drive some residents from the opposing group out of the polity. This can increase his probability of reelection even if the policy is socially inefficient. Note that the only group whose exit can be induced is B. We assume that partisans are too attached to the polity or otherwise invested in the political standoff to leave, and so model only group B's independents as capable of leaving. This has an added technical benefit

Timeline: Exit Model

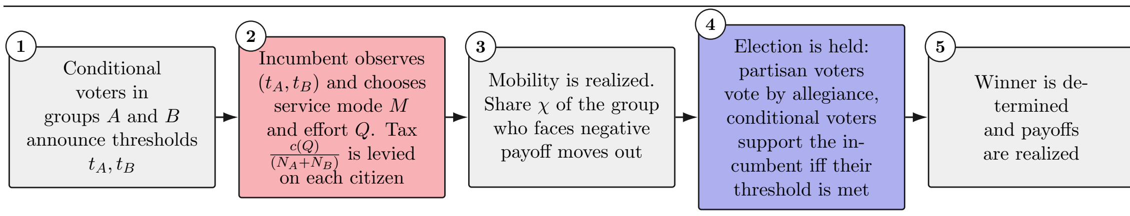

Figure 9: The timeline of the "exit" game. Each cell corresponds to a stage.

of ensuring that the distribution of shocks is consistent with the no-exit case, and enables direct comparisons. We now state and prove our results.

Proposition 10. *Let G_{vrt} be an instance of the restricted tax-and-transfer voice game, and let G_{ver} be its associated exit game. Let $\chi(Q)$ be the exit function with an everywhere non-negative first derivative $\chi'(Q)$, such that $\chi(Q) \in (0,1)\ \forall Q > 0$. Let c be the cost function. Furthermore, let the random shock $J_B - J_A$ be distributed with a strictly increasing cumulative distribution function F_Δ and population $N_A + N_B$ be normalized to 1. If group A receives the quantity $k_A = \frac{1}{N_A} c^{-1}(R[F_\Delta(0.5 - \lambda_A N_A) - F_\Delta(0.5 - N_A)])$ in a subgame perfect equilibrium of the original game, the new game also has targeted provision in SPE. The amount of targeted provision is $k_A^* > k_A$.*

Proof. As usual, define A's bid function:

$$B_{A|voice}(Q) = [1 - F_\Delta(1/2 - N_A)] \times R - c(Q * N_A);$$

We now need to define the bid function in the exit game. Observe that the only relevant thing that changes is the population composition under T_A. A share $\chi(Q)$ of group B's independents leave if they are taxed, according to how much they are taxed. Denote the measure of the leaving group as $L(Q) = N_B * (1 - \lambda_B) * \chi(Q)$. This is a positive number for all $Q > 0$. As such, the probability of victory is:

$$\mathbb{P}[\text{Vote share} \geq \frac{1}{2}] = \mathbb{P}[N_A + [J_B - J_A] \geq \frac{N_A + N_B - L}{2}]$$

$$\mathbb{P}[\text{Vote share} \geq \frac{1}{2}] = \mathbb{P}\left[[J_B - J_A] \geq \frac{N_B}{2} - \frac{N_A}{2} - \frac{L(Q)}{2}\right].$$

We now write the bid function:

$$B_{A|exit}(Q) = \left[1 - F_\Delta(\frac{N_B - N_A - L(Q)}{2})\right] \times R - c(Q * N_A).$$

We see that $F_\Delta(\frac{N_B - N_A - L(Q)}{2})$ is strictly smaller than $F_\Delta(1/2 - N_A)$. That is true because $N_A + N_B = 1$ and $\frac{N_A + N_B}{2} - N_A - \frac{L(Q)}{2}$ is smaller than $\frac{N_A + N_B}{2} - N_A$. As such, the entire first term of the exit value function is larger for any $Q > 0$, and second terms are the same for voice and exit games.

As we remember, $k_A > k_U$ is a necessary and sufficient condition for the existence of a subgame-perfect equilibrium with targeted provision. Since k_A is defined as the intersection of B_A and G_N, we see that the intersection now lies further to the right. As such, $k_A > k_U$ is going to hold in the new game. □

Corollary 1. *Consider an unrestricted voice game with taxation,* $F_\Delta(0.5 - \lambda_A N_A - (1-\lambda_B)N_B) > F_\Delta(0.5 - N_A)$*. The projection value* $k^*_{B|A} : B_B(0) = B_A(k^*_{B|A})$ *in its exit-enabled extension is larger than the corresponding projection value* $k_{B|A}$ *in the original game.*

Corollary 2. *Consider an unrestricted voice game with taxation,* $F_\Delta(0.5 - N_A) > F_\Delta(0.5 - \lambda_A N_A - (1-\lambda_B)N_B)$*. The x-coordinate of the intersection value of the bidding curves* B_A *and* B_B *(*k^{AB}_**) in its exit-enabled extension is smaller than the corresponding intersection value* k^{AB} *in the original game.*

Corollaries 1 and 2 tell us that both the restricted and unrestricted versions of the game change significantly when exit is introduced. The change is not in B's favour. Firstly, in the restricted game $k_A > k_U$ is required to get to a subgame-perfect equilibrium where the provision is limited to A. Since k_A is even higher in the exit-enabled game (and k_U is the same), life gets harder for B. Some of the parameter combinations that would have resulted in uniform provision now support only targeted provision to A.

Turning to the unrestricted case, we see that since $k_{B|A} > k_U$ is a necessary and sufficient condition for targeted provision to be reached by self-interested players, we can see that making $k_{B|A}$ higher makes things worse. Moreover, the change hurts both A and B in the case when $B_B(0) > B_A(0)$ by shifting the agreement point to the left.

The figure 10, generated using the closed-form solver of our proposed model implemented in Python, showcases a cross-section of the restricted game's parameter space along the plane (λ_A, λ_B). We see that the exit-enabled version of the restricted game is more hostile to the anti-incumbent demographic, as there are now more combinations (λ_A, λ_B) that can produce an equilibrium with targeted provision. We see that the bite is concentrated in the region with low λ_A and high λ_B. That makes sense: when there are few independents in B, it makes more sense to drive some of them out cheaply than to pay their hostile compatriots a lot.

Another phase diagram, seen in 11, shows the changes in the unrestricted case. As mentioned in the section on taxes, this is a conservative estimate of the extent of the changes. Once again, we see that the changes are concentrated in the low-λ_A, low-λ_B region. This is necessary for $F_\Delta(0.5 - \lambda_A N_A - (1-\lambda_B)N_B) - F_\Delta(0.5 - N_A)$ to be sufficiently large and for the "pushing out" of $\chi(Q) * (1-\lambda_B)N_B$ voters to have significant effect.

3.6.1 Exit: discussion

As shown above, the balance tips against B when the exit option is introduced. In the restricted case, targeted provision to A now also changes the electorate: when group B receives no service but bears the tax burden, a share of the negatively affected group exits. This raises the incumbent's electoral return from targeted provision without changing the corresponding cost. The extent of the change directly depends on how easy it is to leave.

Formally, the targeted provision bidding curve shifts so that the relevant intersection lies further to the right. Hence k_A increases, while k_U remains the same. This can also be seen through the condition that governs targeted provision in the restricted game $k_A > k_U$. Once exit is introduced, this condition becomes easier to satisfy. Some parameter combinations that would have produced universal provision

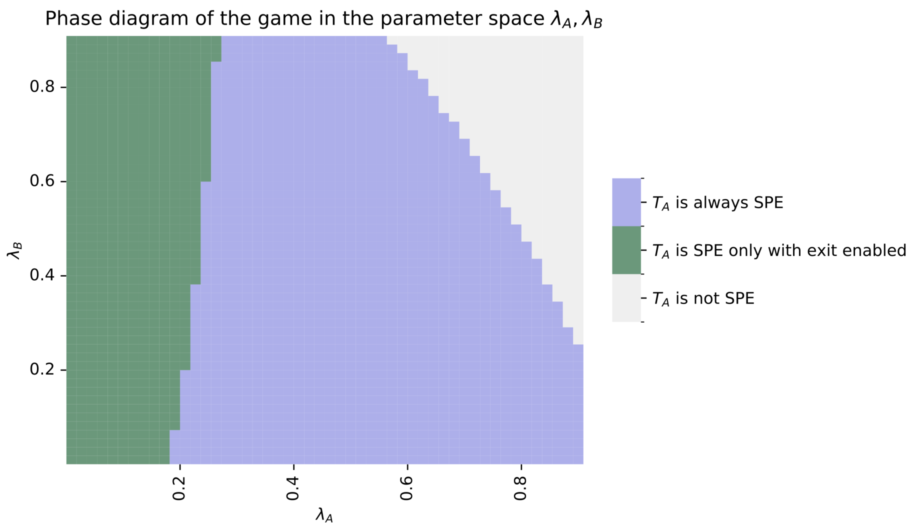

Figure 10: Phase diagram of the restricted game in the parameter space (λ_A, λ_B). The figure compares the regions in which targeted provision to group A is sustained as a subgame-perfect equilibrium with and without exit. Other parameters used: $N_A = 0.5, N_B = 0.5, R = 5.0, \chi(Q) = 0.7, c(Q) = 0.1Q^2$. The distribution of both shocks is uniform, and the distribution of the composite shock is triangular.

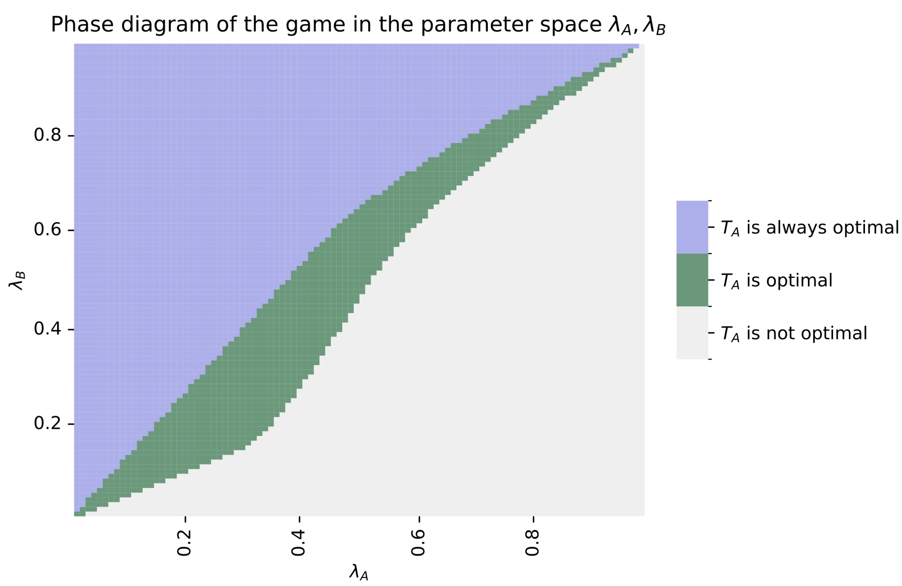

Figure 11: Phase diagram of the unrestricted game in the parameter space (λ_A, λ_B). The figure illustrates how the equilibrium regions change once unrestricted universal provision and exit are allowed. Other parameters used: $N_A = 0.5, N_B = 0.5, R = 5.0, \chi(Q) = 0.7, c(Q) = 0.1Q^2$. The distribution of both shocks is uniform, and the distribution of the composite shock is triangular.

in the tax voice game now produce targeted provision to A. The unrestricted case is not as straightforward, because group B still has the reverse clientelism available to them. Exit does not automatically eliminate this threat, so targeted provision need not appear everywhere. Its effect is instead to move the relevant bargaining objects against B. When $F_\Delta(0.5 - \lambda_A N_A - (1 - \lambda_B)N_B) > F_\Delta(0.5 - N_A)$, the exit-enabled projection value $k^*_{B|A}$ is larger than the original projection value $k_{B|A}$, making the targeted-provision condition easier to satisfy. When $F_{\Delta}(0.5 - N_A) > F_{\Delta}(0.5 - \lambda_A N_A - (1 - \lambda_B) N_B)$, the intersection value k^{AB*} shifts to the left relative to k^{AB}. Then the crossover equilibria become worse for both groups, while the universal equilibria remain unchanged.

Before we move to the conclusion section, two observations are in order. Firstly, both charts in this section assume a constant $\chi(Q)$, which is unrealistic. However, we argue that there must be some non-zero initial probability of leaving, and that our constant function setup provides a lower bound on the effects of exit. Consider e.g. the generalized sigmoid that outputs values starting at some $x \in (0, 1)$ and ending at 1. It has a positive derivative, meaning that A's bidding curve would see both its intercept increase and its slope become flatter, which are both things that push the relevant intersections to the right. Secondly, this paper is about discriminatory provision and its determinants, and we show that the situation where exit is enabled is worse in that it can result in discriminatory provision. For this reason, we do not provide the full solution of the taxation model, although it can make for an interesting follow-up.

4 Conclusion

Having studied the different versions of the model, we are now ready to conclude our paper. First, we are going to recapitulate our results and put them in context of our modeling assumptions. Second, we are going to state our contributions to the literature. Finally, we are going to put forward several questions about our model that have yet to be answered.

Discrimination in the provision of public goods and services can be modeled as arising from competition between demographic groups that use simple retrospective voting rules. In our model, voters setting thresholds on minimal resource or effort provision can be understood as “bargaining” or “bidding” to be selected by the incumbent. In many cases, this bidding process supports normal, non-discriminatory government provision. In other cases, the group that is more electorally valuable to the incumbent obtains disproportionate leverage, and resources are directed to it.

Three primary factors determine a group's success in this game. First, it is important to be less partisan than the opposing group, meaning that the group must contain many conditional voters who are willing to withhold support unless their retrospective threshold is satisfied. We thus replicate a classical result in political economics (Lindbeck and Weibull (1987), Dixit and Londregan (1995)), i.e. that groups with many swing voters that can be "bought" for cheap will more likely "get theirs", all other things equal. We have shown that the same conclusion can be achieved within a game-theoretic setup without probabilistic voting.

Second, incumbent costs determine how much more expensive it is to provide a service to the entire population rather than to a politically relevant subgroup. When universal provision is sufficiently costly relative to targeted provision, discriminatory provision becomes easier to sustain.

Finally, when voters can easily leave the polity for an exogenous outside option, targeted provision becomes more attractive for an office-motivated incumbent. The particular way this pressure is modeled is taxation on many to fund provision to the few, and we show that adding this lever expands the region of the parameter space in which discriminatory provision is optimal.

Many things about our model might not strike the reader as particularly realistic. Still, we believe most of our choices serve to clarify real phenomena rather than simply generate a particular result. For instance, it might appear unintuitive that overt partisans are practically absent from the political process as presented in our model. However, that choice helps us model sections of a “passive” electorate that parties and politicians often have. Moreover, since we explicitly say that the incumbent’s effort is managerial rather than monetary in nature, it can be argued that the partisans represent sections of the electorate that can be satisfied with perfunctory effort whose cost is negligible relative to the potential value of office.

To explain two more of the stronger assumptions that we use, we turn to explicit dialogue with Glaeser and Shleifer (2005), where it is argued that it is exactly the possibility of exit that creates the incentive for discriminatory provision. Our results support that idea in that exit helps sustain more discriminatory provision equilibria in both the restricted and unrestricted versions of the model. The model in Glaeser and Shleifer (2005) presupposes both that provision to the ”opponent bloc” cannot happen and that discriminatory provision cannot be profitable unless population reshuffles. We start out by assuming the former and then partially relax that assumption. The latter is not true in our model for many technical reasons, but mostly because we do not model the ”continuous satisfaction” of voters as relevant to their vote. While this is somewhat extreme, using this binary logic from Kiss (2012) helps us add a lot of game-theoretic structure to the bargaining problem and have it remain tractable.

Another thing that we do differently from Glaeser and Shleifer (2005) is that we model voters as rigid blocs. What that paper presents as an aggregate of many micro-level decisions happens in ours as a result of interaction between just two population-scale players. We understand that this might also not seem particularly realistic. However, political groups that are closer to organized blocs do exist and can exert influence through organized “voice”. The manifestly disorganized “exit” option shown in our model can then be viewed as a hindrance, and blocs might take action to bring down the function χ. For instance, a union leader or an ethnic party candidate might try to signal upcoming positive changes to dissuade people from leaving and thus weakening his group’s bargaining position this period.

Wrapping up the discussion of our assumptions, we want to mention several important limitations follow directly from the simplifying choices made in the paper. First, the analysis is entirely theoretical. We characterize equilibrium mechanisms, but we do not provide empirical evidence that municipalities actually behave according to these mechanisms, nor do we estimate the size of the relevant effects. Second, our blocks are modeled as infinitely organized, which abstracts from within-group heterogeneity, collective-action problems, turnout, and the possibility that different members of the same group may respond differently to the same policy. Third, the voting rule is deliberately binary: conditional voters support the incumbent if and only if their group-specific threshold is met. This keeps the model tractable, but it leaves out more continuous forms of retrospective evaluation. Fourth, the exit mechanism is reduced-form. The model does not explicitly describe where citizens move, how destination municipalities respond, or how exit affects the future tax base and public-service capacity of the original municipality. Fifth, the challenger is not modeled as a strategic player. This is consistent with much of the accountability literature, but it means that the model abstracts from campaign promises, counter-actions from challenger, and other possible elements of the political landscape. Finally, the service provided is modeled as a divisible private benefit rather than a pure public good. This is useful for studying targeted provision, but it limits the direct applicability of the model to services whose benefits are difficult to exclude or assign on a per-capita basis.

Taking stock now of the broader literature on voter mobility, we have to mention Ferejohn (1986). The general retrospective voting model there shows, among other things, that heterogeneous voting groups allow for weaker incumbent accountability. Our model is in direct agreement with that result, using inter-group competition as a direct driver of accountability failures.

Our model adds to the discussion in papers like Tiebout (1956) and Epple et al. (1984). Unlike those papers, we mostly abstract from the spatial element to simplify the game-theoretic logic, but still find that it requires minimal assumptions to model the political process as being downstream of population dynamics.

On the technical side, Kiss (2012) was the main reference we drew from to start building our model. In that paper, a game-theoretical construction serves to establish that sowing political division is beneficial for the incumbent. Of particular interest to us was the sequential logic. In our opinion, the prolonged aspect of the bargaining game that we present is crucial: not only does the incumbent know what is happening and what the voters demands are, but the voters themselves act in accordance with that.

Having discussed the way our results are situated within the political economic literature, we now state our open questions and avenues for further research. First and foremost we would like to mention the lack of moral hazard structure in our model. In models like Ferejohn (1986), the incumbent can fail to satisfy the electorate for reasons outside of his control. If introduced to our model, this change would transform the bidding curves and potentially make for new game-theoretic structure. In particular, a more risk-averse incumbent could be facing higher effort costs, or be subjected to mixed bargaining strategies. It also can be argued that the voters should be risk-averse with regards to election outcomes, as is often the case in the real world.

Envy, ingroup preference, and fairness would also make worthwhile additions to voter preferences. Right now, voters in our model often find themselves indifferent between universal and targeted provision. While envy and altruism can be seen as exogenous and essentially irrational, ingroup bias might arise as an emergent preference when voters in the exit model care about the future and want the incumbent to remain in office.

Another factor that warrants more thorough treatment in future work is taxation. When solving the model, we noticed that while the introduction of taxes leaves targeted-provision equilibria intact, it materially changes players' behavior in many other situations. For instance, group A is no longer indifferent between universal and targeted provision at the same per-capita level, because the former results in a larger tax burden. It is also important to find out which exact form and manner of taxation the model should have.

Finally, we would like to discuss potential ways to test our model empirically. Our model produces two principal results that can be tested. Firstly, having a lot of swing voters influences the group's chances of getting redistribution their way. This result has been tested in e.g. Dahlberg and Johansson (2002), where the authors use data on Swedish temporary grants, where the central government transferred funds, along with significant authority on spending these funds, to municipalities. This is a good setup for testing the first result, and we do not propose any improvements. However, we cannot use grants from the central government to test our second finding, namely that exit makes discriminatory provision more lucrative for the incumbent. For this, we would have to look at closed-loop fiscal setups, where funds are redistributed within the municipality. The ideal setup would feature a polity whose different areas experience a sudden heterogeneous decrease in moving costs. Such a decrease could be compared to us "enabling" the exit option in the game. The extent of the decrease can then be used as a regressor in a model where the amount of redistribution is the dependent variable.

References

Alesina, A. (1988). Credibility and policy convergence in a two-party system with rational voters. *American Economic Review*, 78(4):796–805.

Banks, J. S. and Sundaram, R. K. (1998). Optimal retention in agency problems. *Journal of Economic Theory*, 82:293–323.

Dahlberg, M. and Johansson, E. (2002). On the vote-purchasing behavior of incumbent governments. *American Political Science Review*, 96(1):27–40.

Dixit, A. and Londregan, J. (1995). Redistributive politics and economic efficiency. *American Political Science Review*, 89(4):856–866.

Epple, D., Filimon, R., and Romer, T. (1984). Equilibrium among local jurisdictions: Toward an integrated treatment of voting and residential choice. *Journal of Public Economics*, 24:281–308.

Feddersen, T. J. (2004). Rational choice theory and the paradox of not voting. *Journal of Economic Perspectives*, 18(1):99–112.

Ferejohn, J. (1986). Incumbent performance and electoral control. *Public Choice*, 50(1/3):5–25.

Ferejohn, J. A. and Fiorina, M. P. (1974). The paradox of not voting: A decision theoretic analysis. *American Political Science Review*, 68(2):525–536.

Fernandez, R. and Rogerson, R. (1996). Income distribution, communities, and the quality of public education. *Quarterly Journal of Economics*, 111(1):135–164.

Fieldhouse, E., Green, J., Evans, G., Mellon, J., and Prosser, C. (2020). *Electoral Shocks: The Volatile Voter in a Turbulent World*. Oxford University Press, Oxford.

Glaeser, E. L. and Shleifer, A. (2005). The curley effect: The economics of shaping the electorate. *Journal of Law, Economics, & Organization*, 21(1):1–19.

Hindriks, J. (2001). Mobility and redistributive politics. *Journal of Public Economic Theory*, 3(1):95–120.

Hirschman, A. O. (1970). *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Harvard University Press, Cambridge, MA.

Kiss, A. (2012). Divisive politics and accountability. *European Journal of Political Economy*, 28:208–214.

Lindbeck, A. and Weibull, J. W. (1987). Balanced-budget redistribution as the outcome of political competition. *Public Choice*, 52(3):273–297.

Riker, W. H. and Ordeshook, P. C. (1968). A theory of the calculus of voting. *American Political Science Review*, 62(1):25–42.

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, 64(5):416–424.