UNIVERSITÉ DU LUXEMBOURG

PhD-FDEF-2025-032

The Faculty of Law, Economics and Finance

DISSERTATION

Defence held on 21/11/2025 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG

EN SCIENCES ÉCONOMIQUES

by

SARAI GODO LUQUE

Born on 28 August 1990 in Girona (Spain)

STRATEGIC BEHAVIOUR IN POLITICAL BARGAINING

Dissertation defence committee

Prof. Dr Vincent ANESI, dissertation supervisor

Université du Luxembourg

Prof. Dr Skerdilajda ZANAJ, Chairwoman

Université du Luxembourg

Prof. Dr Pierre PICARD, Vice-Chairman

Université du Luxembourg

Prof. Dr Matthias DAHM

University of Leicester

Prof. Dr Arnaud DELLIS

Université du Québec à Montréal

UNIVERSITÉ DU LUXEMBOURG

Affidavit / Statement of originality

I declare that this thesis:

- is the result of my own work. Any contribution from any other party, and any use of generative artificial intelligence technologies have been duly cited and acknowledged;
- is not substantially the same as any other that I have submitted, and;
- is not being concurrently submitted for a degree, diploma or other qualification at the University of Luxembourg or any other University or similar institution except as specified in the text.

With my approval I furthermore confirm the following:

- I have adhered to the rules set out in the University of Luxembourg's Code of Conduct and the Doctoral Education Agreement (DEA)[1], in particular with regard to Research Integrity.
- I have documented all methods, data, and processes truthfully and fully.
- I have mentioned all the significant contributors to the work.
- I am aware that the work may be screened electronically for originality.

I acknowledge that if any issues are raised regarding good research practices based on the review of the thesis, the examination may be postponed pending the outcome of any investigation of such issues. If a degree was conferred, any such subsequently discovered issues may result in the cancellation of the degree.

Approved on 2025-09-24

[1] If applicable (DEA is compulsory since August 2020).

Acknowledgements

Obtaining a PhD is one of those things in life that cannot be done alone. Not because it would be impossible, but because one would lose their sanity along the way. I am grateful to everyone who has walked with me along this path.

First and foremost, I am deeply grateful to my supervisor, Prof. Vincent Anesi. He gave me the opportunity to pursue this path and supported me at every stage. He has been patient, generous, and always available—no matter how big or small my struggles were. He has been an endless source of ideas and solutions, and our conversations often kept me moving forward. His encouragement, trust, and guidance have been invaluable, and working with him has been a true privilege.

I am grateful to the members of my committee, Professors Skerdi Zanaj and Luisito Bertinelli, for their guidance and for accompanying me throughout these four years. I would also like to thank Professor Pierre Picard for giving me the opportunity to teach and conduct tutorials in his master's course, which has given me valuable teaching experience that will benefit my future in academia.

I am thankful to Marina Legrand, Elisa Gesellchen-Ferreira, Roswitha Glorieux, and Noémie Courtois for their support with the administrative aspects of my PhD, with special appreciation to Noémie for her constant availability and kindness.

To my family: I would never have chosen a different one. My parents, Mingu and Carmen Mary, gave me everything they could and made possible for me what they themselves never had: an easier life full of opportunities and choices. Their unconditional support has allowed me to build a life I truly cherish, and for that I am forever grateful. To my brother Cristian and his wife Carla, thank you for

your encouragement and for reminding me that, no matter the distance, home is never far. Living away has always been easier knowing that our parents can always count on you for whatever they might need.

I am blessed with a large, close-knit extended family on both sides, where love and support are the only languages spoken. I feel fortunate that my grandmother is here to witness this milestone, and I cherish the memory of my grandfather, whom I dearly wish could have been here too. I am especially grateful for my aunt Mari Carme, who has been like a second mother to me, and her daughters Mari and Silvia—my role models, the first economists I knew and my earliest inspiration for my decision to study Economics.

To my friends, who have been my strongest support network: to the old ones in Girona and Tübingen—especially Gabi, Lluc, Núria, Isaac, Alba, Clàudia, Ida, and Natasha—who, despite the distance, have always made me feel close; and to the new ones in Luxembourg, who made life here better than I could ever have imagined. From university, Maria, Camila, Cecilia, and Ari—you have been the best companions on this rollercoaster of ups and downs, and your encouragement and support have been priceless. Especially Maria: sharing this journey with you has been a joy I will always remember. From climbing, Valerie, Karol, Vova, Merry, Brice, Steffi, Joy, Sarah, and Caro—thank you for the long training sessions, the climbing trips, the constant encouragement, and the countless hours of laughter. Climbing has been my therapy and my peace, and without you this journey would have been much harder.

To my partner, Agustín: thank you for making Luxembourg your home so that we could be together. Your patience, love, understanding, and endless belief in me have been my refuge, and your calmness and humor have made this journey lighter and happier. You have stood by me day after day, even when I was hardest on myself. This has always been a team effort, and you are the best companion I could have ever dreamed of for this adventure called life.

Reaching this point is also the result of all the privileges and opportunities I have been fortunate to enjoy. I am deeply aware of this and grateful every day for the life I have.

Contents

Abstract

This dissertation consists of three self-contained chapters, each developing a distinct theoretical framework. The three chapters are unified by a common theme: how private information shapes strategic behavior in political bargaining.

In Chapter 1, I examine how domestic political constraints shape the strategic use of sanctions and rewards in international bargaining under asymmetric information. I develop a sequential game in which a sender policymaker, privately informed about its cost of imposing sanctions, seeks to extract a policy concession from a foreign target by offering a reward while simultaneously threatening sanctions. Without domestic constraints, all offers fail and sanctions are imposed in equilibrium, as signaling frictions prevent credible offers. In contrast, when domestic agents can condition reelection on the size or outcome of transfers, the policymaker's office rents act as a commitment device: pooling equilibrium with agreement becomes sustainable, and inefficient bargaining breakdowns are avoided. The analysis shows that office-seeking motives—often viewed as distortions—can enhance welfare by offsetting informational frictions. An extension with heterogeneous citizens demonstrates that strategic delegation to "weaker" representatives, who bear higher sanction costs, may increase the credibility of transfers and facilitate cooperation. The results underscore how domestic institutions, electoral incentives, and the identity of negotiators critically affect whether international bargaining ends in costly deadlock or mutually beneficial agreement.

In Chapter 2, coauthored with Professor Vincent Anesi, we analyze how the absence of both commitment technology and quasilinear preferences constrains the set of policy agreements that legislators can sustain in the long term. To do so, we develop a dynamic model of legislative bargaining with an infinite horizon. In each period, legislators consider policy reforms under majority rule. Legislators have privately observed, evolving preferences; the policy space is highly dimensional; and the status quo is endogenous. We introduce a notion of dynamic stability for legislative bargaining with incomplete information that we call "cluster-stability", shifting the analysis from single policy points to decision rules. Our main result characterizes a family of perfect Bayesian equilibria for

our framework with incomplete information that exhibit a property of cluster-stability. We establish this result by introducing a sufficient condition, which we call "internal optimality". Our main theorem shows that, for generic policy preferences, every internally optimal outcome is cluster-stable. Internal optimality identifies outcomes that could be supported by incentive-compatible transfers if money and commitment were available. In our setting, such transfers do not exist, but sufficiently patient legislators can replicate them through continuation values, which substitute for money, quasilinearity, and commitment. These results have implications for three classic puzzles in political economy. They address the vote trading paradox by showing that efficient agreements can persist even without monetary transfers; they shed light on how lobbies can engage in vote buying with arbitrarily little money by exploiting continuation values; and they clarify the logic of agenda manipulation, as repeated interaction enables agenda setters to sustain preferred outcomes even under incomplete information.

In Chapter 3, I analyze a principal's choice between full delegation and veto-based delegation in a setting characterized by extreme polarization of preferences and asymmetric information. I develop a principal–agent model in which an informed agent privately observes the state of the world and proposes a reform, while an uninformed principal must decide ex ante whether to retain veto power or delegate full authority. Under full delegation, the agent faces no constraints on her reform choice and therefore always selects the upper bound of the policy space, which is closest to her ideal policy. Under veto delegation, by contrast, the agent must account for the principal's acceptance constraint. Because acceptance depends on the state, the agent uses her private information strategically, but in equilibrium every reform proposal is vetoed. The central contribution of this chapter is to show that, under extreme polarization and incomplete information, veto delegation collapses into inefficient gridlock. The principal thus faces a fundamental trade-off when choosing the delegation regime: retaining veto power insures him against the risk of harmful reforms under full delegation, but at the cost of permanent gridlock. This strong result connects to the broader literature on legislative stalemate, showing how polarization and veto power interact to reinforce one another as drivers of gridlock under incomplete information.

Chapter 1

Policy Interventions under Domestic Political Constraints

1.1 Introduction

International relations are characterized by strategic interdependence: the policy choices of one country often shape the welfare and behavior of others. Foreign policy is therefore not merely the external projection of domestic preferences, but a strategic instrument through which governments seek to influence external actors while constrained by their own domestic agents. Foreign influence takes diverse forms, from voluntary cooperation and diplomacy to coercive measures such as military interventions. Since World War II, policy interventions—defined as *the strategic use of sanctions and/or rewards by a sender country to induce a specific policy concession from a target country* (Aidt et al., 2021)—have become the principal instrument of foreign policy.

Sanctions are coercive policy measures intended to pressure a target country to alter its behavior (Morgan et al., 2009, 2023; Syropoulos et al., 2024). These measures include restrictions on trade and finance, arms embargoes, suspension of military cooperation, or travel prohibitions, and are used to enforce international norms, deter aggression, promote human rights, or reshape political regimes (Syropoulos et al., 2024).[1] Conversely, rewards—often referred to as incentives or positive sanctions—entail the granting of political or economic benefits in exchange for cooperation (Cortright, 1997). These benefits may take the form of direct aid, including development assistance, concessional loans, grants, or technical support, as well as indirect instruments such as preferential market access (PMA) or preferential trade agreements (PTA).

Prominent policy episodes illustrate that the effectiveness of sanctions and rewards depends on domestic political conditions. During the renegotiation of the United States–Mexico–Canada Agreement, the United States administration relied on tariff threats to secure concessions, but the credibility and scope of those threats were shaped by lobbying from the steel and automobile sectors. In Europe, the introduction of the Carbon Border Adjustment Mechanism illustrates how sanction-like measures can be combined with conditional rewards: imports of carbon-intensive

[1] For a detailed explanation on the evolution and consequences of sanctions, see Morgan et al. (2023).

goods face tariffs, while countries adopting comparable climate policies retain equivalent access. Strong public support for environmental regulation made this policy politically feasible. In the security domain, the Joint Comprehensive Plan of Action with Iran demonstrated that domestic realignments can decisively alter the credibility of international bargains. Under the Obama administration, electoral and partisan incentives supported sanctions relief in exchange for nuclear limits, whereas under the Trump administration, the relevant coalition reversed and the agreement was abandoned. More recently, the sanctions imposed on Russia since 2022 have provided a similar lesson: their credibility and effectiveness have hinged not only on international resolve but also on enforcement capacity and the willingness of sender governments to absorb the associated costs.

These episodes illustrate that sanctions and rewards are rarely determined only by international bargaining power. Their credibility and effectiveness are mediated by the domestic incentives facing policymakers, by the constituencies able to support or punish them, and by the stability of those preferences over time. This observation reflects the foundational insight of Putnam (1988) that international negotiations are embedded in domestic political struggles. The large literature since Putnam has demonstrated how domestic institutions, electoral incentives, and lobbying pressures shape the use of sanctions and rewards, confirming the centrality of domestic constraints in both their design and their effectiveness.

Although economic sanctions and rewards are frequently used simultaneously (Hufbauer et al. 2007; Morgan et al. 2014), the game-theoretical models that address them fail to consider their simultaneous use. Focusing on only one of these two intervention modes cannot fully identify the actual causes and consequences of policy interventions, nor can it accurately assess their efficacy in achieving the intended policy concessions or the impact of domestic constraints on these results (Aidt et al. 2021; Brekhov 2022, Peksen 2019; Verdier and Woo 2011).

Motivated by these insights, this chapter analyzes how domestic political constraints shape the strategic use and effectiveness of sanctions and rewards in international bargaining. It asks: How do domestic politics affect the outcomes of policy interventions when sanctions and rewards are used simultaneously?

To answer this question, I develop a sequential game of incomplete information

in which a policymaker in the sender country seeks a costly policy concession from a target country. The policymaker, who faces electoral constraints, can offer a monetary reward—such as a transfer or aid package—to induce acceptance of the concession, while simultaneously issuing a threat of sanctions if the offer is rejected.

The extent of sanctions is determined by the sender's relative economic strength, captured by a payoff-relevant state variable observed only by the policymaker in the sender country. The target government knows only the distribution of this variable, creating uncertainty about the actual costs of disagreement. A stronger sender incurs lower sanctioning costs and can therefore impose harsher disagreement payoffs on the target. This private information generates a signaling problem: the target's willingness to accept the offer depends on its beliefs about the sender's strength, beliefs it must infer from the transfer proposal itself. A formal game-theoretic approach is thus essential, as the credibility of sanctions and the attractiveness of rewards emerge from strategic interaction between the sender and the target.

The policymaker in the sender country is driven by mixed motives: he cares both about citizens' welfare and about staying in office. His utility is a weighted sum of the representative citizen's welfare and his ego rents from office—his personal value of holding power, such as prestige, influence, or other office-related benefits, as is standard in the literature (Persson and Tabellini 2002; Besley 2006). These mixed motives allow me to study how variations in political incentives affect equilibrium outcomes.

Domestic constraints are modeled as an implicit "contract" offered by a domestic political actor who conditions reelection on the policymaker's bargaining behavior or on its outcome. In equilibrium, without loss of generality, attention can be restricted to two types of contracts: (i) reelection is promised only if the policymaker imposes sanctions (i.e., no agreement is reached); or (ii) reelection is promised only if an agreement is secured and the transfer offered exceeds a minimum threshold $\bar{t}$ set by the domestic agent. The domestic agent can be interpreted as a voter, in the electoral accountability tradition (e.g., Persson and Tabellini 2002), or as a lobby, in the interest-group literature where lobbies offer contingent contracts

to policymakers (e.g., Grossman and Helpman 2001). It is well-known from this literature that such contracts can be microfounded as reduced forms of richer dynamic models of political accountability or interest-group bargaining.

I characterize the (pure strategy) perfect Bayesian equilibria (PBE) of the model. I begin with a benchmark scenario in which the policymaker is unconstrained and cares only about maximizing social welfare. Without electoral constraints, incomplete information leads to bargaining failure: sanctions are always imposed, and the sender fails to secure the policy concession from the target.

To understand this result, note that the target's decision to accept the sender's proposal depends on its beliefs about the sender's relative economic strength, which determines the cost of sanctions for both countries. The lower the cost for the sender to impose sanctions, the more aggressively it sanctions in case of disagreement. Consequently, a sender with low sanction costs is less inclined to offer substantial rewards, while a target facing high sanction costs may accept lower rewards. This creates an informational asymmetry: the target does not observe the true cost structure and must infer it from the transfer offer.

In equilibrium, the sender's optimal strategy involves signaling high sanction costs to the target, and he does so by offering a low reward. When the uninformed target government receives a low offer, it infers that the sender is strong and will impose harsh sanctions if no agreement is reached. Such beliefs lead the sender's policymaker to offer rewards that are too low to be accepted, resulting in sanctions.

This signaling spiral is unavoidable: regardless of the sender's true strength, it is always profitable for its policymaker to mimic a stronger type in order to extract a better deal. As a result, no informative signal is credible, and the bargaining process collapses—even when mutually beneficial agreements are feasible. The equilibrium is therefore Pareto inefficient: sanctions are imposed even in states where both countries would benefit from reaching an agreement.

These results prompt a crucial question: Can domestic constraints prevent a policymaker from offering rewards that are too low to be accepted? To address this, I introduce domestic constraints into the model. I first consider the case in which the domestic actor prefers that no agreement be reached. In this case, as in the bench-

mark scenario, no agreement is concluded and sanctions are imposed. This result is consistent with a growing body of literature that argues that sanctions, while costly to both countries, are often imposed not to compel compliance from the target, but to satisfy domestic constituencies within the sender country (Kaempfer and Lowenberg 2007; Renwick 1981).

More interestingly, different outcomes emerge when the domestic agent conditions the policymaker's reelection on successfully offering an acceptable transfer. I show that an equilibrium can exist that involves both rewards and a policy concession, provided that both countries would be worse off under sanctions. For such an equilibrium to arise, the policymaker's valuation of reelection—specifically, the size of the ego rents—must be sufficiently large. In this case, all types of policymakers find it optimal to offer the minimum transfer required by the domestic agent to secure reelection and, obtain the associated personal benefits from it.

To understand this mechanism, note that if the policymaker deviates and offers a transfer below the lower bound set by the domestic agent, he is not reelected and, consequently, loses his ego rents. This trade-off between reducing the reward and forfeiting reelection benefits makes an equilibrium with an agreement possible. In other words, the personal gains from staying in office can offset the sender's incentives to signal strength through low offers. I further show that the lower bound $\bar{t}$ selected by the domestic agent corresponds precisely to the minimum reward the target country is willing to accept under prior beliefs, assuming that both countries benefit from reaching an agreement.

Such an agreement becomes feasible when the sender's expected cost of imposing sanctions is greater than the minimum transfer needed to compensate the target for its policy concession. That transfer depends on both the expected cost of sanctions for the target and the intrinsic cost of the policy change. I also prove that the domestic agent chooses to offer such a contract—i.e., one that induces agreement with a positive transfer—only if doing so maximizes the payoff of the sender country's citizens, and the reward is accepted by the target.

The role of office- and rent-seeking motives is crucial for the results. Existing literature on incumbents' reelection incentives shows that elections mitigate moral hazard between voters and officeholders by reducing informational asymmetries.

Specifically, electoral accountability motivates incumbents to exert greater effort or to limit the diversion of rents for personal benefit (e.g., Ferejohn 1986; Besley and Smart 2007). In my model, however, *it is precisely the presence of ego rents from reelection that allows domestic constraints to improve bargaining outcomes*. When these ego rents are sufficiently large, they create incentives for the policymaker to offer a higher reward, which increases the likelihood of reaching an agreement and raises the payoff of the citizens.

The model's implications for delegated bargaining are explored through an extension in which the sender country's population includes two types of citizens, heterogeneous in their costs of sanctions. I show that the probability of reaching an agreement increases when a "softer" policymaker—one who belongs to the group that faces higher costs from sanctions—is in office. Because this type has higher costs from imposing sanctions, he is more willing to offer a higher reward to secure cooperation. Anticipating this, citizens with lower sanction costs may strategically delegate bargaining authority to a representative from the high-cost group, rather than selecting a policymaker who directly reflects their own preferences. This delegation improves the credibility of the offer and expands the set of feasible agreements, even if it involves misrepresentation of voter preferences.

The remainder of the chapter is organized as follows. Section 1.2 reviews the use of sanctions and rewards in practice and discusses related empirical and theoretical literature. Section 1.3 presents the core model. Section 1.4 characterizes the equilibria, both with and without domestic constraints. Section 1.5 explores the implications of the findings and develops an extension on strategic delegation. Section 1.6 concludes.

1.2 Sanctions, rewards, and related literature

In this section, I provide context on key features of real-world sanctions and rewards, review how the existing literature has addressed these features, and explain how this chapter contributes to that body of work.

Following World War II, the use of economic sanctions as a nonmilitary means of

exerting international pressure expanded significantly. In the early stages of the Cold War (1950–1975), the United States dominated the use of unilateral sanctions, primarily through trade and arms embargoes. In the subsequent stages of the Cold War—despite domestic political challenges such as the Vietnam War—the United States remained the foremost initiator of sanctions. Meanwhile, European countries began to establish themselves as coordinated sanctioners, particularly with the formation of the European Economic Community.

The 1990s, often referred to as the "sanctions decade" (Cortright et al., 2002), saw a surge in financial sanctions and restrictions on arms transfers, alongside a rise in multilateral sanctions spearheaded by the United Nations and the European Union. This upward trend accelerated after 2001, ushering in a distinct phase in which sanctions—especially "targeted" measures directed at individual entities, firms, or persons—were employed with unprecedented frequency. According to the third edition of the Global Sanctions Database (Syropoulos et al., 2024), a total of 1,325 sanctions were imposed between 1950 and 2022, with the rate of imposition now nearly ten times higher than in 1950 (Morgan et al., 2023). Today, the United States, the European Union, and the United Nations remain the principal sanctioning entities, while the United Kingdom and China are playing increasingly active roles. Sanctions currently affect countries that are home to roughly 40% of the global population (Aidt et al., 2021).

The persistent use of sanctions has spurred a substantial literature—mainly in political science—on their effectiveness in altering the behavior of target states to induce compliance with the sender's demands (see, e.g., Baer, 1973; Bapat and Morgan, 2009; Doxey, 1972; Galtung, 1967; Oechslin, 2014; Pape, 1997). This literature generally finds limited success: fewer than 35% of imposed sanctions have succeeded in fulfilling their intended political aims (Morgan et al., 2021; Syropoulos et al., 2024).

Several empirical studies, along with the theoretical findings of Morgan and Miers (1999), have pointed to a crucial selection bias in prior assessments of sanctions' effectiveness. Earlier studies focused primarily on imposed sanctions, neglecting the possible effects of threatened sanctions. A growing body of work (Aidt et al., 2021; Drezner, 2003; Lacy and Niou, 2004; Morgan and Miers, 1999; Smith, 1995)

argues that sanctions often "work" at the threat stage: "successful sanctions might not actually be imposed" (Morgan et al., 2021). The central idea is that target states, anticipating the high cost of non-compliance, may choose to concede as soon as sanctions are threatened—thus preventing imposition altogether.

To evaluate this hypothesis, the Threat and Imposition of Economic Sanctions (TIES) dataset was developed. Its most recent version, presented in Morgan et al. (2014), documents 1,412 iepisodes of sanctions and sanction threats between 1945 and 2005. Notably, threats alone account for 40% of these cases. When the threat stage is included in the analysis, the measured effectiveness of sanctions improves by approximately 25% (Morgan et al., 2021).

The prevailing theoretical approach treats economic sanctions as bargaining instruments, emphasizing the strategic interplay between sender and target states—often modeled through game-theoretic approaches. Building on Nash's foundational work on the role of threats in cooperative bargaining (Nash, 1953), the literature has increasingly incorporated the sanction threat stage into formal models.

Much of this theoretical development relies on the crisis bargaining framework from international relations (Fearon 1994, 1995, 1997; Schultz 1999, 2001; Slantchev 2003, 2010; Wagner 2000). This framework uses a sender–receiver game to identify the incentives and constraints states face in international conflicts, typically war. The sender first decides whether to issue a threat or accept the status quo; if a threat is issued, the target chooses to concede or to resist; if it resists, the sender then decides whether to follow through or back down, accepting the status quo.

When applied to economic sanctions, this framework explains the conditions under which threats are made and whether they are likely to succeed. It aligns with empirical findings suggesting that sanctions may be more effective than commonly assumed, as target states often comply at the threat stage (Drezner, 2003; Krustev, 2010; Lacy and Niou, 2004; Morgan and Miers, 1999; Smith, 1995). Some studies focus on the conditions under which sanctions are imposed (Drezner, 1999; Morgan and Miers, 1999), while others explore their success (Krustev, 2010; Smith, 1995; Tsebelis, 1990; Whang et al., 2013), or analyze the credibility of sanction threats and sender follow-through (Dekker et al., 2020; Guisinger and Smith, 2002; Schultz, 2001).

A core assumption in this literature is the presence of incomplete information. Studies adopting the crisis bargaining framework—whether in the context of war or sanctions—typically conclude that under complete information, neither conflict nor sanctions would occur, due to the inefficiencies involved in both (Aidt et al., 2021; Spaniel and Smith, 2015). However, as I show in Section 1.4.1, sanctions may still arise under complete information when the concession cost exceeds the combined sanction costs for both countries. In this sense, the simultaneous use of threats and rewards generates sanction outcomes even in environments where classic crisis bargaining models would predict agreement.

Uncertainty remains a central explanatory factor in the escalation of international conflict and policy interventions (Eaton and Engers, 1992, 1999; Fearon, 1995, 1997; Hovi et al., 2005; Lacy and Niou, 2004; Morgan and Miers, 1999; Morrow, 1989; Powell, 1990, 1999; Schultz, 2001; Whang et al., 2013). In the sanctions literature, asymmetric information is modeled in several dimensions: over the target's type or cost of compliance (Krustev, 2010; Morgan and Miers, 1999; Spaniel and Smith, 2015); the sender's resolve or sanctioning capacity (Hovi et al., 2005); or both sides simultaneously (Eaton and Engers, 1994; Lacy and Niou, 2004; Whang et al., 2013). In these models, the use of sanctions indicates a breakdown in expectations—showing that either the sender misjudged the target's resolve, the target misread the sender's commitment to act, or both parties failed in their assessments.

The sequential model introduced in this chapter departs from the traditional crisis bargaining framework in both its structure and treatment of domestic politics. I abstract from the initial decision to issue a threat and the final choice of whether to follow through if the target resists. By omitting these nodes, I rule out scenarios where the sender either maintains the status quo or backs down after resistance. This simplification is intentional: it allows for a more focused analysis of the bargaining stage under domestic political constraints and enables the inclusion of heterogeneous actors and alternative policy instruments, which would be difficult to model in the full crisis bargaining sequence.

Excluding the final decision node also removes the mechanism of "audience costs" incurred by sender governments when bluffing—commonly used to model domestic politics in international bargaining while still treating countries as unitary

actors. In this chapter, rather than nesting domestic politics within a unitary rational actor model, I explicitly model internal political dynamics in the sender country. This connects my work to the political economy literature on sanctions, which traces its roots to Galtung (1967) and Schreiber (1973). This tradition emphasizes that the imposition of sanctions is often shaped by domestic political incentives, even when sanctions are costly and unlikely to achieve their stated foreign policy goals (Aidt et al., 2021). Classic studies (Leyton-Brown and Ruggie, 1987; Renwick, 1981; Tsebelis, 1990) argue that sanctions are frequently imposed for symbolic purposes—to appease domestic interest groups or signal resolve—rather than to coerce compliance from target states. Kaempfer and Lowenberg (2007) further formalize this view, showing how domestic actors may value sanctions as a demonstration of commitment or identity. For instance, Renwick (1981) emphasizes that sanctions are often used to consolidate support domestically or internationally, while Leyton-Brown and Ruggie (1987) see them as communicative tools toward both domestic and foreign audiences.

Within this framework, two types of domestic actors in democracies emerge as key players: voters and special interest groups. Using the TIES dataset (1945–2005), McLean and Whang (2014) are among the few to empirically examine the joint influence of both actors. Their findings suggest that sanction strategies are shaped by the preferences of both the electorate and influential interest groups.

A growing literature investigates the electoral motivations behind sanctions. Studies show that voters respond to foreign policy performance (Aldrich et al. 1989; Hurwitz and Peffley 1987; Nincic and Hinckley 1991), and that sanctions can be used to increase domestic political support (Drury, 2005; Smith, 1998). For example, Whang (2011) shows that sanction imposition boosts presidential approval, while Peksen and Peterson (2016) find that United States presidents are more likely to impose sanctions during periods of low popularity. However, the influence of voters depends critically on their awareness. Voters tend to influence foreign policy only when they are well informed, feel directly affected, or when the dispute is highly salient (Grossman and Helpman 1996; Sobel 2001). In line with this, Drury et al. (2014) and Nielsen (2013) find that voters support sanctions against human rights violators—but only when such violations are publicized.

Special interest groups also exert substantial influence. Sanctions create distributional consequences, affecting some economic sectors more than others. This gives rise to lobbying pressures from export industries, humanitarian NGOs, or security-related constituencies (Kaempfer and Lowenberg, 1988; Kaempfer and Lowenberg, 1992). As McGillivray and Stam (2004) and Aidt et al. (2021) note, policymakers are especially sensitive to the pressures from special interests. For example, Lektzian and Souva (2003) argue that democracies, facing diverse constituencies, use sanctions more often—sometimes as a political cover for protectionist trade measures. Conversely, as the size of the export sector increases, the probability of sanction imposition declines. Empirical work supports this logic: McLean and Whang (2014) and Lektzian and Biglaiser (2014) find a negative relationship between United States foreign direct investment and the likelihood of sanctions. Drury et al. (2014), analyzing both threats and impositions, show that greater economic interdependence increases the likelihood of threats but reduces the probability that the sender follows through.

Domestic political influence in this chapter is captured by a policymaker in the sender country with mixed motives: he maximizes a weighted sum of citizens' welfare and ego rents from reelection. The notion of mixed motives in policymaking originates with Wittman (1983), and has been further developed in models of electoral competition by Calvert (1985), Mitchell (1987), and Harrington Jr (1992). Much of this work focuses on equilibrium existence and policy outcomes in probabilistic spatial voting models and hybrid election games (Cowen and Sutter 1998; R. Ball 1999; Indridason 2008; Reny 1999; Saporiti 2008). However, little attention has been devoted to how the relative importance of office versus policy in the policymaker's utility function shapes strategic outcomes in foreign policy settings.

Turning to political agency models, a related strand of literature analyzes how incumbents' reelection incentives shape their behavior. Here, ego rents from remaining in office play a central role. Introduced by Downs (1957), ego rents are defined as the prestige, power, and material benefits associated with holding public office. Their influence on policy distortions is well documented: Congleton (1986) shows that ego rents induce inefficient campaign spending, while Rogoff (1987) demonstrates how they can distort fiscal policy for electoral advantage.

Empirical support includes findings from psychology: Watts et al. (2013) document narcissistic tendencies among United States presidents. More recently, Gustafsson (2019) argues that ego rents can lead incumbents to adopt inefficient policies if such actions help preserve power.

In my model, ego rents serve a novel and counterintuitive function: they improve welfare. When reelection is conditional on securing an agreement, ego rents may provide just enough incentive for the policymaker to overcome the strategic inefficiencies caused by asymmetric information. This mechanism turns the standard critique of rent-seeking on its head: under certain conditions, ego rents restore efficiency in international bargaining.

The model developed in this chapter further departs from the traditional crisis bargaining framework by allowing for the simultaneous use of a reward and a sanction threat. In practice, governments frequently offer conditional aid, loans, and economic privileges to induce policy changes in target countries. These rewards can take both direct forms, such as bilateral transfers and concessional loans, and indirect ones, including preferential market access, preferential trade agreements, or conditional International Monetary Fund (IMF) lending. For example, Aidt et al. (2021) document the rising prevalence of PMA and PTA programs—such as the Generalized System of Preferences and the African Growth and Opportunity Act —used to incentivize cooperation. Similarly, Andersen et al. (2006) develop a principal-agent model showing that while policy-for-aid deals are often intermediated by multilateral institutions, dominant states can manipulate these mechanisms to extract concessions. Their empirical evidence highlights how the United States, through its influence over the IMF, benefits from favorable policy adjustments in borrowing countries as a condition for receiving IMF loans.

Offering rewards alone in exchange for adjusting a specific policy is noncoercive, allowing the target country to refuse the offer. Standard policy-for-aid frameworks are built on a principal-agent setup, where the target country receives compensation precisely equal to the welfare loss it incurs from departing from its preferred, uncoordinated policy stance (see Grossman and Helpman 2001, Chapter 7). The literature on the effects of domestic politics within the sender country on rewards is relatively scarce, and it has mainly analyzed the role of special interest groups.

Keck and Sikkink (1998) argue that NGOs in industrialized countries are able to affect their governments' foreign policies and foreign aid, while Mansfield et al. (2007), covering a sample of 194 countries between 1950 and 1999, demonstrate that an increase in the number of domestic veto players within the sender country is associated with a reduced likelihood of establishing a PTA.

Despite the growing relevance of policy-for-aid strategies, research comparing their effectiveness and efficiency relative to economic sanctions remains limited. A few studies have addressed this question—most notably Dorussen (2001), Haas and O'Sullivan (2000), and Verdier and Woo (2011)—but they typically frame rewards and sanctions as substitute tools, overlooking the possibility of their simultaneous use. To the best of my knowledge, the model developed in this chapter provides the first formal framework in which sanction threats and reward offers are jointly analyzed under incomplete information and domestic political constraints.

1.3 The model

Consider two countries, indexed by $i \in \{1, 2\}$, where Country 1 is the *sender* and Country 2 is the *target*.[2] Without loss of generality, I normalize the population size in both countries to 1, and assume that each population is homogeneous.

The incumbent policymaker in Country 1 offers a per capita transfer t to the government of Country 2 in exchange for a policy concession. Simultaneously, the policymaker threatens to impose sanctions if the target refuses the offer. Both countries are assumed to comply with the outcome: if Country 2 accepts the offer, it makes the policy concession; if it rejects, Country 1 enforces the threatened sanction.

The cost of sanctions for each country depends on a state variable $\theta \in [0, 1]$, distributed according to a cumulative distribution function F with continuous and strictly positive density f on $[0, 1]$. The state variable θ captures the relative sanctioning strength of Country 1. As θ increases, the imposition of sanctions becomes

[2] More specifically, I use the term "policymaker" to refer to Country 1 and "government" to refer to Country 2.

less costly for the sender and more damaging for the target. The realization of θ is observed only by the policymaker in Country 1. The target government and the domestic agent in the sender country do not observe θ; they only know its distribution. This informational asymmetry introduces incomplete information into the game.

To study the role of domestic political constraints in shaping international bargaining outcomes, I introduce a domestic political agent in Country 1—such as a voter or lobby group—who offers the policymaker an implicit "contract" that conditions reelection on the policymaker's proposal and/or the outcome of negotiations.

Generally, a contract assigns a retention decision (i.e., reelect or not) as a function of the offer made and whether there is agreement. It can be shown, however, that I can restrict attention to two types of contracts without loss of generality. In the first one, denoted $o = \bar{t}$, the policymaker is reelected if and only if he proposes a transfer $t \geq \bar{t}$ and the target government accepts the offer; the threshold $\bar{t}$ is chosen by the domestic agent. In the second one, denoted $o = \varnothing$, the policymaker is reelected if and only if the negotiation fails and sanctions are imposed.

Timing:

The model is a sequential game of incomplete information between the incumbent policymaker in Country 1, the domestic political agent in Country 1, and the government of Country 2. The timing is as follows:

- Stage 0. Nature selects the realized value of the state variable $\theta \in [0, 1]$, which is privately observed by the policymaker in Country 1.
- Stage 1. Without observing θ, the domestic political agent in Country 1 offers a contract $o \in O$ that conditions reelection on the policymaker's actions or the outcome of negotiations.
- Stage 2. Given the realization of θ and the contract o, the policymaker in Country 1 proposes a transfer t to the government in Country 2 in exchange for the policy concession, and simultaneously threatens sanctions if the offer is rejected.
- Stage 3. Upon observing the transfer offer t, the government in Country 2 updates its beliefs about θ and decides whether to accept or reject the offer.

If it rejects, Country 1 imposes sanctions (which are not modeled explicitly, for simplicity).

- Stage 4. Payoffs are realized and the policymaker is either reelected or replaced based on the outcome and the contract o.

Individual payoff functions:

Each citizen's utility in either country depends on income, which varies depending on whether an agreement is reached. If the policy concession is accepted, individuals in both countries experience a shift in income based on transfers and concession costs. If no agreement is reached, citizens bear the costs associated with sanctions, which depend on the underlying state θ.

Let $d \in \{a, r\}$ denote Country 2's response to Country 1's transfer proposal, where $d = a$ indicates acceptance and $d = r$ denotes rejection. The payoff function of individuals in Country 1 is:

$$U_1(t, d, \theta) = \begin{cases} y_1 - t, & \text{if } d = a, \\ y_1 - c_1(\theta), & \text{if } d = r, \end{cases} \tag{1.3.1}$$

where y_1 is exogenous and constant income in Country 1, t is the per capita transfer offered to Country 2, and $c_1(\theta) \in [0, 1]$ represents the per capita cost of imposing sanctions. The function $c_1(\cdot)$ is assumed to be strictly decreasing in θ, so that stronger sender types (higher θ) face lower sanction costs. I normalize $c_1(1) = 0$, implying that the strongest type faces no cost from imposing sanctions, while weaker types incur positive costs.

The payoff function of individuals in Country 2 is:

$$U_2(t, d, \theta) = \begin{cases} y_2 + t - \gamma, & \text{if } d = a, \\ y_2 - c_2(\theta), & \text{if } d = r, \end{cases} \tag{1.3.2}$$

where y_2 is exogenous and constant income, t is the per capita reward received from Country 1, and $\gamma \in (0, 1)$ is the utility cost incurred from accepting the policy concession. The function $c_2(\theta)$ denotes the cost of facing sanctions and is assumed to be strictly increasing in θ: the stronger the sender, the harsher the

sanctions imposed on Country 2. I normalize $c_2(0) = 0$, so that when the sender is weakest, the target bears no cost from sanctions. Finally, I assume $c_1(0) < \gamma$, which guarantees that for the weakest type, the costs of imposing sanctions are lower than the cost of reaching an agreement, thereby avoiding situations in which no equilibrium can be sustained.

Policymaker's objective function:

The objective of the policymaker in Country 1 is to maximize a weighted sum of the welfare of the representative citizen and his ego rents. Formally, the policymaker maximizes:

$$\lambda U_1(t, d, \theta) + (1 - \lambda)R, \tag{1.3.3}$$

where $\lambda \in [0, 1]$ reflects the relative weight placed on social welfare compared to ego rents, and $R > 0$ denotes the rents from reelection.

In section 1.4.1 and section 1.4.2, I use as a benchmark the case of an *unconstrained* policymaker, i.e., $\lambda = 1$, who places full weight on citizen welfare and no value on ego rents. This captures a setting without electoral constraints.

Strategies and equilibrium concept:

This setup constitutes a sequential game of incomplete information among three players: (i) the incumbent policymaker in Country 1, (ii) the domestic political agent in Country 1, and (iii) the government in Country 2. The equilibrium concept employed is (pure strategy) Perfect Bayesian Equilibrium (PBE).

A strategy for the domestic political agent is the selection of a contract $o \in O$ to offer the policymaker. The policymaker is reelected if and only if the conditions specified in the contract o are met.

An action for the incumbent policymaker in Country 1 is the offer of a transfer $t \in \mathbb{R}$ to Country 2. Then, a strategy for Country 1 is a function $\sigma_1 : [0, 1] \times o \to \mathbb{R}$. In words, a strategy for the policymaker in Country 1 maps every pair (θ, o) of type and contract to an offer $\sigma_1(\theta, o)$ to Country 2.

An action for Country 2 is a decision $d \in \{a, r\}$, where a stands for "accept" and r stands for "reject". Then, a strategy for Country 2 is a function $\sigma_2 : \mathbb{R} \to \{a, r\}$. In

words, a strategy for the government in Country 2 is a function that maps an offer $t \in \mathbb{R}$ to a choice to accept or reject.

The government of Country 2 does not observe the contract o proposed by the domestic political agent in Country 1. It observes only the transfer offer t and updates its beliefs about θ accordingly. A belief system for Country 2 is a function $\eta : \mathbb{R} \to \Delta([0,1])$, where $\eta(t)$ denotes the posterior belief about the distribution of types conditional on observing offer t. Beliefs are updated according to Bayes' rule whenever possible.

A *Perfect Bayesian Equilibrium* (PBE) in this setting consists of a strategy profile (σ_1, σ_2, o) and a belief system η such that:

- Given σ_2 and η, the strategies σ_1 and o maximize the expected utility of the policymaker and the domestic political agent, respectively.
- Given σ_1, o, and η, the strategy σ_2 maximizes the expected utility of the government in Country 2.
- The belief system η satisfies Bayes' rule, whenever possible.

To rule out implausible equilibria supported by unreasonable off-path beliefs, I apply the D1 refinement introduced by Cho and Kreps (1987). According to this criterion, if the set of responses from Country 2 that would make a type-θ policymaker willing to deviate to some t' is strictly contained within the set of responses that would induce a type-θ' policymaker to do so, then Country 2 must assign zero probability to type-θ and place full belief on type-θ' following the deviation to t'.[3]

From this point forward, I use the term "equilibrium" to refer to a pure strategy PBE that satisfies the D1 criterion. Additionally, to ensure a unique outcome when players are indifferent, I assume that indifferent players choose to induce agreement rather than disagreement.

[3]See Section 1.A for a detailed explanation of Cho and Kreps (1987) D1 criterion.

1.4 Benchmark and equilibrium

This section characterizes the equilibrium outcomes of the model under different assumptions about information and domestic political constraints. I begin with the benchmark case in which the policymaker in Country 1 is unconstrained, i.e., he places full weight on citizen welfare and none on reelection rents ($\lambda = 1$). Section 1.4.1 analyzes this benchmark under complete information. Section 1.4.2 introduces incomplete information about θ. Finally, Section 1.4.3 considers domestic political constraints, where the policymaker must secure reelection to earn ego rents.

1.4.1 Perfect information without electoral constraints

The equilibrium in this benchmark is straightforward. Let $\lambda = 1$, so the policymaker in Country 1 is unconstrained and does not value reelection. He chooses between proposing a transfer that induces agreement, yielding a payoff of $y_1 - t$, or proposing a transfer that is rejected, leading to sanctions and a payoff of $y_1 - c_1(\theta)$.

If the policymaker prefers to reach agreement, he proposes the smallest transfer that makes rejection unprofitable for Country 2. This condition requires $y_2 + t - \gamma \geq y_2 - c_2(\theta)$, which simplifies to $t \geq \gamma - c_2(\theta)$. In equilibrium, the policymaker offers $t = \gamma - c_2(\theta)$, exactly the minimum amount needed for acceptance.

His choice between agreement and sanctions depends on which yields higher payoff. Agreement is preferred whenever $y_1 - (\gamma - c_2(\theta)) \geq y_1 - c_1(\theta)$, which simplifies to $\gamma \leq c_1(\theta) + c_2(\theta)$. If this condition holds, he proposes $t = \gamma - c_2(\theta)$ and agreement is reached. Otherwise, if $\gamma > c_1(\theta) + c_2(\theta)$, the policymaker prefers sanctions and offers $t < \gamma - c_2(\theta)$, leading to rejection and disagreement.

Observation 1.1. *Suppose $\lambda = 1$ and that there is perfect information about the realization of θ. Then, in state $\theta \in [0, 1]$, a disagreement outcome occurs if and only if:*

$$\gamma > c_1(\theta) + c_2(\theta) \tag{1.4.1}$$

The intuition is straightforward. Agreement requires the sender to cover the target's concession cost γ, but this is only worth it if that cost is smaller than the combined sanction costs $c_1(\theta) + c_2(\theta)$. When γ is larger, both countries prefer disagreement, and sanctions are imposed. This result contrasts with the standard crisis bargaining literature discussed in Section 1.2, where sanctions are never observed in equilibrium under complete information. Here, by introducing the possibility of combining threats and rewards, sanctions may emerge in equilibrium as the least costly option, even when all actors are fully informed.

1.4.2 Incomplete information with an unconstrained policymaker

The objective of the policymaker remains to maximize $U_1(t, d, \theta)$, as $\lambda = 1$. However, there is now incomplete information: the realization of θ is observed only by the policymaker in Country 1, while the domestic agent and the government in Country 2 do not observe it and only know its distribution. After receiving the proposed transfer t, the government in Country 2 updates its beliefs about θ. The belief system is a distribution function $\eta(\cdot \mid t)$ that specifies the probability assessment over types θ after observing a proposal t.

The government in Country 2 accepts the proposal if the expected payoff from agreement exceeds that under sanctions. This condition is:

$$t \geq \gamma - \int_0^1 c_2(\theta)\, d\eta(\theta \mid t). \tag{1.4.2}$$

The decision to accept or reject now depends on updated beliefs $\eta(\theta \mid t)$, which reflect the inferred type of Country 1's policymaker based on the proposed transfer. This introduces a signaling problem: transfers are used strategically to influence beliefs. In what follows, I show that this friction leads to an equilibrium in which sanctions are imposed in all states.

Proposition 1.1. *Suppose $\lambda = 1$. Then there exists a unique equilibrium outcome in which every proposal made by Country 1's policymaker is rejected and sanctions are implemented.*

Proof. See Section 1.B.

To understand Proposition 1.1, it is helpful to see why no agreement can be sustained in equilibrium. Consider a type-θ policymaker proposing a transfer t such that the inequality in equation (1.4.2) is satisfied. This type will offer the lowest acceptable transfer, which I denote by t_o. Any higher offer $t' > t_o$ is strictly dominated, as it increases costs without improving the outcome. Moreover, if some type proposes t_o in equilibrium, then $c_1(\theta) > t_o$ must hold—otherwise, deviating would be optimal.

Now suppose t_o is proposed by only one type, say θ_o. In that case, Country 2's beliefs would assign probability 1 to θ_o upon observing t_o. However, due to the continuity of $c_1(\theta)$, nearby types would have an incentive to mimic θ_o and also propose t_o, especially if $c_1(\theta) > t_o$. Thus, this cannot be an equilibrium unless $\theta_o = 0$. But for type $\theta = 0$ to prefer agreement over disagreement, it must hold that $c_1(0) > \gamma$, which contradicts the maintained assumption that, for the strongest policymaker, sanctions are always preferable to agreement, i.e., $c_1(0) < \gamma$.

Next, consider whether multiple types could pool at t_o. Let $\hat{\theta}$ be the highest type in the pooling set that proposes t_o. This type could profitably distinguish itself from lower-type policymakers by deviating to $t_o - \varepsilon$ for some small $\varepsilon > 0$. Upon observing such a deviation, the government of Country 2 would update its beliefs and assign probability 1 to the deviating type being $\hat{\theta}$, and probability 0 to any $\theta < \hat{\theta}$. Formally:

$$t_o \geq \gamma - \int_0^1 c_2(\theta)\, d\eta(\theta \mid t_o) > \gamma - \int_0^1 c_2(\theta)\, d\eta(\theta \mid t_o - \varepsilon). \tag{1.4.3}$$

Thus, for sufficiently small ε, the deviation would induce acceptance by Country 2, making it strictly profitable. This logic implies that pooling equilibria are not sustainable and that types are incentivized to separate—but separation is also incompatible with acceptance.

Intuitively, no agreement can be sustained because of the mimicking incentives across types. If only one type offers the minimal acceptable transfer, nearby types find it profitable to imitate, which rules out separation. If several types pool at a common transfer, the highest-cost type in the pool can profitably deviate by slightly lowering its offer, since such a deviation would still induce acceptance

and reveal its type. These forces make both pooling and separating equilibria with agreement unsustainable, implying that sanctions are the only equilibrium outcome.

In Section 1.B, I show that beliefs supporting these deviations are the only ones consistent with the D1 criterion of Cho and Kreps (1987). I also show that the only separating equilibrium is one in which all types propose transfers that are rejected. These proposals fully reveal the sender's type, but no agreement is ever accepted. Such an equilibrium is supported by beliefs assigning probability 1 to $\theta = 0$ following any off-path proposal.

A normative implication of Proposition 1.1 is that signaling under incomplete information gives rise to systematically inefficient disagreement outcomes. Specifically, for every $\theta \in [0,1]$ satisfying $\gamma - c_2(\theta) \leq c_1(\theta)$, there exists a transfer $t \in (\gamma - c_2(\theta), c_1(\theta))$ that would induce agreement and strictly improve welfare in both countries. However, such offers are not made or accepted in equilibrium due to the strategic constraints imposed by signaling incentives.

1.4.3 Incomplete information under mixed motives

The previous section demonstrated the core problem faced by the policymaker in Country 1: uncertainty over the consequences of sanctions—combined with the resulting signaling process—leads to inevitable disagreement and the imposition of sanctions. This implies inefficiencies: in some states θ, mutually beneficial agreements exist but are not reached. This inefficiency suggests that introducing domestic political constraints may improve the Pareto-efficiency of negotiation outcomes.

In this section, I relax the assumption $\lambda = 1$, allowing for the presence of domestic constraints. These constraints are modeled through a domestic political agent who offers the policymaker a contract that conditions reelection on either the agreement outcome or the proposed transfer. Two types of contracts are possible: (i) the policymaker is reelected if and only if he successfully proposes a transfer $t \geq \bar{t}$, where $\bar{t}$ is chosen by the domestic agent; or (ii) the policymaker is reelected if and only if no agreement is reached and sanctions are imposed (i.e., $o = \varnothing$).

Recall that under these domestic constraints, the policymaker has mixed motives: he maximizes a weighted sum of the representative citizen's welfare and the ego rents R from remaining in power. The key finding under this setup is formally stated in the following proposition:

Proposition 1.2. *Consider the continuation subgame in which Country 1's policymaker is confronted with a contract $o \in O$ offered by the domestic agent. Then:*

1. *If $o = \varnothing$:*
   There is a unique equilibrium in which every proposal made by Country 1's policymaker is rejected, sanctions are implemented, and the policymaker is reelected.
2. *If $o = \bar{t}$, for some $\bar{t} \in \mathbb{R}$:*
   The subgame possesses an equilibrium regardless of the value of $\bar{t}$.
   Specifically, a pooling equilibrium exists in which all types of policymakers in Country 1 propose a successful offer if and only if:

$$\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda} R\right] \tag{1.4.4}$$

   In any such equilibrium, all types of policymakers offer exactly $\bar{t}$, the proposal is accepted by Country 2, and reelection occurs.

Proof. See Section 1.C

In the first case, when the domestic agent conditions reelection on the failure of negotiations ($o = \varnothing$), the outcome replicates that of Proposition 1.1: all proposals are rejected, and sanctions are imposed. The only difference is that the policymaker now earns ego rents from being reelected. The intuition is unchanged—signaling prevents agreement, and domestic political incentives reinforce this outcome.

In the second case, where reelection is conditional on the policymaker successfully offering a transfer $t \geq \bar{t}$, a pooling equilibrium with agreement may arise if $\bar{t}$ lies within a specific range. When this condition is satisfied, all types of policymakers find it optimal to offer $\bar{t}$, as doing so secures both agreement and reelection. In this setting, ego rents serve to counteract the signaling distortions from incomplete information.

The intuition behind this result is as follows: when facing a lower bound $\bar{t}$, the

policymaker may be tempted to deviate to a slightly lower transfer $\bar{t} - \varepsilon$ to reduce costs. However, such a deviation would forfeit reelection and thus the ego rents. This trade-off between saving on the transfer and losing reelection rents is what makes an equilibrium with agreement possible. In this sense, the personal benefits from holding office can counterbalance the signaling incentives created by uncertainty about the cost of sanctions. In Section 1.C, I formally show that under this contract there can be no equilibrium with disagreement, since offering a transfer within the specified range ensures the best possible outcome for both the policymaker and the citizens in Country 1, while rejection is suboptimal for the government of Country 2. This equilibrium is supported by beliefs assigning probability 1 to all types θ proposing $\bar{t}$ in equilibrium, and probability 1 to type $\theta = 1$ for any off-path offer. These beliefs satisfy the D1 criterion.

When $\bar{t}$ lies outside the specified range, no type of policymaker is able to successfully reach an agreement. The resulting equilibrium resembles that in Section 1.4.2, with separating proposals that are always rejected.

Proposition 1.2 characterizes the equilibrium negotiation outcomes that may arise under *any* contract offered by the domestic agent. To complete the full-game equilibrium, the next step is to identify the optimal contract for the domestic agent. To this end, it is useful to consider more closely the range of $\bar{t}$ values that lead to successful agreements.

As shown in Equation (1.4.4), the lower bound of acceptable $\bar{t}$ values is determined by the minimum transfer Country 2 is willing to accept. Let $\tilde{t} \equiv \gamma - \mathbb{E}(c_2(\theta))$ denote this threshold. Since all types of Country 1's policymakers pool on the same offer in equilibrium, the government of Country 2 cannot infer θ and must rely on prior beliefs. The higher the expected cost of sanctions for Country 2, the lower the minimum transfer required to induce agreement. Conversely, the higher the policy adjustment cost γ, the greater the compensation Country 2 demands.

The upper bound for $\bar{t}$ is governed by the policymaker's incentives. A higher ego rent R and a lower weight λ on citizen welfare both raise the policymaker's willingness to offer generous transfers. Thus, the admissible range for $\bar{t}$ is increasing in R and decreasing in λ.

The domestic agent faces a binary choice: either choose $\bar{t} < \tilde{t}$, resulting in disagreement, or choose $\bar{t} = \tilde{t}$, resulting in agreement. To see why these are the only viable options, observe that choosing any $\bar{t} > \tilde{t}$ cannot be sustained in equilibrium. A policymaker would profitably deviate to $t \in (\tilde{t}, \bar{t})$, which would still be accepted but yield higher utility, since $y_1 - t > y_1 - \bar{t}$.

The domestic agent prefers selecting $\bar{t} = \tilde{t}$ only if it leads to a higher expected payoff than disagreement. This occurs if:

$$y_1 - \tilde{t} \geq y_1 - \mathbb{E}(c_1(\theta)), \tag{1.4.5}$$

Corollary 1.1. *The domestic agent induces agreement in equilibrium whenever the following two conditions are met:*

$$\tilde{t} \leq \frac{1-\lambda}{\lambda} R \tag{1.4.6}$$

$$\gamma \geq \mathbb{E}(c_1(\theta) + c_2(\theta)). \tag{1.4.7}$$

In that case, it does so by selecting $\bar{t} = \tilde{t}$.

If either condition in Corollary 1.1 fails to hold, the domestic agent can offer $o = \bar{t}$ outside the equilibrium-supporting range—resulting in rejection and sanctions—or instead $o = \varnothing$, which conditions reelection on disagreement. Since both choices yield the same expected payoff, the domestic agent is indifferent.

Equation (1.4.6) reveals that the desirability of reaching an agreement declines as $\tilde{t}$ increases. As noted earlier, $\tilde{t}$ rises with γ, the cost for Country 2 of implementing the policy change. Thus, high adjustment costs reduce the likelihood of successful negotiation.

Finally, the domestic agent's willingness to induce agreement also depends on its prior beliefs about the distribution of θ. The impact of belief shifts is generally ambiguous. Consider a first-order stochastic dominance shift in the distribution of θ—that is, a shift where higher values of θ, corresponding to stronger sender types, become more likely. On one hand, this increases the expected cost of disagreement for Country 2, thereby reducing the minimum transfer $\tilde{t}$ required to

sustain agreement. On the other hand, it lowers the expected cost of sanctions for Country 1, potentially making disagreement more appealing from the domestic agent's perspective. The overall effect of such a shift depends on the curvature and relative slope of the cost functions $c_1(\theta)$ and $c_2(\theta)$, which jointly determine how belief updates affect the relative attractiveness of sanctions versus agreement.

1.5 Implications

This chapter investigates how uncertainty about the costs of sanctions and domestic political constraints affects the outcomes of international policy interventions. In the absence of domestic constraints—when the policymaker in the sender country cares only about policy outcomes—uncertainty about sanction costs and the associated signaling process leads to systematic disagreement and the imposition of sanctions. A key consequence of incomplete information is the emergence of inefficient disagreements in states where mutually beneficial agreements are feasible, particularly when the social cost of sanctions is high for both countries relative to the cost of the policy concession.

Building on existing research showing that domestic institutions can mitigate the effects of asymmetric information, this chapter examines how domestic political actors in the sender country influence intervention outcomes. Corollary 1.1 demonstrates that when the cost of the policy concession for the target is relatively low compared to the expected sanction costs for both countries, domestic agents can induce agreement by strategically designing the contract offered to the policymaker. The underlying mechanism is that the policymaker, facing reelection constraints, may choose not to deviate to a lower transfer—even when doing so would reduce immediate costs—if such a deviation would lead to losing office. The personal gains from reelection thus counterbalance the signaling distortions caused by incomplete information, restoring Pareto efficiency.

The ego rents and office-seeking motives of the policymaker in the sender country are pivotal for these results. Their impact is captured in the following corollary:

Corollary 1.2. *Suppose condition* (1.4.7) *holds. Then:*

- *All else equal, there exists a threshold* $\bar{\lambda} > 0$ *such that an agreement is reached in equilibrium if and only if* $\lambda \leq \bar{\lambda}$.
- *All else equal, there exists a threshold* $\bar{R} > 0$ *such that an agreement is reached in equilibrium if and only if* $R \geq \bar{R}$.

Consider first the role of $\bar{\lambda}$. The result implies that when the policymaker places sufficiently high weight on reelection relative to citizen welfare, an agreement is more likely to be reached. This leads to the next corollary, which characterizes how the representative citizen's equilibrium utility depends on the policymaker's preferences.

Corollary 1.3. *Let* $U^*(\lambda)$ *denote the equilibrium utility of citizens in Country 1 as a function of* λ. *Then:*

- *If condition* (1.4.7) *is not satisfied,* $U^*(\lambda)$ *is constant.*
- *If condition* (1.4.7) *is satisfied, then* $U^*(\lambda)$ *depends on* λ. *In particular, when* $\lambda \leq \bar{\lambda}$, *an agreement is reached and voter utility is higher than when* $\lambda > \bar{\lambda}$, *in which case disagreement and sanctions occur. Hence, equilibrium citizen welfare decreases in* λ.

This result challenges conventional wisdom and prior findings. Contrary to the common view that closer alignment between citizen and policymaker preferences leads to optimal outcomes, this model shows that a certain degree of misalignment—reflected in strong office-seeking motives—can improve citizen welfare by correcting the distortions introduced by asymmetric information.

A parallel logic applies to the role of ego rents, as captured in Corollary 1.2. Keeping everything else constant, there exists a minimum threshold $\bar{R}$ such that, if the policymaker's rents fall below this level, the incentive to reach an agreement is too weak. The policymaker mimics stronger types, leading to rejection by the target country and the imposition of sanctions. However, if $R \geq \bar{R}$, the policymaker has sufficient private incentive to comply with the transfer requirement set by the domestic agent, facilitating agreement. This finding also contrasts with earlier

literature that emphasized the harmful consequences of ego rents for voters, as discussed in Section 1.2. Here, ego rents act as a corrective device, offsetting informational frictions and enabling efficient outcomes.

1.5.1 Extension: Delegated bargaining

To understand the broader implications of the results for the literature on delegated bargaining, consider a modification to the baseline model. Suppose the population in Country 1 now consists of two distinct citizen groups, denoted by $N = \{L, H\}$, who differ in their individual costs of disagreement.

Let the cost of sanctions for group N be represented by $c_{1,N}(\theta) \in [0, 1]$, strictly decreasing in θ. Group L consists of citizens who incur lower costs from disagreement, while group H bears higher costs. That is, $c_{1,L}(\theta) < c_{1,H}(\theta)$ for all θ.

Suppose now that the policymaker in Country 1 represents only one of the two groups—specifically, the group to which he belongs. His objective becomes:

$$\lambda U_{1,N}(t, d, \theta) + (1 - \lambda)R, \tag{1.5.1}$$

where $U_{1,N}(t, d, \theta)$ denotes the welfare of a citizen in group N. This welfare function is defined as in the previous sections, except that the sanction costs now correspond specifically to the group the policymaker represents.

As in previous sections, a necessary condition for equilibrium with agreement is that the policymaker prefers offering the transfer $\bar{t}$ and being reelected over proposing a lower transfer that leads to sanctions and forfeiting reelection. Put formally,

$$\lambda(y_1 - \bar{t}) + (1 - \lambda)R \geq \lambda(y_1 - c_{1,N}(\theta)), \tag{1.5.2}$$

which can be rearranged as:

$$\bar{t} \leq \frac{1 - \lambda}{\lambda} R + c_{1,N}(\theta). \tag{1.5.3}$$

Combining this upper bound with the minimum transfer the government in

Country 2 is willing to accept yields the feasible range for $\bar{t}$:

$$\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)),\ \frac{1-\lambda}{\lambda} R + c_{1,N}(\theta) \right]. \tag{1.5.4}$$

This formulation makes it clear that the feasibility of an agreement depends on the identity of the policymaker. When the policymaker belongs to a group with high costs of sanctions (i.e., group H), the upper bound of the interval increases, making agreement more likely. Conversely, a policymaker from group L, who bears lower costs, may be unwilling to offer sufficient transfers to induce agreement.

To illustrate, suppose a policymaker representing group L is in office and that $c_{1,L}(\theta)$ is sufficiently low such that:

$$\frac{1-\lambda}{\lambda} R + c_{1,L}(\theta) < \gamma - \mathbb{E}(c_2(\theta)). \tag{1.5.5}$$

In this case, no feasible transfer $\bar{t}$ exists that satisfies both the policymaker's incentive constraint and the target country's acceptance condition. As a result, no agreement is reached. In such a setting, it becomes strategically optimal for the citizens of group L to delegate bargaining authority to a representative from group H, as doing so widens the feasible range for $\bar{t}$ and enables agreement. The welfare of all citizens, including those in group L, would then improve due to the avoidance of inefficient sanctions.

This observation highlights how the choice of the representative can affect negotiation outcomes and welfare. It suggests that heterogeneous citizen groups may find it optimal to delegate bargaining authority to representatives whose preferences reflect higher costs of disagreement. This behavior aligns with the literature on strategic delegation and voting, where voters may select representatives not based solely on ideological alignment but on anticipated negotiation outcomes. Unlike much of the earlier delegation literature, which assumes complete information and shows that voters benefit from delegating to more "conservative" or "hawkish" representatives (e.g., Persson and Tabellini 1992; Segendorff 1998), the mechanism here stems from incomplete information: *delegation to a high-cost policymaker mitigates signaling frictions and makes agreements feasible*.

The presence of incomplete information also alters the mapping between voter preferences and electoral outcomes. The Condorcet winner of the election may not correspond to the median voter's true preference, but rather to the median of the stated preferences, as in strategic voting models. This departs from the predictions of classical median-voter frameworks such as Black (1948) and Downs (1957).

Strategic delegation has been extensively studied in the context of international cooperation and negotiations (Burtraw 1992, 1993; Segendorff 1998; Siqueira 2003). Using two-stage models in which voters first elect a politician who later negotiates on their behalf,[4] previous work has examined strategic delegation in monetary and tax policy, showing that citizens may elect more "conservative" representatives than their own preferences would suggest (Persson and Tabellini 1992; Dolado et al. 1994). In the context of environmental negotiations, by contrast, the direction of the delegation effect is more ambiguous, depending on model assumptions and voter preferences.[5]

Much of this literature develops two-stage complete-information models in which voters anticipate negotiation payoffs and select "tougher" or "softer" representatives to shift outcomes, abstracting from signaling frictions. By contrast, the extension here embeds delegation in an incomplete-information environment: types differ in their sanction costs, and proposals convey information. Delegation to a high-cost representative relaxes the sender's incentive to signal strength via low transfers and can increase the likelihood of an agreement. The mechanism is therefore distinct from commitment-through-type-selection in complete-information models: it operates by mitigating signaling distortions rather than by shifting known payoffs.

1.6 Concluding remarks

The objective of this chapter has been to understand how domestic political constraints shape the strategic use and effectiveness of sanctions and rewards in international bargaining under asymmetric information. While sanctions and

[4] See Persson and Tabellini (2002), Chapter 12.

[5] For a literature review, see Wangler et al. (2013).

rewards are among the most widely used instruments of foreign policy, their credibility and impact depend not only on international power relations but also on the incentives and constraints policymakers face at home. It is therefore crucial to understand how domestic institutions condition the simultaneous use of sanctions and rewards in order to explain the observed dynamics of international policy interventions.

This chapter develops a sequential game of incomplete information in which a policymaker in the sender country seeks to induce a costly policy concession from a target country by combining a reward offer with a threat of sanctions. Incomplete information arises because the policymaker privately observes the costs of imposing sanctions, creating a classic signaling problem.

In the absence of domestic constraints—when the policymaker is purely policy-driven—uncertainty about sanction costs leads to systematic bargaining failure: in equilibrium, all compensation offers are rejected and sanctions are imposed. Intuitively, this occurs because policymakers of all types have incentives to mimic stronger types by offering low transfers, which undermines credibility and drives the target to reject all proposals. Comparing this result to the complete information benchmark highlights that it is precisely asymmetric information that generates inefficient outcomes: under complete information, agreement occurs whenever the required transfer is no larger than the joint sanction costs, while under incomplete information, these mutually beneficial agreements collapse.

Introducing domestic political constraints fundamentally changes this logic. In the model, a domestic political agent can condition the policymaker's reelection on either the failure of negotiations (sanctions) or on the successful offer of a reward above a minimum threshold. When reelection is tied to the sanction outcome, the equilibrium replicates the baseline: sanctions are imposed. By contrast, when reelection is conditioned on a reward above a threshold, ego rents from office act as a commitment device that offsets signaling frictions and restores efficiency. A pooling equilibrium with agreement exists whenever the threshold lies between the target's acceptance floor (based on prior beliefs) and the policymaker's private incentive to deviate. The analysis shows how the likelihood of agreement depends on political primitives: the weight the policymaker places on voter welfare relative

to ego rents, and the magnitude of the rents from office.

Contrary to much of the political economy literature, which typically views office-seeking motives and misalignment between citizens and policymakers as a source of social inefficiency, the results here demonstrate the opposite. A certain degree of misalignment—captured by stronger office-seeking motives—can improve voter welfare by increasing the credibility of transfers and reducing inefficient sanctions. Equivalently, larger ego rents expand the set of states in which agreements can be sustained when threats and rewards are deployed jointly.

The extension on delegated bargaining further illustrates how representation matters. When citizens differ in their costs of sanctions, the choice of who represents them in negotiations becomes crucial. Electing a policymaker from the high-cost group increases the willingness to offer transfers and thus the probability of agreement. This provides a novel rationale for strategic delegation: voters may prefer representatives with preferences different from their own, not because such representatives are "tougher," as in classical models of commitment under complete information, but because under incomplete information they mitigate signaling frictions. This mechanism departs from the traditional delegation literature and offers a new explanation for why voters may sometimes strategically delegate to "weaker" negotiators.

The analysis carries several implications for international bargaining. First, the domestic political environment—the size of ego rents, the structure of electoral incentives, and the identity of representatives—is central to determining whether sanctions produce costly deadlock or whether rewards sustain cooperation. This helps account for why some sanction threats result in agreements while others devolve into prolonged disputes. Second, the results yield testable implications. In particular, increases in office rents or in the electoral salience of foreign policy should raise the incidence of reward-based settlements conditional on observed threats. Similarly, institutional arrangements that empower domestic actors to impose transparent minimum-offer thresholds (such as legislative mandates tied to human rights or governance benchmarks) should increase the probability of agreement at the threat stage.

Two important limitations suggest avenues for future research. First, the analysis

abstracts from the sender's decision of whether to comply with the threat of sanctions if the target refuses concessions. In practice, sanction threats are often not implemented even when the target resists. For example, China repeatedly threatened sanctions against the United States in response to arms sales to Taiwan (2010, 2015, 2019, 2020), but no sanctions were ever imposed. Similarly, during the 1980s and 1990s the United States threatened to revoke China's Most Favored Nation trade status unless human rights conditions improved, but the status was renewed annually despite continued violations. Such cases underscore the importance of modeling the sender's option to back down, and the reputational consequences this has for credibility in subsequent interactions.

Second, the prevalence of multilateral sanctions calls for extending the analysis to sender coalitions. According to the Global Sanctions Database, roughly one third of sanctions are multilateral, imposed by institutions such as the United Nations, the European Union, or coalitions of states. In such cases, the costs of imposing sanctions are shared across senders, lowering per-country sanction costs, while the costs for the target increase because multilateral sanctions close off substitution opportunities and carry greater legitimacy. Coalitions may also enhance the credibility of threats, since backing down carries reputational costs within the coalition. Incorporating multilateral sanctions into theoretical models would therefore refine predictions about when agreements are feasible and when sanction threats are likely to succeed.

Overall, this chapter demonstrates that both domestic political constraints and asymmetric information play a crucial role in shaping the credibility and effectiveness of sanctions and rewards. By highlighting the conditions under which domestic incentives improve rather than hinder bargaining efficiency, the analysis offers a new perspective on the political economy of foreign policy interventions and opens promising directions for both theoretical and empirical research.

Appendix

1.A Reasonable beliefs and criterion D1

To eliminate PBE that rely on implausible off-path beliefs, I restrict attention to those that satisfy the D1 refinement introduced by Cho and Kreps (1987).

Fix an equilibrium, and let $G^*(\theta, t \mid d)$ denote the payoff for a type-θ policymaker in Country 1 under the equilibrium, when he proposes a transfer t. According to the D1 criterion, I ask: which types of Country 1's policymakers could plausibly be thought to choose an off-the-equilibrium offer t'?

Let $MBR(\eta, t')$ denote the set of mixed best responses by the government in Country 2 to offer θ', given that it holds belief η over Country 1's type.

Define $D_\eta(\theta, t')$ as the subset of $MBR(\eta, t')$ by Country 2 that gives type-θ a strict incentive to deviate from the equilibrium strategy to t'. That is, given belief η, the expected payoff to a type-θ policymaker from proposing t', when Country 2 responds with strategy α, exceeds the equilibrium payoff $G^*(\theta, t \mid d)$.

Let:

$$D(\theta, t') = \bigcup_{\eta} D_\eta(\theta, t'), \tag{1.A.1}$$

denote the set of all possible responses by Country 2—across all beliefs η—that would induce type-θ to strictly prefer deviating to t'. In other words, that is the full set of conditions under which the type-θ policymaker would find it profitable to deviate.

Next, let $D^o(\theta, t')$ be the set of responses by Country 2 to t' that leave type-θ exactly indifferent between deviating and following the equilibrium strategy. Formally, these are strategies α that yield an expected payoff equal to $G^*(\theta, t \mid d)$.

Under the D1 criterion, type θ is eliminated as a plausible deviator to the off-equilibrium offer t' if there is a type $\theta' \neq \theta$ for which it holds that:

$$[D(\theta, t') \cup D^o(\theta, t')] \subset D(\theta', t'). \tag{1.A.2}$$

That is, if the set of Country 2's responses that would make type-θ willing (or indifferent) to deviate is strictly smaller than the set of responses that make type-θ' strictly prefer the deviation, then Country 2 should place all belief on θ' being the true type behind the off-equilibrium proposal t'.

1.B Proof of Proposition 1.1

Roadmap. I divide the proof of Proposition 1.1 into two parts. In the first part, using Lemma 1.1, I prove in two claims that (i) there exists an equilibrium in which Country 2 always rejects the agreement, and (ii) this equilibrium satisfies Criterion D1. In the second part, I use Lemma 1.2 and Lemma 1.3 to show that agreement is never a possible equilibrium outcome. First, I demonstrate that if an offer were successful in equilibrium, it must come from a single type θ; then, I show that no such successful offer can exist in equilibrium.

Lemma 1.1. *There exists an equilibrium in which Country 2 always rejects the agreement with Country 1.*

Proof. I proceed in two steps: first, I show the existence of the equilibrium; second, I prove that it satisfies Criterion D1.

Claim 1.1. *An equilibrium in which Country 2 always rejects the agreement with Country 1 exists.*

Let $k > 0$ be an arbitrarily small positive number. The strategy of each type-θ policymaker in Country 1 is to propose a transfer

$$t(\theta) = \gamma - c_2(\theta) - k. \tag{1.B.1}$$

That is, the policymaker offers a transfer that (i) fully reveals his type and (ii) lies strictly below the minimum value that Country 2 would accept under agreement. The strategy of the government in Country 2 is to reject any transfer offer made by Country 1.

This pair of strategies constitutes a PBE when combined with the following belief system: upon receiving an offer $t(\theta)$, Country 2 assigns posterior belief $\eta(\theta, t(\theta)) = 1$ to type θ; for any offer not in the range $t([0,1])$, Country 2 assigns probability 1 to $\theta = 0$.

Note first that when confronted with transfer $t(\theta)$, the optimal strategy for Country 2 is to reject the offer. The reason is that, as stated in Section 1.4.1, the government in Country 2 receives a higher payoff under agreement only if $t > \gamma - c_2(\theta)$.

Consider first a deviation of Country 1 to an off-the-equilibrium-path proposal $t' < t(\theta)$. As this offer is still below the acceptance level for Country 2, the response of Country 2 is to reject the offer. Deviating is therefore not profitable for the policymaker in Country 1.

Now consider a deviation to a transfer $t' > t(\theta)$. Under the defined posterior beliefs $\eta(\theta, t)$, the government in Country 2 assigns probability 1 to type $\theta = 0$ being the one deviating. Given the assumption $c_2(0) = 0$, the government in Country 2 will not accept the offer, as it would not incur any cost from sanctions under disagreement. Thus, the optimal strategy of the government in Country 2 when observing any offer outside $t([0,1])$ is to reject the agreement and opt for sanctions. $\diamondsuit$

Claim 1.2. *The equilibrium in which Country 2 always rejects the agreement satisfies criterion D1.*

Let θ be an arbitrary type and let t' be an off-the-equilibrium-path transfer proposal. Let α denote the mixed-strategy best response to t' under which type-θ in Country

1 is weakly better off deviating than following the equilibrium strategy. This corresponds to the condition:

$$\alpha(y_1 - t') + (1 - \alpha)(y_1 - c_1(\theta)) \geq y_1 - c_1(\theta), \tag{1.B.2}$$

which simplifies to:

$$\alpha(c_1(\theta) - t') \geq 0. \tag{1.B.3}$$

Country 2's mixed-strategy best responses that fulfill this condition is defined as:

$$D(\theta, t') = \left\{\alpha \in [0,1] : \alpha(c_1(\theta) - t') \geq 0\right\}. \tag{1.B.4}$$

Under the belief system defined in Claim 1.1, Criterion D1 requires that, for any $\theta > 0$:

$$[D(0, t') \cup D^o(0, t')] \supseteq D(\theta, t'). \tag{1.B.5}$$

Two cases are relevant:

- If $t' > c_1(0)$, then $c_1(\theta) - t' < 0$ for all $\theta \in [0,1]$, since $c_1(\cdot)$ is strictly decreasing and $c_1(0)$ is maximal. The inequality $\alpha(c_1(\theta) - t') \geq 0$ does not hold for any $\alpha \in [0,1]$, so $D(\theta, t') = \emptyset$ for all types. In this case, the inclusion condition holds trivially.
- If $t' < c_1(0)$, then $D(0, t') \cup D^o(0, t') = [0,1]$. Then, for $\theta > 0$, either $D(\theta, t') \cup D^o(\theta, t') = [0,1]$ or $D(\theta, t') = \emptyset$, as $c_1(\theta)$ decreases in θ.

Therefore, the belief system $\eta(\theta, t)$ satisfies Criterion D1 under the equilibrium in which Country 2 always rejects the agreement. $\diamondsuit$

Together with Claim 1.1, this completes the proof of Lemma 1.1. □

To complete the proof of Proposition 1.1, it remains to show that an agreement is never a possible outcome. Suppose, contrary to the equilibrium, that there exists a transfer that can be successfully agreed upon in equilibrium.

Lemma 1.2. *If an offer is successful in equilibrium, then it must emanate from a single type.*

Proof. The proof of Lemma 1.2 proceeds in three claims.

Claim 1.3. *For each type-θ, only one offer can successfully be made in equilibrium.*

Suppose, for contradiction, that two different transfers t_1 and t_2 lead to successful agreements in equilibrium, made by types θ_1 and θ_2, respectively. Without loss of generality, assume $t_1 < t_2$.

By equilibrium behavior, type θ_2 must find it optimal to offer t_2, so $t_2 \leq y_1 - c_1(\theta_2)$. But this implies that θ_2 could profitably deviate to t_1, since $t_1 < t_2 \leq y_1 - c_1(\theta_2)$. This contradicts the optimality of t_2 in equilibrium. $\diamondsuit$

Suppose now that multiple types make a successful offer, denoted t_o, in equilibrium. From Claim 1.3, all other offers in equilibrium must be rejected.

Define the set $\Theta_o \equiv \{\theta \in [0,1] : \sigma_1(\theta) = t_o\}$, and let $\hat{\theta}$ be the supremum of Θ_o.

Claim 1.4. $\Theta_o = [0, \hat{\theta}]$

Suppose, to the contrary, that there exists a type $\theta < \hat{\theta}$ such that $\sigma_1(\theta) \neq t_o$. By definition of a PBE, $t_o \leq c_1(\hat{\theta}) < c_1(\theta)$. This implies that the type-θ policymaker in Country 1 could profitably deviate by offering t_o, contradicting equilibrium. Thus, it must be that $\Theta_o = [0, \hat{\theta})$.

Now consider whether $\hat{\theta} \in \Theta_o$. Suppose, instead, that $\hat{\theta} \notin \Theta_o$. Then, by definition of a PBE, $c_1(\hat{\theta}) \leq t_o$. By continuity of $c_1(\theta)$, this implies $c_1(\hat{\theta}) = t_o$, since otherwise there exists $\varepsilon > 0$ such that $c_1(\hat{\theta} - \varepsilon) < t_o$, contradicting the definition of $\hat{\theta}$ as the supremum. As $\hat{\theta} \notin \Theta_o$, this violates the maintained assumption that an indifferent policymaker always prefers agreement. This completes the proof of the claim. $\diamondsuit$

Claim 1.5. *For any $t \in (\gamma - c_2(\hat{\theta}), t_o)$ and any $\theta < \hat{\theta}$, the following inclusion holds:*

$$[D(\theta, t) \cup D^o(\theta, t)] \subset D(\hat{\theta}, t). \tag{1.B.6}$$

Fix $\theta \in \Theta_o \setminus \{\hat{\theta}\}$ and consider a deviation to some transfer $t \in (\gamma - c_2(\hat{\theta}), t_o)$.

Let α denote the mixed-strategy best response of Country 2 when receiving the offer t such that type-θ is weakly better off deviating than following the equilibrium play. Then the following incentive constraint must be satisfied:

$$\alpha(y_1 - t) + (1 - \alpha)(y_1 - c_1(\theta)) \geq y_1 - t_o.$$

Rearranging this expression yields:

$$\alpha \geq \frac{c_1(\theta) - t_o}{c_1(\theta) - t} \in (0, 1). \tag{1.B.7}$$

Define $\overline{\alpha}(x) = \frac{x - t_o}{x - t}$, which is strictly increasing in x as long as $t < t_o$. Since $c_1(\cdot)$ is strictly decreasing and $\theta < \hat{\theta}$, we have $c_1(\theta) > c_1(\hat{\theta})$, implying:

$$\alpha(\theta) = \frac{c_1(\theta) - t_o}{c_1(\theta) - t} > \frac{c_1(\hat{\theta}) - t_o}{c_1(\hat{\theta}) - t} = \alpha(\hat{\theta}). \tag{1.B.8}$$

This inequality implies that any response strategy α that induces type θ to weakly prefer the deviation also induces type $\hat{\theta}$ to strictly prefer the same deviation. Therefore:

$$[D(\theta, t) \cup D^o(\theta, t)] \subset D(\hat{\theta}, t). \tag{1.B.9}$$

The strictness of the inclusion follows from the fact that, for any $\alpha \in (0, 1)$, there exists a belief system η under which α constitutes a best response. As a result, for every α within the open interval $(\alpha(\hat{\theta}), \alpha(\theta))$:

$$D(\hat{\theta}, t) \setminus (D(\theta, t) \cup D^o(\theta, t)). \tag{1.B.10}$$

Thus, the proof of the claim is completed. $\diamondsuit$

Now suppose type $\hat{\theta}$ deviates to some transfer $t \in (\gamma - c_2(\hat{\theta}), t_o)$, a nonempty interval established in the preceding claim. By the D1 criterion, Country 2 updates its belief and assigns probability one to the deviator being of type $\hat{\theta}$. Since $t < c_1(\hat{\theta})$, the deviation yields a strictly higher payoff than the equilibrium offer, contradicting incentive compatibility. This concludes the proof of Lemma 1.2. $\square$

Lemma 1.3. *In any equilibrium, an agreement is never a possible outcome.*

Proof. Suppose there exists a unique type $\theta_o \in [0,1]$ that offers transfer t_o and for which agreement is the outcome. By Claim 1.3, this must be the only type making a successful offer in equilibrium.

By Bayes' Rule, the offer t_o perfectly reveals the type of the policymaker in Country 1. For this equilibrium to exist, it must be that $\theta_o = 0$, since otherwise, by continuity of $c_1(\theta)$, types in a neighborhood of θ_o would find it profitable to deviate to t_o.

However, for the policymaker of type $\theta_o = 0$ to prefer an agreement at t_o, it must be that $c_1(0) > \gamma$. This contradicts the assumption that $c_1(0) < \gamma$.

The proof of Lemma 1.3 is completed. □

The proof of Proposition 1.1 is obtained by combining Lemma 1.1 to Lemma 1.3. ■

1.C Proof of Proposition 1.2

Observe first that under the first form of contract, where the domestic agent offers $o = \varnothing$, no lower bound is imposed on the transfer. Hence, the logic of Proposition 1.1 applies directly: the unique equilibrium outcome is a disagreement with sanctions.

Roadmap The results under the second form of contract—where the domestic agent offers reelection if and only if the policymaker in Country 1 successfully secures an agreement with $t \in [\bar{t}, \infty)$—are established in two lemmas.

I show in Lemma 1.4 that if $\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$, then, an equilibrium in which all types of the policymaker in Country 1 propose $\bar{t}$ exists, and the agreement is accepted by Country 2. Moreover, this equilibrium satisfies Criterion D1. I also show that if $\bar{t} \notin \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$, there exists an equilibrium in which Country 1's proposals are always rejected, and that this equilibrium also satisfies Criterion D1.

In Lemma 1.5, I show that when $\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$, no type of policymaker proposes a transfer that is rejected in equilibrium. Conversely, when $\bar{t} \notin \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$, no agreement occurs in equilibrium.

Lemma 1.4. *The following statements are true:*

1. *Suppose $\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$. Then there exists an equilibrium in which all types successfully propose $\bar{t}$.*
2. *Suppose $\bar{t} \notin \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$. Then there exists an equilibrium in which disagreement is the only possible outcome for all types θ.*

Proof. To prove the first statement, suppose the strategy of every type-θ policymaker in Country 1 is to offer the transfer $\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$, i.e., $\sigma_1(\theta) = \bar{t}$ for all θ. Upon receiving $\bar{t}$, the strategy of the government in Country 2 is to accept the transfer and implement the policy concession. Upon receiving any transfer $t' \neq \bar{t}$, the government in Country 2 accepts if and only if $t' \geq \gamma - c_2(1)$.

These strategies constitute a PBE together with a belief function $\eta(\theta, t)$ for Country 2's government, which equals the prior $F(\theta)$ when observing $\bar{t}$, and assigns probability one to type $\theta = 1$ for any off-path proposal $t' \neq \bar{t}$.

Confronted with the transfer $\bar{t}$, the optimal strategy of Country 2 is to accept the transfer. Since all types of Country 1 policymakers offer the same transfer $\bar{t}$, the government in Country 2 cannot update its beliefs according to Bayes' rule on the type of θ and thus accepts an offer t if and only if $t \geq \tilde{t} \equiv \gamma - \mathbb{E}(c_2(\theta))$, which holds by construction.

Now consider a deviation by Country 1's policymaker to some $t' > \bar{t}$. If this offer is accepted, then:

$$\lambda(y_1 - \bar{t}) + (1-\lambda)R > \lambda(y_1 - t') + (1-\lambda)R \tag{1.C.1}$$

so the deviation is not profitable. If t' is rejected, then the policymaker receives $y_1 - c_1(\theta)$ instead of $y_1 - \bar{t}$. Since $\bar{t} \leq c_1(\theta)$ must hold for $\bar{t}$ to be offered in equilibrium, the deviation to t' is also unprofitable.

Next, consider a deviation of the policymaker in Country 1 to $t' < \bar{t}$. For the deviation to be profitable, it must either be accepted or yield higher utility if rejected. A necessary condition for the policymaker to prefer $\bar{t}$ with agreement to a rejected $t' < \bar{t}$ is:

$$\bar{t} \leq \frac{1-\lambda}{\lambda} R + c_1(\theta) \tag{1.C.2}$$

Under the specified beliefs, off-path proposals are interpreted as coming from type $\theta = 1$, so the condition becomes:

$$\bar{t} \leq \frac{1-\lambda}{\lambda} R \tag{1.C.3}$$

which holds by construction. Similarly, to prevent a profitable deviation to $t' < \bar{t}$ that is accepted, it must be that:

$$\bar{t} \leq \gamma - c_2(1) + \frac{1-\lambda}{\lambda} R \tag{1.C.4}$$

which is again satisfied by the definition of the interval for $\bar{t}$.

Finally, consider the strategy of the government in Country 2 when observing $t' \neq \bar{t}$. It assigns probability 1 to type $\theta = 1$ and thus believes its cost of disagreement is $c_2(1)$. Let $\dot{t} \equiv \gamma - c_2(1)$ denote the minimum acceptable transfer under this belief. Then, Country 2 accepts the offer if and only if $t' \geq \dot{t}$.

The system of beliefs assigning probability 1 to $\theta = 1$ upon observing any off-the-equilibrium-path proposal satisfies Criterion D1. To see this, consider the set of mixed strategies α that would make a given type θ strictly prefer deviating to t' over sticking to the equilibrium offer $\bar{t}$:

$$\alpha \left[\lambda U_1(\theta, t'|a)\right] + (1-\alpha) \left[\lambda U_1(\theta, t'|r)\right] \geq \lambda U_1(\theta, \bar{t}|a) + (1-\lambda)R \tag{1.C.5}$$

$$\Leftrightarrow \alpha(y_1 - t') + (1-\alpha)(y_1 - c_1(\theta)) \geq y_1 - \bar{t} \tag{1.C.6}$$

This reduces to:

$$\alpha \geq \frac{c_1(\theta) - \bar{t}}{c_1(\theta) - t'} + \frac{1-\lambda}{\lambda} R \in (0,1) \tag{1.C.7}$$

The function $\bar{\alpha}(x) \equiv \frac{x-\bar{t}}{x-t'}$ is strictly increasing in x, as $\bar{t} > t'$. Together with $c_1(\theta)$

being strictly decreasing in θ, it follows that $c_1(1) < c_1(\theta)$ for all $\theta < 1$. Hence, equation (1.C.7) implies:

$$\hat{\alpha}(\theta) \equiv \frac{c_1(\theta) - \bar{t}}{c_1(\theta) - t'} > \frac{c_1(1) - \bar{t}}{c_1(1) - t'} \equiv \hat{\alpha}(1) \tag{1.C.8}$$

It follows that the incentive to deviate to $t' < \bar{t}$ is strictly higher for type $\theta = 1$ than for any lower θ, so the posterior belief assigning probability 1 to $\theta = 1$ after observing any deviation from $\bar{t}$ satisfies Criterion D1.

This completes the proof that, under the condition of the first statement, an equilibrium exists in which all types of Country 1 successfully propose $\bar{t}$, and that this equilibrium satisfies Criterion D1.

To prove the second statement, consider the two cases: first, when $\bar{t} < \gamma - \mathbb{E}(c_2(\theta))$, and second, when $\bar{t} > \frac{1-\lambda}{\lambda} R$.

Assume first that $\bar{t} < \gamma - \mathbb{E}(c_2(\theta))$. The strategy of each type-θ policymaker in Country 1 is to offer the transfer $\bar{t}$. That is, $\sigma_1(\theta) = \bar{t}$ for all θ. The strategy of the government in Country 2 is to reject any transfer offered by Country 1.

This pair of strategies constitutes a PBE together with a belief function $\eta(\theta, t)$ for the government in Country 2, which assigns the prior $F(\theta)$ upon observing $\bar{t}$ and assigns probability one to type $\theta = 0$ for any off-the-equilibrium-path proposal.

Note first that, when confronted with the transfer $\bar{t}$, the optimal strategy for the government in Country 2 is to reject the offer. Under poling, beliefs are not updated, so Country 2 accepts a transfer t only if $t \geq \tilde{t} \equiv \gamma - \mathbb{E}(c_2(\theta))$, which is not satisfied by $\bar{t}$.

Next, consider deviations by the policymaker in Country 1. First, suppose a deviation to $t' < \bar{t}$. Under the belief system, the government in Country 2 assigns probability one to $\theta = 0$ being the deviating type. Define $\ddot{t} \equiv \gamma - c_2(0)$. Since $\bar{t} < \tilde{t} < \ddot{t}$, the government in Country 2 rejects the offer. Hence, the deviation is unprofitable: the policymaker still receives $y_1 - c_1(\theta)$.

Now consider a deviation to $t' > \bar{t}$. Two subcases arise. If $t' \in (\bar{t}, \ddot{t})$, the offer remains below the acceptance threshold implied by beliefs about type $\theta = 0$.

Therefore, Country 2 still rejects the offer, and the policymaker's payoff remains $y_1 - c_1(\theta) < y_1 - \bar{t}$, so the deviation is unprofitable. If $t' > \ddot{t}$, then the policymaker still earns his ego rents if the offer is accepted. However, by assumption, $c_2(0) = 0$, so the government in Country 2 faces no cost from disagreement and thus rejects the offer. As before, the policymaker receives $y_1 - c_1(\theta)$, and the deviation is unprofitable.

In all cases, the policymaker in Country 1 has no profitable deviation from $\bar{t}$, and the government in Country 2 always rejects. Hence, this constitutes a PBE.

To find the second equilibrium, assume now $\bar{t} > \frac{1-\lambda}{\lambda} R$.

The strategy of each type-θ policymaker in Country 1 is to offer a transfer $t'(\theta) = t(\theta) - \delta$, where $t(\theta) = \gamma - c_2(\theta) - k$ with $k > 0$ as defined in the proofs of Proposition 1.1, and $\delta = \bar{t} - \ddot{t}$. That is, the policymaker offers a type-revealing transfer strictly below $\bar{t}$ and also below the minimum level the government in Country 2 is willing to accept. The strategy of the government in Country 2 is to reject any offer from Country 1, whether $t'(\theta)$ or any other off-the-equilibrium path proposal.

This pair of strategies constitutes a PBE together with a belief function $\eta(\theta, t)$ such that the government in Country 2 assigns probability 1 to type θ when observing $t'(\theta)$, and assigns probability 1 to type $\theta = 0$ when observing any $t \notin \{t'(\theta) : \theta \in [0, 1]\}$.

When confronted with the transfer $t'(\theta)$, the government in Country 2 rejects the offer. By construction, $t'(\theta) < \gamma - c_2(\theta)$, so the expected payoff from acceptance is strictly lower than from disagreement.

Now consider potential deviations by the policymaker in Country 1. First, suppose a deviation to some $t' < \ddot{t}$ with $t' \neq t'(\theta)$. The government in Country 2, under its belief system, assigns probability 1 to $\theta = 0$ and rejects the offer. Hence, the policymaker's payoff remains $y_1 - c_1(\theta)$, and the deviation is not profitable.

Next, suppose a deviation to $t' > \ddot{t}$. The government in Country 2 still rejects the offer. Since $c_2(0) = 0$ by assumption, rejection yields a higher payoff than

acceptance for the target country. Consequently, the deviation does not alter the outcome, and is again not profitable for the policymaker in Country 1.

It remains to prove that the system of beliefs attaching probability 1 to $\theta = 0$ when observing an off-the-equilibrium-path proposal satisfies Criterion D1. To verify this, consider $D(\theta, t')$, the set of mixed strategies α by Country 2 that make a given type-θ policymaker better off deviating to $t' > \bar{t}$.

Fix any type θ, and consider an off-equilibrium transfer proposal t'. Let α denote the mixed-strategy response by Country 2 to t' under which type-θ weakly prefers deviating to t' rather than adhering to the equilibrium transfer. Formally:

$$\alpha(y_1 - t') + (1 - \alpha)(y_1 - c_1(\theta)) \geq y_1 - c_1(\theta) \tag{1.C.9}$$

This inequality simplifies to:

$$\alpha(c_1(\theta) - t') \geq 0 \tag{1.C.10}$$

The set of mixed strategies that would make type-θ strictly prefer deviating from equilibrium is therefore:

$$D(\theta, t') = \{\alpha : \alpha(c_1(\theta) - t') \geq 0\} \tag{1.C.11}$$

Under the belief system described above, which attaches probability 1 to $\theta = 0$ following any off-path offer, it follows that:

$$[D(0, t') \cup D^o(0, t')] \supseteq D(\theta, t') \tag{1.C.12}$$

This shows that the system of beliefs $\eta(\theta, t)$, which assigns probability 1 to $\theta = 0$ upon observing any off-the-equilibrium-path proposal, satisfies Criterion D1. This completes the proof of Lemma 1.4. □

The proof of Proposition 1.2 is concluded by establishing the following lemma:

Lemma 1.5. *The following holds:*

1. *If* $\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$*, then all types of Country 1's policymakers successfully offer* $\bar{t}$ *in any equilibrium.*
2. *If* $\bar{t} \notin \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$*, then all types make unsuccessful offers in any equilibrium.*

Proof. Suppose an offer t is successfully made in equilibrium. From Claim 1.3, it follows that t must be the only successful offer in equilibrium. In what follows, I prove that $t = \bar{t}$.

Claim 1.6. $t = \bar{t}$.

Suppose first that $t > \bar{t}$. Then any policymaker offering t could profitably deviate to some $t' \in (\bar{t}, t)$, which would still be accepted by Country 2 while yielding a strictly higher payoff. This contradicts the assumption that t is successfully proposed in equilibrium.

Now suppose $t < \bar{t}$. Following the logic of Proposition 1.1, the signaling incentives generated by incomplete information would re-emerge, leading to a separating equilibrium in which no offer is accepted. This contradicts the assumption that t is successfully proposed in equilibrium. Therefore, it must be that $t = \bar{t}$.$\diamondsuit$

Claim 1.7. *Suppose* $\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda}R\right]$*. No type makes an unsuccessful offer in equilibrium.*

Suppose that multiple types make a successful offer $\bar{t}$ in equilibrium, which is, by Claims 3 and 6, the only successful offer in equilibrium. Let $\Theta_o \equiv \{\theta \in [0,1] \mid \sigma_1(\theta) = \bar{t}\}$, and denote $\hat{\theta}$ be the supremum of Θ_o.

From Claim 1.4, it follows that $\Theta_o = [0, \hat{\theta}]$, so that every type $\theta < \hat{\theta}$ successfully proposes $\bar{t}$ in equilibrium. Suppose now that there exists some $\theta' \in (\hat{\theta}, 1]$ that makes an unsuccessful offer t', so that $\hat{\theta} < 1$. Then there exists an arbitrarily small $\varepsilon > 0$ such that type $\hat{\theta} + \varepsilon$ makes an unsuccessful offer.

Let $t^c(\theta)$ denote Country 2's reservation transfer when it knows with certainty that the proposing type is θ. Note that at $\hat{\theta}$, the reservation transfer of Country

2 is strictly below $\bar{t}$, i.e., $t^c(\hat{\theta}) < \bar{t}$. To see this, observe that since $t^c(\theta)$ is strictly decreasing, the inequality $t^c(\hat{\theta}) \geq \bar{t}$ would imply $t^c(\theta) \geq \bar{t}$ for all $\theta \in [0, \hat{\theta}]$, which contradicts the fact that Country 2 accepts $\bar{t}$ in equilibrium.

Let type $\hat{\theta} + \varepsilon$ deviate to transfer $\bar{t} - \eta$, where $\eta > 0$. Suppose first that this deviation is off-the-equilibrium-path. The belief system assigns probability 1 to type $\theta = 1$ being the one deviating. Given this belief and the assumption that $c_2(1) = 0$, Country 2 is willing to accept any offer above γ. Since $\bar{t} - \eta > \gamma$, the offer is accepted. Therefore, the deviation is profitable for type $\hat{\theta} + \varepsilon$, contradicting equilibrium.

Suppose instead that the deviation to $\bar{t} - \eta$ is made on the equilibrium path. Then, by Bayes' Rule, Country 2 must assign probability only to types $\theta > \hat{\theta}$. Since $t^c(\theta)$ decreases in θ, it follows that $t^c(\theta) < t^c(\hat{\theta}) < \bar{t}$ for all $\theta > \hat{\theta}$. Hence, the government in Country 2 is willing to accept $\bar{t} - \eta$, and the deviation remains profitable for type $\hat{\theta} + \varepsilon$, again contradicting equilibrium.

This proves that all types θ must successfully offer $\bar{t}$ in equilibrium when

$$\bar{t} \in \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda} R\right] \tag{1.C.13}$$

$\diamondsuit$

Claim 1.8. *Suppose now $\bar{t} \notin \left[\gamma - \mathbb{E}(c_2(\theta)), \frac{1-\lambda}{\lambda} R\right]$. No type makes a successful offer in equilibrium.*

Suppose a transfer t is successfully offered. From Claim 1.6, it must be that $t = \bar{t}$.

Assume first $\bar{t} < \gamma - \mathbb{E}(c_2(\theta))$. Following an argument analogous to Claim 1.6, the equilibrium must be a pooling one, so that the government in Country 2 holds prior beliefs about θ, as beliefs cannot be updated via Bayes' Rule. Under these priors, the government in Country 2 does not accept a transfer $\bar{t} < \gamma - \mathbb{E}(c_2(\theta))$, so an agreement cannot occur in equilibrium.

Now consider $\bar{t} > \frac{1-\lambda}{\lambda} R$. Note that a type $\theta = 1$ policymaker strictly prefers to offer a transfer that is rejected by the government in Country 2 over offering $\bar{t}$ and having it accepted. To see this, note that a rejected offer yields a payoff

of $\lambda(y_1 - c_1(\theta))$, whereas an accepted offer $\bar{t}$ yields $\lambda(y_1 - \bar{t}) + (1 - \lambda)R$. The policymaker prefers rejection if:

$$\lambda(y_1 - c_1(\theta)) > \lambda(y_1 - \bar{t}) + (1 - \lambda)R \tag{1.C.14}$$

which simplifies to:

$$\bar{t} > c_1(\theta) + \frac{1 - \lambda}{\lambda} R \tag{1.C.15}$$

For type $\theta = 1$, with $c_1(1) = 0$ by assumption, this condition becomes $\bar{t} > \frac{1-\lambda}{\lambda} R$. Therefore, type 1 strictly prefers to induce disagreement over accepting $\bar{t}$, implying that such an offer cannot be sustained in equilibrium. $\diamondsuit$

Lemma 1.5 is completed. □

The proof of Proposition 1.2 is completed by Lemmas 1.4 and 1.5. ■

Chapter 2

Buying Votes without a Numeraire: Dynamic Stability with Preference Uncertainty

This chapter is joint work with Professor Vincent Anesi.

2.1 Introduction

Agreements in the context of private contracting typically rely on two fundamental pillars: they specify monetary transfers between the parties involved, and they allow these parties to credibly commit to such transfers conditional on performance. Commitment and the availability of monetary transfers are what enable private actors to reach efficient agreements. This is particularly important in settings with informational asymmetries—when one side privately knows its costs, valuations, or effort—because, absent instruments to align incentives, mutually beneficial trades either fail to occur or take inefficient forms. A transferable numeraire makes it possible to design contingent transfers—such as prices, bonuses, penalties, and deductibles—that induce truthful revelation of private information, while credible commitment allows the parties to bind future payments to verifiable outcomes. Together, transfers and commitment allow contracts to mitigate the effects of adverse selection (and, more broadly, moral hazard), enabling the implementation of second-best allocations that would otherwise be blocked by informational frictions (e.g., P. Bolton and Dewatripont 2004; Laffont and Martimort 2009).

In legislative politics, however, neither of these tools is available. Legislators cannot formally contract on conditional monetary payments in exchange for support on specific bills, nor can they credibly commit to honor such exchanges in the future. Certain practices—such as the allocation of budgetary resources ("pork"), ministerial portfolios, or influential committee assignments—may resemble transfers to some degree, but they are imperfect substitutes. Pork is limited in scope, appointments are indivisible, and their allocation is governed by procedural rules rather than private bargaining. As a result, legislative exchange lacks both the flexibility of monetary contracts and the enforcement mechanisms that make private agreements credible.

A large literature confirms that practices such as vote trading and logrolling, while sometimes feasible, systematically fall short of achieving the efficiency levels that markets with money can generate.[6] When trades are "votes for votes," each bilateral swap shifts pivotal power and imposes externalities on non-participants,

[6] See Casella and Macé (2021) for an extended literature review.

so trading paths can cycle and even end below sincere-voting welfare—the Riker–Brams paradox—because a vote's value comes from how it changes the collective decision rather than from its holder alone. Stability is fragile and efficiency is not robust: individually rational, pairwise trades need not converge to a Condorcet winner (when one exists). Moving to "votes for money" does not restore market logic, because pivotality makes demands interdependent and discontinuous, so competitive equilibrium in decentralized vote markets typically does not exist. Even if legislators could make monetary commitments, legislators' preferences over such payments are not necessarily quasilinear or additively separable from policy and, therefore, cannot serve as direct utility transfers among legislators.

The aim of this chapter is to examine the extent to which the absence of both commitment technology and quasilinear preferences with transfers of a numeraire constrains the set of policy agreements that legislators can durably sustain in the long term. To this end, we develop a dynamic model of legislative bargaining with an infinite horizon. In each discrete period t, legislators meet, deliberate, and consider proposals; a qualified majority vote determines whether a new policy replaces the inherited status quo, and if no proposal is adopted the status quo persists to the next period. The model captures the following important features of real-world legislatures.

(i) Privately observed policy preferences. In many legislative environments, individual legislators possess private information about their preferences, priorities, or ideological positions. These preferences may stem from a range of unobservable factors—such as constituency-specific interests, personal values, electoral incentives, or lobbying pressures—that are not fully known to other legislators. We capture this by assuming that, in each period, a legislator's utility depends on the realization of a privately observed type that evolves over time according to an autonomous Markov chain.

(ii) High-dimensional policy spaces. Legislative policymaking typically involves decisions over a wide array of complex and heterogeneous issues. Legislators do not vote on a single, one-dimensional ideological axis, but rather face a sequence of proposals on diverse topics—ranging from taxation and education to healthcare, infrastructure, and environmental regulation. Each of these domains itself entails

multiple interdependent choices, making individual bills inherently multidimensional. Furthermore, legislators consider and vote on hundreds of such proposals over the course of a legislative session. It is therefore natural to represent the policy space as a high-dimensional region within Euclidean space. While this assumption departs from more traditional one- or two-dimensional models, it better captures the complexity and richness of modern legislative agendas.

(iii) Evolving status quos. A common feature of legislative policymaking is the path dependence of enacted policies: a policy adopted in one period typically remains in place until explicitly amended by future legislation. In the absence of new agreement, the current policy continues to serve as the default option—the "status quo"—for subsequent periods. This reflects both legal and procedural realities in most parliamentary and congressional systems. For instance, once a tax structure or regulatory framework is enacted, it tends to persist until a new bill successfully replaces it (e.g., Bowen et al. 2014). We capture this feature by assuming that the policy implemented in any given period becomes the status quo in the next period and remains in place unless a new agreement is reached.

The first step in our analysis is to overcome a major methodological challenge. Characterizing equilibrium long-run policy outcomes in legislative bargaining models with an endogenous status quo is notoriously difficult, as is widely acknowledged in the literature—see Eraslan et al. (2022) for a detailed discussion. In addition to the usual difficulties associated with such models, our framework introduces asymmetric information, with legislators' unobserved types evolving over time. Under complete information, the literature on dynamic legislative bargaining has typically identified *dynamically stable* policies, that is, policies that sufficiently farsighted legislators do not wish, or fail, to amend once implemented (e.g., Acemoglu et al. 2012; Diermeier and Fong 2011, 2012; Anesi and Duggan 2017, 2018). However, in our incomplete-information framework, stability cannot be defined as the persistence of a single policy point, because legislators' preferences evolve stochastically and remain privately observed.

How can we then define a notion of durable policy outcome in our framework? We begin by noting that a *policy outcome* in our incomplete-information framework

is not merely an element of the policy space, as it is in complete-information policymaking models, but a mapping from legislators' types into a policy. Guided by this logic, we introduce a weaker and more realistic notion of dynamic stability for legislative bargaining with incomplete information, which we call "cluster-stability." Formally, a *policy outcome* is a mapping from type profiles to policies. Then, we say that a policy outcome is *cluster-stable* if there exists a perfect Bayesian equilibrium of our model such that, in every period, the policy chosen by the legislature closely approximates the policy prescribed by the policy outcome, with probability approaching one, regardless of the realized types of legislators in that period. Stability is thus defined in terms of convergence *around* a decision rule, not exact coincidence with a particular policy in every state and period. By relaxing exactness, we recover a notion of stability that is robust to private information and the stochastic evolution of preferences.

Despite the technical difficulties outlined above, our main result characterizes a family of perfect Bayesian equilibria for our dynamic legislative bargaining framework with incomplete information, which exhibits cluster-stability. To do so, we introduce an axiomatic condition on policy outcomes that we call "internal optimality." We define an outcome as *internally optimal* if, for any realization of legislators' types, the policy it prescribes maximizes the (utilitarian) social welfare of the legislature *within its own range*—that is, among the set of all policies that the outcome could possibly prescribe for some realization of the types. Our main theorem establishes that, for generic policy preferences, every internally optimal outcome is cluster-stable. As we impose only extremely mild restrictions on the legislators' policy preferences and the distributions of types, this result provides a particularly simple and tractable characterization of long-run policy outcomes across a broad variety of policy environments, while delivering a powerful result: long-run outcomes in legislative bargaining with private preferences coincide with outcomes that are socially optimal *relative to their own image*.

The intuition is instructive. Suppose, counterfactually, that legislators could commit to type-dependent monetary transfers in each period. Then one could construct Groves–Clarke–style payments that induce truthful revelation of private information, ensuring that the socially optimal policy within the relevant range

is implemented. In our environment, such transfers do not exist, preferences need not be quasilinear, and the legislature cannot commit to future choices. The internal optimality condition plays two essential roles. First, it guarantees that—if transfers were available—incentive-compatible payments exist to support truth-telling. Second, it allows us to adjust equilibrium continuation values so that, *in each period*, the continuation values anticipated by legislators *approximately replicate the incentive-compatible payments* required to induce truth-telling. Thus, along the equilibrium path, legislators' future payoffs are additively and separably incorporated—like a numeraire—into their present utilities, effectively constituting incentive-compatible "payments" *conditional on current actions*. Despite the absence of commitment, these payments are credible because they are generated endogenously by future equilibrium behavior. In short, internal optimality substitutes for monetary transfers, for quasilinear preferences among legislators, and for the ability to commit to future choices—features typically unavailable in real-world legislative environments.

As we illustrate using the standard example of Euclidean preferences, the set of internally optimal outcomes is typically very large—in fact, indeterminate. This suggests that the absence of monetary transfers and commitment technology within legislatures does not constitute a major constraint on the sustainability of desirable policy agreements in the long run. A first corollary of our characterization theorem highlights the point: any socially optimal policy outcome that is one-to-one (as is generically the case) is trivially internally optimal and must therefore be cluster-stable. The social optimality of equilibrium outcomes in our framework stands in sharp contrast with the classic "paradox of vote trading," which argues that in the absence of a transferable numeraire, mutually advantageous agreements among legislators (e.g., logrolling) are unlikely to result in socially efficient outcomes. In our dynamic setting, repeated interaction and a multidimensional policy space allow legislators to approximate the cooperative benefits of markets even without transfers. The paradox disappears once stability is defined in terms of clusters and enforced through continuation values.

We exploit the generality of our equilibrium characterization, as well as the underlying numeraire-through-continuation-values logic, to study two issues that

occupy a central place in political economy: vote buying and agenda manipulation. First, we analyze how a lobbyist who can condition payments on observed legislative behavior—without observing legislators' private preferences—can still implement her ideal policy in every period while spending an arbitrarily small amount in expectation. When legislators are sufficiently patient, promised continuation values (endogenously generated by equilibrium play) substitute for direct transfers, providing a new theoretical rationale for the empirical observation that interest groups often exert persistent influence despite limited monetary expenditures. The mechanism then offers a resolution to the Tullock paradox: there can be "so little money in politics" not because money is irrelevant, but because dynamic incentives do most of the work once small payments align on-path behavior.

We then consider an extension with endogenous agenda formation. In each period, legislators submit proposals to an agenda setter, who chooses the order of pairwise votes. Proposals serve a dual purpose: alongside substantive options, each legislator includes a distinct "tagging proposal" that publicly signals her private type, so the relevant information is revealed on the path of play while the target option is guaranteed a place on the agenda. With repeated interaction, internal optimality and continuation values deter deviations both when proposing and when voting. The takeaway mirrors our main result: every internally optimal outcome remains cluster-stable even when the agenda is controlled endogenously. The direct implication is that incomplete information does not eliminate agenda power—the setter can still steer play toward a preferred cluster within the internally optimal set.

The remainder of the chapter is organized as follows. Section 2.2 situates this chapter within the literature. Section 2.3 introduces the baseline model of dynamic legislative bargaining with privately observed and evolving preferences. Section 2.4 defines cluster-stability and establishes our main equilibrium characterization. Section 2.5 discusses implications of the theory for vote trading, vote buying, and agenda manipulation. Finally, Section 2.6 concludes.

2.2 Related literature

The modern noncooperative literature on legislative bargaining originates with Baron and Ferejohn (1989). In their canonical model, a randomly recognized legislator proposes a division of a fixed surplus that the assembly accepts or rejects under majority rule, with discounting disciplining delay. The key contribution is the existence of equilibrium and—within the stationary class—uniqueness of equilibrium payoffs. By replacing unrestricted pairwise majority comparisons with a simple in-legislature proposal game, this framework offers a tractable response to classical instability: while single-peaked preferences on a line restore stability via Black (1948), the majority-rule core is generically empty in higher dimensions (Plott, 1967), and restricting agendas does not restore predictability because any alternative can be defeated by any other through suitable sequencing (McKelvey 1976, 1979; Schofield 1978). Against this backdrop, the Baron–Ferejohn baseline provides a unified way to study how rules—recognition, voting thresholds, vetoes—shape collective choices, and it has been applied widely to government formation, bicameralism, executive vetoes, and committee organization.[7]

Baron and Ferejohn's work has been followed by a large literature on legislative and multilateral bargaining that extends their model along several dimensions.[8] A first line relaxes symmetry across legislators by allowing heterogeneity in proposer recognition and patience (Eraslan and Merlo 2002; Kalandrakis 2015; Kawamori 2005); by introducing weighted voting and studying when payoffs are or are not proportional to voting weights (Snyder Jr et al. 2005; Eraslan and McLennan 2013; Montero 2006; Montero et al. 2017); and by incorporating heterogeneous preferences to show how risk attitudes and payoff externalities reshape coalition formation (Harrington Jr 1990; Calvert and Dietz 2005).

A second line broadens the feasible set of policies: spatial bargaining places policies in Euclidean spaces and connects stationary outcomes to core-like objects under suitable conditions (Baron 1991; Banks, Duggan, et al. 2006; Duggan and Banks 2000; Jackson and Moselle 2002); stochastic-surplus models examine timing

[7] A review of these applications lies beyond the scope of this chapter; we focus on theoretical extensions.

[8] For a comprehensive survey, see Eraslan and Evdokimov (2019).

and delay (Merlo and Wilson 1995, 1998; Eraslan and Merlo 2002); and mixed environments study bargaining over public goods (Volden and Wiseman 2007, 2008; Cardona and Rubí-Barceló 2014). Finally, a third group of literature analyzes alternative bargaining protocols and shows how procedure redistributes proposal power and affects efficiency: endogenous proposer selection via contests (Yildirim, 2007), demand-bargaining formats (Morelli 1999; Montero and Vidal-Puga 2011), and "rejection-friendly" rules that strengthen countervailing leverage (Kawamori, 2013). Crucially, the vast majority of these contributions are static or "once-and-for-all": a single agreement is reached and implemented, with no endogenous carryover of today's policy into tomorrow's bargaining environment.

This chapter relates most closely to two relatively recent lines in the legislative-bargaining literature: *dynamic* legislative bargaining and legislative bargaining with *incomplete information*. Building on early precursors such as Ingberman (1985) and Epple and Riordan (1987), and on the spatial and distributive formulations in Baron (1996) and Kalandrakis (2004), a growing literature models *bargaining with an endogenous status quo*: in each period a new policy is chosen and, if no agreement is reached, the implemented policy is the one endogenously determined in the previous period.[9]

Within dynamic legislative bargaining, yesterday's policy becomes today's default, and the linkage is implemented along four complementary margins: models with state-dependent stage utilities, letting payoffs vary with economic or political conditions (Riboni and Ruge-Murcia 2008; Dziuda and Loeper 2016; Bowen et al. 2017; Bowen et al. 2014); models with noisy status-quo mapping, in which the default inherits last period's status quo with small perturbations (Duggan and Kalandrakis, 2012); models with stochastic surpluses in the size of the redistributive pie (Merlo and Wilson, 1995); or models with state-dependent winning coalitions (Chen and Eraslan, 2017).

Existence of stationary Markov perfect equilibria is not guaranteed in general infinite-horizon models with endogenous status quo: beyond finite-horizon or finite-policy cases, proposer optimization can create discontinuities when the

[9] This section is based on Eraslan et al. (2022), which the interested reader can refer to for an exhaustive literature review.

default equals last period's policy, so generic stochastic-game results do not directly apply. The broader literature provides conditions and counterexamples (He and Sun 2017; Y. Levy 2013; Y. J. Levy and McLennan 2015) and explains why these results do not fit standard legislative environments (Duggan and Kalandrakis, 2012). Existence is typically secured either by adding smoothness—noisy status-quo transitions and smooth preferences (Duggan and Kalandrakis, 2012), or smooth transitions from a continuous policy space to a countable state space (Duggan, 2017)—or by discretizing the policy space, which guarantees existence of a solution (Anesi 2010; Diermeier and Fong 2012).

Given the fragility of existence, much of the dynamic literature constructs equilibria and studies their properties (Eraslan et al., 2022). In distributive (divide-the-dollar) environments, Kalandrakis (2004) builds a three-player equilibrium in which the randomly chosen proposer captures the entire surplus (the "rotating dictator" outcome), and Kalandrakis (2010) extends agenda-setter dominance to larger odd-sized legislatures with heterogeneous recognition, tracing coalition composition and risk-aversion effects. Bowen and Zahran (2012) show that with strictly risk-averse symmetric legislators and a larger assembly, non–minimum-winning "compromise" coalitions can arise; Richter (2014) shows that allowing waste can sustain equal division. Anesi and Seidmann (2015) provide a significant generalization and show that, as legislators become arbitrarily patient, a large set of outcomes can be supported. Related constructions include Baron and Bowen (2015) and Baron (2019); and work with veto power documents environments in which the veto player asymptotically appropriates the surplus (Nunnari, 2021).

The second dynamic strand constructs equilibria within *one-dimensional spatial* dynamic legislative bargaining. Building on Baron (1996), who analyzes a dynamic public-good scale with an endogenous default, subsequent work with three players and quadratic utilities establishes existence and shows that a non-median legislator may strategically moderate to constrain future proposals, while the median's acceptance set behaves as under myopic median voting (Kalandrakis, 2016; Zápal, 2016). Under unanimity, by contrast, Dziuda and Loeper (2016, 2018) show that endogenous defaults can sustain polarization. Moving beyond three players and using a *multi-dimensional spatial* model, Zápal (2020) provides an algorithmic

construction based on strategic bliss points and proves generic (within a class) uniqueness, though some games remain unresolved.

As in our model, several papers allow dynamic bargaining models with changing preferences, although their models only partially resemble ours: Riboni and Ruge-Murcia (2008) assume a fixed proposer; Duggan and Kalandrakis (2011) compute equilibria numerically; Dziuda and Loeper (2016, 2018) restrict the policy space to two alternatives; Austen-Smith et al. (2019) allow three alternatives; Bowen et al. (2017) and Buisseret and Bernhardt, 2017, 2018 focus on two-period environments.

Our model differs from these existing contributions in its focus on high-dimensional policy spaces (possibly without a distributive dimension), which as we argued above, are a regular feature of real-world legislative policymaking environments. Zápal (2014), Duggan and Kalandrakis (2012) and Anesi and Duggan (2018) also allow for general, high-dimensional policy spaces, but not for privately observed policy preferences—another important feature of legislative politics.

Legislative bargaining with incomplete information remains comparatively thin, and most contributions analyze once-and-for-all (or short-horizon) environments rather than perpetual bargaining with an endogenous status quo. The main themes are private types and communication. With privately known patience, Tsai and Yang (2010) show that majority rule can induce delay and oversized coalitions. In common-value settings, Meirowitz (2007) identifies conditions for full revelation. In three-player models that mix ideological and distributive dimensions, Chen and Eraslan (2013, 2014) analyze cheap talk before proposals, showing that bundling versus separating issues changes what can be learned, and that competition between two potential coalition partners under majority can reduce informativeness and even leave the proposer worse off relative to unanimity. Chen (2021) studies how procedural rules shape communication incentives. Related work with privately known disagreement (breakdown) payoffs compares voting thresholds; extending Miller et al. (2018), Piazolo and Vanberg (2025) find that under unanimity responders are "more expensive," making delay and disagreement more likely than under majority except in narrow cases.

Most contributions in the legislative bargaining literature with incomplete information focus on environments where, although the bargaining process may

be sequential, a single policy is chosen once and for all. They therefore do not capture dynamic contexts in which a policy enacted today may be reconsidered and amended in a future legislative session. A few exceptions include Dziuda and Loeper (2016, 2018), who show how incomplete information under unanimity can sustain polarization, and Anesi and Duggan (2017, 2018), who characterize dynamically stable outcomes in multidimensional policy spaces. The paper most closely related to ours is Anesi (2018), which like the present chapter, considers a dynamic framework with an information structure similar to ours. That paper adopts a mechanism design approach, studying how social choice functions that are strategy-proof with compensatory transfers can be implemented, at least approximately, in dynamic bargaining settings. Our equilibrium construction builds on some of the approximation lemmas established in that paper. Nevertheless, we adopt a more standard approach in the literature on dynamic stability in legislative policymaking, which consists in providing a tractable axiomatic condition for characterizing long-run legislative outcomes (e.g., Anesi 2010; Acemoglu et al. 2012; Diermeier and Fong 2012; Anesi and Seidmann 2015; Anesi and Duggan 2017, 2018). This tractability, in turn, allows us to extend the scope of our result beyond the baseline bargaining framework and to consider richer models, enabling us to address important issues about vote buying and agenda manipulation.

Finally, our main result has strong welfare implications: the dynamically stable policy outcome is efficient. This complements evidence from distributive environments showing that, with linear utilities and common discounting, no-waste policy paths are ex-ante efficient (Anesi and Seidmann, 2015; Kalandrakis, 2004), whereas allowing waste or making outcomes sensitive to the proposer's identity can generate inefficiency (Anesi and Seidmann, 2015; Bowen et al., 2017). In public-goods settings, endogenous defaults can raise provision and welfare (Bowen et al., 2014), while gridlock may serve as a commitment device that improves efficiency in taxation environments (Piguillem and Riboni, 2013; Piguillem and Riboni, 2015). By contrast, in spatial models under unanimity, endogenous defaults can sustain polarization and inefficiency (Dziuda and Loeper, 2016, 2018), and with evolving preferences dynamic linkage can induce inefficient fluctuations (Riboni and Ruge-Murcia, 2008).

2.3 Baseline model

This section presents a dynamic legislative bargaining game in which a legislature meets in every period to make policy decisions. The horizon is infinite, so the same bargaining process is repeated indefinitely: in each round legislators bargain over a new policy, today's outcome becomes tomorrow's status quo, and the shadow of future interaction shapes current behavior. Because there is no terminal date, strategies must specify how legislators act in every possible continuation of the game, and the expectation of future payoffs provides the incentives that sustain agreements. The environment captures the key institutional features of real legislatures: policies persist until amended, legislators' preferences evolve over time and are privately observed, proposals must gather a qualified majority to pass, and while cheap talk communication is possible, commitments might not be enforceable and transfers might not be available, or may not enter policy preferences in a quasilinear way.

Formally, we consider a legislature composed of $n \geq 3$ members, indexed by $i = 1, \ldots, n$, who must choose a policy x^t from the set $X \equiv \mathbb{R}^k$, $k \in \mathbb{N}$, in each of an infinite sequence of discrete time periods $t = 1, 2, \ldots$. Each legislator i has a privately observed type θ_i^t, drawn from a finite set Θ_i and evolving according to a Markov process (p_i, M_i): at $t = 1$ the type is drawn from the initial distribution p_i, and in each subsequent period it transitions according to the matrix M_i. Type processes are independent across legislators. Let $\Theta \equiv \prod_{i=1}^{n} \Theta_i$ denote the set of type profiles, and let (p, M) denote the aggregate type process, which we assume is irreducible so that every type profile is eventually reachable.

Following Coughlan (2000) and Austen-Smith and Feddersen (2006), we include a deliberation stage in which legislators can communicate through simultaneous cheap talk before any policy proposal is made.[10] Messages are costless and non-binding, and we assume that the communication language is sufficiently rich so that the set of available messages for legislator i is at least as large as her type set

[10] The fact that deliberation occurs before proposals is purely for expositional ease and entails no loss of generality. All results would remain unchanged if deliberation took place during the bargaining stage.

Θ_i. We focus on equilibria in which legislators truthfully signal their types, so we lose no generality by restricting messages to Θ_i.

Policy change occurs through sequential proposals. At the start of each period, the legislature inherits the status quo x^{t-1} from the preceding period. A proposer order ι is drawn randomly from the set of all permutations of $\{1, \ldots, n\}$, with each order having strictly positive probability. The first legislator in the proposing order, $\iota(1)$, submits a proposal of a policy $y \in X$. Legislators then vote sequentially according to some arbitrary order on whether to accept the proposal. A proposal is adopted if it receives at least q votes in favor, where $\frac{n}{2} < q < n$; otherwise, it is rejected. If adopted, y becomes the new status quo x^t and payoffs (described below) are realized. If rejected, the next proposer $\iota(2)$ makes a new proposal, and the process repeats until either some proposal is adopted or all proposers fail, in which case the status quo x^{t-1} prevails for period t and carries forward into period $t+1$. The game begins with an exogenously given status quo policy x^0.

Payoffs:

At the end of each period t, each legislator i obtains stage utility of $u_i(x^t, \theta_i^t)$, where $u_i(\cdot, \theta_i) : X \to \mathbb{R}$ is bounded and twice continuously differentiable for all $\theta_i \in \Theta_i$. We denote by $\mathcal{U}$ the set of all such utility functions. The game repeats indefinitely, and each player discounts her stage payoffs using a discount factor $\delta \in [0, 1)$. Each legislator aims to maximize the expected discounted sum of these payoffs over time.

Timing:
Each period t unfolds as follows:

1. *Type realization.* Legislator i privately observes $\theta_i^t \in \Theta_i$, drawn according to (p_i, M_i).
2. *Deliberation.* Legislators simultaneously send cheap-talk messages.
3. *Bargaining.* A proposer order ι is drawn. Proposals $y \in X$ are made sequentially until one secures at least q votes, or all are rejected. If a proposal is accepted, it is implemented and becomes the new status quo x^t. If all are rejected, x^{t-1} persists.
4. *Payoffs.* Each legislator i receives $u_i(x^t, \theta_i^t)$. The game advances to $t+1$.

Strategies and equilibrium:

In each period t, an action for legislator i consists of sending a message in the deliberation stage, proposing $y \in X$ if recognized as proposer, and casting a vote in favor or against each proposal. A strategy for legislator i specifies a complete contingent plan mapping her private type θ_i^t and the public history h^t into actions at every decision point. A strategy profile is denoted by $\sigma = (\sigma_1, \ldots, \sigma_n)$ and describes the strategies of all legislators.

We study Perfect Bayesian Equilibria (PBE) of this dynamic game, allowing for mixed strategies. Strategies must be sequentially rational given beliefs, and beliefs must be updated consistently with Bayes' rule whenever possible.

Policy space dimensionality:

As emphasized in the introduction, we focus on environments in which the policy space is highly dimensional. To ensure that our results hold generically in $\mathcal{U}$, we assume that the number of dimensions of the policy space is at least as large as the total number of possible types across legislators,

$$k \geq \sum_{i=1}^{n} |\Theta_i|. \tag{2.3.1}$$

2.4 Policy outcome clusters

In this section, we address the central question of this chapter: whether dynamically stable outcomes can arise in legislative bargaining when legislators' preferences are private and evolving, and if so, how such outcomes can be characterized.

The central difficulty is that the standard notion of dynamic stability developed for complete-information models does not directly apply: under incomplete information, what matters is not whether a single policy is stable, but whether a collective decision rule can persist over time. Our goal is therefore to provide a definition of stability suited to this environment, and to identify the conditions under which such stable outcomes exist. To this end, Section 2.4.1 introduces the concept of *cluster-stability* as a new notion of dynamic stability for legislative bargaining settings with adverse selection. Section 2.4.2 then defines a simple condition, *internal*

optimality, and establishes our main theorem: any policy outcome that satisfies internal optimality is cluster-stable. Section 2.4.3 develops the intuition behind this result, showing how continuation values can substitute for monetary transfers—in effect, how legislators can "buy votes without a numeraire."

2.4.1 Cluster-Stability: Definition

As discussed in the introduction, much of the political economy literature on dynamic collective choice under complete information has focused on the characterization of dynamically stable policies—i.e., policies that can be supported as absorbing outcomes of equilibria in dynamic bargaining games when legislators are sufficiently patient. The central idea is that certain policies, once adopted, persist because no legislator has the incentive to change them—they are dynamically "stable" in the sense that they can be supported as long-run equilibrium outcomes. This is the notion developed, for example, by Acemoglu et al. (2012), Anesi and Duggan (2017), and others. From a conceptual standpoint, one might be tempted to apply exactly the same criterion to legislative environments with incomplete information.

However, in settings where legislators' preferences evolve stochastically over time and remain privately observed, stability cannot meaningfully be defined in terms of a single policy $x \in X$. What matters is not whether a fixed policy persists, but whether a collective decision rule persists. If in each period the legislature consistently applies the same mapping from type profiles to policies, then observed changes in outcomes simply reflect the evolution of preferences. In that case it would be odd to describe the process as dynamically "unstable"—what changes are the types, not the decision rule.

Following the standard approach in collective choice theory, we define a *policy outcome* as a mapping $\wp : \Theta \to X$, where Θ is the set of type profiles. We denote by $\mathcal{P}$ the set of all such functions. Since Θ is finite, $\mathcal{P}$ can be identified with a subset of Euclidean space $\mathbb{R}^{k|\Theta|}$. Following the language of social choice theory, a policy outcome is thus nothing more than a "social choice function"—a collective decision rule that prescribes a policy for each possible configuration of legislators'

private types.

Given a strategy profile σ in the bargaining game, let $\{x_\sigma^t\} \in X^\infty$ denote the (possibly stochastic) sequence of policies implemented over time, where each x_σ^t represents the policy chosen by the legislature in period t. A very demanding criterion for stability of policy outcome $\wp$ would require that in every period t, for every realization of the type profile $\theta^t \in \Theta$, the legislature chooses exactly $\wp(\theta^t)$, so that $x_\sigma^t = \wp(\theta^t)$ holds. This criterion proves to be extremely demanding: outside constant outcomes, it is unlikely that such exact implementation can be sustained once private information and strategic incentives are taken into account.

We therefore propose a weaker notion of dynamic stability, which we call *cluster-stability*. Instead of requiring that the policy choices exactly coincide with those prescribed by the policy outcome $\wp$, cluster-stability requires that the policies chosen each period, x_σ^t, lie within an arbitrarily small neighborhood of $\wp(\theta^t)$ with probability arbitrarily close to one. Intuitively, in a cluster-stable equilibrium the long-run play of the legislature "converges" around the prescriptions of $\wp$—outcomes are not exact, but they concentrate in a tight neighborhood around the intended decision rule. This relaxation is precisely what makes the notion applicable under incomplete information.

Definition 1. A policy outcome $\wp \in \mathcal{P}$ is *cluster-stable* (or satisfies *cluster-stability*) if, for every $\varepsilon \in (0, 1)$, there exists a strategy profile σ for the legislative bargaining game such that the following hold:

(i) σ is a Perfect Bayesian Equilibrium (PBE) of the bargaining game for sufficiently high discount factor $\delta < 1$; and

(ii) for every period $t \in \mathbb{N}$, the probability that x_σ^t belongs to the ε-neighborhood of $\wp(\theta^t)$ exceeds $1 - \varepsilon$.

Condition (i) is a standard dynamic stability requirement—see Acemoglu et al. (2012). It requires that the strategy profile supporting the outcome be an equilibrium for legislators who are sufficiently farsighted. Condition (ii) requires that, conditional on the stochastic evolution of types and strategies, equilibrium policy choices concentrate with high probability in neighborhoods of the prescribed outcomes. The probability in condition (ii) is computed with respect to both the

realizations of the type process $\{\theta^t\}$ and the (possibly mixed) equilibrium strategies of legislators. It is evaluated at the start of the game, before any types are realized.[11]

Cluster-stability thus extends the complete-information notion of absorbing policies to environments with private information. The focus shifts from individual policies to policy outcomes, and from exact implementation to approximate implementation with high probability. The lesson is the same as in the vote-trading literature: insisting on exact stability is too strong a requirement, but relaxing the notion allows for a meaningful characterization of long-run outcomes. In this weaker, probabilistic sense, stability captures how legislatures can sustain collective decision rules even in the absence of money or enforceable contracts.

2.4.2 Internal optimality and Cluster-stability

To characterize dynamically stable outcomes, we use the standard approach in the literature: first introduce an axiomatic solution concept and then show that it coincides with the long-run policy outcome in an equilibrium of the legislative bargaining game—see, for example, Acemoglu et al. (2012), Diermeier and Fong (2012), and Anesi and Duggan (2017, 2018). The logic of this approach is straightforward: if equilibria are often multiple and difficult to characterize directly, then it is useful to start from an abstract property that a "stable" outcome should satisfy, and then to verify that such outcomes can in fact be supported in equilibrium when legislators are sufficiently patient. We follow this approach here by introducing the notion of *internal optimality* as our solution concept.

Definition 2. A policy outcome $\wp$ is *internally optimal* if, for all $\theta \in \Theta$, the policy $\wp(\theta)$ is the unique solution to the following problem:

$$\max_{x} \sum_{i=1}^{n} u_i(x, \theta_i) \text{ , subject to } x \in \wp(\Theta) \,. \tag{2.4.1}$$

[11] An alternative formulation would compute $\Pr\{\|x_\sigma^t - \wp(\theta^t)\| < \varepsilon\}$ at each period-s history on the equilibrium path and require that it exceed $1-\varepsilon$ for all sufficiently large $t > s$. This alternative definition would yield exactly the same results.

Internal optimality requires that for each realization of the type profile $\theta = (\theta_1, \ldots, \theta_n)$ in Θ, the policy prescribed by $\wp$ maximizes aggregate welfare *among the set of policies in the range of* $\wp$. In other words, once we restrict attention to the set of policies that $\wp$ itself ever prescribes, $\wp(\theta)$ must be the unique social optimum. Two observations follow immediately. First, that the existence of internally optimal policy outcomes is not an issue: constant outcomes $\wp$, where $\wp(\Theta)$ is a singleton, are trivially internally optimal. Second, non-constant outcomes also frequently satisfy the conditions of Definition 2 for standard policy preferences, as the following example shows.

Example 1: Euclidean preferences. Suppose legislators' preferences take the standard quadratic Euclidean form, commonly used in the political economy literature. For all i and $\theta_i \in \Theta_i$, we have

$$u_i(x, \theta_i) = -\left\| x - \hat{x}_i(\theta_i) \right\|^2, \tag{2.4.2}$$

where $\hat{x}_i(\theta_i) = (\hat{x}_{i1}(\theta_i), \ldots, \hat{x}_{ik}(\theta_i))$ is legislator i's (exogenously given) ideal policy when her type is θ_i.

It is easy to see that, in this case, the socially optimal policy outcome is the average of legislators' ideal points:

$$\wp^*(\theta) = \frac{1}{n} \sum_{i=1}^{n} \hat{x}_i(\theta_i), \qquad \text{for all } \theta \in \Theta. \tag{2.4.3}$$

Moreover, by strict concavity of $u_i(\cdot, \theta_i)$, this optimum is unique for all $\theta \in \Theta$:

$$\sum_{i=1}^{n} u_i(x, \theta_i) < \sum_{i=1}^{n} u_i\big(\wp^*(\theta), \theta_i\big), \tag{2.4.4}$$

for all $\hat{\theta}_i : \wp(\hat{\theta}_i) \neq \wp(\theta_i)$.[12] It follows that, in the (generic) case where $\wp^*(\theta) \neq \wp^*(\theta')$ for all distinct $\theta, \theta' \in \Theta$, *the (utilitarian) social optimum is internally optimal*.

However, $\wp^*$ is certainly not the only policy outcome that is internally optimal. This is easy to see by considering small perturbations of $\wp^*$. By continuity of

[12] Trivially, it holds with equality if $x = \wp(\hat{\theta}_i) = \wp(\theta_i)$

the utility functions $u_i(\cdot, \theta_i)$, there exists a sufficiently small $\varepsilon > 0$ such that if, for each $\theta \in \Theta$, one selects any policy x_θ within the ε-neighborhood of $\wp^*(\theta)$, then the policy outcome defined by $\wp^\varepsilon(\theta) \equiv x_\theta$, for all $\theta \in \Theta$, is also internally optimal. Intuitively, small perturbations of $\wp^*$ preserve the welfare ranking of policies within $\wp^*(\Theta)$, and therefore the internal optimality condition continues to hold. We conclude that there exists a continuum of non-constant policy outcomes that satisfy internal optimality. $\Diamond$

We can now state the main result of this chapter. Recall that $\mathcal{U}$ denotes the set of bounded, twice continuously differentiable utility functions on X, endowed with the Whitney topology.[13] The following theorem establishes that for generic policy-preference profiles in $\mathcal{U}$, every internally optimal policy outcome satisfies cluster-stability.

Theorem 2.1. *There exists a dense subset $\mathcal{U}^* \subseteq \mathcal{U}$ such that, whenever $u_i(\cdot, \theta_i) \in \mathcal{U}^*$ for all $i \in N$ and all $\theta_i \in \Theta_i$, every internally optimal policy outcome is cluster-stable.*

Proof. See Section 2.A

An immediate consequence of Theorem 2.1, together with the observation that all constant policy outcomes are internally optimal, is that, generically, *there exist infinitely many cluster-stable outcomes.*

The theorem is non-trivial. It establishes that the axiomatic notion of internal optimality is not just internally consistent, but also provides a sufficient condition for cluster-stability. In other words, every outcome that satisfies Definition 2 can be sustained in equilibrium when legislators are sufficiently patient. This provides a powerful characterization of long-run outcomes in legislative bargaining with important implications for several strands of the political economy literature, as we discuss in Section 2.5. To fully appreciate those implications, it is essential to first understand the logic underpinning the cluster-stability of internally optimal policy outcomes. This is the focus of the next subsection.

[13] The Whitney topology is the standard mathematical device for studying "genericity" of functional spaces: it ensures that statements such as "for a dense subset of $\mathcal{U}$" have a rigorous meaning.

2.4.3 The underlying logic: Buying votes without a numeraire

The proof of Theorem 2.1 is constructive. Starting from an internally optimal policy outcome $\wp$ and an $\varepsilon \in (0,1)$, we show that for sufficiently high δ, one can construct a PBE σ in which legislators implement policies that lie within ε of $\wp(\theta^t)$ with probability arbitrarily close to one in every period $t \in \mathbb{N}$.[14] The logic of this construction can be understood by comparing three environments: a complete-information benchmark, the private-information setting of our model, and a hypothetical world in which transfers are available.

Suppose first that legislators' types were publicly observable, or equivalently, that they were required to truthfully reveal their types in every deliberation stage. This yields a dynamic legislative-bargaining game with complete information and, in this case, any deviation from equilibrium prescriptions would be observable by all players and hence "off-path"—in the language of dynamic games, it would occur outside of the histories that equilibrium strategies specify. A result by Anesi and Duggan (2018) shows that in such games, when the policy space is sufficiently high-dimensional, as in our model, it is generically possible to sustain any desired policy outcome $\wp$ (or at least some outcome within any small neighborhood of $\wp$) as a Markovian equilibrium, with probability one in every period. The construction works by punishing legislators who fail to propose, or to support, $\wp(\theta^t)$ in period t, through small perturbations of $\wp$ that lower their continuation payoffs in all future periods. Deviations are deterred because the discounted loss in continuation value outweighs any short-term gain.

Thus, if it is guaranteed that in each period legislators reveal their true types $\theta^t = (\theta_1^t, \ldots, \theta_n^t)$ (at least with probability $1 - \varepsilon$), then we know precisely what actions they are supposed to take—i.e., propose and vote in favor of $\wp(\theta^t)$ (or of another prescribed policy within ε of $\wp(\theta^t)$). Any deviation from such behavior is then off-path and can be deterred using, off-the-shelf, the construction of Anesi and Duggan (2018) described above. To establish the cluster-stability of $\wp$, therefore, it remains to ensure that legislators truthfully reveal their types with probability

[14] Whenever two types are prescribed the same policy, we treat them as the same type. Henceforth, we focus on $\mathcal{P}$ that prescribe different policies to different types - "one-to-one" - for expositional ease.

close to one.

In our baseline model, however, legislators' types are private. Deviations from truthful reporting in the deliberation stage are therefore "on-path"—they occur along the equilibrium path of play, since misrepresentation of preferences cannot be directly observed. Punishments of the type used in Anesi and Duggan (2018) are no longer sufficient: if a legislator misreports her type, the deviation cannot be distinguished from truthful play. Establishing cluster-stability in this setting therefore requires a mechanism to induce truth-telling in the deliberation stage.

How can legislators be incentivized to truthfully reveal their types during the deliberation stage? A natural way to induce truthful reporting under asymmetric information, familiar from the mechanism design literature, is to use monetary payments designed so that agents find it in their own interest to tell the truth. Suppose, as a thought experiment, that a constitution-maker or "equilibrium designer" could use unbounded monetary payments and that the per-period preferences of a type-θ_i legislator i take the quasilinear form:

$$v_i(x, m_i \mid \theta_i) \equiv u_i(x, \theta_i) + m_i, \tag{2.4.5}$$

where $u_i(\cdot, \theta_i)\colon X \to \mathbb{R}$ is interpreted as before, and m_i denotes the payment received by i.

As the internal optimality of the policy outcome $\wp$ implies that for every realization $\theta = (\theta_1, \ldots, \theta_n)$ of legislators' types, the policy $\wp(\theta) \in X$ is socially optimal *among those in the subset* $\wp(\Theta)$, it is straightforward to show that a standard Groves-Clarke-style payment function $m^* = (m_1^*, \ldots, m_n^*)\colon \Theta \to \mathbb{R}^n$ can be used to *strictly* deter untruthful reports $\hat{\theta}_i \neq \theta_i$ (e.g., Mas-Colell et al. 1995). Formally,

$$v_i\big(\wp(\hat{\theta}_i, \theta_{-i}), m_i^*(\hat{\theta}_i, \theta_{-i}) \mid \theta_i\big) < v_i\big(\wp(\theta), m_i^*(\theta) \mid \theta_i\big)\ , \tag{2.4.6}$$

for all i, $\hat{\theta}_i \in \Theta_i \setminus \{\theta_i\}$, and $\theta \in \Theta$. In this way, the designer can implicitly buy legislators' proposals and votes in favor of the policies prescribed by $\wp(\Theta)$, regardless of their true policy preferences. As explained above, there is no need to worry about off-path deviations leading to policies outside $\wp(\Theta)$, as these are already deterred by repeated legislative interactions.

In actual legislatures and in our framework, however, unbounded payments between legislators are typically not feasible; preferences over such payments may not be linear and additively separable from other variables; and in any case, there is no technology allowing the legislature to commit to implementing a given social function $\big(\wp(\cdot), m^*(\cdot)\big)$. However, the dimensionality of the policy space provides a substitute. With the high dimensionality of the policy space X, which ensures that for a generic set of policy preferences $\mathcal{U}^*$, it is possible to perturb each policy $\wp(\theta)$ (or at least a policy arbitrarily close to $\wp(\theta)$) in such a way as to replicate any sufficiently small monetary transfers.

Formally, if each component of a payment vector $m = (m_1, \ldots, m_n) \in \mathbb{R}^n$ is sufficiently small in absolute value, then one can construct a policy x in an ε-neighborhood of $\wp(\theta)$ such that:

$$u_i(x, \theta_i) = u_i(\wp(\theta), \theta_i) + m_i, \tag{2.4.7}$$

for every legislator i and type $\theta_i \in \Theta_i$. Thus, by slightly adjusting the desired policy outcome $\wp$, legislators can replicate monetary transfers, so long as these remain sufficiently small.

A final difficulty remains. While the legislature may be able to simulate small monetary transfers among its members, the Groves–Clarke transfers $m^*(\theta)$ required to induce truth-telling in the deliberation stages may be large. That is, the payments required may well exceed what can be replicated through arbitrarily small perturbations of $\wp(\theta)$. To resolve this, note that since $m^*(\cdot)$ strictly induces truth-telling, so does $m^*(\cdot) - c$ for any constant vector $c \in \mathbb{R}^n$. In particular, let $c = \mathbb{E}[m^*(\theta)]$, where the expectation is computed under the invariant distribution of the irreducible Markov chain (p, M). Then, by the strong law of large numbers for Markov chains, we have that the sequence

$$\frac{1}{k} \sum_{t=1}^{k} m_i^*(\theta^t) \to \mathbb{E}[m_i^*(\theta)] \quad \text{a.s. as } k \to \infty, \tag{2.4.8}$$

for each legislator i. This means that by subtracting c from $m^*(\cdot)$, the net transfers owed to each legislator vanish in the long run, with probability one. In practice,

over sufficiently long blocks of periods, the residual transfers become negligible. Combining the perturbation argument with this averaging trick, the legislature can approximate the transfers needed to induce truth-telling. By applying small policy perturbations every k periods, with k large, legislators replicate the payment scheme $m^*(\cdot) - c$ with near certainty. Since they are very patient—i.e., δ is close to one—legislators are willing to wait these k periods to receive their "implicit" transfers through continuation payoffs.

The upshot is that internal optimality guarantees the existence of such constructions. It ensures that the continuation values of repeated bargaining play the role of implicit side payments. In effect, legislators can "buy votes without a numeraire": the high-dimensional policy space substitutes for money, and continuation values substitute for enforceable contracts. Our construction shows that even without money, stable and consistent outcomes can be sustained when policies are multidimensional and legislators are sufficiently patient.

2.5 Implications: Politics without money

The results of the previous sections establish that internal optimality and cluster-stability provide a general characterization of long-run outcomes in dynamic bargaining games with private information. Internal optimality captures the idea that policies must be consistent with the set of prescriptions generated by the outcome itself, while cluster-stability ensures that these prescriptions can be sustained in equilibrium when legislators are sufficiently patient. Together, these results imply that legislatures can support a wide range of stable outcomes even in the absence of money or enforceable contracts, provided that policies are multidimensional and legislators interact repeatedly. The continuation values of repeated bargaining act as implicit transfers, substituting for explicit side payments.

This has important implications for the literature on vote trading, vote buying, and agenda manipulation. In what follows, we return to some of the classic paradoxes that have motivated skepticism about the efficiency of legislative exchange, and show how our framework sheds new light on them.

Throughout this section, we assume that the utility functions of all player types belong to the set $\mathcal{U}^*$, as defined in Theorem 2.1. As this set is generic, the assumption is made with virtually no loss of generality.

2.5.1 The vote trading paradox resolved

The idea that legislators may exchange votes across issues—logrolling—is as old as modern public choice theory. In their classic book, Buchanan and Tullock (1965) argued that logrolling could be efficiency-enhancing. After Arrow (1951) had shown that majority voting may not yield a Condorcet winner, they emphasized that voters rarely face a single decision but rather multiple decisions at once, and that logrolling is the means through which voters can overcome the cycles and instability created by majority rule. By trading votes, legislators could in principle mimic markets: giving up support on issues of low personal value in exchange for support on issues of high personal value, so that influence is reallocated to where preferences are most intense. From this perspective, logrolling held the promise of transforming legislatures into arenas of efficient exchange—delivering the Condorcet winner when it exists, and preventing majority cycles when it does not.

The Buchanan and Tullock (1965) book launched a large literature, but doubts soon emerged, as their models were informal and provided little formal support for the conjectures they advanced. Park (1967) provides one of the first systematic analyses of vote trading and showed that efficiency could not be guaranteed, since introducing vote trading does not ensure convergence to a Condorcet winner. Shortly thereafter, Riker and Brams (1973) published their famous article, *The Paradox of Vote Trading*. They model vote trading as a process of pairwise, enforceable, myopic trades across multiple issues, with each trade benefiting the participants at the moment of exchange, and demonstrated that such trades could converge to outcomes that were Pareto-inferior to sincere voting—final allocations could leave everyone worse off relative to no trading. In their well-known three-legislator, six-issue example, each trade benefited the two parties involved, but imposed negative externalities on the third. As trades accumulated, all legislators ended up

worse off. What appeared to be a cooperative mechanism to internalize preference intensity instead generated instability and inefficiency. The underlying reason is that vote trades impose externalities on those not party to the exchange, and therefore there is no presumption that trading will improve welfare.

The Riker–Brams paradox sparked a wave of skepticism about vote trading. Ferejohn (1974) shows that instability arises even with strictly pairwise trades, and his conclusion that "we really know very little theoretically about vote trading" effectively closed the literature for decades. More recent work has not brought more optimistic results. Casella and Palfrey (2019) formally analyze explicit vote trading and showed that stability requires very specific and restrictive preference assumptions, and even then the equilibrium path need not converge to the Condorcet winner. Laboratory experiments by Casella and Palfrey (2020) and the theoretical survey in Casella and Macé (2021) confirm that inefficiency and instability are generic features of explicit logrolling. The consensus view is that the analogy with markets is misleading: votes are not private goods, and individually rational exchanges fail to aggregate into collectively rational outcomes.

Against this background, our framework provides a resolution to the paradox. In the environment introduced in Section 2.3, there is no explicit vote-trading stage. It would, however, be possible to incorporate, at each period, multiple rounds of communication between two or more legislators. The multidimensionality of the policy space can be interpreted as distinct issues, and legislators could in principle make nonbinding deals of the form: "I will compromise on issue 1, which matters more to you, if you compromise on issue 2, which matters more to me." Because such agreements are nonbinding, they must be supported in equilibrium by the continuation values of repeated interaction. This means that the possibility of logrolling is already embedded in our construction of cluster-stable outcomes, and that introducing explicit trading rounds would not alter the set of stable equilibria.

The difficulties identified by Riker and Brams (1973) are present in our dynamic framework as well—externalities and cycles. To these are added the challenges of asymmetric information, since legislators' preferences evolve stochastically and remain private, and the absence of commitment, as nothing obliges legislators to abide by agreements made during the bargaining process, whether in the current

period or in the future. However, our main result shows that these difficulties do not preclude efficiency. Internally optimal outcomes can be sustained as cluster-stable equilibria. As illustrated in Example 1, legislators can enforce the socially optimal outcome—or at least an outcome arbitrarily close to it—in each period with probability arbitrarily close to one. Moreover, the argument in Example 1 is not limited to Euclidean preferences. More generally, whenever the socially optimal outcome prescribes different policies for different type profiles in Θ—that is, whenever the outcome is one-to-one—it satisfies the internal optimality condition and is therefore cluster-stable. This yields the following corollary.

Corollary 2.1. *Every one-to-one socially optimal outcome $\wp^* \in \mathcal{P}$ is cluster-stable.*

This result stands in sharp contrast with the paradox identified by Riker and Brams. Even in the absence of money and enforceable contracts, repeated legislative interaction in a multidimensional policy space allows legislators to approximate the cooperative benefits that logrolling was supposed to provide. Rather than producing instability and inefficiency, as in the classical paradox, vote trading in our framework converges to stable and socially desirable outcomes. The paradox disappears once stability is defined in terms of clusters and enforced through continuation values, rather than through exact, enforceable trades.

2.5.2 Vote buying: Money in legislative politics

The study of vote buying occupies a central place in the formal literature on legislative politics. Models in this tradition analyze how outside actors—typically interest groups or lobbies—use monetary incentives to influence collective decisions. The basic insight is that financial resources can be deployed to secure political support, shifting policy outcomes closer to the lobbyist's ideal point. A seminal contribution is Snyder Jr (1991), which introduces a game-theoretic model in which a lobby distributes resources to secure majority support in a legislature. Rather than targeting staunch allies or opponents, optimal strategies often involve persuading pivotal or marginal legislators whose support can be obtained at relatively low cost.

This line of research is developed most prominently by Groseclose and Snyder Jr (1996), which shows that sequential competition between rival vote buyers can generate "supermajority" coalitions that are strictly larger than minimal winning ones, while actually minimizing total spending. The result challenges the earlier assumption that minimal winning coalitions are always the cost-effective outcome. Subsequent work extends the framework to finite legislatures (Banks, 2000), bicameral institutions (Diermeier and Myerson, 1999), and sequential lobbying games (Le Breton et al., 2012), reinforcing the finding that equilibrium coalitions tend to be oversized whenever strategic interaction between competing buyers is taken into account.

However, this theoretical tradition sits uneasily with the empirical evidence. Interest groups routinely obtain policies worth billions, but lobbying expenditures and campaign contributions remain strikingly small by comparison—a puzzle known as the *Tullock paradox* (Tullock, 1967, 1997). Ansolabehere et al. (2003) famously ask, "Why is there so little money in U.S. politics?" and argue that much political giving resembles consumption rather than investment: donations reflect ideology, identity, or access-seeking, rather than expected policy returns. They also emphasize that political influence is often achieved by information provision rather than financial incentives, and that institutional constraints—separation of powers, electoral competition, and transparency—limit the effectiveness of monetary influence. Similarly, De Figueiredo and Richter (2014) survey lobbying expenditures and show that, relative to the scale of potential policy rents, the amounts actually spent are surprisingly modest. Their conclusion is that money is only one channel of influence, and often a relatively weak one.

In this subsection we develop a simple variant of the baseline model in Section 2.3 that allows us to revisit the logic of vote buying in a dynamic legislative setting with incomplete information. The extension does not rely on quasilinear preferences with respect to money, a standard assumption in the existing literature, and it delivers a stark implication for the Tullock paradox.

We now extend the baseline model of Section 2.3 to incorporate a lobbyist. As in the formal literature on vote buying, the lobby can commit at the beginning of the game to offering legislators monetary payments in exchange for their votes and

policy proposals. A key feature of our dynamic framework is that the lobby can condition its offers on the entire public history of votes and proposals made by legislators in all previous periods. A new challenge, relative to existing models, is the presence of incomplete information about legislators' policy preferences, which makes it more difficult for the lobby to determine what offers can effectively incentivize legislators.

Formally, consider the mechanism-design problem faced by a lobby that can offer each legislator i a contribution function $C_i(h^t)$ before the beginning of the game. This function specifies the amount the lobby commits to pay to i—or to receive from i if the contribution is negative—in each period t, as a function of the public history h^t of observed actions up to time t. Each legislator must then decide simultaneously whether to accept or reject the lobby's offer. If all legislators accept, then they receive the payments prescribed by the functions $C_i(\cdot)$; otherwise, they receive no payment, regardless of what occurs in the remainder of the game. Once the legislators have accepted or rejected the lobby's offers, events unfold exactly as in the baseline model.

The stage payoff to legislator i, when policy x is implemented and she receives a contribution c_i, is given by

$$v_i(x, c_i \mid \theta_i) \;\equiv\; u_i(x, \theta_i) + c_i, \tag{2.5.1}$$

where $u_i : X \to \mathbb{R}$ is the same policy utility function introduced in Section 2.3. The lobby's stage payoff is given by

$$v_0(x, c_1, \ldots, c_n) \;\equiv\; u_0(x) - \sum_{i=1}^{n} c_i, \tag{2.5.2}$$

which is assumed to have a maximizer, denoted by $\hat{x}_0$. Intuitively, the lobby values the policy x according to $u_0(x)$ but subtracts the cost of the transfers it must pay to legislators.

Stage payoffs are discounted as in the baseline model. We allow the lobby's discount factor to differ from that of the legislators, and we do not require the lobby to be farsighted. For expositional convenience, we assume that the lobby's

preferences are independent of any type. If we instead assumed that the lobby's preferences depended on a type θ_0, privately observed or not, it would suffice to suppose that the set of lobby types, say Θ_0, satisfies $k|\Theta_0| \geq \sum_{i=1}^{n} |\Theta_i|$. Our result below would still hold.

We are interested in the minimum expenditure required for the lobby to induce the legislature to implement its ideal policy $\hat{x}_0$ in every period. For every $\varepsilon > 0$, we say that the lobby *can achieve its first best with less than* ε if, for sufficiently high legislator discount factor δ, there exists a PBE in which, in every period t: (i)the policy $\hat{x}_0$ is implemented by the legislature; and (ii) the expected sum of the lobby's contributions to all legislators is smaller than ε.

To see how the lobby can do so, we can use the equilibrium construction described in Section 2.4.3 (and formally established in Section 2.A), but now replace the "equilibrium designer" of that subsection with the lobby. Since constant policy outcomes are necessarily internally optimal, we can apply the logic of that construction to implement the lobby's ideal policy outcome $\hat{\wp}_0$, which prescribes $\hat{x}_0$ regardless of the realization of the legislators' type profile $\theta \in \Theta$.

In the baseline model (Section 2.3), there are no actual payments to legislators. Rather, transfers are generated fictitiously through perturbations of the desired policy outcome. In contrast, in the vote-buying model considered here, it is no longer necessary to perturb $\hat{\wp}_0$, since the lobby can make actual monetary payments. Moreover, a simple normalization ensures that these transfers can be made arbitrarily close to $\varepsilon/2$ with probability arbitrarily close to one. Denote by $C^{\varepsilon}(\cdot)$ the contribution function that prescribes such payments. We conclude that, *if legislators accept* $C^{\varepsilon}(\cdot)$, then the lobby can induce the legislature to adopt its ideal policy in every period while ensuring that its expected contributions remain smaller than ε.

To show that the lobby can achieve its first best for less than ε, for any $\varepsilon > 0$, it remains to establish that legislators are indeed willing to accept the lobby's offer. By Theorem 2.1, there exists a PBE of the game without monetary payments (and with sufficiently high discount factors) in which each legislator i's expected payoff is arbitrarily close to

$$\mathbb{E}\left[u_i(\hat{x}_0, \theta_i)\right], \tag{2.5.3}$$

since a policy arbitrarily close to $\hat{x}_0$ is implemented by the legislature in every period. Suppose that this equilibrium is precisely the one played in the continuation game of the vote-buying model when one or more legislators reject the lobby's offer $C^\varepsilon(\cdot)$. These legislators then face a simple choice: accept $C^\varepsilon(\cdot)$, thereby receiving the per-period policy benefit $\mathbb{E}[u_i(\hat{x}_0, \theta_i)]$ plus a transfer arbitrarily close to $\varepsilon/2 > 0$; or reject $C^\varepsilon(\cdot)$ and receive only the policy benefit arbitrarily close to $\mathbb{E}[u_i(\hat{x}_0, \theta_i)]$. It follows that it is strictly optimal for all legislators to accept $C^\varepsilon(\cdot)$.

Corollary 2.2. *In the vote-buying model, for every $\varepsilon > 0$, the lobby can achieve its first best with less than ε.*

The corollary shows that, even under a simple bribery logic, a lobby can sustain its ideal policy while spending arbitrarily little. This provides a novel theoretical resolution of the Tullock paradox. Classic empirical studies emphasize that contributions are surprisingly small compared to the value of policy rents (Ansolabehere et al., 2003; De Figueiredo and Richter, 2014). Explanations include that donations are often consumption-like rather than investment-driven, that influence operates through information rather than money, and that institutional constraints limit the effectiveness of financial transfers.

Theoretical work has also sought to rationalize these patterns. Sequential vote-buying models (e.g., Chen and Zápal, 2022; Le Breton et al., 2012) show that leaders can secure passage at very low cost because many legislators are dispensable at the margin: equilibrium payments reflect pivotality under strategic substitution, not willingness to pay for the policy. Related arguments arise in simultaneous contracting with externalities (e.g., Chowdhury and Sengupta, 2012; Dal Bó, 2007; Genicot and Ray, 2006), where "pivotal bribes" or contingent transfers induce equilibria in which no voter is pivotal, collapsing payments toward zero. These mechanisms provide a theoretical foundation for the paradox: money appears scarce not because it is irrelevant, but because substitutability and strategic structure imply that only negligible transfers are needed.

Against this backdrop, Corollary 2.2 highlights a distinct but complementary mechanism. In our dynamic model with incomplete information, the structure of repeated interaction implies that legislators' incentives are already shaped by

continuation values. Because these dynamic payoffs provide discipline, even tiny monetary contributions are sufficient to shift outcomes toward the lobby's preferred policy. The logic is therefore not that money is irrelevant, but that repeated bargaining renders large transfers unnecessary.

Thus, while previous theoretical models emphasize pivotality and substitutability, and empirical work highlights ideology, information, and institutions, our framework complements these perspectives: even in a dynamic legislative environment with incomplete information, where no institutional barriers or ideological motives are assumed, the paradox persists. Monetary influence remains possible, but requires only negligible resources.

2.5.3 Agenda manipulation under incomplete information

The manipulation of legislative outcomes through control of the agenda is a central topic in the study of collective decision-making. This literature has long recognized that majority voting over multidimensional spaces can produce intransitive collective references, leading to cycles and instability in outcomes. The foundational results by McKelvey (1976, 1979) demonstrated that when preferences are sufficiently diverse across multiple dimensions, majority rule can fail to yield a stable social choice and that any point in the space can be reached given the right sequence of proposals. In the often referred to as *chaos theorem*, McKelvey showed that, by structuring the order of pairwise votes, the agenda setters can gain substantial power to manipulate outcomes, effectively steering collective decision toward their preferred alternatives. Schofield (1978) extended these insights to broader classes of differentiable preferences, reinforcing the conclusion that instability is not an anomaly but a generic feature of multidimensional majority rule. These results have been foundational in political economy, highlighting the centrality of agenda control in shaping policy outcomes.

Building on this work, a large literature emerged on how institutional features constrain or amplify agenda power (Banks 1985; Denzau and Mackay 1983; Krehbiel 1998; Romer and Rosenthal 1978, among many). An important dimension of this literature concerns information. Most of the early literature assumes that the

agenda setter had full knowledge of legislators' preferences, a strong assumption that grants the agenda setter the power to predict voting outcomes and strategically sequence votes to achieve preferred policies. However, in many legislative contexts, preferences are private or only imperfectly observed. A relatively small strand of research analyzes agenda manipulation under incomplete information. Banks (1990), Barberà and Gerber (2017), Casella (2011), and Ordeshook and Palfrey (1988) show that uncertainty about preferences or reversion policies limits the ability of agenda setters to dictate outcomes, but rarely eliminates their power altogether. Instead, incomplete information tends to constrain extreme manipulation, generate more nuanced equilibria, and sometimes turn agenda order into a vehicle for information transmission. The general conclusion is that informational frictions reduce the scope of agenda power but also make agenda-setting more strategically important.

To study this question in our setting, we consider a variant of the baseline model that incorporates endogenous agenda formation, following the spirit of Banks and Gasmi (1987). We refer to this extension as the *agenda-formation model*. The policy space and legislators' preferences are as in Section 2.3, assuming now that the number of legislators, $n \geq 3$, is odd. Deliberation and proposal stages are replaced by the following procedure: in each period t, all legislators simultaneously submit one or more proposals to legislator 1, who acts as the agenda setter. The setter then chooses the order in which proposals will be voted upon by the legislature. This ordering defines the amendment agenda through which the period-t policy will be selected: the first proposal is put to a simple majority vote against the second; the winner of that pairwise vote is then voted on against the third proposal in the agenda, and so on, until the final proposal is reached. The winner of the last pairwise vote constitutes the legislature's policy choice for period t. As before, the legislators' types are privately observed, while all the actions they take, including their proposal choices, are publicly observable.

Cluster-stability can be defined exactly as in Definition 1, and the equilibrium construction in Section 2.4.3 adapts naturally to this new framework. Specifically, policy proposals submitted by legislators can be used as messages to convey information about their type in each period. To see this, fix an internally optimal

policy outcome $\wp$. In order to allow each legislator i to signal their type, we assign a distinct "tagging proposal" $\bar{x}_i(\theta_i)$ outside $\wp(\Theta)$, for each legislator i and type $\theta_i \in \Theta_i$. This ensures that all such policies are distinct—i.e., that $\left|\bar{x}_i(\Theta_i)\right| = |\Theta_i|$. In each period t, the type-θ_i^t legislator is then prescribed to submit a set of proposals that includes $\bar{x}_i(\theta_i^t)$. This guarantees that, after the submission stage, all legislators have signaled their preference types to the others. Deviations in the proposal stages are discouraged in the same manner as in the baseline model—that is, through variations in future policy choices that negatively affect deviators.

It remains to guarantee that, in each period t, the policy $y(h^{t-1}, \theta^t)$ prescribed by the equilibrium designer of Section 2.4.3 to approximate $\wp(\theta^t)$ is included in the period-t endogenously formed agenda. Upon observing the public history h^{t-1} at the beginning of period t, legislators know the policy $y(h^{t-1}, \theta^t)$ that the designer would want implemented for any realization of types $\theta^t = (\theta_1, \ldots, \theta_n^t)$ on the equilibrium path. To do so, each legislator i submits the set of proposals

$$\left\{y(h^{t-1}, \theta) \colon \theta \in \Theta\right\} \cup \left\{\bar{x}_i(\theta_i^t)\right\}. \tag{2.5.4}$$

This ensures both that the designer's target policy $y(h^{t-1}, \theta^t)$ is always on the agenda in period t, and that legislators have signaled their types.

The remainder of the equilibrium construction proceeds as in the baseline model: internal optimality guarantees that no player has an incentive to deviate, regardless of the agenda setter's sequencing choices. We therefore have the following variant of Theorem 2.1.

Corollary 2.3 (Theorem 1')**.** *Every internally optimal policy outcome is cluster-stable in the agenda-formation model.*

Two implications follow. First, the agenda setter's most-preferred policy outcome among those that are internally optimal is cluster-stable. This suggests that asymmetric information about legislators' preferences only marginally limits the agenda setter's ability to secure favorable long-run outcomes, given the flexibility of internal optimality. In this sense, Corollary 2.3 aligns with the classic insights of McKelvey's chaos theorem, but it also resonates with later work showing that incomplete information alters but does not eliminate agenda power. For instance,

Banks (1985) shows that uncertainty about reversion points constrains the extent of manipulation, while Ordeshook and Palfrey (1988) emphasize how incomplete information about legislators' preferences undermines standard majority-rule results but still leaves room for agenda influence through procedural design. Similarly, Casella (2011) demonstrates that agenda control can operate as a form of cheap talk, transmitting limited information with only minor welfare effects, and Barberà and Gerber (2017) show how sequential voting structures allow agenda setters to extract information and maintain influence despite hidden preferences. Our contribution differs from these approaches: rather than focusing on one-shot or static voting games, we embed agenda manipulation in a dynamic framework with repeated interaction and cluster-stability. This allows us to show that even when preferences are private and evolving, the agenda setter retains robust long-run power, highlighting the resilience of agenda control in environments of incomplete information.

Second, the result also implies that *any* internally optimal outcome—including those least favorable to the setter—can be sustained in equilibrium. This highlights an issue of equilibrium multiplicity and selection familiar from the literature on dynamic legislative bargaining (e.g., Anesi and Duggan 2018). Determining how equilibria are selected, and under what conditions agenda setters can reliably capture their most-preferred cluster-stable outcome, remains an open question and lies beyond the scope of the present chapter.

2.6 Concluding remarks

In legislative bargaining, money-like transfers are scarce and procedurally rationed, credible commitment is limited, preferences are private and evolve over time, and policy is chosen in high-dimensional spaces where today's decision becomes tomorrow's status quo. Our objective in this chapter is to develop a stability notion and an accompanying equilibrium characterization for dynamic legislative bargaining with privately observed, evolving preferences and endogenous status quo, without relying on quasilinearity and commitment. We address this objective by shifting the unit of analysis from single policy points to decision rules and by

showing how repeated interaction supplies—through continuation values—the incentives that quasilinearity and commitment would otherwise provide.

We model an infinite-horizon legislature that meets in discrete periods. In each period, legislators privately observe their current preferences, propose policies, and then vote under a qualified-majority rule. If a proposal meets the threshold it is implemented; otherwise, status quo remains in place. The implemented policy becomes next period's status quo. Payoffs reflect discounted policy utilities over time. We study equilibria in which strategies depend on legislators' private information and the public history, and are optimal given beliefs, while beliefs are consistent with observed behavior—that is, perfect Bayesian equilibrium in a nontechnical sense: plans are sequentially rational and learning follows from what the legislature does on the path of play.

We formalize a stability notion tailored to this environment, which we call *cluster-stability*. Rather than require exact implementation period by period, we ask that equilibrium play concentrates, with high probability, in arbitrarily small neighborhoods of the policy prescribed by a decision rule. The focus is thus on the persistence of a mapping from private types to collective choices, not on a single policy point. This perspective adapts the complete-information idea of absorbing policies to settings where preferences are private and evolving, while remaining tractable in high-dimensional spaces.

Our main result provides a simple sufficient condition. We introduce *internal optimality*—utilitarian welfare maximization within an outcome's own range for every realized type profile—and show that, for generic preferences and sufficiently patient legislators, every internally optimal outcome is cluster-stable. Conceptually, the persistent objects in incomplete-information legislatures are internally optimal social choice mappings, and, generically, any such mapping can be supported as a long-run equilibrium rule.

The mechanism behind this result can be summarized as follows. If monetary transfers were available, Groves–Clarke payments would induce truthful revelation and implement welfare-maximizing policies on the relevant range. In our setting, repeated interaction substitutes for those missing instruments: continuation values operate as "shadow transfers" that replicate, period by period, the

incentive-compatible payments needed for truth-telling. Along the equilibrium path, future payoffs enter additively—like a numeraire—into current utilities, and history-contingent punishments calibrated by internal optimality deter deviations. Patience and continuation values supply the credibility that money and commitment would otherwise provide.

Our main result has strong welfare implications. Under Euclidean preferences, the set of internally optimal outcomes is typically very large—often indeterminate—so the absence of monetary transfers and commitment is not, by itself, a barrier to sustaining desirable agreements. A corollary underscores the point: any socially optimal, one-to-one outcome is trivially internally optimal and therefore cluster-stable. Recasting stability around decision rules also reframes the classic "paradox of vote trading": once exact implementation is relaxed and continuation values substitute for transfers, the standard inefficiency of logrolling without a numeraire no longer constrains long-run outcomes.

We apply the framework to two canonical problems. In a vote-buying extension, a lobby that precommits to contribution schedules contingent on public histories—without observing private preferences—implements its ideal policy every period while spending an arbitrarily small amount in expectation. Small transfers coordinate play on the lobby's preferred cluster because dynamic incentives do most of the work, offering a distinct resolution of the Tullock paradox consistent with modest spending and persistent influence. In an agenda-formation extension, proposals serve a dual role: alongside substantive options, each legislator includes a distinct 'tagging proposal' that publicly signals type, which reveals the relevant information on the path of play and guarantees that the target option appears on the amendment agenda. Every internally optimal outcome remains cluster-stable, so incomplete information restricts but does not eliminate agenda power.

Our characterization is deliberately one-sided: the main theorem establishes a *sufficiency* result only. In particular, we show that internal optimality—welfare maximization within an outcome's own range for every realized type profile—is enough to guarantee cluster-stability under the primitives of our model. We do not claim the converse. So an open question remains: Is cluster-stability possible without internal optimality?

Appendix

2.A Proof of Theorem 2.1

Roadmap. The objective is to prove that, generically, every internally optimal policy outcome is cluster-stable. The proof begins with preliminaries: we define the generic set $\mathcal{U}^*$ and introduce a notion of regular policy outcomes that guarantee local control of utilities via small perturbations of policies. The argument then proceeds in four steps. Step 1 shows that internal optimality on the restricted range of policies implies strict implementability through Groves-type transfers. Step 2 establishes that such outcomes can be implemented in a static mechanism with a numeraire. Step 3 extends this to a dynamic mechanism with blocks and statistical tests à la Escobar–Toikka, ensuring that implicit transfers remain uniformly small and mean-zero across blocks. Step 4 embeds this mechanism into the original legislative-bargaining game without commitment, using public punishments that adjust the implicit transfer vector; when δ is sufficiently close to one, these punishments are deterrent.

Preliminaries and genericity. We begin by defining the set $\mathcal{U}^*$ in the statement of Theorem 2.1. To do so, we first establish some notation.

Let $N \equiv \{1, \ldots, n\}$ denote the set of legislators, and recall that the set of policy outcomes $\mathcal{P}$ can be identified with a subset of Euclidean space $\mathbb{R}^{k|\Theta|}$. Define the

utility map

$$U : X \to \mathbb{R}^{\sum_i |\Theta_i|}, \qquad U(x) \equiv \left(u_i(x, \theta_i)\right)_{i \in N,\, \theta_i \in \Theta_i} \tag{2.A.1}$$

That is, $U(x)$ collects all legislators' utilities from policy x across all possible realizations of their types.

By Schofield (1980)'s Singularity Theorem A, together with our dimensionality assumption $k \geq \sum_{i=1}^{n} |\Theta_i|$, there exists a dense subset $\mathcal{U}^* \subseteq \mathcal{U}$ of policy preferences $\left(u_i(\cdot, \theta_i)\right)_{i \in N,\, \theta_i \in \Theta_i}$ such that, for all utilities in $\mathcal{U}^*$ and outside a closed, measure-zero set of policy outcomes $\wp$, the Jacobian of U has full row rank at every policy $\wp(\hat{\theta})$, $\hat{\theta} \in \Theta$. We refer to such policy outcomes $\wp$ as *regular*. Intuitively, regularity guarantees that near each prescribed policy $\wp(\hat{\theta})$ we can move utilities *locally and independently* (one degree of freedom per component of U) by small policy perturbations.

Suppose legislators' utility functions belong to $\mathcal{U}^*$, irrespective of their types, and take any internally optimal policy outcome $\bar{\wp}$. By density, for any $\eta > 0$ there exists a regular policy outcome $\wp_\eta$ within the η-neighborhood of $\bar{\wp}$ in $\mathcal{P}$. To establish the theorem, it suffices to prove that, for η sufficiently small, $\wp_\eta$ is cluster-stable.

To this end, fix $\varepsilon \in (0, 1)$. We construct a PBE σ of the legislative bargaining game for $\delta < 1$ arbitrarily high such that, for every period $t \in \mathbb{N}$, the probability that x_σ^t belongs to the ε-neighborhood of $\wp_\eta(\theta^t)$ exceeds $1 - \varepsilon$ — see Definition 1. We proceed in four steps.

Step 1. Simulating quasilinear preferences: local replication of small transfers.

Fix $\eta > 0$ and consider the regular policy outcome $\wp_\eta$. Regularity means that at every $\hat{\theta} \in \Theta$, the Jacobian of U at $\wp_\eta(\hat{\theta})$ has full row rank. By the local submersion theorem, this implies that the image of a neighborhood of $\wp_\eta(\hat{\theta})$ is an open neighborhood of $U(\wp_\eta(\hat{\theta}))$.

Hence, there exists a sufficiently small number $M > 0$ such that any utility shifts $m = (m_1, \ldots, m_n) \in \mathbb{R}^n$ with $|m_i| < M$ can be approximated through small perturbations of the policy. Formally, for each $\hat{\theta} \in \Theta$ and each m in that

neighborhood, there is a policy outcome $\wp_\eta^m$ close to $\wp_\eta$ such that

$$u_i\big(\wp_\eta^m(\hat{\theta}), \theta_i\big) = u_i\big(\wp_\eta(\hat{\theta}), \theta_i\big) + m_i, \qquad \forall i \in N,\ \theta_i \in \Theta_i. \tag{2.A.2}$$

Thus, in a sufficiently high-dimensional policy space, small policy moves replicate small "numeraire-like" utility adjustments simultaneously across players and types. This allows us to treat local perturbations of policies as if quasilinear transfers were available. $\diamondsuit$

Step 2. Implementation in $\wp_\eta(\Theta)$ with a numeraire:

The next step is to show that internal optimality implies strict implementability by Groves-type transfers when we restrict attention to the range of the outcome.

Consider the mechanism design problem of implementing the social choice function $\wp_\eta$ using message-dependent monetary payments $(m_1, \ldots, m_n)\colon \Theta \to \mathbb{R}^n$. Since $\bar{\wp}$ is internally optimal and Θ (and thus $\wp(\Theta)$) is finite, we have that

$$\sup_{\theta \in \Theta} \left\{ \sum_{i=1}^{n} u_i\big(\wp(\theta), \theta_i\big) - \max_{x \in \wp(\Theta) \setminus \{\wp(\theta)\}} \sum_{i=1}^{n} u_i(x, \theta_i) \right\} \tag{2.A.3}$$

is bounded away from zero. In words, for every type profile θ, the policy $\wp(\theta)$ delivers strictly higher utilitarian welfare than any other policy in the range $\wp(\Theta)$. By continuity of the utility functions $u_i(\cdot, \theta_i)$, there exists a sufficiently small $\bar{\eta} > 0$ such that $\wp_\eta$ inherits this property for all $\eta < \bar{\eta}$. From this point forward, we restrict attention to such values of η.

The internal optimality of $\wp_\eta$ then guarantees that it can be *strictly* implemented using Groves-type monetary transfers. Define $\tau : \Theta \to \mathbb{R}^n$ by

$$\tau_i(\theta) \ \equiv \ \sum_{j \in N \setminus \{i\}} u_j\big(\wp_\eta(\theta), \theta_j\big), \qquad \forall (i, \theta) \in N \times \Theta. \tag{2.A.4}$$

Suppose, toward a contradiction, that the direct revelation mechanism $(\wp_\eta(\cdot), \tau(\cdot))$ is not strictly incentive compatible. Then there exists a legislator $j \in N$, a profile

$\theta \in \Theta$, and some misreport $\hat{\theta}_j \in \Theta_j \setminus \{\theta_j\}$ such that $\wp_\eta(\hat{\theta}_j, \theta_{-j}) \neq \wp_\eta(\theta)$ and

$$u_j\big(\wp_\eta(\theta), \theta_j\big) + \tau_j(\theta) \;\leq\; u_j\big(\wp_\eta(\hat{\theta}_j, \theta_{-j}), \theta_j\big) + \tau_j(\hat{\theta}_j, \theta_{-j}). \tag{2.A.5}$$

Summing utilities across all legislators and applying the definition of τ, we obtain:

$$\sum_{i=1}^{n} u_i\big(\wp_\eta(\theta), \theta_i\big) \;\leq\; \sum_{i=1}^{n} u_i\big(\wp_\eta(\hat{\theta}_j, \theta_{-j}), \theta_i\big), \tag{2.A.6}$$

which implies that $\wp_\eta(\theta)$ is *not* the unique maximizer of the utilitarian social welfare function in $\wp(\Theta)$, contradicting the internal optimality of $\wp_\eta(\theta)$. Hence, the direct revelation mechanism $(\wp_\eta(\cdot), \tau(\cdot))$ is strictly incentive compatible. $\diamondsuit$

Step 3. A dynamic mechanism design problem:

We now use Steps 1 and 2 to construct a PBE that supports the policy outcome $\wp_\eta$. The idea is to build a dynamic mechanism (with commitment), following Anesi and Duggan (2018), in which the same social choice function is applied to legislators' reports in every period, but fictitious transfers are tracked across blocks of play in order to keep incentives aligned and perturbations small.

Formally, time is divided into blocks of T periods, where $T \in \mathbb{N}$. Each block $b = 1, 2, \ldots$ begins with a vector of fictitious monetary payments $m^{b-1} = (m_1^{b-1}, \ldots, m_n^{b-1})$ that the mechanism designer is understood to "owe" the legislators, based on the history of play in block $b-1$. We initialize with $m^0 \equiv (0, \ldots, 0)$.

Within each period $t = (b-1)T + 1, \ldots, bT$, all legislators simultaneously submit reports $\hat{\theta}^t = (\hat{\theta}_1^t, \ldots, \hat{\theta}_n^t) \in \Theta$. These reports are filtered through the statistical test of Escobar and Toikka (2013): if all of legislator i's previous reports during the current block have passed the test so far, then her report is accepted, $\vartheta_i^t = \hat{\theta}_i^t$; otherwise, ϑ_i^t is redrawn from the invariant distribution of i's type. The resulting generated profile of filtered reports is denoted $\vartheta^t = (\vartheta_1^t, \ldots, \vartheta_n^t)$. The implemented policy in period t is then

$$x^t \;=\; \wp^{m^{b-1}}(\vartheta^t), \tag{2.A.7}$$

where $\wp^{m^{b-1}}$ is a policy outcome determined by the fictitious payment vector m^{b-1},

defined below. Legislator i's collected payoff in period t is therefore

$$u_i\big(\wp^{m^{b-1}}(\vartheta^t), \theta_i^t\big). \tag{2.A.8}$$

Definition of the adjusted outcome. The construction of $\wp^{m^{b-1}}$ applied in each block b relies on Steps 1 and 2. First, for every $i \in N$ and $\theta \in \Theta$, let $\bar{\tau}_i \colon \Theta \to \mathbb{R}$ be defined as:

$$\bar{\tau}_i(\theta) \;\equiv\; \tau_i(\theta) - \mathbb{E}_{\bar{p}}\big[\tau_i(\tilde{\theta})\big], \tag{2.A.9}$$

where τ is as in Step 2, and the expectation is calculated with respect to $\bar{p}$, which denotes the unique invariant distribution of the irreducible Markov chain (p, M). (Existence and uniqueness of $\bar{p}$ is guaranteed by our assumption that (p, M) is irreducible.) That is, $\bar{\tau}_i(\theta)$ is legislator i's Groves payment in state θ, centered so that its expectation under the stationary distribution is zero.

Then, let the monetary payments inherited from block b be defined as:

$$m_i^b \;\equiv\; \begin{cases} \frac{1}{T}\sum_{t=(b-1)T+1}^{bT} \bar{\tau}_i(\vartheta^t) & \text{if } \left|\frac{1}{T}\sum_{t=(b-1)T+1}^{bT} \bar{\tau}_i(\vartheta^t)\right| < M/3, \\ -M/3 & \text{otherwise,} \end{cases} \tag{2.A.10}$$

for each legislator $i \in N$, where $M > 0$ is the bound identified in Step 1. (For expositional simplicity, we omit writing the dependence of m_i^b on the finite sequence $\{\vartheta^t\}_{t=(b-1)T+1}^{bT}$).

This definition guarantees that $|m_i^b| < M$ for all i, so that Step 1 can always be applied and therefore used to define $\wp^{m^{b-1}}$ as the perturbed policy outcome that satisfies:

$$u_i\big(\wp^{m^{b-1}}(\hat{\theta}), \theta_i\big) \;=\; u_i\big(\wp_\eta(\hat{\theta}), \theta_i\big) + m_i^{b-1}, \qquad \forall \hat{\theta} \in \Theta,\ \forall i \in N,\ \forall \theta_i \in \Theta_i. \tag{2.A.11}$$

That is, the adjusted outcome $\wp^{m^{b-1}}$ replicates the target outcome $\wp_\eta$ while embedding the fictitious transfer vector m^{b-1} directly into policy perturbations.

Incentives under the mechanism. Step 2 established that $(\wp_\eta(\cdot), \tau(\cdot))$ is strictly incentive compatible. Hence, so is $(\wp_\eta(\cdot), \bar{\tau}(\cdot))$, since centering by the stationary mean does not affect strict inequalities. Combining this with the Escobar–Toikka

test yields the following: by Lemmata A2–A3 in Anesi and Duggan (2018), if T is sufficiently large and δ sufficiently close to one, then truth-telling passes the statistical test with probability arbitrarily close to one, and the implemented policy remains within any prescribed neighborhood of $\wp_\eta(\theta^t)$ with probability at least $1-\varepsilon$ in every period. In particular, for every $\varepsilon \in (0,1)$ there exist such T and δ for which

$$\Pr\Big(\| \wp^{m^{b-1}}(\vartheta^t) - \wp_\eta(\theta^t) \| < \varepsilon \Big) \;\geq\; 1-\varepsilon \quad \text{for all } t \in \mathbb{N}, \tag{2.A.12}$$

as required by Definition 1. Moreover, because the centered transfers $\bar{\tau}_i$ have mean zero under the invariant distribution, the strong law of large numbers for Markov chains implies that block averages converge almost surely. Together with the truncation rule, this ensures m^b remains uniformly bounded (and in particular $|m_i^b| < M$ for all i and b), so only arbitrarily small policy perturbations are ever required, consistent with Step 1.

Altogether, we have constructed a dynamic direct mechanism (with commitment) in which legislators truthfully reveal their types, policies remain arbitrarily close to $\wp_\eta$, and fictitious transfers remain bounded and mean-zero across blocks. By Lemmata A2–A3 in Anesi and Duggan (2018), for every $\varepsilon \in (0,1)$ there exist T large enough and δ close enough to one such that the mechanism yields a PBE in which the policy implemented in each period lies in the ε-neighborhood of $\wp_\eta(\theta^t)$ with probability at least $1-\varepsilon$, as required by Definition 1. $\diamondsuit$

Step 4. Equilibrium of the legislative-bargaining game:

To establish that $\wp_\eta$ is cluster-stable, it remains to show that for sufficiently high values of $\delta < 1$, the legislative bargaining game introduced in Section 2.3 admits a PBE in which the legislators, despite the absence of any commitment technology, implement the mechanism described in Step 3. We proceed by construction.

Define the assessment (σ, μ) as follows:

- *Deliberation stages:* Legislators send the same reports and update their beliefs exactly as in the equilibrium of the mechanism from Step 3, for every possible message history.

- *Proposal stages:* If no deviation from the prescribed actions has been observed so far, legislators propose the policy that the mechanism would have selected for every message history. If a deviation has been observed, the proposer still submits the mechanism's prescribed policy, but with the fictitious payment vector m^b permanently adjusted: the component corresponding to the last legislator who deviated is reduced by $M/3$, while all others remain unchanged. Because $|m_i^b| < M$ by construction, this adjustment can be carried out using the construction described in Step 1.
- *Voting:* Legislators vote to accept the latest proposal if it is the one dictated by σ; otherwise, they reject.

It follows from Step 3 that deviations from σ in the deliberation stage cannot be profitable. Deviations in proposal and voting stages are publicly observable, since they depend only on message histories and not on private types. Such deviations are punished by permanently reducing the continuation payoff of the deviator by $M/3$ forever, a sanction that can be implemented through the perturbation device of Step 1 since $|m_i^b| < M$ is always maintained. For δ sufficiently close to one, the present value of this punishment strictly exceeds any short-run gain, and it is enforced collectively, as each legislator has an incentive to carry it out to avoid being sanctioned herself. $\diamondsuit$

Together, Steps 1 to 4 complete the proof of Theorem 2.1.■

Chapter 3

Extreme Polarization and the Limits of Delegation

3.1 Introduction

Political polarization has become a defining feature of modern democracies. McCarty (2018) defines it as "a process whereby the normal multiplicity of differences in a society increasingly align along a single dimension, and people increasingly perceive and describe politics and society in terms of "Us" versus "Them". In other words, polarization turns diverse political and social differences into a rigid divide, heightening antagonism between camps. The United States remains the paradigmatic case, with congressional voting records showing a steady widening of the partisan gap since the 1970s (Poole and Rosenthal 2000; McCarty et al. 2016), but similar dynamics are visible in Europe, Latin America, Africa, and Asia (Dalton 2021; Norris and Inglehart 2019; Agarwal 2024; Golder 2016). At the mass level, partisan identities increasingly shape views on economic, cultural, and social issues, producing electorates that are more consistent within parties but further apart across them, and generating growing animosity toward political opponents (Abramowitz 2010; Webster and Abramowitz 2017; Iyengar et al. 2012).

The causes of polarization are complex and mutually reinforcing. Rising inequality politicizes distributive conflict, while cultural and identity-based divisions harden partisan loyalties (Norris and Inglehart 2018; Gidron et al. 2018). Media transformations—from partisan outlets to social networks—create echo chambers, spread misinformation, and raise the electoral costs of compromise (Prior 2013; Bail et al. 2018). At the same time, populist leaders nationalize grievances into identity conflicts (Mudde 2007; Golder 2016). These dynamics converge in what McCoy et al. (2018) call "severe polarization": when ideological, social, and affective divisions align so strongly that adversaries are seen as existential enemies. Building on this, Pierson and Schickler (2020) describe polarization as a developmental process: Initial partisan divergence reshapes the environment of parties, interest groups, and media in ways that intensify polarization rather than correct it. This perspective suggests that current trends are unlikely to stabilize and may escalate toward extreme polarization.

A vast literature documents the consequences of polarization. In politics, it reduces incentives to compromise, lowers legislative productivity, and produces volatile

outcomes as governments undo rather than improve prior reforms (Binder 2004; Hetherington 2001; Aldrich and Rohde 2000; Sinclair 2016; McCarty 2011, 2019). Polarization reshapes electoral behavior (Abramowitz and Saunders 2008), party systems (Dalton 2008; Golder 2016), and coalition bargaining (Reiljan 2020). It undermines democratic legitimacy by reducing trust in institutions, weakening interpersonal trust, and lowering satisfaction with democracy (Fiorina and Abrams 2008; McCoy et al. 2018; Reiljan 2020), while citizens increasingly tolerate norm violations and even political violence (Iyengar et al. 2012; Graham and Svolik 2020; Mason 2018). Beyond politics, polarization weakens cooperation across groups, reduces public goods provision and organizational performance (Dimant 2024; Bursztyn et al. 2020), strains social relations (Nicholson et al. 2016), and distorts markets, consumer behavior, and labor productivity (Michelitch 2015; McConnell et al. 2018; Gift and Gift 2015).

The consequences of polarization extend across all domains of society, from politics to organizations and markets. Delegation is a central mechanism in each of these spheres: legislatures delegate to committees and bureaucracies, executives to agencies, and organizations to managers and boards. As polarization becomes more severe worldwide, it is crucial to understand how these delegation relationships function when trust erodes and preferences diverge sharply. Can delegation between two polarized actors still generate desirable outcomes? Despite the central role of delegation in governance and organizational design, its interaction with polarization remains largely unexplored. This chapter addresses this gap by analyzing the consequences of extreme polarization for delegation, focusing on veto-based delegation.

Veto-based delegation is a widely employed institutional arrangement across both public and private decision-making contexts. In political institutions, it most commonly appears as a closed rule procedure in legislative bargaining, wherein a committee proposes a bill that the full legislature may either approve without amendment or entirely reject. In bureaucratic settings, elected reformers frequently delegate reform-making authority to specialized agencies while retaining the right to overturn agency decisions. For example, the U.S. Congress exercises control over executive agencies through legislative vetoes, disapproval resolutions, and

procedural oversight. Comparable arrangements exist in other political systems, such as parliamentary negative procedures or regulatory review committees in the European Union. For instance, the Council and the European Parliament may object to delegated regulations proposed by the European Commission. In private organizations, veto delegation is equally pervasive: shareholders of publicly traded firms retain veto power over fundamental corporate changes, while boards of directors may ratify executive decisions but lack unilateral reform authority.

All of these examples involve what is referred to as expertise-based delegation: a setting in which a better-informed agent acts on behalf of a less-informed principal. These relationships are formally modeled through principal-agent frameworks with asymmetric information, divergent preferences, and the absence of monetary transfers. The agent holds superior knowledge about a relevant state variable—such as the consequences of reform or the quality of a project—while the principal designs the delegation mechanism. Because the principal cannot perfectly monitor the agent or contract on outcomes, she must choose an institutional arrangement that mitigates incentive problems while leveraging the agent's informational advantage. This leads to a foundational trade-off: granting the agent greater discretion improves the use of private information but increases the risk of reform drift, while restricting discretion reduces misalignment risks at the cost of informational efficiency. Under severe or extreme polarization, preference misalignment becomes so pronounced that the need to control drift may outweigh the benefits of discretion. Analyzing delegation in such environments is therefore essential to understanding how governance operates in polarized societies.

Cases of extreme polarization between principals and agents are common in practice. Far-right and far-left parties in coalition governments often legislate within constitutional bounds far from their ideal policies, creating a structural gap between what they want and what is feasible (Tsebelis 2002; Golder 2016). In the United States, legislators at ideological extremes regularly clash with bureaucracies whose professional cultures pull in the opposite direction (McCarty 2019). A particularly clear mechanism is "agency burrowing," whereby outgoing administrations entrench loyalists and policy changes within agencies to constrain their successors. For example, Mendelson (2003) and Lewis (2010) document how

presidents use appointments to shape long-term bureaucratic preferences. Historical and contemporary examples illustrate the point: from John Adams' "midnight judges" to Trump's efforts to convert political appointees into career posts and to lock in regulatory changes (Emerson and Michaels 2020; Shafie 2020). Similar dynamics appear in parliamentary and organizational settings, where boards and managers may be locked into systematically opposed positions.

Delegation can take many forms. At one extreme lies communication or cheap talk, in which the agent sends a costless, non-binding message and the principal retains full authority over the final decision. In legislative settings, this corresponds to the open rule procedure, where a committee provides information through its proposal, but the floor may amend freely. At the other extreme is full delegation, often described as rubber-stamping, whereby the agent unilaterally selects the reform and the principal automatically implements it. Intermediate between these extremes are several hybrid arrangements. One is interval delegation, where the principal defines a delegation set—a permissible set of actions from which the agent can choose. Another is veto-based delegation, where the agent makes a reform proposal that the principal can either accept or veto, in favor of a predetermined fallback option, typically the status quo.

A core advantage of veto-based delegation is its institutional simplicity and relatively low administrative cost. Constructing a detailed delegation set or specifying a discretion window may require extensive ex ante knowledge of the reform space, which principals often lack. Conversely, ex post oversight of the agent's decisions—through monitoring and sanctions—is costly and frequently infeasible. Veto-based arrangements avoid both of these difficulties by requiring the principal only to define a reversion point and retain binary approval authority. This simplicity has made veto delegation attractive in legislative oversight of bureaucracies and in corporate governance. In addition, some models interpret veto delegation as a reflection of the principal's limited commitment power (Callander et al. 2008). When the principal cannot credibly commit to fully delegate decision-making authority, veto power becomes a second-best device.

Despite its ubiquity, veto-based delegation has received comparatively less theoretical attention than other forms of delegation. One of the key challenges lies in

the richness of the equilibrium set: veto games often admit multiple equilibria. As emphasized by Lubensky and Schmidbauer (2018), equilibrium multiplicity and the absence of clear selection criteria complicate theoretical predictions. One could expect that keeping veto authority would weakly outperform full delegation, as the principal is not obligated to exercise it; however, theoretical models often lead to contradictory outcomes. For instance, Marino (2007) and Mylovanov (2008) show that veto mechanisms can strictly outperform full delegation, particularly when the principal benefits from screening out low-quality proposals. In contrast, Dessein (2002) demonstrates that full delegation may dominate veto, especially when the agent's bias is small and the principal gains more from truthful revelation than from screening.

In this chapter, I develop a sequential principal-agent model with asymmetric information in which an uninformed principal delegates authority to a better-informed agent. The agent privately observes the state of the world and decides whether to propose a reform. Before this, the principal selects an institutional regime: either full delegation, under which any reform proposed by the agent is automatically implemented, or veto-based delegation, under which the principal can approve or reject the agent's proposal in favor of a fixed default—the status quo. In either regime, the agent can abstain from proposing any reform. In such a case, the status quo prevails.

Both players evaluate reforms based on how far they are from their preferred outcomes. Each player has a state-dependent ideal point: the principal's preferred reform exactly matches the true state of the world, while the agent's preferred reform is consistently biased in one direction. This bias, or drift, is always positive but decreases as the state increases. In other words, the agent always favors more extreme reforms than the principal, but the gap between their preferences narrows in higher states. When a reform is implemented, both players receive a baseline benefit from enacting a reform and incur a cost proportional to the distance between the implemented reform and their ideal point. If no reform is implemented—either because the agent chooses not to propose one, or because the principal vetoes it—both players receive a fixed, exogenously given payoff from maintaining the status quo. This fallback utility is constant across all states

and reflects the default value of inaction.

The distinctive feature of the model is the assumption of extreme polarization between the principal and the agent. Formally, the policy space is restricted so that the principal's ideal point lies strictly below its lower bound, while the agent's ideal point lies strictly above its upper bound. This implies that both players' most preferred reforms are unattainable within the feasible set and that conflict over reforms persists in all states of the world.

The analysis focuses on pure strategy perfect Bayesian equilibria that satisfy the D1 refinement criterion of Cho and Kreps (1987). In the complete information setting, both the agent and the principal observe the true state before any proposal is made. Under veto-based delegation, the agent internalizes the principal's approval constraint and adjusts her proposal accordingly. She offers the closest reform the principal is willing to accept, provided that doing so yields her a higher payoff than sticking with the status quo. A reform can therefore be sustained when it is mutually beneficial. However, the agent captures all the benefits from the reform, pushing it as close as possible to her ideal point subject to the incentive compatibility constraint. As a consequence, the principal's utility is no greater than under the status quo. Reform drift is even larger under full delegation. In this case, the agent does not need to take into account the preferences of the principal, since she can implement any reform she wishes from the feasible set. She therefore chooses the upper bound of the reform set, as this maximizes her utility. A direct implication is that the agent always benefits from implementing this reform, whereas the utility of the principal depends on whether the net benefit from reform is sufficiently large relative to the status quo payoff.

In the main part of the analysis, I consider the case in which the agent privately observes the state of the world, while the principal only knows its prior distribution. The first observation is that under full delegation, the outcomes of the game are identical to those under complete information, since the agent continues being informed of the state of the world and can select her most convenient policy. However, under veto-based delegation, the principal retains the right to accept or reject any reform proposed by the agent, who privately observes the state and decides which reform to put forward. Upon receiving a proposal, the principal

updates his beliefs about the state and accepts the reform only if the expected utility exceeds the payoff from maintaining the status quo. This creates a signaling mechanism: the principal relies on the proposal to infer the underlying state and to decide whether to approve it. Anticipating this, the agent has an incentive to use her proposal strategically—not only to implement her preferred reform, but also to signal the state in a way that persuades the principal to accept it.

The main result of this chapter shows that these signaling incentives lead to complete gridlock: the only equilibrium that survives the D1 criterion is one in which all proposals are vetoed and the status quo always prevails. To understand this result, note first that both players' preferred reforms increase with the state of the world, but the agent's bias implies that she always prefers reforms further to the right than those favored by the principal. As a consequence, the agent has an incentive to signal a high state by proposing large reforms. Low-state agents thus find it profitable to mimic the behavior of high-state agents, knowing that the principal will interpret larger proposals as evidence of a high state. However, this incentive leads all agent types to push proposals upward, resulting in reforms that exceed the principal's acceptance threshold. As a result, all reforms are vetoed, and the only sustainable equilibrium involves gridlock.

These results show that extreme polarization, when combined with incomplete information, makes gridlock the inevitable outcome of veto delegation. An important observation is that the presence of incomplete information is essential for reform gridlock: under complete information, extreme polarization and veto delegation can produce policy reform. A related observation is that although incomplete information produces complete gridlock, it does not alter the principal's payoffs: under full delegation, the outcome is identical to complete information, and under veto delegation, the principal always earns a payoff at least as high as the status quo.

Comparing both regimes, whether the principal prefers to keep veto rights depends on the net benefit from implementing a reform, the state, and the upper bound of the policy space. Since under full delegation the principal earns at most the same payoff as under the status quo, veto rights act as a form of "insurance": they shield the principal from reforms that would leave him worse off than in-

action. The price of this insurance, however, is high, as reforms can never be implemented.

These results speak to an extensive literature on gridlock that identifies polarization of preferences as a main driver of stalemate. This literature also emphasizes that gridlock often results from the interaction of multiple forces, among which the existence of veto players and incomplete information are particularly important. Finally, the results also relate to the large literature on veto delegation, which analyzes the effect of bias or policy drift between principals and agents as the main determinant of whether veto delegation can be beneficial relative to other forms of delegation. This literature provides contradictory results on whether drift makes veto delegation superior or inferior to other arrangements. The setup analyzed here pushes the idea of policy drift to the extreme, with a biased agent and ideal policies outside the policy set. This chapter is closely related to Marino (2007), who also compares full delegation with veto delegation in a setting *without* extreme polarization and in which full policy gridlock is never an equilibrium outcome.

The remainder of the chapter is organized as follows. Section 3.2 reviews the relevant literature on delegation. section 3.3 introduces the model and outlines its key components. Section 3.4 characterizes the equilibria under both the benchmark specification and the incomplete information environment, with and without veto power. In Section 3.5 I discuss the results and its implications. Section 3.6 concludes with final observations and remarks.

3.2 Related literature

This chapter contributes to the literature on delegation, specifically the strand that examines how much decision-making authority an uninformed principal should grant to an informed agent, when contingent transfers between both players are infeasible.[15] This question arises in a variety of institutional settings. In public administration, it is central to the design of regulatory frameworks in which elected officials delegate reform-making authority to bureaucratic agencies (Epstein and O'Halloran 1999; Huber and Shipan 2002). In legislative politics, it concerns the

allocation of proposal and amendment powers to legislative committees under procedural rules such as open and closed rules (Gilligan and Krehbiel 1987; Baron 2000). Similar problems of authority allocation appear in organizational economics, where researchers study the optimal delegation of decision rights in hierarchical or decentralized organizations (Melumad and Shibano 1991; Dessein 2002).

When a principal chooses to delegate, she must determine how much discretion to grant the agent. The core trade-off is that broader discretion allows the agent to leverage superior information to improve reform outcomes, but also increases the risk of reform drift due to misaligned preferences. The formal analysis of this problem was initiated by Holmström (1977, 1984), who model a principal choosing a set of permissible actions from which an informed agent selects after privately observing the state. The delegation set—termed also "discretionary window"—is defined as a connected subset of the policy space.[16] Interval delegation, he argues, is both analytically tractable and widely observed in practice, and he demonstrates that discretion expands as the preferences of the principal and agent converge, a result commonly referred to as the *ally principle*. He establishes conditions for the existence of an optimal solution to the delegation problem using a series of examples with quadratic preferences and uniform uncertainty, under the restriction that the delegation set takes the form of a single interval.

Building on Holmström's framework, discretionary delegation models have been applied to analyze statutory delegation in the United States. Epstein and O'Halloran (1994) develop a formal model of bureaucratic discretion in which legislators specify an interval of policies within which bureaucrats may choose. They show how discretion expands with policy uncertainty, and analyze how institutional design adapts to expected shifts in future coalitions. Epstein and O'Halloran (1999) demonstrate empirically that discretion is broader under unified government than under divided government, consistent with the ally principle. Huber and Shipan (2002) expand the model to introduce presidents, veto players, and courts, and analyze how noncompliance risk affects the optimal level of delegation, while

[15]The delegation literature is extensive and cannot be fully reviewed here. For a comprehensive survey, see Bendor et al. (2001).

[16] Full delegation is typically modeled as a special case of discretionary delegation in which the discretionary window spans the entire policy space—see Gailmard and Patty (2019).

Volden (2002) shows that discretion is harder to retract when both Congress and the president must approve it. Beyond politics, Dessein (2002) applies Holmström's framework to organizations, showing that firms delegate authority to informed subordinates when preference divergence is limited, whereas with large conflicts communication[17] is preferable to discretion.[18]

These models treated the discretionary window as exogenously given. A central goal of the subsequent literature has been to characterize the optimal (endogenous) delegation set among all incentive-compatible mechanisms. The earliest approach in this strand was to identify conditions under which an interval delegation set is optimal. Such an interval may take the form of a cap on the actions of an upward-biased agent, a floor on the actions of a downward-biased agent, or a cap together with a floor when the agent's actions are sometimes too high and sometimes too low. Melumad and Shibano (1991) show that interval delegation is optimal with quadratic preferences, unless the agent's bias is highly sensitive to the state, in which case a two-point set may emerge. Using state-independent biases, Martimort and Semenov (2006) provide sufficient conditions under which the optimal delegation mechanism is continuous, while Gailmard (2009) demonstrates that the optimal choice for the principal is always a compact and connected set of actions. Alonso and Matouschek (2008) offer a general characterization of optimal delegation in quadratic settings and establish necessary and sufficient conditions for interval delegation to be optimal. They further show that constrained delegation strictly dominates both full delegation and cheap talk whenever feasible. More recently, Amador and Bagwell (2013) present the most general sufficient and

[17] A parallel strand of the literature studies whether the principal should delegate at all or instead rely on communication to extract the agent's private information. This work builds on the cheap-talk framework of Crawford and Sobel (1982), where the informed agent sends a non-binding message and the principal retains full control. Dessein (2002) shows that delegation dominates communication when preference divergence is small, while Gailmard (2009) finds that delegation strictly outperforms communication for any degree of conflict, due to the agent's ability to internalize the consequences of her choices. As communication is not the focus of this chapter, I do not survey this literature in detail.

[18] Extensions to these models include: repeated interaction to address the principal's commitment problem (Alonso and Matouschek, 2007); the role of bureaucratic capacity (Huber and McCarty, 2004); subversion risks outside the delegated set (Gailmard, 2002); delegation under separation of powers (McCarty 2004; Crombez et al. 2006); interest group influence (Boehmke et al. 2006; Gailmard and Patty 2013); and multi-agent or partial delegation environments (Gailmard and Patty, 2019).

necessary conditions to date for interval delegation to be optimal, extending the analysis to non-quadratic preferences and state-dependent biases. They apply their results to the design of international trade agreements, where the optimal contract often takes the form of a tariff cap.

While early contributions established the optimality of continuous interval delegation, subsequent research has shown that this need not always be the case. Krähmer and Kováč (2016) are the first to demonstrate formally that optimal delegation sets can be disconnected. In a model of sequential delegation, they show that when sequential screening is valuable, the principal may optimally exclude intermediate actions, producing a set with gaps that permits only extreme choices (for example, a union of a low singleton and a high interval). Amador et al. (2018) extend the analysis of Amador and Bagwell (2013) to allow for degenerate intervals—single-point delegation sets in which the agent implements the same action regardless of the state—and provide sufficient conditions under which such degenerate allocations are optimal. More recently, Ball and Gao (2024) study a setting in which the agent can acquire costly information. They show that the optimal delegation set may take several non-standard forms: a "hollow set" featuring a gap around the agent's ex ante preferred action, a standard continuous interval, or a "high-point set" consisting of an interval combined with a singleton at the top. Saran (2022) develops a dynamic optimization approach to delegation with participation constraints, proving that the optimal allocation may involve a finite menu of discrete actions rather than a continuous set, with applications to volunteer contracts and donation schemes. Finally, Amador et al. (2025) provide the most general characterization to date, offering sufficient conditions for a wide class of utility functions under which the optimal delegation set may take the form of an interval, a finite menu, or a hybrid that combines connected intervals with

distinct singleton actions.[19]

Looking at the case of full delegation, a growing body of literature has focused on veto-based delegation. It departs from the optimal delegation literature by focusing on a specific game form in which an informed agent can propose any reform within the reform space and the principal may either accept the proposal or veto it in favor of an outside option, typically the status quo.

The role of veto power in delegation was first analyzed in the legislative bargaining literature, specifically through the study of closed rule procedures: a legislative committee proposes a bill that is then subject to an up-or-down vote by the full legislature, with the status quo prevailing in the event of rejection. Gilligan and Krehbiel (1987, 1989) introduced formal models comparing open and closed rule procedures,[20] using both homogeneous and heterogeneous committees. Alongside Krishna and Morgan (2001) and Martin (1997), they show that the closed rule may outperform the open rule only when the committee's bias relative to the floor is sufficiently small. Baron (2000) adopts a mechanism design approach to compare open rule with deference, a procedure under which the legislature may or may not exercise veto power. He develops a screening model in which deference frequently yields superior outcomes to the open rule, even when the committee is highly biased.[21]

The use of veto delegation has been extended to general principal-agent frame-

[19] Extensions of the optimal delegation framework include: stochastic mechanisms that allow for lotteries over policies (Kováč and Mylovanov, 2009); multidimensional policy spaces (Kleiner, 2022; Koessler and Martimort, 2012); sequential or dynamic delegation (Krähmer and Kováč, 2016; Saran, 2022); money burning and non-monetary incentives (Amador and Bagwell, 2013; Ambrus and Egorov, 2017); the role of capture and endogenous information (Lefebvre and Martimort, 2022); connections between delegation and Bayesian persuasion (Kolotilin and Zapechelnyuk, 2025); costly information acquisition (Ball and Gao, 2024); applications to fiscal and monetary policy rules (Amador and Bagwell, 2022; Halac and Yared, 2014, 2018, 2020, 2022a, 2022b; Sublet, 2023); to tariff caps (Amador and Bagwell, 2013); to monopoly regulation (Amador and Bagwell, 2022; Kolotilin and Zapechelnyuk, 2025); and to quality certification (Zapechelnyuk, 2020).

[20] Under an open rule, the full legislature may freely amend the committee's proposal during floor consideration. Models of open rule closely resemble cheap talk delegation models, as they follow the Crawford–Sobel framework.

[21] Extensions of this line of work include dynamic legislative bargaining with overlapping generations (Diermeier, 1995); the committee's gatekeeping power (Epstein and O'Halloran, 1994); closed rules with oversight committees (Van Gestel and Crombez, 2014); lobbying access to committees (Ambrus et al., 2013); and delegation in supranational settings such as the European Union (Franchino, 2004).

works, in which a principal delegates decision-making authority to a better-informed agent while retaining the right to veto the agent's proposal. Dessein (2002) and Marino (2007) compare full delegation with and without veto rights, in a setup closely related to that analyzed in this chapter. Dessein shows that full delegation outperforms veto delegation when the agent's bias is sufficiently small, primarily because the possibility of rejection introduces noise into the reform outcome. In contrast, Marino (2007) challenges this conclusion, demonstrating that veto power can yield superior outcomes by inducing the agent to screen out low-quality proposals. The divergence between these results has been attributed to differences in the assumed distribution of the state variable and the alignment of the reversion point with the principal's preferences, as emphasized by Mylovanov (2008).

Mylovanov adopts a mechanism design approach in which the reversion point is endogenous and develops what he terms the veto-power principle. He shows that the principal can replicate optimal delegation outcomes through veto-based mechanisms by appropriately selecting the default. By adjusting the fallback option, veto-based delegation can implement outcomes equivalent to those under full delegation, interval delegation, or cheap talk. Building on this logic, Hu and Lei (2025) introduce a stochastic veto mechanism in a model with a state-independent agent. They show that when the principal is more risk-averse than the agent with respect to non-status quo outcomes, stochastic veto becomes the optimal delegation mechanism.

Lubensky and Schmidbauer (2018) argue that the limited use of veto-based delegation in formal economic modeling stems, in part, from the complexity of the equilibrium set and the challenges of equilibrium selection. Much of the legislative bargaining literature has focused on identifying the most informative equilibrium—i.e., the one that transmits the most precise information to the principal.[22] The debate over equilibrium selection originated with Krishna and Morgan (2001)'s critique of Gilligan and Krehbiel (1989), where Krishna and Morgan constructed alternative equilibria that yielded significantly greater informational efficiency. Krehbiel (2001) accepts the theoretical validity of these refinements but raises concerns about their empirical relevance and plausibility. To advance this debate,

Lubensky and Schmidbauer (2018) develop a method to fully characterize the equilibrium set in a veto game and identify the most informative equilibrium, confirming that Krishna and Morgan's proposal is, in fact, maximally informative.

Other contributions in the veto delegation literature move away from information transmission as the primary objective and instead focus on maximizing the principal's expected utility. In this respect, the present chapter is closely related to Marino (2007), who departs from the information-centric approach and develops a welfare comparison between full delegation and veto delegation. The results of this paper are particularly relevant to my analysis, and I return to them in Section 3.5, where I compare Marino's findings to the results of this chapter.

3.3 The model

This section presents a delegation game between a less-informed *principal*, denoted p (he), and a better-informed *agent*, denoted a (she). The model captures a common institutional setting in which decision-making authority is fully delegated to an informed agent, with the possibility that the principal retains veto power. The environment is one of incomplete information: the agent privately observes a payoff-relevant state of the world before deciding whether to propose a reform. The central focus is on the effects of preferences' polarization between both players and the principal's veto rights in shaping the strategic interaction and reform outcomes.

Formally, the principal first chooses the delegation regime it offers to the agent, for her to make a *reform decision*[23] after she has observed the state $\omega \in [0,1]$. The agent may either propose a reform $x \in [\underline{x}, \overline{x}]$, or refrain from proposing any reform, in which case the status quo prevails. If the agent proposes a reform

[22]Formally, an equilibrium is more informative if its posterior beliefs are a mean-preserving contraction of those in another equilibrium. For instance, a fully separating equilibrium is more informative than any pooling equilibrium.

[23] The term "reform" is used generically to denote a delegated decision. It may correspond to a policy change in bureaucratic or legislative settings, such as regulatory delegation to agencies (Epstein and O'Halloran, 1999; Huber and McCarty, 2004), a committee's bill proposal under a closed rule (Gilligan and Krehbiel, 1987; Krishna and Morgan, 2001), or a project or strategic decision in organizational economics (Dessein, 2002).

under *full delegation*, it is automatically enacted. Under *veto delegation*, however, the reform is implemented only if the principal accepts it; otherwise, the status quo is maintained.

Each player $i \in \{a, p\}$ receives a constant and strictly positive utility $\bar{u}_i$ when the status quo is in place, independent of the realized state. If a reform $x \in [\underline{x}, \overline{x}]$ is implemented, player i receives utility:

$$u_i(x, \omega) = \beta - |\hat{x}_i(\omega) - x|, \tag{3.3.1}$$

where $\beta > 0$ is a baseline benefit from implementing a reform,[24] and $\hat{x}_i(\omega)$ is player i's ideal point, which depends on the realization of the state variable ω. The state ω represents implementation uncertainty and is drawn from a distribution with cumulative distribution function F and a strictly positive, continuous density f over $[0, 1]$.

A central feature of the model is the asymmetric information structure: the state ω is observed privately by the agent but not by the principal. The principal only knows the prior distribution of ω.

Payoffs.

The utility functions are given by:

$$u_p(x, \omega) = \begin{cases} \beta - |\omega - x| & \text{if a reform } x \text{ is implemented,} \\ \bar{u}_p & \text{if the status quo prevails,} \end{cases} \tag{3.3.2}$$

$$u_a(x, \omega) = \begin{cases} \beta - |\omega + b(\omega) - x| & \text{if a reform } x \text{ is implemented,} \\ \bar{u}_a & \text{if the status quo prevails.} \end{cases} \tag{3.3.3}$$

The principal's ideal point is given by $\hat{x}_p(\omega) = \omega$, indicating that his preferred reform perfectly aligns with the true state. In contrast, the agent's ideal point is defined as $\hat{x}_a(\omega) = \omega + b(\omega)$, where $b(\omega) > 0$ represents a state-dependent bias or drift. The function $b(\cdot)$ is continuous and strictly decreasing, with derivative

[24] Assuming that β is the same for both agents is without loss of generality and purely to simplify the exposition.

$b'(\omega) \in (-1, 0)$ for all $\omega \in [0, 1]$. This formulation captures the idea that while the agent always favors more extreme reforms than the principal, the extent of this misalignment decreases as the state increases.

To ensure extreme polarization between players, the reform space is constrained in a way that the two players' ideal points are permanently outside of the reform space, on opposite sides. This is, I impose $\underline{x} > 1$ and $b(0) > \overline{x}$, which guarantee that both $\hat{x}_a$ and $\hat{x}_p$ lie strictly outside of the policy space $x \in (\underline{x}, \overline{x})$, for all possible values ω can take. This guarantees that even in extreme states, the principal and the agent never fully agree on the optimal reform, ensuring persistent policy conflict in all states of the world.

I assume $b(0) > 2\beta - \bar{u}_a - \bar{u}_p$, which ensures that the most biased agent does not benefit from any reform that would be acceptable to the principal. This condition rules out borderline cases and avoids equilibrium existence problems.

Two additional assumptions guarantee that, taken individually, each player could benefit from implementing some reform within the reform space, regardless of the realized state. Specifically, I impose that $\beta - \bar{u}_p > 1 - \omega$ and $\beta - \bar{u}_a > 1 + b(1) - b(0)$. Intuitively, these conditions ensure that the net benefit from reform relative to the status quo is sufficiently large, so that for every ω, there exist at least a reform in the interval $[\underline{x}, \overline{x}]$ that is strictly preferred to the status quo, for each player.

Timing.

The sequence of moves is as follows:

- Stage 1. The principal chooses whether to retain veto rights (veto delegation, *vd*) or to delegate without veto (full delegation, *fd*).
- Stage 2. Nature draws the state $\omega \in [0, 1]$, which is privately observed by the agent.
- Stage 3. The agent proposes a reform $x \in [\underline{x}, \overline{x}]$. Alternatively, it may choose not to propose any reform.
- Stage 4. Under veto delegation, the principal observes x, forms beliefs about ω, and decides whether to accept or reject the proposal. Under full delegation, the proposed reform is implemented automatically.

- Stage 5. If the reform is accepted or automatically implemented, x is adopted. Otherwise, the status quo prevails. Payoffs are then realized.

Strategies and Equilibrium.

The model is a sequential game of incomplete information between an informed agent and an uninformed principal. The solution concept is Perfect Bayesian Equilibrium (PBE), restricted to pure strategies and refined using the D1 criterion.

An action for the agent is the choice of a reform to propose or the decision to refrain from any proposal. Formally, the agent's strategy is a function $\sigma_a : [0,1] \to [\underline{x}, \overline{x}] \cup \{\emptyset\}$, where $\emptyset$ denotes no proposal. In words, a strategy for the agent is a function that maps every state of the world ω to either a reform proposal within the reform space or a decision to propose no reform and let the status quo prevail.

An action for the principal at Stage 1 is a choice $\psi \in \{fd, vd\}$, where fd stands for "full delegation" and vd stands for "veto delegation". If $\psi = vd$, then at Stage 4, after observing the agent's proposed reform $x \in [\underline{x}, \overline{x}]$, the principal chooses $d \in \{a, r\}$, where a stands for "accept" and r stands for "reject". Thus, a strategy for the principal at Stage 4 is a function $\sigma_p^4 : [\underline{x}, \overline{x}] \to \{a, r\}$. In words, a strategy for the principal at Stage 4 is a function that maps a proposed reform to a decision to accept or reject the reform. Altogether, the principal's full strategy is a pair $\sigma_p = (\psi, \sigma_p^4)$.

A perfect Bayesian equilibrium (PBE) in this game is a combination of strategies of the agent and the principal, and a belief system for the principal, such that:

- Given the principal's strategy and the belief system, the agent's strategy σ_a maximizes her payoff at every state $\omega \in [0,1]$.
- Given the agent's strategy, the principal's strategy $\sigma_p = (\psi, \sigma_p^4)$ maximizes his expected utility at each of his decision nodes.
- The belief system of the principal, denoted by η, assigns a posterior belief over ω conditional on each proposal x from the agent and satisfies Bayes' Rule whenever possible.

I apply the D1 refinement developed by Cho and Kreps (1987) to rule out equilibria supported by unreasonable off-path beliefs. This refinement operates as follows:

suppose two different agent types, ω and ω', could plausibly be behind a deviation to some off-equilibrium reform x'. If the set of responses by the principal that would make type ω willing to deviate is strictly contained within the set that would induce type ω' to deviate, then the principal assigns full belief to ω' as the likely deviator. In other words, types for whom the deviation is more attractive are considered infinitely more likely to be responsible for it.[25]

Unless otherwise stated, the term "equilibrium" refers to a pure strategy Perfect Bayesian Equilibrium satisfying D1. Finally, I assume that indifferent players induce reforms rather than the status quo.

3.4 Benchmark and delegation regimes

This section characterizes the outcomes of the game under both complete and incomplete information. The first part establishes a benchmark by analyzing the case in which both the agent and the principal observe the true state of the world. The second part introduces the institutional regimes that apply when the agent is privately informed.

3.4.1 Perfect information

Suppose that the realization of the state $\omega \in [0, 1]$ is commonly observed by both players. The reform game then reduces to a complete information setting, allowing for a clear benchmark comparison across institutional regimes.

Consider first the case of veto delegation, in which the principal has veto rights over the reform proposed by the agent. In this regime, the agent must internalize the principal's preferences when choosing a reform. She can either propose a reform x that the principal accepts—yielding her $u_a^{vd} = \beta - |\omega + b(\omega) - x|$—or refrain from proposing any reform and receive $\bar{u}_a$.

The agent proposes a reform only if it is both acceptable to the principal and yields her at least as much utility as the status quo. The principal, in turn, accepts any

[25] See Section 3.A for a full statement and discussion of the D1 criterion.

reform that leaves him weakly better off than maintaining the status quo. That is, he accepts any reform satisfying:

$$\beta - |\omega - x| \geq \bar{u}_p, \tag{3.4.1}$$

which is equivalent to:

$$x \leq \omega + \beta - \bar{u}_p. \tag{3.4.2}$$

Let $x_o(\omega) \equiv \omega + \beta - \bar{u}_p$ denote the largest reform the principal is willing to accept for a given ω. For the agent, any reform $x < x_o(\omega)$ is dominated by $x_o(\omega)$, as she prefers larger reforms. She proposes $x(\omega) = x_o(\omega)$ only if it gives her at least as much utility as the status quo. That is:

$$\beta - |\omega + b(\omega) - x_o(\omega)| \geq \bar{u}_a. \tag{3.4.3}$$

Observation 3.1. *With complete information and veto delegation, the agent successfully proposes a reform $x(\omega) = x_{max}(\omega) = \omega + \beta - \bar{u}_p$ in state $\omega \in [0,1]$ if and only if:*

$$b(\omega) \leq 2\beta - \bar{u}_p - \bar{u}_a. \tag{3.4.4}$$

In such a case, the agent receives a payoff $u_a^{vd} = 2\beta - b(\omega) - \bar{u}_p$, and the principal receives $u_p^{vd} = \bar{u}_p$, leaving him indifferent between this reform and the status quo.

Consider now the case of full delegation, in which the principal does not have veto power. Since the agent seeks to maximize her utility and faces no constraints, she proposes the largest reform possible, $x = \overline{x}$, as this is the closest she can get to her ideal reform $\hat{x}_a(\omega) = \omega + b(\omega)$.

Her payoff is then $u_a^{fd}(\omega) = \beta - |\omega + b(\omega) - \overline{x}|$. That is, her utility depends on how far her ideal reform lies from the upper bound of the reform space. The agent prefers implementing the reform $\overline{x}$ if the utility she receives exceeds her status quo utility. This condition is:

$$\beta - |\omega + b(\omega) - \overline{x}| \geq \bar{u}_a. \tag{3.4.5}$$

Under the assumption that the net benefit from reform relative to the status quo is sufficiently large, and that $\overline{x}$ lies within the interval of policies that the agent strictly prefers over the status quo for all ω, this condition is satisfied. Therefore, the agent always proposes $x = \overline{x}$, and this reform is automatically implemented.

From the principal's perspective, whether he prefers this reform to the status quo depends on whether the reform falls within his acceptable range. The principal benefits from $x = \overline{x}$ being implemented if and only if:

$$\beta - |\omega - \overline{x}| \geq \bar{u}_p, \tag{3.4.6}$$

which simplifies to:

$$\omega + \beta - \bar{u}_p \geq \overline{x}. \tag{3.4.7}$$

Observation 3.2. *With complete information and full delegation, the agent successfully proposes a reform $x = \overline{x}$ and receives $u_a^{fd} = \beta - \omega - b(\omega) + \overline{x}$. Under $\overline{x}$, the principal receives $u_p^{fd} = \beta + \omega - \overline{x}$, which is preferred to the status quo if and only if:*

$$\overline{x} \leq \omega + \beta - \bar{u}_p. \tag{3.4.8}$$

Comparing both regimes is straightforward. Under veto delegation, a reform is implemented only when it weakly benefits both players; otherwise, the status quo prevails. Since the agent sets the reform as close as possible to her own ideal point, subject to the principal's approval, she captures all the net gains from reform, while the principal is left exactly as well off as under the status quo. In contrast, under full delegation, the agent always proposes the most extreme reform in her preferred direction, $x = \overline{x}$, which is automatically enacted. Whether the principal is better off in this case depends on whether the enacted reform yields higher utility than his fallback from the status quo.

Finally, note that if the principal does not delegate and instead retains authority to choose the reform directly, he is strictly better off than under either delegation regime. Since his ideal reform is $\hat{x}_p(\omega) = \omega$, which lies outside the feasible reform space, he would implement the closest available option — namely, $x = \underline{x}$.

3.4.2 Delegation regimes with incomplete information

I now turn to the case of incomplete information. The realization of the state variable $\omega \in [0,1]$ is privately observed by the agent before selecting a reform x, while the principal only knows its distribution. Upon observing a proposed reform x, the principal forms beliefs about the underlying state by updating via Bayes' rule. Let $\eta(\cdot \mid x)$ denote the posterior distribution over ω conditional on the observed proposal x.

Suppose first that the principal has chosen to relinquish his veto rights, so that the regime is full delegation. In this case, the agent does not internalize the principal's preferences and holds full information about the state of the world. Therefore, the game is identical to the complete information benchmark, and the results from Observation 3.2 continue to hold.

Suppose now that the principal has chosen veto delegation. Upon receiving a reform proposal x, the principal must decide whether to accept or veto it based on a comparison between the expected utility under the proposed reform and the status quo. The principal accepts the reform if and only if:

$$\mathbb{E}_{\eta(\omega|x)}\left[\beta - |\omega - x|\right] \geq \bar{u}_p. \tag{3.4.9}$$

Equivalently:

$$\int_0^1 |\omega - x| \, d\eta(\omega \mid x) \leq \beta - \bar{u}_p. \tag{3.4.10}$$

Since, by definition, $\underline{x} < 1$, this condition simplifies to:

$$x \leq \beta - \bar{u}_p + \int_0^1 \omega \, d\eta(\omega \mid x). \tag{3.4.11}$$

Hence, the principal's willingness to accept the reform depends not only on the proposal itself but also on his updated beliefs about the underlying state, $\eta(\omega \mid x)$. These beliefs allow the principal to infer the likely state of the world based on the proposal x. This opens the door to strategic signaling: the agent may use her reform choice to credibly convey private information. However, as the next result shows, this signaling incentive introduces a friction that undermines the feasibility

of any accepted reform in equilibrium.

Proposition 3.1. *There exists a unique equilibrium outcome in which the principal vetoes all reform proposals, resulting in the status quo being maintained in every state.*

Proof. See Section 3.B

To understand Proposition 3.1, it is helpful to see why no reform can be sustained in equilibrium under incomplete information and veto powers. Suppose, toward a contradiction, that there exists a state ω such that the agent proposes a reform x that is accepted by the principal. Let x_o denote again the highest proposal satisfying the principal's acceptance constraint. Since proposing any $x' < x_o$ yields lower utility for the agent and is strictly dominated, the agent optimally chooses x_o.

If type-ω proposes x_o in equilibrium, the agent must find x_o weakly preferable to the status quo. That is:

$$\beta - |\omega + b(\omega) - x_o| \geq \bar{u}_a. \tag{3.4.12}$$

Given that $x_o < \omega + b(\omega)$ by assumption, this implies:

$$x_o \geq \bar{u}_a + \omega + b(\omega) - \beta. \tag{3.4.13}$$

Now suppose x_o is proposed in equilibrium by a single type ω_o. Then, by Bayes' rule, the principal updates her belief with certainty and assigns probability 1 to $\omega = \omega_o$ being the unique type proposing x_o. But this perfect inference makes x_o a separating signal, allowing nearby types $\omega' < \omega_o$ to profitably mimic ω_o by proposing x_o. This contradicts the premise that x_o is only offered by a single type. Then, by continuity, this implies that the only possible type that could remain unchallenged is $\omega_o = 0$.

However, this leads to a contradiction. For type $\omega = 0$ to successfully propose x_o, it must be that:

$$\bar{u}_p \leq \beta - |b(0) - x_o|. \tag{3.4.14}$$

Since $x_o < b(0)$, this inequality becomes:

$$\beta - \bar{u}_p \geq \bar{u}_a + b(0) - \beta, \tag{3.4.15}$$

or, equivalently:

$$b(0) \leq 2\beta - \bar{u}_p - \bar{u}_a. \tag{3.4.16}$$

But this contradicts the maintained assumption that $b(0) > 2\beta - \bar{u}_p - \bar{u}_a$. Therefore, type $\omega = 0$ strictly prefers the status quo, and no single type can sustain x_o as an equilibrium proposal.

Suppose now that, in equilibrium, several agent types propose the same reform x_o. Denote $\hat{\omega}$ the highest type among them. In that case, type $\hat{\omega}$ has a profitable deviation available by offering a marginally larger reform, say $x' = x_o + \varepsilon$. Upon observing x', the principal updates her belief and assigns probability 1 to type $\hat{\omega}$ being the deviator, and 0 to all $\omega < \hat{\omega}$. Because posterior beliefs shift in the direction of higher types, the acceptance constraint becomes more permissive:

Suppose instead that multiple types offer x_o in equilibrium. Let $\hat{\omega}$ be the highest such type. The type-$\hat{\omega}$ agent can profitably distinguish itself from the lower-state types by deviating to a reform $x' = x_o + \varepsilon$. Upon observing x', the principal updates her beliefs and assigns probability 1 to $\hat{\omega}$ being the type deviating to x', and probability 0 to all $\omega < \hat{\omega}$. Since x' lies in a neighborhood of x_o:

$$x_o \leq \beta - \bar{u}_p + \int_0^1 \omega, d\eta(\omega \mid x_o) < \beta - \bar{u}_p + \int_0^1 \omega, d\eta(\omega \mid x_o + \varepsilon). \tag{3.4.17}$$

Hence, for ε sufficiently small, $x' = x_o + \varepsilon$ is also accepted, and the deviation is profitable. This shows that pooling at x_o is unsustainable in equilibrium.

The key insight is that under asymmetric information, higher types cannot separate from lower types without attracting profitable mimicry. I formally establish in Section 3.B that the beliefs used in this argument are the only ones consistent with Cho and Kreps (1987) D1 criterion.

In the same appendix, I also show the existence of a separating equilibrium in which no reform is ever implemented. In this equilibrium, each type-ω agent proposes a distinct reform $x(\omega)$ that perfectly reveals the true state to the principal.

However, the principal always rejects $x(\omega)$, so the status quo prevails in every state. The equilibrium is sustained by off-path beliefs assigning probability 1 to $\omega = 0$ after any deviation from equilibrium play. These beliefs imply that any out-of-equilibrium proposal is interpreted as coming from the most biased type, making it unacceptable to the principal. Anticipating this, the agent has no incentive to deviate, and gridlock becomes the only possible outcome.

To summarize, under incomplete information and full delegation, the agent always proposes the extreme reform $\overline{x}$, which is automatically implemented regardless of the principal's preferences. By contrast, under veto delegation, no reform can be sustained in equilibrium: the principal rejects all proposals, and the status quo prevails in every state. Whether the principal is better off under one regime or the other depends on how the payoff from the extreme reform $\overline{x}$ compares to his fallback utility from the status quo. The agent, in turn, strictly prefers full delegation, as it guarantees implementation of her most extreme reform.

3.5 Implications and discussion

The objective of this section is to analyze the results obtained in the previous section and to discuss their implications for the literature on gridlock and delegation.

This chapter examines how extreme polarization shapes the outcomes of delegation when the state of the world is privately observed by the agent. To understand the implications of incomplete information, a natural starting point is to compare the results with complete information.

Under full delegation—where the reform proposed by the agent is automatically enacted—the principal's information about the state does not affect the outcome. Because the agent is always informed and does not internalize the principal's preferences, she always enacts the most extreme reform available in the feasible set, $x = \overline{x}$. Importantly, this choice is independent of the realized state. This has two main implications. The first implication is that the reform proposal conveys no information to the principal about the state of the world, so he only learns the true state once payoffs are realized. The second one is that the payoffs of both

players depend on the realization of ω. Even though the agent's bias decreases in ω, this effect is dominated by the direct effect of the state, so the principal's utility decreases in ω for a given reform. The logic is that, for a given $\overline{x}$, larger realizations of the state increase the distance between the enacted reform and the ideal one. This distance can be interpreted as the *extreme preferences gap*: the more extreme the agent is, the larger the distance between her favorite reform and the reform set. The principal's payoff is not affected by the size of the agent's drift, since the agent's choice is fixed at $\overline{x}$ regardless of ω. As the ideal reform of the principal lies at the other extreme of the reform set, the intuition of his utility is the opposite as for the agent: the larger the upper bound of the reform set, the lower the utility of the principal, which can only be partially compensated with increases in the state.

Under veto delegation, incomplete information has profound consequences for the outcomes of delegation. In this regime, the agent must internalize the preferences of the principal through his approval constraint. With complete information, she proposes the largest reform that the principal is willing to accept, provided that doing so gives her at least as much utility as the status quo. Importantly, this reform is tailored to the realized state, generating separating equilibria. When a reform is implemented, the utility of each player is equal to their respective status quo utility if:

$$b(\omega) = 2\beta - \bar{u}_a - \bar{u}_p \tag{3.5.1}$$

Whenever the agent's bias is larger than this threshold, she appropriates the entire surplus created by the reform, while the principal remains exactly indifferent. Thus, under veto delegation, the principal always receives $u_p = \overline{u}_p$, regardless of whether or not a reform is implemented, while the agent receives at least $u_a = \overline{u}_a$.

Under incomplete information, the only possible outcome is gridlock. The first implication of this result is summarized in the following observation.

Observation 3.3. ***Normative implication of Proposition 3.1:***
Asymmetric information and signaling frictions lead to inefficient gridlock in some states. Specifically, in states ω such that

$$b(\omega) < 2\beta - \overline{u}_a - \overline{u}_p, \tag{3.5.2}$$

the outcome of the game would strictly improve if a reform were implemented. In these states, there exists a feasible reform

$$x \in \big(\omega + \beta - \overline{u}_p,\, \omega + b(\omega) + \overline{u}_a - \beta\big) \tag{3.5.3}$$

that constitutes a Pareto-improvement over the status quo.

The interpretation of this result is straightforward. In the states identified above, both players could benefit from reform relative to the status quo. An important observation is that the gains from reform are not shared: the principal always remains indifferent, obtaining exactly $\overline{u}_p$, while the agent captures the entire improvement in utility.

A second implication is that the payoffs of the principal are not affected by his information about the state of the world: he always receives $u_p = \overline{u}_p$, regardless of whether he observes ω or not. Whereas moving from complete to incomplete information leaves the agent weakly worse off, the principal remains indifferent. One might be tempted to see this as a positive result—since lack of information does not make the principal worse off. But it has an important consequence: gridlock. Under incomplete information, no reform is ever implemented, which undermines the very purpose of delegation. Holding veto rights under extreme polarization makes gridlock the inevitable outcome.

Before turning to the discussion on gridlock, it is useful to consider which delegation regime the principal prefers under incomplete information. The principal decides the regime before the realization of the state, so he evaluates the expected payoff of each regime. Under full delegation, his expected utility is:

$$\mathbb{E}[u_p^{fd}] \;=\; \int_0^1 \big[\beta - |\omega - \overline{x}|\big] f(\omega)\, d\omega \;=\; \beta - \overline{x} + \int_0^1 \omega f(\omega)\, d\omega, \tag{3.5.4}$$

Under veto delegation, his expected utility is simply $\mathbb{E}[u_p^{fd}] = \overline{u}_p$. Therefore, the principal prefers full delegation to veto delegation if:

$$\beta - \overline{u}_p > \overline{x} - \int_0^1 \omega f(\omega)\, d\omega \tag{3.5.5}$$

Comparative statics follow from this condition. First, an increase in the net benefit from reform, $\beta - \overline{u}_p$, makes full delegation more attractive. Second, the larger the upper bound of the reform space, the greater distance between the principal's ideal reform and the implemented one, making delegation with veto rights more likely to be preferred. Finally, if the distribution f places more probability mass on large values of ω, the expected distance between the principal's ideal reform and the implemented one becomes shorter, making full delegation relatively more attractive to the principal.

Two observations follow. First, the principal's choice does not change with information: his payoffs are exactly the same whether or not he observes the state of the world. Second, veto rights can be interpreted as a form of "insurance": while they completely block the possibility of reform, they guarantee that the principal is never worse off than under the status quo, shielding him from reforms that could lower his payoff.

This discussion also highlights the role of extreme polarization. Because the ideal policies of both players lie outside the feasible reform set, the relevant trade-off is not between reaching one's ideal point, but between accepting the best feasible reform or remaining with the status quo. In this context, the net benefit from reform determines whether a reform can dominate the fallback option. When this net benefit is small, the insurance value of veto dominates, and gridlock becomes the better choice. When it is large, even extreme reforms may deliver higher payoffs than the status quo, making full delegation more attractive

3.5.1 On gridlock and delegation

The results in this chapter yield an important observation: veto delegation under extreme polarization of preferences and incomplete information results in complete reform gridlock. However, it is important to note that veto delegation and extreme polarization alone do not produce gridlock; it is the combination of the two with incomplete information that makes gridlock the inevitable outcome of delegation.

Policy gridlock—also referred to as legislative stalemate, or policy deadlock—

is defined as the inability of political systems to adapt to new circumstances or to update existing policies.[26] Although a minority view argues that under certain conditions gridlock may improve the quality of legislative output,[27] the prevailing perspective emphasizes its negative consequences. A large body of research documents that gridlock reduces the likelihood of policy change, lowers legislative productivity, delays adaptation to new economic and social conditions, and undermines democratic responsiveness (Binder 1999, 2004; Brady and Volden 2006; Chiou and Rothenberg 2008; Krehbiel 2010; D. R. Mayhew 2005; Tsebelis 1999).

As gridlock is a common feature of contemporary policymaking with significant negative consequences, a large literature has aimed at understanding its causes. Explanations point to a variety of factors, including divided government, multiple veto players, inter-branch and inter-chamber conflict, institutional design, actor preferences and polarization, the influence of interest groups and coalitions, as well as electoral incentives and shifts in public mood.[28]

Polarization has been widely identified as a central cause of gridlock across political science and political economy literatures. Classic studies link partisan polarization to declining legislative productivity and bipartisan agreement in the United States Congress (Jones 2001; Binder 2004). More recent work argues that polarization by itself can be sufficient to produce legislative stalemate. Barber and McCarty (2015) and Carmines and Fowler (2017) show that rising partisan distance entrenches stalemate and weakens governing capacity, while McCarty (2018) demonstrates that polarization alone widens the "gridlock interval" and generates policy drift and uncertainty. Other contributions measure the expansion of the interval under rising polarization (Chiou and Rothenberg 2003), examine its impact on congressional partisanship and productivity in state legislatures

[26] See Krehbiel 1996; Chiou and Rothenberg 2003; Beckmann and McGann 2008; Saeki 2009; Bowling and Ferguson 2001; Jones 2001.

[27] See Bradbury and Crain 2001; Congleton 2003; Lupia and McCubbins 1994; Riker 1992; Rogers 2001.

[28] The literature analyzed here focuses on polarization and veto players. Other explanations of gridlock emphasize divided government (D. Mayhew 1991; Coleman 1999; Saeki 2009); inter-branch and inter-chamber conflict (Peterson and Greene 1994; Taylor 1998; Binder 1999), interest groups and coalitions (Dodge and Lee 2017); and electoral incentives and shifts in public mood (Schattschneider 1942).

(Gelman and Wilson 2022; Hicks 2015), or highlight its role in committee conflict and representation (Calvo and Sagarzazu 2016; Kang 2017). Agarwal (2024) develops this sufficiency perspective further, arguing that polarization on its own explains gridlock through mechanisms such as social identity dynamics, elite manipulation, media and social media echo chambers, identity politics, and populist governments.[29]

Another main driver of gridlock is veto players. Classic veto player theory shows that additional veto points—such as bicameralism, presidential vetoes, or supermajority requirements—expand the region in which the status quo cannot be displaced, reducing the likelihood of policy change (Tsebelis 1995, 2002; Tsebelis 1995, 1999; Krehbiel 2010). Empirical studies confirm that multiple veto players lower legislative productivity (Binder 1999; Brady and Volden 1998; Crombez and Hix 2015). Saeki (2009) shows that the ideological configuration of veto pivots, rather than the mere presence of divided government, determines the size of the winset, and hence the probability of policy change.

Veto power has also been analyzed through the lens of bargaining: Cameron (2000) and Cameron and McCarty (2004) highlight that incomplete information and the strategic use of presidential vetoes and veto threats can delay policymaking, induce concessions, or generate repeated veto chains ending in stalemate. Consistent with the results of this chapter, they show that under complete information vetoes rarely occur, whereas incomplete information generates them. Building on this insight, subsequent models show how informational asymmetries can sustain bargaining failure and legislative deadlock. Matthews (1989) models veto threats as cheap talk; McCarty (1997) highlights reputational incentives across issues; and Groseclose and McCarty (2001) develop a blame-game framework in which electoral signaling, combined with voter uncertainty, produces vetoes and policy stalemate. More recent formal work demonstrates that incomplete information can severely limit proposers and generate inefficient, deadlock-like outcomes in sequential veto bargaining (Ali et al. 2023), while also showing that appropriately designed delegation can sometimes mitigate these inefficiencies (Kartik et al. 2021).

[29] See Agarwal (2024) for an extensive discussion of theoretical frameworks linking polarization to gridlock.

The study of polarization and gridlock on delegation has focused on separation of powers and bureaucratic agencies. There is consensus that polarization and gridlock have increased the power of the executive and bureaucratic agencies, undermining Congress (Binder 2004; Barber and McCarty 2015; Carmines and Fowler 2017). Presidents and bureaucrats exploit this weakness by acting unilaterally (Howell 2015) and resisting legislative oversight (Potter 2019). A. Bolton (2022) shows that Congress adapts to this environment by expanding its reliance on nonstatutory tools—most notably appropriations committee reports—which guide agency behavior through repeated interactions, reputational sanctions, and credible budgetary threats. Farhang (2021) similarly argues that Congress adjusts its delegation strategies within statutes: it enacts fewer but more detailed laws, constrains agencies through deadlines and reporting requirements (Epstein and O'Halloran 1999; Lewis 2004), fragments delegation across multiple jurisdictions to reduce presidential control (Farhang and Yaver 2016; Freeman and Rossi 2012), and supplements administrative enforcement with private rights of action (Farhang 2018; Burbank and Farhang 2021).

This literature highlights how polarization and gridlock reshape delegation outside of veto-based frameworks, while this chapter focuses on how extreme polarization and incomplete information collapse veto delegation into gridlock. Previous veto-delegation literature, introduced in Section 3.2, has emphasized the trade-off between bias, information and the choice of the delegation regime.

Starting with the literature on legislative bargaining, a main component in the analysis of veto delegation has been the policy drift or bias between the principal and the agent. Gilligan and Krehbiel (1987, 1989) and Krishna and Morgan (2001) show that closed rules may outperform open rules only when drift is sufficiently small, while Baron (2000), using a mechanism design approach, shows that institutional design can overturn the negative effects of drift and that veto delegation can dominate open rules. In the delegation literature in firms, Dessein (2002) shows that full delegation outperforms veto delegation when the agent's bias is sufficiently small, and that veto delegation only becomes attractive once bias is large. In all these contributions, drift is the central mechanism. My contribution pushes this logic further by introducing conditions of extreme polarization with

unattainable ideal policies. In this context, drift becomes so severe that no reform can be sustained in equilibrium.

Mylovanov (2008) takes a different approach by emphasizing the choice of the reversion point. His "veto-power principle" shows that if the default can be set endogenously, veto delegation can replicate the outcomes of optimal delegation, cheap talk, or interval delegation. This result highlights how the welfare ranking between veto and delegation depends crucially on the location of the status quo. In earlier models such as Dessein (2002) and Marino (2007), the reversion point is fixed and exogenous, which drives their conclusions. In my setting, the default option is the status quo, exogenously given and state-independent. Unlike in those papers, the signaling spiral and resulting gridlock in my model does not depend on the magnitude of the reversion point, as long as there exist reforms within the policy space that generate sufficient benefit relative to this fallback. This assumption is mild and consistent with policy environments.

Finally, this chapter is closely related to Marino (2007), who compares full delegation and veto delegation with risk-neutral players and an upward, state-independent bias. His results show that under full delegation the agent implements projects at a higher scale than the principal would prefer, which yields negative payoffs for the principal when the state is low. Under veto delegation, and using the Intuitive Criterion as a refinement concept, equilibrium features three regions: (1) agents in low states refrain from proposing any reform, anticipating rejection and leading to the status quo; (2) agents in intermediate states pool and propose the same reform, equal to the highest proposal among them; and (3) agents in high states separate and propose truthfully. Because agents in low and intermediate states have incentives to mimic lower types, full separation is infeasible, although full gridlock is never sustained in equilibrium. Marino concludes that veto delegation dominates full delegation under a broad range of parameters, particularly when the distribution of states places significant weight on low types. In my chapter, by contrast, the only outcome under veto delegation is gridlock. Whether the principal prefers veto delegation to full delegation depends not only on the distribution of states but also on the upper bound of the policy space and the net benefit of implementing a reform relative to the status quo.

3.6 Concluding remarks

The objective of this chapter has been to analyze the effect of extreme polarization of preferences on the outcomes of veto-based delegation. Delegation is a central mechanism across all domains of society, from politics to private organizations. Veto delegation is especially pervasive: it appears as the closed rule in legislative bargaining; in bureaucratic settings, where elected politicians delegate reform-making authority while retaining the right to overturn agency decisions; and in private organizations and firms, where, for example, boards of directors retain the right to ratify executive decisions. Polarization has become a defining feature of modern politics, both among political elites and across individuals at all levels of society. This observation has generated a large literature on the causes and consequences of polarization, as well as theories predicting that the severe polarization already present in some spheres of society is likely to intensify, thereby exacerbating its negative effects on politics and beyond.

As the consequences of extreme polarization for delegation remain understudied, this chapter develops a sequential principal–agent model with asymmetric information and extremely polarized players. The agent privately observes the state of the world and decides whether to propose a reform. The principal, who knows only the distribution of the state, retains the power to veto the reform proposed by the agent if he has chosen veto delegation, whereas the reform is implemented automatically if he has chosen full delegation. A distinctive feature of the model is that the ideal policies of both players lie outside the policy space, on opposite sides, representing extreme polarization.

The main results of the chapter establish the outcomes of delegation under both full delegation and veto delegation. Under full delegation, the agent always chooses the upper bound of the policy space, as it is the closest to her ideal policy. Under veto delegation, by contrast, the only possible outcome of the delegation game is gridlock. Comparing the two regimes, the principal's decision on the regime depends on whether the net benefit from implementing a reform is larger than the distance between his ideal policy and the upper bound of the policy space. The more the principal is polarized toward lower reforms, the less attractive it becomes

for him the reform under full delegation. Conversely, larger realizations of the state make full delegation more attractive. In choosing the delegation regime, the principal therefore faces a trade-off: retaining veto rights provides insurance against harmful reformsbut at the cost of blocking any possibility of reform—even when such a reform would constitute a Pareto-improvement. The central contribution of this chapter is therefore to show that under extreme polarization and incomplete information, veto delegation collapses into inefficient gridlock.

Gridlock is a common feature of contemporary policymaking, and a large literature has analyzed its causes and consequences. There is broad consensus that polarization is a central driver of gridlock. Although some authors argue that polarization alone is sufficient to generate legislative stalemate (Barber and McCarty 2015; Carmines and Fowler 2017; Agarwal 2024), much of the literature considers gridlock as the product of multiple forces operating simultaneously. Among these, veto players are another main driver of stalemate. In this sense, the results of this chapter speak directly to this logic, showing that gridlock is the consequence of a combination of extreme polarization, veto power, and incomplete information. Importantly, without incomplete information, extreme polarization and veto delegation alone do not produce reform gridlock.

This chapter also contributes to the delegation literature by pushing the logic of policy drift to its extreme. Previous research has debated whether drift makes veto delegation superior or inferior to other regimes, producing contradictory results depending on the distribution of states, the location of the reversion point, or the refinement concept employed. By contrast, this chapter demonstrates that when polarization places both players' ideal points outside the feasible set, drift becomes so severe that no reform can be sustained in equilibrium.

An important direction for future research is to extend the analysis of extreme polarization and delegation to other delegation mechanisms. Previous studies on polarization and delegation have focused on separation of powers and bureaucratic agencies. In this context, A. Bolton (2022) and Farhang (2021) show that Congress reacts to polarization and gridlock by relying more heavily on nonstatutory tools to constrain agency behavior, and by drafting statutes that are denser, more detailed, and fragmented across agencies. These findings suggest that a

natural response to polarization and gridlock is the design of delegation regimes that constrain the set of actions available to the agent. In principal–agent models, this takes the form of optimal delegation, where the principal limits the agent's discretion by establishing a delegation set or discretion window, but without holding veto power over the agent's choice. An extensive literature on optimal delegation shows that the larger the policy drift between principal and agent, the more constrained the delegation set becomes. A natural extension of this chapter is therefore to introduce optimal delegation into the analysis: is optimal delegation capable of sustaining reform under extreme polarization of preferences?

Appendix

3.A Reasonable beliefs and criterion D1

To eliminate PBE that rely on implausible off-path beliefs, I restrict attention to those that satisfy the D1 refinement introduced by Cho and Kreps (1987).

Fix an equilibrium, and let $G^*(\omega, x)$ denote the payoff for a type-ω agent under the equilibrium, when she proposes reform x. According to the D1 criterion, I ask: which agent types ω could plausibly be thought to choose an off-the-equilibrium reform x'?

Let $MBR(\eta, x')$ denote the set of mixed best responses by the principal to a reform proposal x', when he holds belief η over the agent's type ω.

Let $D_\eta(\omega, x')$ be the set of mixed responses $d \in MBR(\eta, x')$ by the principal that make the agent of type ω strictly prefer deviating to the off-path proposal x', rather than following her equilibrium strategy. That is, under belief η, the expected utility from proposing x' and triggering response d exceeds the equilibrium utility $G^*(\omega, x)$.

Let:

$$D(\omega, x') = \bigcup_{\eta} D_\eta(\omega, x'), \tag{3.A.1}$$

denote the set of all principal responses—across all belief systems η—that would make a type-ω agent strictly prefer deviating to x'. This set represents all conditions under which a type-ω agent finds it profitable to deviate.

Let $D^o(\omega, x')$ denote the set of mixed strategies under which a type-ω agent is indifferent between deviating to x' and adhering to her equilibrium action. Formally, these are the strategies d such that the expected utility from x' matches the equilibrium payoff $G^*(\omega, x)$.

Under the D1 criterion, the agent type ω is eliminated as a plausible deviator to x' if there exists another type $\omega' \neq \omega$ such that:

$$[D(\omega, x') \cup D^o(\omega, x')] \subset D(\omega', x'). \tag{3.A.2}$$

In words, if the set of principal responses that could make type ω willing or indifferent to deviate is strictly smaller than the set that makes type ω' strictly prefer the deviation, then the principal should place all belief on ω' being the true type behind the off-equilibrium proposal x'.

3.B Proof of Proposition 3.1

Roadmap. I divide the proof of Proposition 3.1 into two parts. In the first part, using Lemma 3.1 I prove—via two claims—that there exists an equilibrium in which the principal always vetoes the agent's reform proposal, and that this equilibrium satisfies Criterion D1. In the second part, I use Lemmas 3.2 and 3.3 to show that no reform can be sustained in equilibrium. Specifically, I first show that if a reform x were successfully implemented in equilibrium, it would be proposed by a unique type ω, and then prove that even such proposals are not sustainable.

Lemma 3.1. *There exists an equilibrium in which the principal always vetoes the agent's proposal.*

Proof. The proof proceeds in two steps: first, I show that such an equilibrium exists; second, I verify that it satisfies Criterion D1.

Claim 3.1. *An equilibrium in which the principal vetoes every reform x proposed by the agent exists.*

Let $k > 0$ be an arbitrarily small constant. Define the strategy of a type-ω agent as proposing a reform $x(\omega) = \omega + \beta - \bar{u}_p + k$. That is, each agent proposes a reform that fully reveals the state and lies strictly above the maximum reform the principal is willing to accept. The strategy of the principal is to veto any reform proposed by the agent.

This pair of strategies constitutes a Perfect Bayesian Equilibrium when combined with a belief function $\eta(\omega \mid x)$ such that: Upon observing an on-path proposal $x(\omega)$, the principal assigns probability 1 to the corresponding state ω. Upon observing any off-path reform $x' \notin x([0,1])$, the principal assigns probability 1 to state $\omega = 0$.

On the equilibrium path, the principal's strategy to veto any reform is optimal. A reform is accepted if and only if $x \leq \omega + \beta - \bar{u}_p$, but under the strategy $x(\omega) = \omega + \beta - \bar{u}_p + k$, this condition is violated for all ω. Thus, the principal prefers to veto.

Consider first a deviation by the agent to $x' > x(\omega)$. Since x' lies even further above the acceptance threshold, the principal vetoes it. The agent receives the status quo payoff $\bar{u}_a$, which is identical to her equilibrium payoff. The deviation is not profitable for the agent.

Now consider a deviation to a reform $x' < x(\omega)$. Under the defined belief system $\eta(\omega, x)$, the principal attaches probability 1 to a type-$\omega = 0$ agent being the one deviating upon observing x'. However, the principal does not accept the reform as, by assumption, $b(0) > 2\beta - \bar{u}_p - \bar{u}_a$: the principal strictly prefers the status quo when $\omega = 0$. Again, the agent receives $\bar{u}_a$ and the deviation is not profitable.

Hence, no profitable deviation exists and the principal's veto strategy is optimal both on and off the equilibrium path. $\diamondsuit$

Claim 3.2. *The equilibrium in which the principal always vetoes the reform satisfies criterion D1.*

Take an arbitrary type ω and let x' be an off-equilibrium reform proposal. Let α denote the principal's mixed strategy best response to x', such that the agent of

type ω weakly prefers proposing x' over her equilibrium proposal $x(\omega)$. That is:

$$\alpha\left(\beta - |\omega + b(\omega) - x'|\right) + (1-\alpha)\bar{u}_a \geq \bar{u}_a.$$

Under the assumption that $x' < \omega + b(\omega)$, this simplifies to:

$$\alpha\left(\beta - (\omega + b(\omega) - x')\right) + (1-\alpha)\bar{u}_a \geq \bar{u}_a,$$

which reduces to:

$$\alpha\left(\beta + x' - \omega - b(\omega) - \bar{u}_a\right) \geq 0. \tag{3.B.1}$$

Let:

$$D(\omega, x') = \left\{\alpha \in [0,1] : \alpha\left(\beta + x' - \omega - b(\omega) - \bar{u}_a\right) \geq 0\right\}$$

denote the set of mixed strategies that make the agent of type ω weakly prefer x' over $x(\omega)$. To apply Criterion D1, I compare these sets across types.

Under the belief system described earlier, the principal assigns probability 1 to type $\omega = 0$ upon observing any off-path reform x'. According to D1, it suffices to show:

$$[D(0, x') \cup D^o(0, x')] \supseteq D(\omega, x') \quad \text{for all } \omega \in (0, 1]. \tag{3.B.2}$$

Suppose first that $\beta + x' < \omega + b(\omega) + \bar{u}_a$. In this case, the left-hand side of the inequality defining $D(\omega, x')$ is negative for all $\alpha \in (0, 1]$, and therefore $D(\omega, x') = \emptyset$. Since the empty set is trivially a subset of any other set, the D1 condition is satisfied.

Contrary, suppose $\beta + x' > \omega + b(\omega) + \bar{u}_p$. Then the set $D(\omega, x')$ includes all $\alpha \in [0, 1]$ for which the inequality holds. Note that $\omega + b(\omega)$ is strictly increasing in ω, since $b'(\omega) \in (-1, 0)$. Hence, for any $\omega < 1$, the expression $\beta + x' - \omega - b(\omega) - \bar{u}_a$ is strictly greater than the same expression evaluated at $\omega = 1$. This implies:

$$D(1, x') \subset D(\omega, x') \quad \text{for all } \omega < 1.$$

Consequently:

$$D(\omega, x') \subset D(0, x') \cup D^o(0, x') \quad \text{for all } \omega.$$

This establishes that types $\omega > 0$ are less likely to deviate to x' than type $\omega = 0$, which is the only type not eliminated under D1. Thus, the belief system $\eta(\omega, x)$ assigning probability 1 to $\omega = 0$ upon observing any off-path x' is consistent with Criterion D1. $\diamondsuit$

Together with Claim 3.1, this completes the proof of Lemma 3.1. $\square$

It remains to show that no reform can be successfully implemented in equilibrium. Suppose, to the contrary, that there exists a reform that is successfully accepted in equilibrium.

Lemma 3.2. *If a reform is successful in equilibrium, then it must emanate from a single type.*

Proof. I divide the proof of Lemma 3.2 into three claims.

Claim 3.3. *In each state of the world ω, only one reform x can be successfully implemented in equilibrium.*

Suppose, for contradiction, that there exist two different reforms x_1 and x_2, both successfully implemented in equilibrium, corresponding to two distinct states of the world $\omega_1 < \omega_2$, with $x_1 < x_2$.

Consider the agent of type ω_1. Since x_1 is implemented in equilibrium, this agent must find it optimal to propose x_1. This implies:

$$\beta - |\omega_1 + b(\omega_1) - x_1| \geq \bar{u}_a.$$

Given the assumption that $x_1 < \omega_1 + b(\omega_1)$, this simplifies to:

$$\beta - \omega_1 - b(\omega_1) + x_1 \geq \bar{u}_a.$$

Now consider a deviation by the agent of type ω_1 to the alternative reform x_2. The resulting utility would be:

$$\beta - \omega_1 - b(\omega_1) + x_2.$$

Since $x_2 > x_1$, it follows that:

$$\beta - \omega_1 - b(\omega_1) + x_2 > \beta - \omega_1 - b(\omega_1) + x_1,$$

Therefore, the type-ω_1 agent strictly prefers to deviate to x_2, contradicting the assumption that x_1 is an optimal equilibrium proposal. Hence, in each state of the world ω, only one reform x can be successfully implemented in equilibrium. $\diamondsuit$

Let $\Omega_o \equiv \{\omega \in [0,1] : \sigma_a(\omega) = x_o\}$, and let $\hat{\omega}$ be the supremum of Ω_o.

Suppose now that agents from multiple states of the world ω successfully propose a reform, say x_o, in equilibrium. From Claim 3.3, all other equilibrium offers must be unsuccessful.

Let $\Omega_o \equiv \{\omega \in [0,1]; \sigma_p(\omega) = x_o\}$, and let $\hat{\omega}$ be the supremum of Ω_o.

Claim 3.4. $\Omega_o = [0, \hat{\omega}]$.

Suppose, for contradiction, that there exists a state $\omega < \hat{\omega}$ such that the agent proposes a reform $x \neq x_o$. By the definition of a Perfect Bayesian Equilibrium, if type $\hat{\omega}$ offers x_o in equilibrium, it must be that:

$$\beta - |\hat{x}_a(\hat{\omega}) - x_o| \geq \bar{u}_a$$

Using the assumption that $\hat{x}_a(\omega) = \omega + b(\omega)$:

$$x_o \geq \bar{u}_a + \hat{\omega} + b(\hat{\omega}) - \beta$$

Since the right-hand side of the inequality increases with ω, this implies:

$$x_o \geq \bar{u}_a + \hat{\omega} + b(\hat{\omega}) - \beta > \bar{u}_a + \omega + b(\omega) - \beta$$

Thus, the agent in state ω would strictly prefer to deviate to x_o, which contradicts the equilibrium assumption. This establishes that $\Omega_o = [0, \hat{\omega})$.

To confirm $\hat{\omega} \in \Omega_o$, suppose it is not. Then, by the definition of equilibrium, it must be that:

$$\beta - |\hat{x}_a(\hat{\omega}) - x_o| < \bar{u}_a.$$

By continuity of ω and $b(\omega)$, this implies that:

$$x_o = \bar{u}_a + \hat{\omega} + b(\hat{\omega}) - \beta,$$

otherwise, for sufficiently small $\varepsilon > 0$, it would be the case that:

$$x_o > \bar{u}_a + (\hat{\omega} - \varepsilon) + b(\hat{\omega} - \varepsilon) - \beta,$$

which would contradict the fact that type $\hat{\omega} - \varepsilon \in \Omega_o$ proposes x_o.

As $\hat{\omega} \notin \Omega_o$, this contradicts the assumption that indifferent agents strictly prefer proposing a reform over maintaining the status quo. This completes the proof of Claim 3.4. $\diamondsuit$

Claim 3.5. *For every $x \in (x_o, \omega + \beta - \bar{u}_p)$ and every $\omega < \hat{\omega}$, $[D(\omega, x) \cup D^o(\omega, x)] \subset D(\hat{\omega}, x)$.*
Take an arbitrary $\omega \in \Omega_o \setminus \{\hat{\omega}\}$, so that $\omega < \hat{\omega}$, and suppose that the agent of type ω deviates from its equilibrium proposal x_o to some reform $x \in (x_o, \hat{\omega} + \beta - \bar{u}_p)$.

Let $\alpha \in [0, 1]$ denote the mixed strategy best response of the principal that makes the agent of type ω weakly prefer deviating to x over sticking to x_o:

$$\alpha\,(\beta - |\omega + b(\omega) - x|) + (1 - \alpha)\bar{u}_a \geq \beta - |\omega + b(\omega) - x_o|.$$

Under the assumption $x < \omega + b(\omega)$, this inequality becomes:

$$\alpha\,(\beta - \omega - b(\omega) + x) + (1 - \alpha)\bar{u}_a \geq \beta - \omega - b(\omega) + x_o.$$

Subtracting $\bar{u}_a$ from both sides and rearranging yields:

$$\alpha\,(\beta - \omega - b(\omega) + x - \bar{u}_a) \geq \beta - \omega - b(\omega) + x_o - \bar{u}_a.$$

Solving for α, the threshold value is:

$$\alpha \geq \frac{\beta - \omega - b(\omega) + x_o - \bar{u}_a}{\beta - \omega - b(\omega) + x - \bar{u}_a} \equiv \alpha(\omega).$$

Define $\overline{\alpha}(a) = \frac{a+x_o}{a+x}$, which is strictly increasing in a as long as $x_o < x$. Let $a(\omega) \equiv \beta - \omega - b(\omega) - \bar{u}_a$, so that $\alpha(\omega) = \overline{\alpha}(a(\omega))$. Since $b(\cdot)$ is strictly decreasing with $b'(\omega) \in (-1, 0)$, it follows that $a(\omega)$ is strictly decreasing in ω. Hence, as $\omega < \hat{\omega}$, then $a(\omega) > a(\hat{\omega})$, implying:

$$\alpha(\omega) = \overline{\alpha}(a(\omega)) > \overline{\alpha}(a(\hat{\omega})) = \alpha(\hat{\omega}). \tag{3.B.3}$$

Hence, the set of principal responses that support a profitable deviation is strictly larger for type $\hat{\omega}$ than for type ω, which proves:

$$[D(\omega, x) \cup D^o(\omega, x)] \subset D(\hat{\omega}, x).$$

The inclusion is strict. To see this, observe that any $\alpha \in (\alpha(\hat{\omega}), \alpha(\omega))$ is a mixed response that satisfies the incentive constraint for $\hat{\omega}$ but not for ω. Therefore:

$$D(\hat{\omega}, x) \setminus [D(\omega, x) \cup D^o(\omega, x)] \neq \emptyset.$$

This completes the proof of Claim 3.5. $\diamondsuit$

Now, consider a deviation by the agent of type $\hat{\omega}$ to some reform $x \in (x_o, \hat{\omega} + \beta - \bar{u}_p)$, which is a nonempty interval by Claim 3.4. According to Claim 3.5 and the D1 criterion, the principal assigns probability 1 to the deviation coming from a type in the interval $[\hat{\omega}, 1]$, and accepts the reform.

Given this belief update, the deviation is profitable. The agent of type $\hat{\omega}$ strictly prefers x to x_o because $x > x_o$ and the utility function is increasing in x when the proposal lies below the agent's ideal point. Therefore, type $\hat{\omega}$ prefers the deviating reform to the equilibrium one, contradicting the assumption that x_o is successfully implemented in equilibrium.

This contradiction implies that no reform can be successfully implemented by multiple types in equilibrium. $\diamondsuit$

Together with Claim 3.3 and Claim 3.4, this completes the proof of Lemma 3.2. $\square$

Lemma 3.3. *In any equilibrium, a reform is never a possible outcome.*

Proof. Suppose, for contradiction, that there exists a unique state of the world $\omega_o \in [0,1]$ such that the agent in state ω_o proposes a reform x_o, and that this reform is accepted in equilibrium. By Claim 3.3, x_o must be the only proposal that is successfully implemented.

Bayesian updating then implies that the proposal x_o fully reveals the state of the world to the principal.

By Lemma 3.2, such an equilibrium can only exist if no neighboring type $\omega' \neq \omega_o$ finds it profitable to deviate to x_o. This is only possible if $\omega_o = 0$, i.e., the lowest possible type. Otherwise, by continuity of the utility function and the mapping from ω to the agent's strategy, a nearby type $\omega' < \omega_o$ (e.g., $\omega_o - \varepsilon$ for $\varepsilon > 0$ sufficiently small) would strictly prefer to mimic type ω_o and offer x_o, contradicting uniqueness. Therefore, it must be:

$$\omega_o = 0.$$

For the agent of type $\omega = 0$ to prefer x_o to the status quo, it must be that:

$$\beta - |0 + b(0) - x_o| \geq \bar{u}_a,$$

which implies:

$$x_o \geq \bar{u}_a + b(0) - \beta.$$

However, for the principal to accept x_o, it must also be that:

$$\beta - |0 - x_o| \geq \bar{u}_p,$$

which implies:

$$x_o \leq \beta - \bar{u}_p.$$

Combining these two inequalities:

$$\bar{u}_a + b(0) - \beta \leq x_o \leq \beta - \bar{u}_p,$$

which in turn implies:

$$b(0) \leq 2\beta - \bar{u}_a - \bar{u}_p.$$

But this contradicts the maintained assumption that $b(0) > 2\beta - \bar{u}_a - \bar{u}_p$. Hence, the type-0 agent strictly prefers the status quo and would never propose a reform. This contradiction implies that no reform can ever be implemented in equilibrium. □

I obtain Proposition 3.1 combining Lemmas 3.1 to 3.3 ■

Bibliography

Abramowitz, A. I. (2010). The disappearing center: Engaged citizens, polarization and aerican democracy. *Polarization, and*, *220*.

Abramowitz, A. I., & Saunders, K. L. (2008). Is polarization a myth? *The Journal of politics*, *70*(2), 542–555.

Acemoglu, D., Egorov, G., & Sonin, K. (2012). Dynamics and stability of constitutions, coalitions, and clubs. *American Economic Review*, *102*(4), 1446–1476.

Agarwal, K. (2024). How does political polarization impact legislative gridlock and policy-making processes. *IOSR Journal of Humanities and Social Science*, *29*(10), 53–64.

Aidt, T. S., Albornoz, F., & Hauk, E. (2021). Foreign influence and domestic policy. *Journal of Economic Literature*, *59*(2), 426–487.

Aldrich, J. H., & Rohde, D. W. (2000). The republican revolution and the house appropriations committee. *The Journal of Politics*, *62*(1), 1–33.

Aldrich, J. H., Sullivan, J. L., & Borgida, E. (1989). Foreign affairs and issue voting: Do presidential candidates "waltz before a blind audience?" *American Political Science Review*, *83*(1), 123–141.

Ali, S. N., Kartik, N., & Kleiner, A. (2023). Sequential veto bargaining with incomplete information. *Econometrica*, *91*(4), 1527–1562.

Alonso, R., & Matouschek, N. (2007). Relational delegation. *The RAND Journal of Economics*, *38*(4), 1070–1089.

Alonso, R., & Matouschek, N. (2008). Optimal delegation. *The Review of Economic Studies*, *75*(1), 259–293.

Amador, M., & Bagwell, K. (2013). The theory of optimal delegation with an application to tariff caps. *Econometrica*, *81*(4), 1541–1599.

Amador, M., & Bagwell, K. (2022). Regulating a monopolist with uncertain costs without transfers. *Theoretical Economics*, *17*(4), 1719–1760.

Amador, M., Bagwell, K., & Carpizo, L. (2025). *Generalized delegation* [Working Paper].

Amador, M., Bagwell, K., & Frankel, A. (2018). A note on interval delegation. *Economic Theory Bulletin*, *6*(2), 239–249.

Ambrus, A., Azevedo, E. M., Kamada, Y., & Takagi, Y. (2013). Legislative committees as information intermediaries: A unified theory of committee selection and amendment rules. *Journal of Economic Behavior & Organization*, *94*, 103–115.

Ambrus, A., & Egorov, G. (2017). Delegation and nonmonetary incentives. *Journal of Economic Theory*, *171*, 101–135.

Andersen, T. B., Harr, T., & Tarp, F. (2006). On us politics and imf lending. *European Economic Review*, *50*(7), 1843–1862.

Anesi, V. (2010). Noncooperative foundations of stable sets in voting games. *Games and Economic Behavior*, *70*(2), 488–493.

Anesi, V. (2018). *Dynamic legislative policy making under adverse selection* (CeDEx Discussion Paper No. 2018-08). Centre for Decision Research and Experimental Economics (CeDEx), University of Nottingham.

Anesi, V., & Duggan, J. (2017). Dynamic bargaining and stability with veto players. *Games and Economic Behavior*, *103*, 30–40.

Anesi, V., & Duggan, J. (2018). Existence and indeterminacy of markovian equilibria in dynamic bargaining games. *Theoretical Economics*, *13*(2), 505–525.

Anesi, V., & Seidmann, D. J. (2015). Bargaining in standing committees with an endogenous default. *Review of Economic Studies*, *82*(3), 825–867.

Ansolabehere, S., De Figueiredo, J. M., & Snyder, J. M. (2003). Why is there so little money in us politics? *Journal of Economic Perspectives*, *17*(1), 105–130.

Arrow, K. J. (1951). *Social choice and individual values* (Vol. 1). Wiley New York.

Austen-Smith, D., Dziuda, W., Harstad, B., & Loeper, A. (2019). Gridlock and inefficient policy instruments. *Theoretical Economics*, *14*(4), 1483–1534.

Austen-Smith, D., & Feddersen, T. J. (2006). Deliberation, preference uncertainty, and voting rules. *American Political Science Review*, *100*(2), 209–217.

Baer, G. W. (1973). Sanctions and security: The league of nations and the italian–ethiopian war, 1935–1936. *International Organization*, *27*(2), 165–179.

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, *115*(37), 9216–9221.

Ball & Gao, X. (2024). Benefiting from bias: Delegating to encourage information acquisition. *Journal of Economic Theory*, *217*, 105816.

Ball, R. (1999). Discontinuity and non-existence of equilibrium in the probabilistic spatial voting model. *Social Choice and Welfare*, *16*(4), 533–555.

Banks, J. S. (1985). Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, *1*(4), 295–306.

Banks, J. S. (1990). Monopoly agenda control and asymmetric information. *Quarterly Journal of Economics*, *105*(2), 445–464.

Banks, J. S. (2000). Buying supermajorities in finite legislatures. *American Political Science Review*, *94*(3), 677–681.

Banks, J. S., Duggan, J., et al. (2006). A general bargaining model of legislative policymaking. *Quarterly Journal of Political Science*, *1*(1), 49–85.

Banks, J. S., & Gasmi, F. (1987). Endogenous agenda formation in three-person committees. *Social Choice and Welfare*, *4*(2), 133–152.

Bapat, N. A., & Morgan, T. C. (2009). Multilateral versus unilateral sanctions reconsidered: A test using new data. *International Studies Quarterly*, *53*(4), 1075–1094.

Barber, M. J., & McCarty, N. (2015). Causes and consequences of polarization. *Solutions to political polarization in America*, *15*, 50.

Barberà, S., & Gerber, A. (2017). Sequential voting and agenda manipulation. *Theoretical Economics*, *12*(1), 211–247.

Baron, D. P. (1991). A spatial bargaining theory of government formation in parliamentary systems. *American Political Science Review*, *85*(1), 137–164.

Baron, D. P. (1996). A dynamic theory of collective goods programs. *American Political Science Review*, *90*(2), 316–330.

Baron, D. P. (2000). Legislative organization with informational committees. *American Journal of Political Science*, 485–505.

Baron, D. P. (2019). Simple dynamics of legislative bargaining: Coalitions and proposal power. *Economic Theory*, *67*(1), 319–344.

Baron, D. P., & Bowen, T. R. (2015). Dynamic coalitions.

Baron, D. P., & Ferejohn, J. (1989). Bargaining in legislatures. *American political science review*, *83*(4), 1181–1206.

Beckmann, M. N., & McGann, A. J. (2008). Navigating the legislative divide: Polarization, presidents, and policymaking in the united states. *Journal of Theoretical Politics*, *20*(2), 201–220.

Bendor, J., Glazer, A., & Hammond, T. (2001). Theories of delegation. *Annual review of political science*, *4*(1), 235–269.

Besley, T. (2006). *Principled agents?: The political economy of good government*. Oxford University Press, USA.

Besley, T., & Smart, M. (2007). Fiscal restraints and voter welfare. *Journal of public Economics*, *91*(3-4), 755–773.

Binder, S. A. (1999). The dynamics of legislative gridlock, 1947–96. *American Political Science Review*, *93*(3), 519–533.

Binder, S. A. (2004). *Stalemate: Causes and consequences of legislative gridlock*. Rowman & Littlefield.

Black, D. (1948). On the rationale of group decision-making. *Journal of political economy*, *56*(1), 23–34.

Boehmke, F. J., Gailmard, S., Patty, J. W., et al. (2006). Whose ear to bend? information sources and venue choice in policy making. *Quarterly Journal of Political Science*, *1*(2), 139–169.

Bolton, A. (2022). Gridlock, bureaucratic control, and nonstatutory policymaking in congress. *American Journal of Political Science*, *66*(1), 238–254.

Bolton, P., & Dewatripont, M. (2004). *Contract theory*. MIT Press.

Bowen, T. R., Chen, Y., Eraslan, H., & Zápal, J. (2017). Efficiency of flexible budgetary institutions. *Journal of economic theory*, *167*, 148–176.

Bowen, T. R., Chen, Y., & Eraslan, H. (2014). Mandatory versus discretionary spending: The status quo effect. *American Economic Review*, *104*(10), 2941–2974.

Bowen, T. R., & Zahran, R. (2012). On dynamic compromise. *Games and Economic Behavior*, *76*(1), 391–419.

Bowling, C. J., & Ferguson, M. R. (2001). Divided government, interest representation, and policy differences: Competing explanations of gridlock in the fifty states. *The Journal of Politics*, *63*(1), 182–206.

Bradbury, J. C., & Crain, W. M. (2001). Legislative organization and government spending: Cross-country evidence. *Journal of public Economics*, *82*(3), 309–325.

Brady, D. W., & Volden, C. (1998). Revolving gridlock: Politics and policy from carter to clinton. *(No Title)*.

Brady, D. W., & Volden, C. (2006). *Revolving gridlock: Politics and policy from jimmy carter to george w*. Westview Press.

Brekhov, B. (2022). Rewards versus sanctions in international relations: A game-theoretic analysis of bluffing. *International Interactions*, *48*(1), 75–109.

Buchanan, J. M., & Tullock, G. (1965). *The calculus of consent: Logical foundations of constitutional democracy* (Vol. 100). University of Michigan press.

Buisseret, P., & Bernhardt, D. (2017). Dynamics of policymaking: Stepping back to leap forward, stepping forward to keep back. *American Journal of Political Science*, *61*(4), 820–835.

Buisseret, P., & Bernhardt, D. (2018). Reelection and renegotiation: International agreements in the shadow of the polls. *American Political Science Review*, *112*(4), 1016–1035.

Burbank, S. B., & Farhang, S. (2021). Class certification in the us courts of appeals: A longitudinal study. *Law & Contemp. Probs.*, *84*, 73.

Bursztyn, L., Egorov, G., & Fiorin, S. (2020). From extreme to mainstream: The erosion of social norms. *American economic review*, *110*(11), 3522–3548.

Burtraw, D. (1992). Strategic delegation in bargaining. *Economics Letters*, *38*(2), 181–185.

Burtraw, D. (1993). Bargaining with noisy delegation. *The RAND Journal of Economics*, 40–57.

Callander, S., et al. (2008). A theory of policy expertise. *Quarterly Journal of Political Science*, *3*(2), 123–140.

Calvert, R. L. (1985). Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science*, 69–95.

Calvert, R. L., & Dietz, N. (2005). Legislative coalitions in a bargaining model with externalities. In *Social choice and strategic decisions: Essays in honor of jeffrey s. banks* (pp. 227–247). Springer.

Calvo, E., & Sagarzazu, I. (2016). Presidential agenda authority in plurality-led congresses. *Legislative institutions and lawmaking in Latin America*, 32.

Cameron, C. (2000). *Veto bargaining: Presidents and the politics of negative power*. Cambridge University Press.

Cameron, C., & McCarty, N. (2004). Models of vetoes and veto bargaining. *Annu. Rev. Polit. Sci.*, *7*(1), 409–435.

Cardona, D., & Rubí-Barceló, A. (2014). On the efficiency of equilibria in a legislative bargaining model with particularistic and collective goods. *Public Choice*, *161*(3), 345–366.

Carmines, E. G., & Fowler, M. (2017). The temptation of executive authority: How increased polarization and the decline in legislative capacity have contributed to

the expansion of presidential power. *Indiana Journal of Global Legal Studies*, *24*(2), 369–398.

Casella, A. (2011). Agenda control as a cheap talk game: Theory and experiments with storable votes. *Games and Economic Behavior*, *72*(1), 46–76.

Casella, A., & Macé, A. (2021). Does vote trading improve welfare? *Annual Review of Economics*, *13*, 57–86.

Casella, A., & Palfrey, T. (2019). Trading votes for votes. a dynamic theory. *Econometrica*, *87*(2), 631–652.

Casella, A., & Palfrey, T. R. (2020). Trading votes for votes: A laboratory study. *Games and Economic Behavior*, *125*, 1–26.

Chen, Y., & Eraslan, H. (2013). Informational loss in bundled bargaining. *Journal of Theoretical Politics*, *25*(3), 338–362.

Chen, Y., & Eraslan, H. (2014). Rhetoric in legislative bargaining with asymmetric information. *Theoretical Economics*, *9*(2), 483–513.

Chen, Y., & Eraslan, H. (2017). Dynamic agenda setting. *American Economic Journal: Microeconomics*, *9*(2), 1–32.

Chen, Y., & Zápal, J. (2022). Sequential vote buying. *Journal of Economic Theory*, *205*, 105529.

Chiou, F.-Y., & Rothenberg, L. S. (2003). When pivotal politics meets partisan politics. *American Journal of Political Science*, *47*(3), 503–522.

Chiou, F.-Y., & Rothenberg, L. S. (2008). Comparing legislators and legislatures: The dynamics of legislative gridlock reconsidered. *Political Analysis*, *16*(2), 197–212.

Cho, I.-K., & Kreps, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, *102*(2), 179–221.

Chowdhury, P. R., & Sengupta, K. (2012). Transparency, complementarity and holdout. *Games and Economic Behavior*, *75*(2), 598–612.

Coleman, J. J. (1999). Unified government, divided government, and party responsiveness. *American Political Science Review*, *93*(4), 821–835.

Congleton, R. D. (1986). Rent-seeking aspects of political advertising. *Public choice*, 249–263.

Congleton, R. D. (2003). On the merits of bicameral legislatures: Policy stability within partisan polities. *Jahrbuch für Neue Politische Ökonomie*, *22*, 29–49.

Cortright, D. (1997). *The price of peace: Incentives and international conflict prevention*. Rowman & Littlefield.

Cortright, D., Lopez, G. A., & Gerber, L. (2002). *Sanctions and the search for security: Challenges to un action*. Lynne Rienner Publishers.

Coughlan, P. J. (2000). In defense of unanimous jury verdicts: Mistrials, communication and strategic voting. *American Political Science Review*, *94*(2), 375–393.

Cowen, T., & Sutter, D. (1998). Why only nixon could go to china. *Public Choice*, *97*, 605–615.

Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431–1451.

Crombez, C., Groseclose, T., & Krehbiel, K. (2006). Gatekeeping. *The Journal of Politics*, *68*(2), 322–334.

Crombez, C., & Hix, S. (2015). Legislative activity and gridlock in the european union. *British Journal of Political Science*, *45*(3), 477–499.

Dal Bó, E. (2007). Bribing voters. *American Journal of Political Science*, *51*(4), 789–803.

Dalton, R. J. (2008). The quantity and the quality of party systems: Party system polarization, its measurement, and its consequences. *Comparative political studies*, *41*(7), 899–920.

Dalton, R. J. (2021). Modeling ideological polarization in democratic party systems. *Electoral Studies*, *72*, 102346.

De Figueiredo, J. M., & Richter, B. K. (2014). Advancing the empirical research on lobbying. *Annual Review of Political Science*, *17*(1), 163–185.

Dekker, M., Walentek, D., Haslbeck, J., & Broere, J. (2020). A general framework for international conflict under incomplete information.

Denzau, A. T., & Mackay, R. J. (1983). Structure-induced equilibria and legislative choice. *Public Choice*, *41*(1), 35–57.

Dessein, W. (2002). Authority and communication in organizations. *The Review of Economic Studies*, *69*(4), 811–838.

Diermeier, D. (1995). Commitment, deference, and legislative institutions. *American Political Science Review*, *89*(2), 344–355.

Diermeier, D., & Fong, P. (2011). Legislative bargaining with reconsideration. *Quarterly Journal of Economics*, *126*(3), 947–985.

Diermeier, D., & Fong, P. (2012). Characterization of the von neumann-morgenstern stable set in a non-cooperative model of dynamic policy-making with a persistent agenda setter. *Games and Economic Behavior*, *76*(1), 349–353.

Diermeier, D., & Myerson, R. B. (1999). Bicameralism and its consequences for the internal organization of legislatures. *American Economic Review*, *89*(5), 1182–1196.

Dimant, E. (2024). Hate trumps love: The impact of political polarization on social preferences. *Management Science*, *70*(1), 1–31.

Dodge, J., & Lee, J. (2017). Framing dynamics and political gridlock: The curious case of hydraulic fracturing in new york. *Journal of environmental policy & planning*, *19*(1), 14–34.

Dolado, J. J., Griffiths, M., & Padilla, A. J. (1994). Delegation in international monetary policy games∗. *European Economic Review*, *38*(5), 1057–1069.

Dorussen, H. (2001). Mixing carrots with sticks: Evaluating the effectiveness of positive incentives. *Journal of Peace Research*, *38*(2), 251–262.

Downs, A. (1957). An economic theory of political action in a democracy. *Journal of political economy*, *65*(2), 135–150.

Doxey, M. (1972). International sanctions: A framework for analysis with special reference to the un and southern africa. *International organization*, *26*(3), 527–550.

Drezner, D. W. (1999). *The sanctions paradox: Economic statecraft and international relations*. Cambridge University Press.

Drezner, D. W. (2003). The hidden hand of economic coercion. *International Organization*, *57*(3), 643–659.

Drury, A. C. (2005). *Economic sanctions and presidential decisions*. Springer.

Drury, A. C., James, P., & Peksen, D. (2014). Neo-kantianism and coercive diplomacy: The complex case of economic sanctions. *International Interactions*, *40*(1), 25–51.

Duggan, J. (2017). Existence of stationary bargaining equilibria. *Games and Economic Behavior*, *102*, 111–126.

Duggan, J., & Banks, J. (2000). A bargaining model of collective choice. *American Political Science Review*, *94*(1), 73–88.

Duggan, J., & Kalandrakis, T. (2011). A newton collocation method for solving dynamic bargaining games. *Social Choice and Welfare*, *36*(3), 611–650.

Duggan, J., & Kalandrakis, T. (2012). Dynamic legislative policy making. *Journal of Economic Theory*, *147*(5), 1653–1688.

Dziuda, W., & Loeper, A. (2016). Dynamic collective choice with endogenous status quo. *Journal of Political Economy*, *124*(4), 1148–1186.

Dziuda, W., & Loeper, A. (2018). Dynamic pivotal politics. *American Political Science Review*, *112*(3), 580–601.

Eaton, J., & Engers, M. (1992). Sanctions. *Journal of political economy*, *100*(5), 899–928.

Eaton, J., & Engers, M. (1994). Threats and promises.

Eaton, J., & Engers, M. (1999). Sanctions: Some simple analytics. *American Economic Review*, *89*(2), 409–414.

Emerson, B., & Michaels, J. D. (2020). Abandoning presidential administration: A civic governance agenda to promote democratic equality and guard against creeping authoritarianism. *UCLA L. Rev. Discourse*, *68*, 418.

Epple, D., & Riordan, M. H. (1987). Cooperation and punishment under repeated majority voting. *Public Choice*, *55*(1), 41–73.

Epstein, D., & O'Halloran, S. (1994). Administrative procedures, information, and agency discretion. *American Journal of Political Science*, 697–722.

Epstein, D., & O'Halloran, S. (1999). Measuring the electoral and policy impact of majority-minority voting districts. *American Journal of Political Science*, 367–395.

Eraslan, H., & Evdokimov, K. S. (2019). Legislative and multilateral bargaining. *Annual Review of Economics*, *11*(1), 443–472.

Eraslan, H., Evdokimov, K. S., & Zápal, J. (2022). Dynamic legislative bargaining. *Bargaining: Current research and future directions*, 151–175.

Eraslan, H., & McLennan, A. (2013). Uniqueness of stationary equilibrium payoffs in coalitional bargaining. *Journal of Economic Theory*, *148*(6), 2195–2222.

Eraslan, H., & Merlo, A. (2002). Majority rule in a stochastic model of bargaining. *Journal of Economic Theory*, *103*(1), 31–48.

Escobar, J. F., & Toikka, J. (2013). Efficiency in games with markovian private information. *Econometrica*, *81*(5), 1887–1934.

Farhang, S. (2018). Legislating for litigation: Delegation, public policy, and democracy. *Calif. L. Rev.*, *106*, 1529.

Farhang, S. (2021). Legislative capacity & administrative power under divided polarization. *Daedalus*, *150*(3), 49–67.

Farhang, S., & Yaver, M. (2016). Divided government and the fragmentation of american law. *American Journal of Political Science*, *60*(2), 401–417.

Fearon, J. D. (1994). Domestic political audiences and the escalation of international disputes. *American political science review*, *88*(3), 577–592.

Fearon, J. D. (1995). Rationalist explanations for war. *International organization*, *49*(3), 379–414.

Fearon, J. D. (1997). Signaling foreign policy interests: Tying hands versus sinking costs. *Journal of conflict resolution*, *41*(1), 68–90.

Ferejohn, J. (1974). *Sour notes on the theory of vote trading. social science.* [Working paper No 41, California Institute of Technology].

Ferejohn, J. (1986). Incumbent performance and electoral control. *Public choice*, 5–25.

Fiorina, M. P., & Abrams, S. J. (2008). Political polarization in the american public. *Annu. Rev. Polit. Sci.*, *11*(1), 563–588.

Franchino, F. (2004). Delegating powers in the european community. *British Journal of Political Science*, *34*(2), 269–293.

Freeman, J., & Rossi, J. (2012). Improving interagency coordination in shared regulatory space. *Admin. & Reg. L. News*, *38*, 11.

Gailmard, S. (2002). Expertise, subversion, and bureaucratic discretion. *Journal of Law, Economics, and Organization*, *18*(2), 536–555.

Gailmard, S. (2009). Discretion rather than rules: Choice of instruments to control bureaucratic policy making. *Political Analysis*, *17*(1), 25–44.

Gailmard, S., & Patty, J. W. (2013). Stovepiping. *Journal of Theoretical Politics*, *25*(3), 388–411.

Gailmard, S., & Patty, J. W. (2019). Giving advice versus making decisions: Transparency, information, and delegation. *Political Science Research and Methods*, *7*(3), 471–488.

Galtung, J. (1967). On the effects of international economic sanctions, with examples from the case of rhodesia. *World politics*, *19*(3), 378–416.

Gelman, J., & Wilson, S. L. (2022). Measuring congressional partisanship and its consequences. *Legislative studies quarterly*, *47*(1), 225–256.

Genicot, G., & Ray, D. (2006). Contracts and externalities: How things fall apart. *Journal of Economic Theory*, *131*(1), 71–100.

Gidron, N., Adams, J., & Horne, W. (2018). How ideology, economics and institutions shape affective polarization in democratic polities. *Annual conference of the American political science association*.

Gift, K., & Gift, T. (2015). Does politics influence hiring? evidence from a randomized experiment. *Political Behavior*, *37*(3), 653–675.

Gilligan, T. W., & Krehbiel, K. (1987). Collective decisionmaking and standing committees: An informational rationale for restrictive amendment procedures. *The Journal of Law, Economics, and Organization*, *3*(2), 287–335.

Gilligan, T. W., & Krehbiel, K. (1989). Asymmetric information and legislative rules with a heterogeneous committee. *American journal of political science*, 459–490.

Golder, M. (2016). Far right parties in europe. *Annual review of political science*, *19*(1), 477–497.

Graham, M. H., & Svolik, M. W. (2020). Democracy in america? partisanship, polarization, and the robustness of support for democracy in the united states. *American Political Science Review*, *114*(2), 392–409.

Groseclose, T., & McCarty, N. (2001). The politics of blame: Bargaining before an audience. *American Journal of Political Science*, 100–119.

Groseclose, T., & Snyder Jr, J. M. (1996). Buying supermajorities. *American Political Science Review*, *90*(2), 303–315.

Grossman, G. M., & Helpman, E. (1996). Electoral competition and special interest politics. *The review of economic studies*, *63*(2), 265–286.

Grossman, G. M., & Helpman, E. (2001). *Special interest politics*. MIT press.

Guisinger, A., & Smith, A. (2002). Honest threats: The interaction of reputation and political institutions in international crises. *Journal of Conflict Resolution*, *46*(2), 175–200.

Gustafsson, A. (2019). Busy doing nothing: Why politicians implement inefficient policies. *Constitutional Political Economy*, *30*(3), 282–299.

Haas, R., & O'Sullivan, M. L. (2000). Terms of engagement: Alternatives to. *Survival*, *42*(2), 113–35.

Halac, M., & Yared, P. (2014). Fiscal rules and discretion under persistent shocks. *Econometrica*, *82*(5), 1557–1614.

Halac, M., & Yared, P. (2018). Fiscal rules and discretion in a world economy. *American Economic Review*, *108*(8), 2305–2334.

Halac, M., & Yared, P. (2020). Commitment versus flexibility with costly verification. *Journal of Political Economy*, *128*(12), 4523–4573.

Halac, M., & Yared, P. (2022a). Fiscal rules and discretion under limited enforcement. *Econometrica*, *90*(5), 2093–2127.

Halac, M., & Yared, P. (2022b). Instrument-based versus target-based rules. *The Review of Economic Studies*, *89*(1), 312–345.

Harrington Jr, J. E. (1990). The role of risk preferences in bargaining when acceptance of a proposal requires less than unanimous approval. *Journal of risk and uncertainty*, *3*(2), 135–154.

Harrington Jr, J. E. (1992). The revelation of information through the electoral process: An exploratory analysis. *Economics & Politics*, *4*(3), 255–276.

He, W., & Sun, Y. (2017). Stationary markov perfect equilibria in discounted stochastic games. *Journal of Economic Theory*, *169*, 35–61.

Hetherington, M. J. (2001). Resurgent mass partisanship: The role of elite polarization. *American political science review*, *95*(3), 619–631.

Hicks, W. D. (2015). Partisan competition and the efficiency of lawmaking in american state legislatures, 1991-2009. *American Politics Research*, *43*(5), 743–770.

Holmström, B. (1977). *On incentives and control in organizations* [Ph.D. thesis]. Stanford University.

Holmström, B. (1984). On the theory of delegation. In M. Boyer & R. Kihlstrom (Eds.), *Bayesian models in economic theory* (pp. 115–141). North-Holland.

Hovi, J., Huseby, R., & Sprinz, D. F. (2005). When do (imposed) economic sanctions work? *World politics*, *57*(4), 479–499.

Howell, W. G. (2015). Power without persuasion: The politics of direct presidential action.

Hu, X., & Lei, H. (2025). The optimality of (stochastic) veto delegation. *Games and Economic Behavior*, *150*, 215–234.

Huber, J. D., & McCarty, N. (2004). Bureaucratic capacity, delegation, and political reform. *American Political Science Review*, *98*(3), 481–494.

Huber, J. D., & Shipan, C. R. (2002). *Deliberate discretion?: The institutional foundations of bureaucratic autonomy*. Cambridge University Press.

Hufbauer, G., Schott, J., Elliott, K., & Oegg, B. (2007). Economic sanctions reconsidered. washington, dc. *Peterson Institute for International Economics*.

Hurwitz, J., & Peffley, M. (1987). The means and ends of foreign policy as determinants of presidential support. *American Journal of Political Science*, 236–258.

Indridason, I. H. (2008). When to run and when to hide: Electoral coordination and exit. *Economics & Politics*, *20*(1), 80–105.

Ingberman, D. E. (1985). Running against the status quo: Institutions for direct democracy referenda and allocations over time. *Public Choice*, *46*(1), 19–43.

Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology: A social identity perspective on polarization. *Public opinion quarterly*, *76*(3), 405–431.

Jackson, M. O., & Moselle, B. (2002). Coalition and party formation in a legislative voting game. *Journal of Economic theory*, *103*(1), 49–87.

Jones, D. R. (2001). Party polarization and legislative gridlock. *Political Research Quarterly*, *54*(1), 125–141.

Kaempfer, W. H., & Lowenberg, A. D. (1988). The theory of international economic sanctions: A public choice approach. *The American Economic Review*, *78*(4), 786–793.

Kaempfer, W. H., & Lowenberg, A. D. (2007). The political economy of economic sanctions. *Handbook of defense economics*, *2*, 867–911.

Kaempfer, W. H., & Lowenberg, A. D. (1992). International economic sanctions: A public choice perspective. *(No Title)*.

Kalandrakis, T. (2004). A three-player dynamic majoritarian bargaining game. *Journal of Economic Theory*, *116*(2), 294–322.

Kalandrakis, T. (2010). Minimum winning coalitions and endogenous status quo. *International Journal of Game Theory*, *39*(4), 617–643.

Kalandrakis, T. (2015). Computation of equilibrium values in the baron and ferejohn bargaining model. *Games and Economic Behavior*, *94*, 29–38.

Kalandrakis, T. (2016). Pareto efficiency in the dynamic one-dimensional bargaining model. *Journal of Theoretical Politics*, *28*(3), 525–536.

Kang, M. (2017). Representation, sophisticated voting, and the size of the gridlock region. *Journal of Theoretical Politics*, *29*(4), 623–646.

Kartik, N., Kleiner, A., & Van Weelden, R. (2021). Delegation in veto bargaining. *American Economic Review*, *111*(12), 4046–4087.

Kawamori, T. (2005). Players' patience and equilibrium payoffs in the baron-ferejohn model. *Economics Bulletin*, *3*(43), 1–5.

Kawamori, T. (2013). Rejecter-proposer legislative bargaining with heterogeneous time and risk preferences. *Social Choice and Welfare*, *40*(1), 27–40.

Keck, M. E., & Sikkink, K. (1998). *Activists beyond borders: Advocacy networks in international politics*. Cornell University Press.

Kleiner, A. (2022). Optimal delegation in a multidimensional world. *arXiv preprint arXiv:2208.11835*.

Koessler, F., & Martimort, D. (2012). Optimal delegation with multi-dimensional decisions. *Journal of Economic Theory*, *147*(5), 1850–1881.

Kolotilin, A., & Zapechelnyuk, A. (2025). Persuasion meets delegation. *Econometrica*, *93*(1), 195–228.

Kováč, E., & Mylovanov, T. (2009). Stochastic mechanisms in settings without monetary transfers: The regular case. *Journal of Economic Theory*, *144*(4), 1373–1395.

Krähmer, D., & Kováč, E. (2016). Optimal sequential delegation. *Journal of Economic Theory*, *163*, 849–888.

Krehbiel, K. (1996). Institutional and partisan sources of gridlock: A theory of divided and unified government. *Journal of Theoretical Politics*, *8*(1), 7–40.

Krehbiel, K. (1998). *Pivotal politics: A theory of u.s. lawmaking*. University of Chicago Press.

Krehbiel, K. (2001). Plausibility of signals by a heterogeneous committee. *American Political Science Review*, *95*(2), 453–457.

Krehbiel, K. (2010). *Pivotal politics: A theory of us lawmaking*. University of Chicago Press.

Krishna, V., & Morgan, J. (2001). A model of expertise. *The Quarterly Journal of Economics*, *116*(2), 747–775.

Krustev, V. L. (2010). Strategic demands, credible threats, and economic coercion outcomes. *International Studies Quarterly*, *54*(1), 147–174.

Lacy, D., & Niou, E. M. (2004). A theory of economic sanctions and issue linkage: The roles of preferences, information, and threats. *The journal of politics*, *66*(1), 25–42.

Laffont, J.-J., & Martimort, D. (2009). *The theory of incentives: The principal-agent model*. Princeton University Press.

Le Breton, M., Sudhölter, P., & Zaporozhets, V. (2012). Sequential legislative lobbying. *Social Choice and Welfare*, *39*(2), 491–520.

Lefebvre, P., & Martimort, D. (2022). Delegation, capture and endogenous information structures. *Journal of Theoretical Politics*, *34*(3), 357–414.

Lektzian, D., & Biglaiser, G. (2014). The effect of foreign direct investment on the use and success of us sanctions. *Conflict Management and Peace Science*, *31*(1), 70–93.

Lektzian, D., & Souva, M. (2003). The economic peace between democracies: Economic sanctions and domestic institutions. *Journal of Peace Research*, *40*(6), 641–660.

Levy, Y. (2013). Discounted stochastic games with no stationary nash equilibrium: Two examples. *Econometrica*, *81*(5), 1973–2007.

Levy, Y. J., & McLennan, A. (2015). Corrigendum to "discounted stochastic games with no stationary nash equilibrium: Two examples". *Econometrica*, *83*(3), 1237–1252.

Lewis, D. E. (2004). *Presidents and the politics of agency design: Political insulation in the united states government bureaucracy, 1946-1997*. Stanford University Press.

Lewis, D. E. (2010). The politics of presidential appointments: Political control and bureaucratic performance.

Leyton-Brown, D., & Ruggie, J. G. (1987). The north american political economy in the global context: An analytical framework. *International Journal*, *42*(1), 3–24.

Lubensky, D., & Schmidbauer, E. (2018). Equilibrium informativeness in veto games. *Games and Economic Behavior*, *109*, 104–125.

Lupia, A., & McCubbins, M. D. (1994). Who controls? information and the structure of legislative decision making. *Legislative Studies Quarterly*, 361–384.

Mansfield, E. D., Milner, H. V., & Pevehouse, J. C. (2007). Vetoing co-operation: The impact of veto players on preferential trading arrangements. *British Journal of Political Science*, *37*(3), 403–432.

Marino, A. M. (2007). Delegation versus veto in organizational games of strategic communication. *Journal of Public Economic Theory*, *9*(6), 979–992.

Martimort, D., & Semenov, A. (2006). Continuity in mechanism design without transfers. *Economics Letters*, *93*(2), 182–189.

Martin, E. M. (1997). An informational theory of the legislative veto. *The Journal of Law, Economics, and Organization*, *13*(2), 319–343.

Mas-Colell, A., Whinston, M. D., Green, J. R., et al. (1995). *Microeconomic theory* (Vol. 1). Oxford university press New York.

Mason, L. (2018). *Uncivil agreement: How politics became our identity*. University of Chicago Press.

Matthews, S. A. (1989). Veto threats: Rhetoric in a bargaining game. *The Quarterly Journal of Economics*, *104*(2), 347–369.

Mayhew, D. (1991). Divided we govern new haven. *CT: Yale University*.

Mayhew, D. R. (2005). *Divided we govern: Party control, lawmaking and investigations, 1946-2002*. Yale university press.

McCarty, N. (1997). Presidential reputation and the veto. *Economics & Politics*, *9*(1), 1–26.

McCarty, N. (2004). The appointments dilemma. *American Journal of Political Science*, *48*(3), 413–428.

McCarty, N. (2011). The policy effects of political polarization. In *The transformation of american politics* (pp. 223–255). Princeton University Press.

McCarty, N. (2018). Pivotal politics, partisan polarization, and policy predictability. *The Journal of Politics*, *80*(3), 1076–1081.

McCarty, N. (2019). *Polarization: What everyone needs to know®*. Oxford University Press.

McCarty, N., Poole, K. T., & Rosenthal, H. (2016). *Polarized america: The dance of ideology and unequal riches*. mit Press.

McConnell, C., Margalit, Y., Malhotra, N., & Levendusky, M. (2018). The economic consequences of partisanship in a polarized era. *American Journal of Political Science*, *62*(1), 5–18.

McCoy, J., Rahman, T., & Somer, M. (2018). Polarization and the global crisis of democracy: Common patterns, dynamics, and pernicious consequences for democratic polities. *American behavioral scientist*, *62*(1), 16–42.

McGillivray, F., & Stam, A. C. (2004). Political institutions, coercive diplomacy, and the duration of economic sanctions. *Journal of Conflict Resolution*, *48*(2), 154–172.

McKelvey, R. D. (1976). Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory*, *12*(3), 472–482.

McKelvey, R. D. (1979). General conditions for global intransitivities in formal voting models. *Econometrica*, *47*(5), 755–767.

McLean, E. V., & Whang, T. (2014). Designing foreign policy: Voters, special interest groups, and economic sanctions. *Journal of Peace Research*, *51*(5), 589–602.

Melumad, N. D., & Shibano, T. (1991). Communication in settings with no transfers. *The RAND Journal of Economics*, 173–198.

Mendelson, N. A. (2003). Agency burrowing: Entrenching policies and personnel before a new president arrives. *NYUL Rev.*, *78*, 557.

Merlo, A., & Wilson, C. (1995). A stochastic model of sequential bargaining with complete information. *Econometrica: Journal of the Econometric Society*, 371–399.

Merlo, A., & Wilson, C. (1998). Efficient delays in a stochastic model of bargaining. *Economic Theory*, *11*(1), 39–55.

Michelitch, K. (2015). Does electoral competition exacerbate interethnic or interpartisan economic discrimination? evidence from a field experiment in market price bargaining. *American Political Science Review*, *109*(1), 43–61.

Miller, L., Montero, M., & Vanberg, C. (2018). Legislative bargaining with heterogeneous disagreement values: Theory and experiments. *Games and Economic Behavior*, *107*, 60–92.

Mitchell, D. W. (1987). Candidate behavior under mixed motives. *Social Choice and Welfare*, *4*(2), 153–160.

Montero, M. (2006). Noncooperative foundations of the nucleolus in majority games. *Games and Economic Behavior*, *54*(2), 380–397.

Montero, M., et al. (2017). Proportional payoffs in legislative bargaining with weighted voting: A characterization. *Quarterly Journal of Political Science*, *12*(3), 325–346.

Montero, M., & Vidal-Puga, J. J. (2011). Demand bargaining and proportional payoffs in majority games. *Games and Economic Behavior*, *71*(2), 395–408.

Morelli, M. (1999). Demand competition and policy compromise in legislative bargaining. *American Political Science Review*, *93*(4), 809–820.

Morgan, T. C., Bapat, N., & Kobayashi, Y. (2014). Threat and imposition of economic sanctions 1945–2005: Updating the ties dataset. *Conflict Management and Peace Science*, *31*(5), 541–558.

Morgan, T. C., Bapat, N., & Krustev, V. (2009). The threat and imposition of economic sanctions, 1971—2000. *Conflict Management and Peace Science*, *26*(1), 92–110.

Morgan, T. C., Bapat, N. A., & Kobayashi, Y. (2021). The threat and imposition of economic sanctions data project: A retrospective. *Research Handbook on Economic Sanctions*, 44–61.

Morgan, T. C., & Miers, A. C. (1999). When threats succeed: A formal model of the threat and use of economic sanctions. *Annual Meeting of the American Political Science Association, Atlanta, GA*, *264*.

Morgan, T. C., Syropoulos, C., & Yotov, Y. V. (2023). Economic sanctions: Evolution, consequences, and challenges. *Journal of Economic Perspectives*, *37*(1), 3–29.

Morrow, J. D. (1989). Capabilities, uncertainty, and resolve: A limited information model of crisis bargaining. *American Journal of Political Science*, 941–972.

Mudde, C. (2007). *Populist radical right parties in europe* (Vol. 22). Cambridge university press Cambridge.

Mylovanov, T. (2008). Veto-based delegation. *Journal of Economic Theory*, *138*(1), 297–307.

Nash, J. (1953). Two-person cooperative games. *Econometrica: Journal of the Econometric Society*, 128–140.

Nicholson, S. P., Coe, C. M., Emory, J., & Song, A. V. (2016). The politics of beauty: The effects of partisan bias on physical attractiveness. *Political Behavior*, *38*(4), 883–898.

Nielsen, R. A. (2013). Rewarding human rights? selective aid sanctions against repressive states. *International Studies Quarterly*, *57*(4), 791–803.

Nincic, M., & Hinckley, B. (1991). Foreign policy and the evaluation of presidential candidates. *Journal of Conflict Resolution*, *35*(2), 333–355.

Norris, P., & Inglehart, R. (2019). *Cultural backlash: Trump, brexit, and authoritarian populism*. Cambridge University Press.

Norris, P., & Inglehart, R. F. (2018). Understanding brexit: Cultural resentment versus economic grievances. *Paper for presentation at the Panel on 'Populism in Advanced Capitalist Democracies', Thursday*, *30*, 4–00.

Nunnari, S. (2021). Dynamic legislative bargaining with veto power: Theory and experiments. *Games and Economic Behavior*, *126*, 186–230.

Oechslin, M. (2014). Targeting autocrats: Economic sanctions and regime change. *European Journal of Political Economy*, *36*, 24–40.

Ordeshook, P. C., & Palfrey, T. R. (1988). Agendas, strategic voting, and signaling with incomplete information. *American Journal of Political Science*, 441–466.

Pape, R. A. (1997). Why economic sanctions do not work. *International security*, *22*(2), 90–136.

Park, R. E. (1967). The possibility of a social welfare function: Comment. *The American Economic Review*, *57*(5), 1300–1304.

Peksen, D. (2019). When do imposed economic sanctions work? a critical review of the sanctions effectiveness literature. *Defence and Peace Economics*, *30*(6), 635–647.

Peksen, D., & Peterson, T. M. (2016). Sanctions and alternate markets: How trade and alliances affect the onset of economic coercion. *Political Research Quarterly*, *69*(1), 4–16.

Persson, T., & Tabellini, G. (1992). The politics of 1992: Fiscal policy and european integration. *The review of economic studies*, *59*(4), 689–701.

Persson, T., & Tabellini, G. (2002). *Political economics: Explaining economic policy*. MIT press.

Peterson, P. E., & Greene, J. P. (1994). Why executive-legislative conflict in the united states is dwindling. *British Journal of Political Science*, *24*(1), 33–55.

Piazolo, D., & Vanberg, C. (2025). Legislative bargaining with private information: A comparison of majority and unanimity rule. *Games and Economic Behavior*.

Pierson, P., & Schickler, E. (2020). Madison's constitution under stress: A developmental analysis of political polarization. *Annual Review of Political Science*, *23*(1), 37–58.

Plott, C. R. (1967). A notion of equilibrium and its possibility under majority rule. *The American Economic Review*, *57*(4), 787–806.

Poole, K. T., & Rosenthal, H. (2000). *Congress: A political-economic history of roll call voting*. Oxford University Press, USA.

Potter, R. A. (2019). *Bending the rules: Procedural politicking in the bureaucracy*. University of Chicago Press.

Powell, R. (1990). *Nuclear deterrence theory: The search for credibility*. Cambridge University Press.

Powell, R. (1999). *In the shadow of power: States and strategies in international politics*. Princeton University Press.

Prior, M. (2013). Media and political polarization. *Annual review of political science*, *16*(1), 101–127.

Putnam, R. D. (1988). Diplomacy and domestic politics: The logic of two-level games. *International Organization*, *42*(3), 427–460.

Reiljan, A. (2020). 'fear and loathing across party lines'(also) in europe: Affective polarisation in european party systems. *European journal of political research*, *59*(2), 376–396.

Renwick, R. (1981). *Economic sanctions*. Center for International Affairs, Harvard University.

Reny, P. J. (1999). On the existence of pure and mixed strategy nash equilibria in discontinuous games. *Econometrica*, *67*(5), 1029–1056.

Riboni, A., & Ruge-Murcia, F. J. (2008). The dynamic (in) efficiency of monetary policy by committee. *Journal of Money, Credit and Banking*, *40*(5), 1001–1032.

Richter, M. (2014). Fully absorbing dynamic compromise. *Journal of Economic Theory*, *152*, 92–104.

Riker, W. H. (1992). The justification of bicameralism. *International Political Science Review*, *13*(1), 101–116.

Riker, W. H., & Brams, S. J. (1973). The paradox of vote trading. *American Political Science Review*, *67*(4), 1235–1247.

Rogers, J. R. (2001). An informational rationale for congruent bicameralism. *Journal of Theoretical Politics*, *13*(2), 123–151.

Rogoff, K. S. (1987). Equilibrium political budget cycles.

Romer, T., & Rosenthal, H. (1978). Political resource allocation, controlled agendas, and the status quo. *Public Choice*, *33*(4), 27–43.

Saeki, M. (2009). Gridlock in the government of the united states: Influence of divided government and veto players. *British Journal of Political Science*, *39*(3), 587–607.

Saporiti, A. (2008). Existence and uniqueness of nash equilibrium in electoral competition games: The hybrid case. *Journal of Public Economic Theory*, *10*(5), 827–857.

Saran, R. (2022). A dynamic optimization approach to delegation with an application to volunteer contracts. *Available at SSRN 4912403*.

Schattschneider, E. E. (1942). Party government. new york: Farrar and rinehart. *Inc. SchattschneiderParty Government1942*.

Schofield, N. (1978). Instability of simple dynamic games. *Review of Economic Studies*, *45*(3), 575–594.

Schofield, N. (1980). Generic properties of simple bergson-samuelson welfare functions. *Journal of Mathematical Economics*, *7*(2), 175–192.

Schreiber, A. P. (1973). Economic coercion as an instrument of foreign policy: Us economic measures against cuba and the dominican republic. *World Politics*, *25*(3), 387–413.

Schultz, K. A. (1999). Do democratic institutions constrain or inform? contrasting two institutional perspectives on democracy and war. *International Organization*, *53*(2), 233–266.

Schultz, K. A. (2001). *Democracy and coercive diplomacy* (Vol. 76). Cambridge University Press.

Segendorff, B. (1998). Delegation and threat in bargaining. *Games and Economic Behavior*, *23*(2), 266–283.

Shafie, D. M. (2020). *The administrative presidency and the environment: Policy leadership and retrenchment from clinton to trump*. Routledge.

Sinclair, B. (2016). *Unorthodox lawmaking: New legislative processes in the us congress*. CQ Press.

Siqueira, K. (2003). International externalities, strategic interaction, and domestic politics. *Journal of Environmental Economics and Management*, *45*(3), 674–691.

Slantchev, B. L. (2003). The principle of convergence in wartime negotiations. *American Political Science Review*, *97*(4), 621–632.

Slantchev, B. L. (2010). Feigning weakness. *International Organization*, *64*(3), 357–388.

Smith, A. (1995). The success and use of economic sanctions. *International Interactions*, *21*(3), 229–245.

Smith, A. (1998). International crises and domestic politics. *American Political Science Review*, *92*(3), 623–638.

Snyder Jr, J. M. (1991). On buying legislatures. *Economics and Politics*, *3*(2), 93–109.

Snyder Jr, J. M., Ting, M. M., & Ansolabehere, S. (2005). Legislative bargaining under weighted voting. *American Economic Review*, *95*(4), 981–1004.

Sobel, R. (2001). *The impact of public opinion on us foreign policy since vietnam*. Oxford: Oxford University Press.

Spaniel, W., & Smith, B. C. (2015). Sanctions, uncertainty, and leader tenure. *International Studies Quarterly*, *59*(4), 735–749.

Sublet, G. (2023). The optimal degree of discretion in fiscal policy.

Syropoulos, C., Felbermayr, G., Kirilakha, A., Yalcin, E., & Yotov, Y. V. (2024). The global sanctions data base–release 3: Covid-19, russia, and multilateral sanctions. *Review of International Economics*, *32*(1), 12–48.

Taylor, A. J. (1998). Explaining government productivity. *American Politics Quarterly*, *26*(4), 439–458.

Tsai, T.-S., & Yang, C.-C. (2010). On majoritarian bargaining with incomplete information. *International Economic Review*, *51*(3), 959–979.

Tsebelis, G. (1990). Are sanctions effective? a game-theoretic analysis. *Journal of Conflict Resolution*, *34*(1), 3–28.

Tsebelis, G. (1995). Decision making in political systems: Veto players in presidentialism, parliamentarism, multicameralism and multipartyism. *British journal of political science*, *25*(3), 289–325.

Tsebelis, G. (1999). Veto players and law production in parliamentary democracies: An empirical analysis. *American political science review*, *93*(3), 591–608.

Tsebelis, G. (2002). *Veto players: How political institutions work*. Princeton University Press.

Tullock, G. (1967). The welfare costs of tariffs, monopolies, and theft. *Western Economic Journal*, *5*(3), 224–232.

Tullock, G. (1997). Where is the rectangle? *Public Choice*, *91*(2), 149–159.

Van Gestel, W., & Crombez, C. (2014). The role of oversight committees in closed rule legislation. *Journal of Theoretical Politics*, *26*(4), 521–547.

Verdier, D., & Woo, B. (2011). Why rewards are better than sanctions. *Economics & Politics*, *23*(2), 220–238.

Volden, C. (2002). A formal model of the politics of delegation in a separation of powers system. *American Journal of Political Science*, 111–133.

Volden, C., & Wiseman, A. E. (2007). Bargaining in legislatures over particularistic and collective goods. *American Political science review*, *101*(1), 79–92.

Volden, C., & Wiseman, A. E. (2008). Erratum to "bargaining in legislatures over particularistic and collective goods" american political science review 101 (1): 79–92. *American Political Science Review*, *102*(3), 385–386.

Wagner, R. H. (2000). Bargaining and war. *American Journal of Political Science*, 469–484.

Wangler, L., Altamirano-Cabrera, J.-C., & Weikard, H.-P. (2013). The political economy of international environmental agreements: A survey. *International Environmental Agreements: Politics, Law and Economics*, *13*, 387–403.

Watts, A. L., Lilienfeld, S. O., Smith, S. F., Miller, J. D., Campbell, W. K., Waldman, I. D., Rubenzer, S. J., & Faschingbauer, T. J. (2013). The double-edged sword of grandiose narcissism: Implications for successful and unsuccessful leadership among us presidents. *Psychological science*, *24*(12), 2379–2389.

Webster, S. W., & Abramowitz, A. I. (2017). The ideological foundations of affective polarization in the us electorate. *American politics research*, *45*(4), 621–647.

Whang, T. (2011). Playing to the home crowd? symbolic use of economic sanctions in the united states. *International Studies Quarterly*, *55*(3), 787–801.

Whang, T., McLean, E. V., & Kuberski, D. W. (2013). Coercion, information, and the success of sanction threats. *American journal of political science*, *57*(1), 65–81.

Wittman, D. (1983). Candidate motivation: A synthesis of alternative theories. *American Political science review*, *77*(1), 142–157.

Yildirim, H. (2007). Proposal power and majority rule in multilateral bargaining with costly recognition. *Journal of Economic Theory*, *136*(1), 167–196.

Zápal, J. (2014). Simple markovian equilibria in dynamic spatial legislative bargaining. *CERGE-EI Working Paper Series*, (515).

Zápal, J. (2016). Markovian equilibria in dynamic spatial legislative bargaining: Existence with three players. *Games and Economic Behavior*, *98*, 235–242.

Zápal, J. (2020). Simple markovian equilibria in dynamic spatial legislative bargaining. *European Journal of Political Economy*, *63*, 101816.

Zapechelnyuk, A. (2020). Optimal quality certification. *American Economic Review: Insights*, *2*(2), 161–176.