COORDINATING COMPROMISE: INFORMATION MANIPULATION AND BICAMERAL BARGAINING IN THE EUROPEAN UNION

BY

DANIEL B. PEMSTEIN

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Political Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2010

Urbana, Illinois

Doctoral Committee:

Professor William T. Bernhard, Chair
Associate Professor Brian J. Gaines
Professor Robert D. Pahre
Associate Professor Tracy E. Sulkin

Abstract

The European Union relies on a complicated—some would say arcane—set of institutions and legislative rules to develop public policy that affects millions of Europeans. Actors within these institutions must effectively navigate this convoluted institutional structure in order to legislate. Crucially, government ministers in the Council of the European Union and members of the European Parliament seek to forge bicameral bargains in a complex information environment. This study examines how European politicians construct such compromises and explores how political elites coordinate around particular proposals when crafting policy. It highlights the ways in which European lawmakers manage and share information to encourage—and hamper—legislative coordination, and emphasizes the role that the European Commission—the Union's bureaucratic arm—plays in transmitting information between lawmakers, modulating legislative efficiency.

Acknowledgements

My committee—Bill Bernhard, Brian Gaines, Bob Pahre, and Tracy Sulkin—have all been wonderful teachers and mentors over the past few years and I would like to thank them for making my graduate school experience so positive. Many people discussed aspects of this research with me but Steve Meserve deserves special recognition for enduring the bulk of my ranting. Finally, my wife Laura provided terrific moral support, especially during the hectic, and difficult, final months of this project. She put up with me and deserves thanks.

Table of Contents

Chapter 1

Introduction

This project is about the art of political compromise. Or, more specifically, it is about how politicians manipulate information to forge compromises—how they figure out what is possible, how and when they coordinate around a particular political possibility, and how they influence the public record to highlight, or obscure, their roles in striking political bargains. This study focuses on political compromises constructed in the context of lawmaking in the European Union (EU) and explores how political elites—most notably members of the European Parliament (MEPs)—coordinate around particular policy proposals when crafting European legislation. A primary goal of this project is to improve our understanding of how day-to-day policymaking gets done in the EU, and to highlight the centrality of information transmission in this process.

The popular history of the EU is largely written in terms of large-scale negotiations between member states. The introduction of the Euro, treaty revisions like those accomplished at Maastricht, Amsterdam, and Lisbon, and the difficulty surrounding the ratification of the ill-fated constitution, all captured the attention of the world, or at least of Europe. Furthermore, much of the literature on the EU—starting with the foundational theories of neofunctionalism (Haas 1958, Lindberg 1963) and intergovernmentalism (Hoffman 1966, Taylor 1982) and extending to their modern incarnations in liberal intergovernmentalism (Moravcsik 1993, Moravcsik 1998) and theories of supranational governance (Stone Sweet & Sandholtz 1998)—treats the EU, first and foremost, as an international regime evolving from treaty to treaty, focusing scholars' attention on the broad sweep of European integration. But, as Thomson, Stockman, Achen & König (2006) argue, everyday

politics in the EU are fundamentally important, and regularly affect the lives of millions of Europeans. Legislation resulting from the sorts of political compromises considered in this project regulate everything from cross-border trade in services, to the security procedures governing aviation in the Union, and even the particulars of windscreen defrosting and demisting systems in motor vehicles. In short, everyday EU lawmaking is affecting the daily lives of all Europeans. Just as national legislatures routinely pass laws that matter to their constituents, the EU has for some time now enacted directives and regulations—as a matter of standard operating procedure, and without recourse to lengthy intergovernmental negotiations—that affect millions.

Indeed, the reality of routine policy-making in modern Europe has led to an explosion in work that sees the Union not simply as a traditional international organization, dominated by high-level interstate bargaining, but as a policymaking institution that is amenable to analysis with general tools developed primarily to help explain legislating at the national level (Hix 2005). Although certain policy domains—such as citizenship, social security, tax harmonization, and common foreign and security policy—remain exclusively subject to intergovernmental bargaining between the 27 EU member states, policy creation in the modern EU is, largely, an exercise in bicameral lawmaking between the Council of the European Union (Council) and the European Parliament (EP). Thus while many European lawmakers—namely the ministers serving as agents of their national governments in the Council—represent the preferences of entire states (or at least their governments), those preferences are channeled through institutions that are largely analogous to those responsible for lawmaking in nations around the world. Therefore, the ways in which information flows affect the ability of European politicians to coordinate around policy compromises provide lessons that are potentially applicable to similar political systems; notably, polities sporting bicameral legislatures.

Precisely because they serve such a varied constituency, European institutions—like most federal structures—prioritize minority protection over agile lawmaking. As the Union has

evolved into a truly bicameral system, with legislative power distributed reasonably equally between the Council and the Parliament, the pace of lawmaking has, understandably, slowed (Golub 2002, Schulz & König 2000, Golub 2007, König 2007, König 2008). Starting with Madison and Montesquieu, theorists have traditionally understood bicameralism as a way to improve governance by forcing legislation to clear multiple, asymmetric, hurdles on its path to becoming policy (Heller 2007). Modern versions of this thesis are built largely on the idea that bicameralism adds an additional veto player to the mix (Tsebelis 2002), biasing outcomes towards the status quo, but authors nonetheless render such arguments in a positive light, pointing out bicameralism's tendency to generate moderate, or Pareto efficient, policies (Hammond & Miller 1987, Tsebelis & Money 1997). Conversely, one can depict bicameral caution, not as a moderating influence on potentially unfettered policymaking, but as a cause of legislative gridlock (Alt & Lowry 1994) and inefficient cross-chamber logrolling (Heller 1997).

Information management is crucial to the ability of the Union to overcome an institutional bias towards glacial lawmaking so that it may quickly react to policy priorities and efficiently produce legislation to meet Europeans' diverse needs. During bicameral negotiations the representatives of two legislative chambers must balance the priorities of both chambers to hammer out a text that is acceptable to all. At the same time, the proxies of each house must strive to protect those aspects of the bill that are most important to their own legislative body, political parties, and constituencies, and each set of negotiators will prefer to make compromises only when absolutely necessary to placate voters in the other house. Yet this process plays out in a complicated informational environment. On the one hand, bicameralism may serve to overcome informational hurdles and produce better policy by bringing lawmakers with access to varying sources of knowledge to the table (Rogers 2001). On the other hand, information asymmetry across legislative chambers can exacerbate bicameralism's tendency towards delay in policymaking (Tsebelis & Money 1997). Legislative houses are not unified actors but collections of individual legislators representing varying

constituencies and sometimes belonging to differing political parties. These legislators differ not only in their policy preferences, but in their priorities, and in their receptiveness to compromise on particular topics. It may be very difficult for negotiators to predict how the individual houses will react to specific compromise language, and it may be especially difficult for representatives from one house to predict the positions of members in the other chamber.

An influential strand of literature on everyday decision-making in the EU uses game theory to examine how the Union's rules help to determine the relative policy influence wielded by the EU's major lawmaking institutions, the European Commission, Council, and Parliament (see e.g. Steunenberg 1994, Garrett 1995, Garrett & Tsebelis 1996, Tsebelis 1997, Crombez 1997, Tsebelis & Garrett 2000, Schackleton 2000, Crombez 2001, Thomson et al. 2006). This literature, which is based predominantly on procedural spatial models of legislative bargaining, has highlighted the growth of the EP into a full-fledged co-legislator of the Council, documents the Commission's transition from a powerful agenda setter and veto player into a—largely bureaucratic—supporting player, and provides the underpinnings for recent work conceptualizing the EU as a bicameral system (Hagemann & Høyland 2010). The application of rational choice institutionalism to the EU has been instrumental in producing clear and testable hypotheses about who should be influential in European lawmaking, and what sorts of policies should survive the Union's legislative process, but it has been criticized for the relative simplicity of its assumptions (Hörl, Warntjen & Wonka 2005). Furthermore, its focus on de jure institutions paints a picture of European politics that is often at odds with the on-the-ground impressions of practitioners and scholars alike (Burns 2004, Thomson & Hosli 2006).

Taking these criticisms seriously, I build upon the foundations laid by current formal theories of European lawmaking to construct a model that addresses questions about how European actors cope with their complex institutional, and informational, environment. I open my examination of how EU politicians coordinate around bicameral bargains in chapter

2, highlighting one mechanism by which MEPs may identify acceptable policy compromises when haggling over legislation with the Council. I argue that, while the Commission no longer holds strong formal powers over the majority of EU legislation, it serves as an informational link in the bargaining game between the Council and the Parliament. The model relaxes a standard, but untenable, assumption commonly held in the existing literature.

Specifically, rather than treating EU policymakers as perfectly informed, I focus on the fact that information asymmetries across the Council and Parliament can lead to costly bargaining delay. More specifically, because the Council conducts most of its internal negotiations in secret, while the Parliament is a model of institutional transparency, it is easier for Council ministers to gauge the Parliament's bargaining position than it is for MEPs to measure resolve within the Council. Therefore, the Council has a distinct advantage in inter-chamber negotiations with the Parliament because its bargaining strength—the cohesiveness of the internal coalition supporting its position and the willingness of its members to incur delay costs—is obscured, while Parliament's weaknesses and divisions are on full display. But, because the Commission—in its capacity as the bureaucratic engine of the Union—has privileged access to internal Council negotiations, it can uncover and expose Council weakness, and strength, to the Parliament. The Council, on the other hand, cannot credibly signal its resolve to the Parliament because its protestations of strength will be seen as cheap talk (Crawford & Sobel 1982, Austen-Smith & Banks 2000).

This story is, at its heart, about how politicians figure out what bargains are possible and it explores how legislators decide to settle on particular compromises. If the Commission were not in the picture, MEPs would be forced to court delay to sound out the depths of the Council's resolve. Balancing its own appetite for extended negotiation, the Parliament would open inter-chamber negotiations with a hard line, slowly moderating its position to zero in on a compromise acceptable to both parties. The Commission can short-circuit this process, modulating the pace of policymaking. Of course, Commissioners have their own goals, and will use this power strategically. Indeed, the model shows that the Commission can only

credibly transmit information when unbiased, or when its preferences coincide with the Parliament's. Thus, I treat the Commission as a strategic mediator, importing techniques from the study of interstate conflict resolution (Kydd 2003, Rauchhaus 2006) to a complex institutional environment not traditionally found in international relations. When the Commission is willing to spill the beans, it acts as a coordinator of comprises, speedily guiding the Council and Parliament to policy bargains. Therefore, the model helps to explain a perception, common among practitioners of EU politics, that the Commission exerts impressive power in the modern Union by forging agreement between the co-legislators (Rasmussen 2003, Thomson & Hosli 2006). Moreover, it helps to reconcile the stark image of EU lawmaking rendered by procedural spatial models with the rich picture of conciliatory interactions, and technocratic compromise often described—if perhaps somewhat optimistically—by policymakers in Brussels and Strasbourg.

The model adds to a growing literature that seeks to explain how European politicians have adapted to the Union's institutional complexity and daunting information environment, both by changing their patterns of internal organization and by adopting new modes of communication across institutions (Farrell & Hèritier 2004, Høyland 2006). Moreover, it extends our understanding of how the preferences of actors within the key EU institutions may interact to regulate the pace of lawmaking in Europe. Indeed, it shows that the relationship between the Commission's policy priorities, and those of the other EU institutions, should continue to affect the rate of policy production in the Union, even in the current bicameral era. Similarly, it describes an important mechanism through which the rules governing the selection of European Commissioners may influence Union policies. More generally, the European example demonstrates how informational asymmetries can interact with institutional factors to determine the likelihood of gridlock in bicameral legislatures and highlights the role that bureaucratic actors can play in bridging knowledge gaps between policymakers.

Finding bicameral bargains does not simply require cross-chamber coordination. Actors within each of the EU institutions must also coordinate with one another to support a

particular bargaining position in multi-institutional negotiations. Indeed, another common criticism of formal models of EU policymaking is that they tend to treat the three European institutions as unitary actors, rather than exploring how each policy organ channels its members' preferences. The model that I present in chapter 2, while relaxing common informational assumptions, does little to address this critique. Chapter 3, on the other hand, investigates how MEPs react to policy recommendations from the Commission, and balance their ideological preferences against the need to coordinate around particular compromises in order to create public policy.

A primary focus of this project is that politicians need tools to help them to coordinate around particular policies in order to legislate efficiently. Furthermore, chapter 2 highlights the Commission's ability to provide information that can help MEPs to quickly coalesce around positions that will reduce the time and effort necessary to produce European public policy. Yet, for this coordinating device to work, MEPs must be willing to moderate their own positions and support second-best outcomes, rejecting their immediate preferences in favor of proposals that are likely to produce viable bicameral bargains. In other words, MEPs need to compromise their own ideological leanings and those of their constituents—be they national party leaderships or voters in their home states—in order to effectively engage with the Council in policy debates.

The European treaties provide a highly institutionalized mechanism for the Commission to transmit coordinating information to the Parliament. Specifically, before the Parliament votes on legislation, the Commission provides MEPs with official opinions on each of the amendments that MEPs have tabled and will consider in plenary. By examining how MEPs react to these recommendations, I am able to delve into mechanisms that are necessary to support the inter-institutional dynamic postulated by chapter 2's formal model and investigate whether or not MEPs actually alter their legislative behavior in the face of expected bargaining constraints. Chapter 3 therefore asks the question, do MEPs alter their voting behavior in anticipation of bicameral bargaining and do they incorporate information

transmitted by the Commission into their voting calculus?

To answer these questions I take advantage of the variety of contexts within which MEPs vote in the legislature. Specifically, while MEPs spend much of their voting time dealing with actual legislation that is subject to negotiations with the Council and for which the Commission renders opinions, they also vote on a variety of internal non-legislative resolutions that do not directly concern the Council and upon which the Commission does not comment. I leverage the extensibility of the Bayesian statistical roll call voting model (Clinton, Jackman & Rivers 2004) to exploit this natural experiment and identify the extent to which MEPs alter their votes in the face of Commission messages and expected Council responses. I build a statistical framework that assumes that MEPs are utility maximizers that balance ideological concerns against pressure to quickly generate policy and use it to identify which MEPs incorporate coordinating information into their voting decisions and break from their typical voting patterns in light of bargaining constraints.

Using these tools, I find that a subset of MEPs behave differently on bicameral legislation than on other measures and that they vote in a way that varies systematically with Commission recommendations. Specifically, MEPs from national governing parties, and who have ideological predispositions that do not completely preclude compromise, moderate their positions in face of coordinating information. On the other hand, MEPs from opposition parties either ignore Commission recommendations or actively agitate against them, sometimes even sacrificing their ideological purity to stand apart from the inter-institutional consensus. Moreover, Commission opinions appear to play a crucial role in this story, providing information that does not simply mirror MEPs' expectations about Council positions.

These findings paint a micro-level picture of how legislators within one house coordinate around compromises when engaging in bicameral bargaining. Furthermore, they provide a link between the high-level story of inter-institutional information transmission, told by chapter 2, and individual decisions with the EP. Indeed, these findings should force researchers to think carefully about how voting coalitions form in the Parliament and, more

generally, have implications for our understanding of the underlying causes of legislative voting decisions across bicameral systems. Moreover, the theoretical and statistical models that I develop to examine this story represent a useful for framework for understanding previous results describing a government-opposition "dimension" to voting in the EP (Hix, Noury & Roland 2007). Of course, the statistical tools that I describe here also have potential applications in other contexts, and can help to elucidate a wide range of external influences on legislative voting behavior. Similarly, they speak to existing tools that use voting patterns to measure legislative ideology (e.g. Poole & Rosenthal 1985, Clinton, Jackman & Rivers 2004, Poole 2005) and provide an example of one way to disentangle the variety of factors that drive voting behavior (Hall 1992).

Forging compromises among political elites is one thing; selling these bargains to constituents is quite another. Furthermore, while party heads—and in the case of the EP, party group leaderships—may often wish to build durable coalitions around viable bicameral bargains, it may be difficult to bring the rank and file into line. In parliaments that record votes, legislators' policy compromises become an indelible artifact of the public record. This reality can constrain both legislators' ability to sacrifice principles in the name of expediency and to escape detection when they ignore their political masters' attempts to forge voting coalitions. Chapter 4 extends my discussion of the role of information in compromise coordination in the EU to examine one way in which MEPs massage the public record of their behavior in the Parliament. Specifically, the chapter investigates the determinants of roll call vote publication in the EP. In the process I engage an important methodological debate about roll call vote analysis that asks how strategic considerations within parliaments influence the picture of legislative behavior that is painted by the voting record (Carrubba, Gabel, Murrah, Clough, Montgomery & Schambach 2006, Clinton & Lapinski 2008).

Like many legislatures, the EP records and publishes only a subset of legislative votes, with party group leaderships requesting roll call when it suits them. Of course, this raises the specter of selection bias in studies that rely on roll call votes in the EP. And, from a

substantive point of view, it forces us to ask why politicians choose to publicize certain types of votes and not others. Using techniques imported from machine learning and computer science, I demonstrate how one can use legislative text—namely the transcripts of speeches delivered during debates on the EP floor—to predict when party groups will request roll call votes. In turn, I use measures derived from the predictive model to explore when, and why, EP party groups choose to emphasize the voting record. The technology that I develop to examine roll call votes in the EP is broadly applicable to other legislatures that publish only a portion of their voting records and can serve as a foundation for building tools for modeling selection bias in roll call votes that will help to improve the inferences we draw from cross-national studies of legislative voting (Carey 2008).

Not surprisingly, EP groups request roll on especially controversial issues and use public votes to highlight disagreements within the legislature. Most notably, party groups call roll when they are unified in opposition to a policy for which they are not responsible. This finding demonstrates one way in which roadblocks to building consensus around particular policies in the Parliament may form. Effective, and efficient, legislating—and especially bicameral bargaining—often requires all parties to moderate their positions in ways that violate ideological orthodoxy, and may elicit the wrath of external interests. By publicizing controversial votes, factions within the EP can make compromise costly for their opponents. To a lesser extent, party groups also call roll to advertise their support for legislation on which their members took active leadership roles. Thus, the overall picture that emerges is one of blaming and credit claiming through public votes. On the other hand, I find only modest support for the idea that party leaderships use public votes to enforce voting discipline among their members.

Chapter 4 emphasize the fact that patterns of parliamentary compromise are functions, not only of legislative institutions and inter-chamber information transmission, but of the structures that allow external actors to monitor legislative behavior. At the beginning of this introduction I argued that day-to-day lawmaking in Europe matters. Nonetheless, the

citizens of Europe are still largely oblivious to what goes on within the EP. Of course, interest groups, lobbyists, and national party leaderships all take an interest in what MEPs do in the legislature, providing an audience for public records of MEP behavior, such as roll call records. In its quest to reduce its "democratic deficit," the EU has taken great strides to promote transparency in lawmaking and strives to increase public interest in European policymaking. While institutions play a pivotal role in determining the form of EU policy, and the efficiency with which it produces legislative product, long-term changes in the external information environment also have the potential to factor prominently in how European politicians coordinate around compromise, by altering the patterns of pressure that drive politicians' decisions.

Chapter 2

Bureaucratic Mediation and European Lawmaking

We have long understood that bureaucrats can exert substantial influence over both the content and implementation of laws in democratic political systems. Similarly, just as governments represent their states' interests within inter-governmental organizations (IGOs) in a manner that is often largely analogous to the way in which lawmakers represent constituents within legislatures, so too can the bureaucrats that populate such organizations influence supranational policy creation and application. Thus, while elected officials maintain monopolies on lawmaking in most democracies and states' representatives hold the reigns of IGOs, appointed civil servants, nonetheless, may substantially determine the policies pursued by national governments and international organizations.

On the one hand, because bureaucracies house stores of policy expertise that are often unrivaled among elected officials and interest groups, politicians are commonly forced to rely on bureaucratic recommendations when formulating policy. Bureaucrats, therefore, may take advantage of their informational assets to ensure that policy-makers' understanding of issues leads them to make choices that the relevant civil servants—the staff of a government ministry, for example, or perhaps the permanent employees of an IGO—prefer to other possible outcomes. Furthermore, bureaucracies are often largely responsible for implementing policy, providing appointed civil servants with substantial leeway in interpreting exactly what laws and international agreements mean in practice. Indeed, a large literature examines the circumstances under which national government agencies may subvert the policy intentions of elected officials (Brehm & Gates 1997) and investigates how lawmakers tailor their legislation—and bureaucratic oversight mechanisms—to channel civil

servants' discretionary tendencies to their best advantage (Epstein & O'Halloran 1999, Huber & Shipan 2002). A similar literature explores analogous questions within international organizations, using principal-agent models to describe the extent to which states can maintain control over international institutions and to probe the circumstances under which international bureaucracies may take advantage of delegated powers to exert influence over international policy-making and implementation (Hawkins, Lake, Nielson & Tierney 2006).

This essay discusses another avenue of bureaucratic influence, highlighting the role that government agencies and similar institutions may play as informational conduits linking groups of policy-makers. Specifically, I argue that appointed civil servants may influence policy-outcomes—and, furthermore, the dynamics of policy-making—by acting as mediators between those actors—usually elected politicians—with the power to choose policy. Indeed, because bureaucrats possess information and skills that are indispensable to policy-makers, they often engage directly with all of the actors that have a say in a given policy-making institution. For example, in separation-of-powers systems where both the executive and legislature influence policy—or in a bicameral system where two legislative houses must both approve legislation—it is common for civil servants from government ministries, owing to their singular access to expertise on a given issue, to consult directly with all the relevant decision-makers. Similarly, the secretariat of an IGO will typically make itself available to all the member states of the organization. Even in unitary political systems, where a single lawmaking body decides policy, bureaucrats will tend to interact with all those factions that have a agenda or veto power within the policy-making process. These contacts represent both a way for policy-makers to learn about the issue at hand and an opportunity for bureaucrats to obtain unique insights into the preferences, and bargaining positions, of policy-makers. I argue that bureaucrats can leverage their access, strategically revealing information about policy-makers to one another. Thus, civil servants are often in a position to adopt a mediating role, acting much like a third party arbiter in other bargaining environments. Furthermore, under some circumstances, bureaucrats can take advantage of their

mediating capacity to guide policy-making in a direction that they prefer. Moreover, by regulating the information environment within which policy-makers negotiate, bureaucrats may sometimes modulate the rate at which policy-making gets done, and therefore may fundamentally alter the efficiency of a given political institution.

I explore the circumstances under which bureaucratic mediation may influence policy-making outcomes and efficiency[1] in the context of lawmaking in the European Union (EU). In particular, I develop a model of lawmaking under the codecision procedure in the EU that illuminates the mediating role that the European Commission plays in inter-institutional bargaining between the two house of the EU's bicameral legislature, the Council of the European Union and the European Parliament. Building on a standard sequential bargaining model (Fudenberg & Tirole 1983, Sobel & Takahashi 1983), I argue that the Commission can influence European lawmaking by relaying information about the Council's bargaining resolve to the Parliament, but only when the Commission and the Parliament share similar preferences. While the Parliament is a largely transparent body, the Council makes most of its internal decisions in secret and it is difficult for other actors to gauge the preferences of its members or the strength of its bargaining resolve. This asymmetry in information disadvantages the Parliament and has the potential to hamper lawmaking efficiency. Indeed, the model shows that, lacking other sources of information, the Parliament has incentives to float aggressive early offers to evaluate the Council's bargaining strength, only moderating its position if the Council rebuffs its initial proposal. The Commission can sometimes play an important role in mitigating this asymmetry, relaying information to the Parliament that allows it to forgo a costly search for compromise.

Specifically, in line with the above discussion, Commission representatives sit in on all levels of internal negotiations within the Council (Cini 1996, Nugent 2001), yielding the Commission particular insight into the Council's negotiating position. The model shows that,

[1]Here, I use the term (in)efficiency to refer to the degree to which policy makers exert effort and time when bargaining over policy. These are dead-weight costs that actors can avoid, in principle, when they perfectly understand each others' negotiating positions.

when the Commission's interests and the Parliament's coincide, the Commission has incentives to relay this information to the Parliament. On the other hand, when the two actors disagree, the Parliament has reason to be wary of the Commission's information. Furthermore, the Commission's ability to signal truthfully interacts with the Council's bargaining strength to determine lawmaking duration. When the Council holds a weak bargaining position—when its procedural costs to delay are large, it places little importance on the manner at hand, or when disagreement within the Council hampers its ability to maintain a hard-line position—lawmaking will take little time. On the other hand, when the Council's delay costs are low and its bargaining position is strong, delay may be substantial; but this inefficiency is mitigated when the Parliament trusts the Commission's estimate of the Council's bargaining strength. Thus, the Commission plays a pivotal role in regulating the EU's legislative efficiency. Not surprisingly, the Parliament also achieves better bargains, on average, when the Commission provides it with accurate information about the Council's bargaining strength. Therefore, while it has no direct say over the policies decided under codecision, the Commission may collude with the Parliament to push outcomes in its preferred direction.

The model that I present here bridges a gap between theoretical accounts of codecision lawmaking in the EU that treat the Commission as a bit part player in codecision (Tsebelis & Garrett 2000, Crombez 2001) and empirical research that ascribes substantial lawmaking influence to the European bureaucracy, even after the wide-spread application of codecision to Union decisions (Rasmussen 2003, Burns 2004). It demonstrates the Commission's ability to wield influence in the absence of strong formal agenda-setting or veto powers, but also helps to delineate the circumstances under which it may be influential. Furthermore, it provides a prime example of the mediation game that can arise between policy-makers and bureaucrats, and exposes an important relationship between bureaucratic bias and policy-making efficiency. Indeed, because the EU straddles the line between a traditional IGO and an out-and-out government, this case highlights the wide array of situations in which

such a dynamic may take hold. Moreover, the relationship between legislative efficiency and bureaucratic mediation that I discuss here is, in principle, quite general. Indeed, this mechanism may help explain why political institutions that look similarly prone to deadlock and delay—perhaps two states with comparable bicameral legislative setups and similar political cleavages—may behave quite differently in practice, expending varying degrees of time and effort to produce effective policy. Whenever veto-power over policy is spread across institutional actors, how the biases of bureaucratic mediators interact with what each player knows about the others' bargaining strengths can be crucial.

2.1 Codecision: Power, Information, and Efficiency

Lawmaking in the European Union is dominated by three bodies, the Council of the European Union, the European Commission, and the European Parliament. Prior to the Treaty of Maastricht in 1993, the Parliament played third fiddle to the Commission—effectively the government executive and bureaucratic arm of the Union—and the Council,[2] which is composed of the ministers of the Union's member states.[3] Maastricht ushered in the codecision procedure, and with it, a new era of Parliamentary ascendency in the EU (Crombez 1997, Crombez, Steunenberg & Corbett 2000, Tsebelis & Garrett 2000, Schackleton 2000). The Treaty of Amsterdam in 1999 simplified the procedure somewhat, and expanded the scope of codecision to cover over half of EU legislation. Counter-balancing the wide extension of this complicated—and often very lengthy procedure—the Treaty also made it possible for the Council and Parliament to conclude bargaining after their first reading of a Commission proposal, should the Council agree with the Parliament's initial opinion on the legislation. 2001's Treaty of Nice did little to change codecision beyond tinkering with the weighting of states' votes in the Council and expanding its purview to yet a wider swath of legislative

[2]The exact composition of the Council varies depending on the topic at hand: ministers represent their nations in the Council with respect to their portfolios.

[3]But see Tsebelis (1994) for an argument about the Parliament's conditional agenda setting power at this time.

topics (Galloway 2001). Similarly, provisions of the Lisbon Treaty that went into effect on December 1, 2009 extended codecision's reach still further to cover all but a few exclusive areas of policy-making, re-naming codecision the "ordinary legislative procedure" in the process, while doing little to alter the function and form of the institution. Thus, the Amsterdam incarnation of codecision remains largely intact to this day. And, while applicable to the majority of EU legislation since the turn of the century, with the expansions of Nice and Lisbon, codecision now truly represents the standard lawmaking procedure in the Union. Therefore, to explore the dynamics of law-making in modern Europe one must necessarily engage the—often arcane—details of this complicated procedure.

The literature on lawmaking in the EU, and on codecision more specifically, has largely focused on the relative power that the Commission, Council, and Parliament have to obtain policy outcomes that they, or their pivotal members, most prefer when crafting legislation. This literature largely revolves around a series of procedural spatial models, based on non-cooperative game theory, that describe the ability of the EU institutions to affect policy choices as a function of institutional preferences and the decision-making rules embedded in the Union's various lawmaking procedures.[4] Initially, arguments revolved around whether or not codecision advantaged the Parliament relative to the earlier consultation and cooperation procedures. Indeed, one stated goal of the Maastricht treaty was to improve the Union's democratic legitimacy by putting the directly elected Parliament on more even institutional footing with the indirectly elected Council, and especially, the appointed Commission (Crombez 2001). Furthermore the Parliament itself maintained that the Maastricht version of codecision failed to put it on equal legislative ground with the Council because the Council could act unilaterally should bargaining between the two bodies break down during the penultimate stage of the procedure (European Parliament 1992). The opinion that codecision failed to increase the Parliament's powers was initially maintained by

[4]See Hörl, Warntjen & Wonka (2005) for a critical review of this literature. There is also a tradition of using cooperative bargaining models to explain European lawmaking; Boekhoorn, Van Deemen & Hosli (2006) is a recent example of work in this mode.

the academic literature (Curtin 1993, Steunenberg 1994). Furthermore, some authors argued that codecision actually undermined the conditional agenda-setting prerogatives that the Parliament enjoyed under the older cooperation procedure (Garrett 1995, Garrett & Tsebelis 1996, Tsebelis 1997). These arguments rested largely on analyses of the procedure that focused only on the final few stages of the legislative game and Crombez (1997) showed that, if one analyzed the game from beginning to end, one could show that Maastricht substantially improved the Parliament's ability to obtain its preferred policy outcomes with respect to previous decision-making procedures. Similarly, considering the final steps of the procedure, Steunenberg (2000) argued that the Amsterdam version of codecision would drive policy outcomes towards the Parliament's preferred point, weakening both the Commission and the Council. Furthermore, the argument that the Parliament has enjoyed enhanced policy-making power under codecision has found substantial empirical support (Kasack 2004, König & Pöter 2001, Selck & Steunenberg 2004, Steunenberg & Selck 2006, Thomson & Hosli 2006).

One issue on which commentators have generally agreed is that codecision weakens the European Commission relative to previous procedures. This argument rests on the fact that the Commission has no formal proposal or veto powers in the final stages of codecision. Specifically, while the Commission initiates legislation under the procedure, plays an informal role throughout deliberations, and has some formal power to influence the rules under which the Council takes decisions in the earlier stages of the process, the Parliament and Council may hold off on making a decision until the final possible stage of codecision—dubbed conciliation—where the two co-legislators haggle over a joint text in a setting that closely resembles conference committee in the US Congress (European Commission 2010). Furthermore, both institutions may amend the Commission's initial proposal, under an open rule, in the stages preceding conciliation. Therefore, given sufficient information and patience, the Parliament and the Council can anticipate the conciliation outcome and refuse to accept any early agreement that is not as good for both of them as the expected outcome of

the entire process. As Garrett (1995, pp. 305) maintains, "...under co-decision, the Commission is effectively taken out of the game before the real bargaining over policy begins." Crombez (2001, pp. 101) modifies this argument somewhat, drawing a distinction between the Maastricht and Amsterdam versions of the institution. He argues that, while the original Maastricht version of codecision actually provided the Commission with considerable agenda-setting power, Amsterdam "...renders the Commission irrelevant." In sum, under the modern codecision procedure, the Commission relinquishes its former role as a first-class policy-maker with biting agenda-setting and veto rights; rather, its involvement in the process is predominantly bureaucratic and its role is largely analogous to that of a government ministry in national politics. That is, while it is active in every part of the legislative process, the Commission provides logistical and technical support to decision-makers, rather than rendering decisions itself.

Nonetheless, there is an enduring perception, and some empirical evidence, that the Commission can be influential in post-Amsterdam codecision (Rasmussen 2003, Burns 2004).[5] For example, Thomson & Hosli (2006) conducted a survey of "practitioners of European affairs," that asked officials from the Commission and Council secretariat to rate the relative influence of the European institutions—the Commission, Council, and Parliament—over a variety of issue areas under both the consultation and codecision procedures. On average, the practitioners rated the Commission's influence under codecision as virtually equal to the Council's and above the Parliament's. In fact, the respondents deemed the Commission more influential during codecision than when operating under the consultation procedure, even though consultation provides greater formal power to the Commission than codecision.[6] Furthermore, many respondents cited the Commission's skill at "...forging political deals

[5]Additionally, while Kasack (2004) concludes that the Commission's codecision influence declined after Amsterdam at second reading, she nonetheless finds that Commission support for Parliament's amendments strongly predicts their likelihood of adoption by the Council, a finding that is at least consistent with continued Commission influence.

[6]I should note that Thomson & Hosli's (2006) findings in the same paper—based on a quantitative analysis of expert ratings of European lawmaking outcomes and institutional preferences—are at odds with practitioners' perceptions.

between other actors, both within the Council and between the Council and EP," when justifying their judgements (Thomson & Hosli 2006, pp. 398). In other words, respondents identified *mediation* as key to the Commission's perceived influence. Similarly, in interviews with a variety of EU officials, Rasmussen (2003, pp. 4) finds strong support for the picture of the Commission as a mediator, or forger of compromises. In the words of one respondent, "Sometimes the positions of the Council and Parliament are so far apart that only the Commission can find common ground." Thus, as Tsebelis & Garrett (2000, pp. 26) conclude, with the advent of codecision, "[t]he remaining influence of the Commission over legislation is thus likely to rely more on informal channels—asymmetries of information, persuasion, deal-brokering—than on formal roles written into various procedures."

In fact, the Commission is well-situated to use its informational assets to its advantage, adopting a mediating role. Unlike members of the Parliament, who are not privy to the internal discussions of the Council, Commission representatives are directly involved in internal Council deliberations. Indeed, from the time that the Commission transmits its initial proposal to the Council and Parliament at first reading, Commission staff sit in on all intra-Council negotiations, ranging from working group discussions and meetings of the Council's committee of permanent representatives (COREPER) to ministerial meetings (Nugent 2001).[7] The Commission also maintains contacts with ministerial staff in the member states themselves (Cini 1996). Similarly, Commission representatives sit in on Parliamentary committee meetings, and maintain close contact with rapporteurs. As Nugent (2001, pp. 253) remarks, "[The Commission] thus has excellent knowledge of what the EP and Council want and may settle for. This knowledge can be used, often in an informal manner, to promote and broker compromises and settlements on difficult points." The Commission's unique access to the range of decisive policy-makers in the Union puts it in a

[7]Ministerial level meetings of the Council dealing with codecision are now open to the public but are, largely, stale affairs that provide little information about internal Council negotiations. The Council does most of its work within COREPER meetings and working groups; these sessions are held behind closed doors.

position to play an influential mediating role.

Policy-making is, at its core, about bargaining; agents with veto power over policies negotiate with one another to find outcomes that all decisive actors can agree upon. While it may be relatively easy for the relevant factions to find a set of policies that every veto-wielding policy-maker prefers to the status quo in a given issue area,[8] policy-makers may, nonetheless, vehemently disagree about exactly what policy to pursue from the alternatives within this set. Exactly what policy prevails, therefore, will rest heavily on the bargaining positions of the decisive policy-makers. Those veto-players with greater bargaining resolve—those policy-makers who are more willing, or able, to drag out negotiations to obtain their most preferred policy, or who find the issue under consideration particularly salient—will often be able to extract concessions from other policy-makers (Muthoo 1999). But, of course, the difficulty inherent in evaluating opponents' resolve is exactly what makes bargaining interesting, and negotiations would end instantly if every party at the table knew the resolve of every other bargainer (Rubenstein 1982). Furthermore, knowing that the most resolute bargainer generally wins the day, each negotiator will do its utmost to project an aura of infinite resolve, hoping to convince its counterparts that capitulation is the least costly course of action.

In the context of codecision, this means that, whenever there is substantial disagreement between the pivotal coalition in the Council and the decisive voter in the Parliament, the nature of the final policy outcome will rest, largely, on the relative bargaining strength of these pivotal actors. Amsterdam made it possible for codecision to conclude at first reading and, in response, the Council and Parliament have greatly increased their inter-institutional contacts during the early stages of codecision, in the form of informal "trilogues" between representatives of the three core EU institutions. As others have pointed out, these early-stage points of contact have altered the dynamics of codecision since Amsterdam (Farrell & Hèritier 2004). Indeed, as figure 2.1 shows, codecision files increasingly end early on in

[8]Indeed, if they cannot, then policy movement on the issue at hand will be very unlikely.

Figure 2.1: Codecision since Amsterdam.

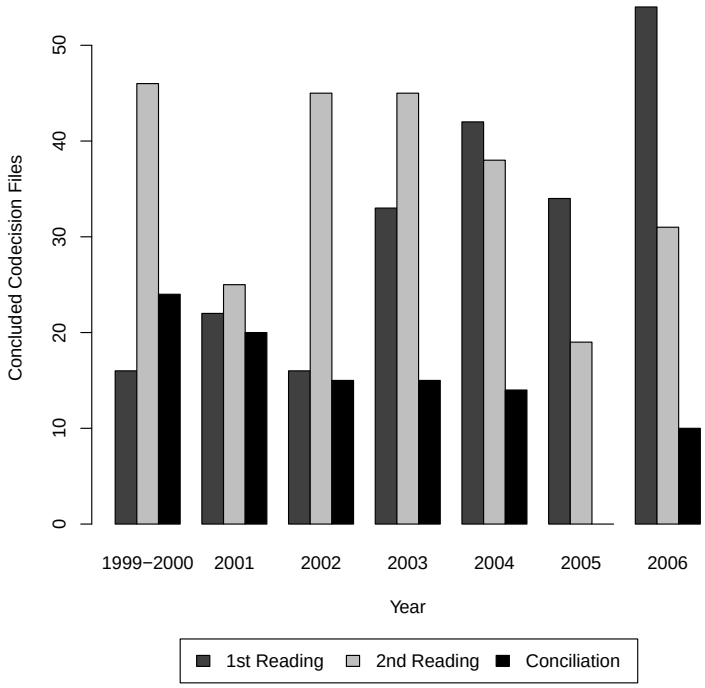

the process and first-reading agreements have become the norm. Nonetheless, the legislative bodies have little incentive to reveal their bargaining strengths to one another during informal negotiations, when their preferences diverge. These exchanges are a form of cheap talk, where the Council and Parliament send costless messages to one another about their positions, the importance they place on the issue at hand, and their bargaining resolve. The core result of models of cheap talk is that such communication is informative only when the relevant parties share similar preferences (Crawford & Sobel 1982, Austen-Smith & Banks 2000).

Therefore, while first reading trilogues may help EU institutions to resolve bargaining quickly on "easy" cases where inter-institutional preferences largely match, they are unlikely to ameliorate informational asymmetries when the negotiators' preferences diverge. Instead, each side will want to talk tough and will largely disregard the other side's protestations about an inability to compromise.[9] Under such circumstances, bargaining in bicameral legislatures can result in costly delay (Fukumoto 2008). It is exactly in such conditions that a well-connected bureaucracy can leverage its role as a mediator to influence outcomes and potentially improve policy-making efficiency. While the mediator's communications to policy-makers are also a form of cheap talk, the mediator's biases may cut across those of the pivotal policy-makers, sometimes making effective communication possible. Exactly when effective mediation occurs is a function of the negotiating parties' preferences, the bias—if any—of the mediator, and the institutional rules governing bargaining. Thus, it is necessary to carefully model the bargaining environment when evaluating a given bureaucracies' mediative capacity. The model that I present here does this for the case of the Commission and codecision.

Below, I first provide a thorough description of codecision, for the uninitiated. I then introduce a model that that describes sequential bargaining with asymmetric information

[9]Of course, this assumes that the pivotal actors on both sides are represented by agents that share their preferences in these early negotiations. Høyland (2006) develops a model that examines how the party of the Parliament's rapporteur affects codecision dynamics and integrating rapporteur influence into the information game that I present in this paper would be an interesting avenue for future research.

between the Council and Parliament during codecision and derive its equilibrium properties, notably a tendency towards inefficient delay. I then augment the model, demonstrating how the addition of a mediator—in the form of the Commission—changes the equilibrium policy outcomes, and expected duration, of codecision negotiations. I conclude with a discussion of the empirical predictions of the model.

2.2 How Codecision Works

Figure 2.2 provides a graphical depiction of lawmaking under codecision. Codecision is a complicated, and potentially lengthy, process.[10] The Commission holds the right to initiate legislation under the procedure and the process begins when the Commission forwards its initial proposal to the Parliament for a first reading.[11] Within the Parliament, the committee responsible for the proposal, party groups, and any group of 40 or more MEPs may table amendments to the legislation. Next, during the debate on the bill that precedes voting in the Parliament, the responsible Commissioner delivers the Commission's opinions on any tabled amendments, along with explanations of the Commission's positions. Subsequently, the EP holds floor votes on each amendment and the whole bill as amended, all under a simple majority rule. After the bill has cleared the Parliament the Commission drafts a new version of the legislation known as the "amended proposal" which may or may not include each of the amendments voted on by the parliament. Next, the legislation proceeds to the Council, which may either adopt the Parliament's proposal or draft a new version of the legislation, known as the "common position." When deciding whether or not to adopt the Parliament's proposal, the Council votes by qualified majority if each Parliamentary amendment to the initial proposal also features in the Commission's amended proposal, and by unanimity

[10]See European Commission (2010) for an exhaustive explanation of the current version of the ordinary legislative procedure (codecision).

[11]To foster efficiency, the Commission actually forwards its initial proposal the Parliament and Council simultaneously. Nonetheless, the Council must wait for the Parliament to table a first reading proposal, and for the Commission to render its opinion on that proposal, before formally acting in the first reading.

Figure 2.2: The Codecision procedure.

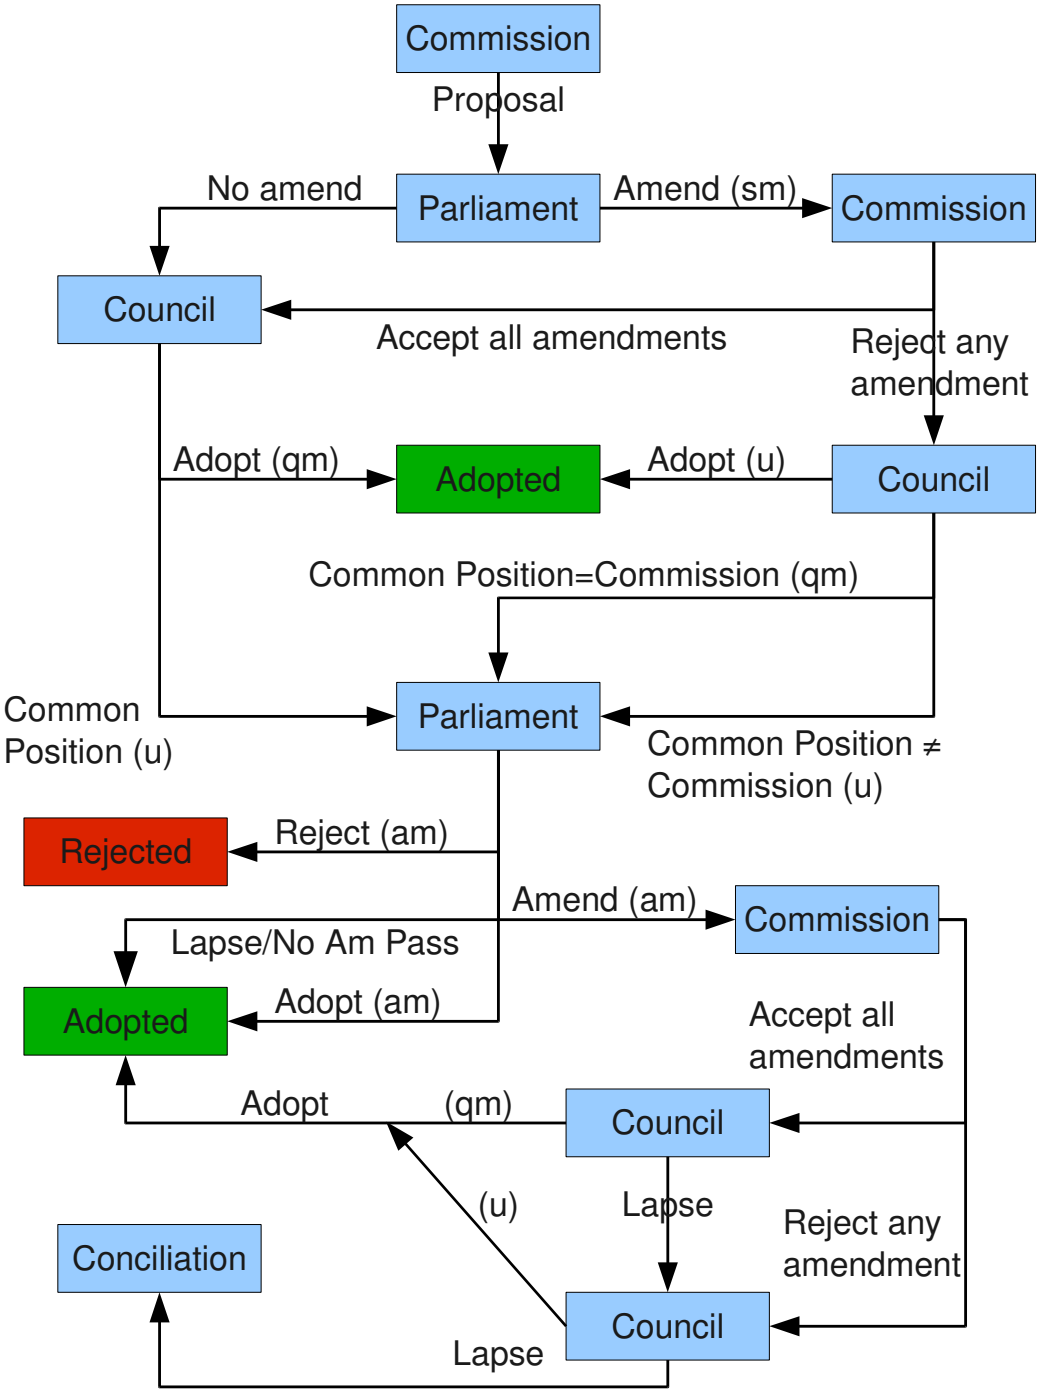

if the Parliament's proposal contains language that does not appear in the Commission's version.[12] Thus, the Commission sets the threshold for adoption of Parliament's proposals. If the Council accepts every provision in the EP text, and successfully votes to adopt the proposal, the process ends and the proposal becomes law. Otherwise the Council must propose a common position containing any mix of Commission text, EP amendments, and new Council-generated amendments. Again, the Commission sets the threshold for adoption of the Council's counter-proposal. The Council may only adopt a common position that differs from the Commission's amended proposal if it garners unanimous support from the member states; when the Council's common position matches the amended proposal it may adopt the position by qualified majority.[13]

The second reading starts largely like the first, with the EP first considering and tabling amendments to the common position, the Commission next giving its opinion on each amendment prior to the EP vote on the amendments and bill as a whole, and the Council subsequently voting on each aspect of the EP text. Nonetheless, there are a number of differences between the two stages. First, the EP conducts all votes under an absolute—rather than simple—majority rule during this reading. Second, the EP can choose to either accept or reject the common position outright and end the procedure, rather than amending the text, and indeed, the common position is adopted unless the Parliament can find an absolute majority that supports amending or rejecting the Council's text. Third, the Commission's opinions on EP amendments, as voiced by the commissioner in plenary, directly set the voting rule used by the Council when voting on the amendments, rather than indirectly through an amended Commission proposal. Finally, both the Parliament and the Council face three

[12] In a subset of issue areas, the number of which has dwindled over time, the Council must act unanimously in all codecision decisions.

[13] While the Treaty appears to give the Commission the right to force a unanimous vote at first reading, the Commission's ability to use this power is somewhat unclear in practice. Indeed, in a personal communication with the author, representatives from Europe Direct, the Commission's public information office, referred to this issue as "controversial" and refused to provide a straightforward interpretation of the Commission's powers in this context. The academic literature tends to assume that the Commission will be unwilling to exert its legal prerogative if push comes to shove, and assumes that the Council takes its decisions by qualified majority. I largely adopt this convention in this paper.

to four month time limits during the second reading.

Specifically, if the Council fails to adopt the act in the second round the conciliation committee—composed of representatives from the EP and the Council[14]—forms. This committee works to hammer out a compromise bill that is acceptable to both legislative bodies, while the Commission maintains a purely advisory role. The bill is rejected if the committee fails to reach an agreement on a joint text in a timely fashion, or if the agreed-upon text fails to garner an absolute majority of votes in the EP and a qualified majority of votes in the Council.[15]

2.3 The Model

The Parliament, P, and the Council, $C = \{\underline{C}, \overline{C}\}$, bargain to set a policy $x \in \mathbb{R}$ on some issue governed by the codecision procedure. Decisions in the Parliament are decided by (absolute) majority; therefore P represents the median voter in the Parliament, with ideal point θ_P.[16] On the other hand, the Council takes its decisions by qualified majority vote, thus both the left qualified majority pivot (q-pivot), $\underline{C}$, and right q-pivot, $\overline{C}$, where $\theta_{\underline{C}} < \theta_{\overline{C}}$, may, in general, be decisive when the Council takes a position. The status quo on the dimension of interest is q before bargaining begins. The status quo is stable when $q \in [\min\{\theta_{\underline{C}}, \theta_P\}, \max\{\theta_{\overline{C}}, \theta_P\}]$ because at least one actor will resist any change to the policy on the given dimension and bargaining is futile. Therefore, because the Commission will have no reason to waste resources by proposing policy change in such circumstances, I restrict attention to cases where q falls outside the support of the Pareto set. Furthermore,

[14]The EP delegation is composed of a representative (in terms of party group membership) of EP members, typically drawn from the committee responsible for the bill. Each Council minister is represented in the negotiations.

[15]Again, certain issue areas require unanimity in the Council.

[16]I assume full attendance in the Parliament and thus absolute and simple majority rule are identical in the model. In reality, attendance in the Parliament can be spotty and, as Hagemann & Høyland (2010) point out, this can endow the Council with a degree of conditional agenda power in the bicameral game. Incorporating this wrinkle into the model I present here would represent an interesting extension, although I leave this consideration to future research.

I assume that players have single-peaked utility functions that are strictly increasing in the proximity of their ideal points and a given policy, and that $q < \theta_{\underline{C}} < \theta_P$.[17] Given this setup, all potential bargains—or points in the core of the bargaining game—will fall on the range $[\theta_{\underline{C}}, \min\{2\theta_{\underline{C}} - q, \theta_P, \theta_{\overline{C}}\}]$. Nonetheless, after making some further assumptions, and outlining the stages of the game, I will be able to limit attention to a more restrictive set of possible proposals and outcomes.

Throughout the game, actors' choices will hinge on what bargain they believe will obtain should proceedings extend all the way into conciliation. Note, nonetheless, that predicting this outcome is difficult without precisely modeling the bargaining procedure during conciliation. Moreover, because the Treaty does not specify a restrictive bargaining protocol, the literature is divided on this question (Crombez 1997, Steunenberg & Dimitrova 1999, Tsebelis & Garrett 2000, Napel & Widgrén 2006). Rather than commit to an explicit conciliation bargaining structure, I simply assume that all players in the game share an ex ante expectation about the likely conciliation outcome, x_{cc}. Furthermore, following Napel & Widgrén (2006), I maintain that the Council q-pivot closer to the status quo, $\underline{C}$, will exert substantial influence over conciliation bargaining. More specifically, I assume that $\theta_{\underline{C}} \leq x_{cc} < \min\{\theta_P, 2\theta_{\underline{C}}\}$, as predicted by Nash's (1950) bargaining solution and Napel & Widgrén (2006).[18] Thus, in this model, I treat conciliation as a lottery over possible outcomes with expected value x_{cc}, rather than specifying strategic actions by players at the penultimate stage of codecision. Note, furthermore, that because x_{cc} is closer to all pivotal actors' ideal points than the status quo, both houses will accept the conciliation bargain when it is put to a vote at third reading.

The core Council-Parliament lawmaking game, which I outline here, but explicitly de-

[17] This ordering assumption rules out a number of symmetric cases that are qualitatively equivalent to the circumstances that I examine here, particularly cases where the Parliament's preferences are better aligned with the status quo than the most conservative member state.

[18] This logic stems from the fact that $\underline{C}$ has the least to lose from a status quo outcome, should its obstinacy lead to such an undesirable result. Furthermore, player's ideal points are common knowledge in this game and questions of bargaining patience are largely obviated in conciliation by the strict negotiating time limit specified by the Treaties.

scribe below, consists of two stages, corresponding to Codecision readings, in which the Parliament makes a proposal to the Council and the Council decides whether or not to accept the proposal or to hold out for a better outcome at each reading,[19] and where each player pays some delay cost, $0 \leq k_i \leq \theta_P - x_{cc}$,[20] when the Council rejects Parliament's proposal at a given reading, extending the game either into the second reading or triggering conciliation. I assume that, while the Parliament median's cost to delay, k_P, is common knowledge, members of Parliament have difficulty discerning how sensitive members of the Council are to delay, and thus $k_{\underline{C}}$ and $k_{\overline{C}}$ are known only to the Council members themselves.

This asymmetry in information makes good empirical sense. Parliamentary debates are conducted in public, reports written by Parliament's rapporteurs are generally available, and members' codecision voting records—both on final votes and intermediary amendments—are, largely, a matter of common knowledge. The Council, on the other hand, conducts its internal negotiations behind closed doors and publicly deliberates and votes on codecision legislation only once member states reach an agreement that all parties know will pass muster. Indeed, internal divisions in the Council are largely hidden from the public, and Parliament's, eye. Even when taking final votes, the Council is notoriously fond of maintaining a unified front. Member states rarely vote against Council positions; rather they abstain on contentious votes, and even abstentions are rare. To some extent, this tendency is a result of the member states' interest in promulgating belief in the Union's resilience, demonstrating the strength of their commitment to the Treaties, and projecting an image of European unity. It may also represent an explicit strategy by the Council to limit the Commission and Parliament's ability to exploit its internal divisions (Tsebelis 1994). Furthermore, the intra-Council bargaining procedure is a bit of a black box. As in conciliation, there are no formal institutions that restrict bargaining within the Council. Therefore, it may be difficult

[19]The model does not allow the Council to make a counter-proposal at first reading, contrary to reality. As I argue below, this restriction, while unrealistic, greatly simplifies the exposition without substantively altering the model's predictions.

[20]Thus, I assume that players' bargaining costs are non-negative and not so large that they dwarf the bargaining space.

to predict Council bargaining outcomes, even for Council members. Nonetheless, as was the case in conciliation, $\underline{C}$'s preference-proximity to the status quo gives it substantial leverage over $\overline{C}$ under common bargaining assumptions (Nash 1950, Rubenstein 1982). In fact, this dynamic is intensified here because there is no time limit to negotiations during the Council's first reading.

The above discussion implies that it is reasonable to model the Council's responses to Parliamentary proposals as if it were a unitary, yet unpredictable, actor. This is clearly the case when $\theta_{\overline{C}} < x_{cc}$. As long as Council pivots accept proposals to which they are indifferent, the Council will always accept the proposal $x = x_{cc}$ at any point in the game, knowing that the Parliament will never have any incentive to propose $x < x_{cc}$. On the other hand, $\underline{C}$ and $\overline{C}$ might be willing to accept $x > x_{cc}$ at a given point in time, depending on their delay costs, and one pivot may be more or less willing to compromise than the other. So there will, in general, be some division in the Council about the circumstances under which capitulation makes sense, but neither pivot has any incentive to help the Parliament figure out exactly where that sweet spot is and would prefer the Parliament believe that $k_{\underline{C}} = k_{\overline{C}} = 0$, thereby forcing it to propose $x = x_{cc}$ at first reading. Furthermore, the assumption that the Council qualified majority pivots both hold preferences substantially closer to the status quo than the Parliament median is common in models of EU politics (Garrett 1995, Tsebelis & Garrett 2000).[21] On the other hand, when $\theta_{\overline{C}} > x_{cc}$, $\overline{C}$'s loyalties may be mixed. Indeed, if $\theta_{\overline{C}} \geq \theta_P$, then $\overline{C}$ will wish to tell P everything it knows about the likelihood that a given Parliamentary proposal will be accepted. When $x_{cc} < \theta_{\overline{C}} < \theta_P$, $\overline{C}$ would not wish to be quite so forthcoming, but would nonetheless like to transmit some information about Council internals—information that P would find incomplete, but credible. This is a potentially interesting dynamic, but not one that I wish to focus on

[21] Of course, others have criticized this approach (Hörl, Warntjen & Wonka 2005). In general, this assumption seems warranted when talking about policy changes that affect the extent of European integration, or the relative power of European and national institutions. It seems less justified when considering issues that deal with overarching political cleavages, such as the traditional left-right divide.

here. Therefore, drawing on the above paragraph, I assume that the norm of presenting a unified front within the Council eliminates $\overline{C}$'s ability to send direct signals to P about likely Council bargaining outcomes, without $\underline{C}$'s consent.[22] Furthermore, I assume that the relative bargaining advantage that $\underline{C}$'s proximity to the status quo affords it is such that $\overline{C}$ is never in a position to force a rejection of a first reading offer, simply to signal to P that it had been too accommodating and could strike a harder subsequent bargain. Nonetheless, $\overline{C}$ is able to increase $\underline{C}$'s bargaining costs by putting up a fight in the Council when the two pivots' preferences diverge with respect to delay.

Given these assumptions, $\underline{C}$ is the pivotal decision maker in the Council, although $\overline{C}$'s preferences and bargaining resilience help determine $\underline{C}$'s cost to delay. Given the expected conciliation outcome, all potential bargains will fall on the range $[x_{cc}, \theta_P]$ because P will never have any incentive to make a proposal that is more accommodating than x_{cc}, nor more hard-line than its own ideal point. Therefore, I restrict attention to proposals in this range and, without loss of generality, scale the space such that $[x_{cc}, \theta_P] = [0, 1]$. The Parliament, therefore, prefers policies closer to 1, while the left Council q-pivot prefers bargains closer to 0. Up to this point I have specified only vague restrictions on player utilities, but I now adopt specific utility functions. Specifically, I assume linear Euclidean utility functions for both players, such that

$$u_P(x, t) = x - k_P(t - 1) \tag{2.1}$$

and

$$u_{\underline{C}}(x, t) = 1 - x - k_{\underline{C}}(t - 1) \tag{2.2}$$

where $t = 1, 2, 3$ is the stage in the game at which bargaining concludes, and k_i is the cost of delay to player i, and $x \in [0, 1]$.

[22] This argument is most persuasive if one conceives of each codecision file in the context of a larger repeated game. $\overline{C}$ may wish to reveal information to P in this case, but may prefer that the Council appear united on many other files. The same holds true for every other potential q-pivot on the Council. Furthermore, publicizing divisions within the Council may undermine the image of the Union more broadly. Therefore, appealing to the Folk Theorem, the unanimity norm may be self-reinforcing.

Furthermore, as I noted above, $\underline{C}$ knows the value of k_P while P is uncertain about $\underline{C}$'s delay costs. Specifically, P knows that $\underline{C}$'s cost falls on a range, $k_{\underline{C}} \in [0, \omega]$, where $0 < \omega \leq \theta_P - x_{cc}$, but believes that $k_{\underline{C}}$ is equally likely to take any value along the given interval. This means, of course, that P's prior belief is that $k_{\underline{C}}$ is uniformly distributed, such that initially

$$k_{\underline{C}} \sim U[0, \omega], \tag{2.3}$$

from P's point of view. Note that I have not directly addressed $\overline{C}$'s costs in this formalization, but have rather focused on $\underline{C}$'s cost distribution, an approach that requires some further explanation. Note that the simple linear costs that feature in the model may capture a variety of issues that modulate players' willingness to prolong negotiations. On the one hand, delay costs include basic issues like the time dedicated to evaluating proposals, deliberation and bargaining over prospective policies, and voting on positions. But opportunity costs also play an important role here, as time spent on the proposal in question restricts an actor's ability to consider other issues on its agenda. Therefore, each k_i can be thought of, perhaps predominantly, as a measure of the relative salience that actor i places on the bill under consideration. Furthermore, with respect to bargaining within the Council, $k_{\underline{C}}$ may incorporate the degree to which the more conservative q-pivot is willing to contest positions with the more progressive q-pivot, and therefore captures the net effect of each pivot's bargaining constraints. In sum, we can interpret ω as a measure of how uncertain the Parliament is, not only about the Council's raw bargaining costs, but of the importance the Council places on the legislation in question, and the relative intra-institutional bargaining strengths of the Council's pivotal actors. To formalize the above logic one might represent $k_{\underline{C}}$ as a linear function of its constituent parts, such that

$$k_{\underline{C}} = k^d_{\underline{C}} + k^o_{\underline{C}} + k^b_{\underline{C}} \tag{2.4}$$

where $k^d_{\underline{C}}$ is $\underline{C}$'s fixed cost to delay, $k^o_{\underline{C}}$ represents $\underline{C}$'s opportunity cost to putting more re-

sources into the current bill rather than deploying them on other legislation, and $k_{\underline{C}}^b$ describes the amount of effort that $\underline{C}$ must put into intra-Council bargaining to build coalitions around rejecting the Parliament's proposal when the more liberal q-pivot would rather accept the proposal.

2.3.1 Sequential Codecision Bargaining

I begin the analysis by examining how the game plays out between the Council and Parliament when the Commission is not involved. As I mentioned above, the game has multiple stages. First, Nature chooses $\underline{C}$'s type (i.e. delay cost), $k_{\underline{C}}$, drawing it randomly according to equation 2.3. Subsequently, the first reading commences ($t = 1$) and P makes a policy proposal, $x_{P,1} \in [0, 1]$. Upon receiving P's proposal, $\underline{C}$ may either accept it, ending the game with payoffs $(x_{P,1}, 1-x_{P,1})$, or reject the proposal, kicking off the second stage. At this point, if the Council rejects Parliament's offer, each player i pays cost to delay k_i. Next, Parliament chooses its second reading proposal, $x_{P,2}$, and transmits it to the Council. Again $\underline{C}$ may choose to accept or reject P's proposal, resulting in net payoffs $(x_{P,2} - k_P, x_{P,2} - k_{\underline{C}})$.[23] If $\underline{C}$ accepts, the game ends. Otherwise, the conciliation committee convenes. Should conciliation be necessary, each player again pays cost k_i. As I described above, neither player formally performs any action after the second stage of the game; rather, in this model, conciliation represents a fixed—but probabilistic—payoff-generating mechanism that takes effect should the Council reject the Parliament's second proposal. This means that, at the start of the game, the players expect a net payoff from conciliation of $(x_{cc} - 2k_P, 1 - x_{cc} - 2k_{\underline{C}})$.[24]

[23] Again, note that I make a rather large simplification here by ignoring the Council's counter-proposal and modeling the process as if the Parliament makes all the proposals. This choice is largely a matter of convenience as it simplifies the exposition considerably, especially when I expand the discussion to include the Commission's role, in the following sub-section. Nonetheless, as I will argue below, this assumption does not fundamentally alter the logic of the game, and the many equilibria of the more more realistic game are qualitatively similar to the single equilibrium that characterizes this, simplified, model.

[24] This model is a variant of a well-established bargaining model (Sobel & Takahashi 1983, Fudenberg & Tirole 1983) that economists have applied to a wide variety of topics, most notably union wage negotiations and strikes. Gibbons (1992, pp. 218–224) provides an introduction to the canonical formulation. Within political science, Cameron & Elmes (1994) and Cameron (2000) use a similar approach to model veto bargaining between Congress and the President in the United States.

This is a sequential game with incomplete information. Therefore, I focus on identifying perfect Bayesian equilibria for the model. As is often the case in dynamic games, it is helpful to consider players' behavior in the last stage of the game first. At second reading $\underline{C}$ will accept Parliament's second reading proposal, $x_{P,2}$, if it provides it with at least as much utility as it can expect to get from rejecting the proposal, paying its cost of delay, and obtaining the conciliation bargain. Therefore, $\underline{C}$ accepts $x_{P,2}$ if

$$1 - x_{P,2} \geq 1 - x_{cc} - k_{\underline{C}} \implies k_{\underline{C}} \geq x_{P,2} \tag{2.5}$$

and rejects P's proposal otherwise.

Now, look at things from the Parliament's point of view. P takes the condition in equations 2.5 into account when formulating its course of action at second reading. Furthermore, as I will subsequently show, the sequence of play leading up to the second reading—most notably the fact that the Council chose to reject the Parliament's initial offer—allows P to update its beliefs about the cost that $\underline{C}$ pays for delaying bargaining. Note that $k_{\underline{C}}$ is a measure of $\underline{C}$'s bargaining strength. An especially strong type of $\underline{C}$ incurs no cost from delay—$k_{\underline{C}} = 0$—while the weakest possible type of Council q-pivot pays $k_{\underline{C}} = \omega$. Stronger types are more able to delay; therefore, as I will further describe below, finding itself at second reading allows P to update its estimate of $\underline{C}$'s strength. Specifically, the history of the game leading up to P's move at second reading allows P to put a new lower bound on $\underline{C}$'s strength—or, equivalently, to update its estimate of the upper bound on $\underline{C}$'s delay cost—inducing P to update its beliefs such that, at this point in the game, P believes

$$k_{\underline{C}} \sim U[0, \omega_2] \tag{2.6}$$

where $\omega_2 \leq \omega$. Therefore, using equations 2.5 and 2.6 to calculate the probability that $\underline{C}$ will accept a given offer at second reading, P's conditional expected utility from an optimal

second reading proposal, $x^*_{P,2}(\omega_2)$, is

$$u_P(\omega_2) = x^*_{P,2}(\omega_2) \cdot \frac{\omega_2 - x^*_{P,2}(\omega_2)}{\omega_2} - k_P \cdot \frac{x^*_{P,2}(\omega_2)}{\omega_2}, \tag{2.7}$$

where $\frac{\omega_2 - x^*_{P,2}(\omega_2) + x_{cc}}{\omega_2}$ is the probability, given P's updated beliefs, that $\underline{C}$ accepts proposal $x^*_{P,2}(\omega_2)$, $\frac{x^*_{P,2}(\omega_2) - x_{cc}}{\omega_2}$ is the probability that $\underline{C}$ rejects the second reading proposal, $x^*_{P,2}(\omega_2)$ is P's payoff if $\underline{C}$ accepts its optimal second reading proposal, and $-k_P$ is P's expected payoff from conciliation. Maximizing, P's optimal proposal at second reading is

$$x^*_{P,2}(\omega_2) = \begin{cases} \frac{\omega_2 - k_P}{2} & \text{if} \quad \omega_2 > k_P, \text{ and} \\ x_{cc} = 0 & \text{otherwise.} \end{cases} \tag{2.8}$$

Interpreting equation 2.8, we can see that the Parliament makes a second reading proposal that balances its own distaste for delay, k_P, with its updated belief about the Council's bargaining strength. When the Parliament believes that the Council is relatively weak (i.e. ω_2 is comparatively large), it makes tough proposals that benefit it more than the likely conciliation outcome, expecting that the Council will accept a Parliament-biased bargain because it will not find delay worth the added cost of in time, effort, and lost opportunity to pursue other issues. When the Council is especially weak, Parliament will demand allocations close to its own ideal point; when facing a particularly strong Council, the Parliament will offer the expected conciliation policy.[25]

Now consider the first reading. If P offers $x_{P,1}$ at first reading, and $\underline{C}$ expects that it will offer $x_{P,2}$ at second reading, $\underline{C}$ accepts the first reading offer when it believes that it will give it a payoff that is at least as good as what it could get by rejecting initially and accepting the second offer, or by rejecting both offers and holding out for conciliation. That

[25]Note that the assumption that $\omega \leq \theta_P - x_{cc}$ bounds Parliament's optimal proposal above, such that Parliament always proposes compromise positions that are no more extreme than its own ideal point.

is, the Council accepts the first reading offer if, and only if,

$$k_{\underline{C}} \geq \max\{x_{P,1} - x_{P,2}, \frac{x_{P,1}}{2}\} = \omega_2. \tag{2.9}$$

Furthermore, should the Council q-pivot reject the initial offer, it can anticipate $x_{P,2}$, as identified by equation 2.8. This means that $\underline{C}$ follows a simple cutoff rule, as a function of its type, when deciding whether to accept or reject the Parliament's first offer. There is a type of $\underline{C}$—with delay cost ω_2—that is indifferent between accepting $x_{P,1}$ now, and rejecting it either to accept the second offer or to hold out all the way to conciliation. All types of $\underline{C}$ with $k_{\underline{C}} \geq \omega_2$ accept $x_{P,1}$ while stronger types of $\underline{C}$ with $k_{\underline{C}} < \omega_2$ reject the initial offer. Therefore, the model exhibits a screening property that, as I mentioned above, allows the Parliament to eliminate weaker Council types with its first round offer. Thus, should the game proceed to the second reading, the Parliament will update its beliefs about $k_{\underline{C}}$ as I described in equation 2.6.

The Parliament chooses a first reading offer that maximizes its overall expected utility, or

$$\max_{x_{P,1},\omega_2} x_{P,1} \cdot \frac{\omega - \omega_2}{\omega} + (u_P(\omega_2) - k_P) \cdot \frac{\omega_2}{\omega}. \tag{2.10}$$

Moreover, in the lemma 1 in appendix A I show that, in any equilibrium, $x_{P,1} - x_{P,2} \geq \frac{x_{P,1}}{2}$. Therefore, substituting Parliament's optimal second reading proposal, $x^*_{P,2}(\omega_2)$, from equation 2.8 into equation 2.9 and solving for ω_2 yields an implicit equation for ω_2, in terms of $x_{P,1}$. In particular,

$$\omega_2(x_{P,1}) = \begin{cases} \frac{2x_{P,1}+k_P}{3} & \text{if} \quad \omega_2 > k_P, \text{ and} \\ x_{P,1} & \text{otherwise.} \end{cases} \tag{2.11}$$

First, assume that $\omega_2 > k_P$ so that, from equation 2.11, $\omega_2(x_{P,1}) = \frac{2x_{P,1}+k_P}{3}$. Substituting $\omega_2(x_{P,1})$ and the optimal second round proposal from equation 2.8 for this case into equation

2.10 makes it possible to express equation 2.10 completely in terms of ω and $x_{P,1}$. Maximizing this expression with respect to $x_{P,1}$ and solving for $x_{P,1}$ yields the optimal first offer:

$$x^*_{P,1} = \frac{9\omega - 11k_P}{10}. \tag{2.12}$$

Substituting this result back into equation 2.8, the optimal second period offer is

$$x^*_{P,2} = \frac{3\omega - 7k_P}{10}, \tag{2.13}$$

while, plugging the optimal first offer in equation 2.12 into equation 2.11, the cutoff point for Council types that reject the first offer in equilibrium is

$$\omega^*_2 = \frac{3\omega - 2k_P}{5}. \tag{2.14}$$

Following analogous logic, the equilibrium offers and cutoff for the case where $\omega_2 \leq k_P$ are

$$x^*_{P,1} = \max\{0, \frac{\omega - k_P}{2}\}, \quad x^*_{P,2} = 0, \quad \omega^*_2 = x^*_{P,1}. \tag{2.15}$$

While the details differ, the equilibrium proposals and cutoffs are qualitatively similar across both cases. During each reading the Parliament makes a proposal that balances its beliefs about the range of possible Council types with its own costs to delay. In the case where $\omega_2 > k_P$—or, plugging in the equilibrium values, $\omega > \frac{7}{3}k_P$—the Parliament makes an aggressive first offer, followed by a more accommodating offer should the game extend into second reading. The same is true when $\omega_2 \leq k_P \implies \omega \leq \frac{7}{3}k_P$, except, in this case, the Parliament will certainly offer the conciliation outcome at second reading. Furthermore, if its own delay costs are too high—if $k_P \geq \omega$ and the risk of delay outweighs the expected benefit of any aggressive offer—then the Parliament proposes $x_{cc} = 0$ at first reading, knowing that the Council is bound to accept it, ending bargaining.

In general, when its own costs are sufficiently low, the Parliament makes a series of offers that maximize the bargaining surplus that it expects it can extract from the Council. The Parliament's offer at first reading is at least as aggressive as its second reading offer; indeed, if the Parliament does not offer the conciliation outcome at first reading, it will certainly propose a policy at first reading that is closer to its own ideal point than the policy it suggests subsequently. As I have already noted, the Parliament's offer at first reading will satisfy weaker types of Council q-pivots, and they will accede to the Parliament's demands. But stronger types will stay in the game. The details of this dynamic varies, depending on the difference between ω and k_P. When ω is sufficiently above k_P, the Parliament's proposals track an optimal balance between short-term utility maximization, and longer-term information acquisition. Knowing that only stronger Council types will reject the initial proposal, the Parliament softens its demands at second reading in hopes of concluding negotiations without investing in conciliation.

Yet, depending on the circumstances, the Parliament may nonetheless attempt to take advantage of Council weakness at second reading, and propose an allocation that benefits the Parliament (Council) more (less) than the conciliation outcome. This is, of course, a gamble, and especially strong Council q-pivots will reject even the more accommodating second reading offer. On the other hand, when k_P is relatively large, the Parliament simply maximizes its single-shot expected utility at first reading, knowing that Council types that will reject the proposal and force a second reading will accept nothing short of the conciliation outcome. Under such circumstances bargaining never extends into conciliation because the Parliament will not take the risk of making an aggressive offer in the second stage.

One notable aspect of the model is that the median MEP's cost to delay must be substantially lower than that of the Council q-pivot to make extended bargaining worthwhile for the Parliament. Nonetheless, it would make sense that the average MEP's cost to delay would be significantly below that of the member state representatives on the Council. While the member states have government bureaucracies behind them, and have reasonably large

staffs dedicated to considering European legislation, they must nonetheless commit substantial resources when considering each and every piece of legislation that the Commission proposes. On the other hand, while individual MEPs have limited staffs, and few personal resources to draw upon, the work-load of the Parliament is distributed across upwards of 700 MEPs. While the rapporteur assigned to the file, and indeed the Parliament's permanent secretariat, would probably prefer to complete the case quickly, the marginal cost of extending a given codecision file is likely negligible for any given MEP. Therefore, the Parliament as a whole, and the median MEP in particular, should be relatively unencumbered by delay costs in most circumstances.

It is worth emphasizing that delay is possible in this game and, depending on the values of $k_{\underline{C}}$ and k_P, bargaining may extend all the way to conciliation. That is, when it is possible that $\underline{C}$'s costs are quite high relative to k_P, Parliament will attempt to extract concessions from the Council at both readings. Yet, when presented with these offers, some especially strong types of $\underline{C}$ will still prefer waiting until conciliation to accepting P's second reading compromise proposal. Therefore, unlike full information models of codecision, the sequential bargaining game allows for the multi-reading codecision files that we observe in real life.

On the other hand, this model does make one assumption—that the Council does not counter-offer, but rather only accepts or rejects Parliament's proposals—that is clearly untrue. This assumption greatly simplifies the analysis because, given the Council's severely restricted signaling language, the Parliament can never be put into a situation where it has to update its beliefs, knowing that it is off of the equilibrium path.[26] When the Council can counter-propose, its messages become much more detailed, off-path belief updating becomes a serious issue, the number of possible equilibria multiplies exponentially, and finding closed-form solutions becomes difficult. Thus, such models provide poor foundations for building more complicated theories. Nonetheless, the equilibria of asymmetric bargaining models in

[26]That is, unless the Council were to reject an offer of x_{cc} at the first reading. In this case, one can simply assume that the Parliament maintains its belief that $k_{\underline{C}} = 0$, as Bayes' rule provides no guidance for what the Parliament's beliefs should be off the equilibrium path.

which both sides can propose retain much of the flavor of the model that I present here (Grossman & Perry 1986, Ausubel, Cramton & Deneckere 2001). Indeed, in such models, the uninformed player makes sequentially more accommodating offers, screening out weaker types of the informed player as the game progresses, just as the Parliament does in this model. When there is no hope of extracting an offer from the uninformed player that will compensate the informed actor for dragging things out further, the informed player makes an acceptable offer to its counterpart. Otherwise, a cheap-talk dynamic maintains and the uninformed bargainer makes offers that it knows the uninformed actor will reject, resulting in a situation that looks quite similar to bargaining with one-sided offers. Clearly, the simplified model that I develop here cannot help explain whether or not Parliament will choose to make amendments at second reading. Similarly, it tells us nothing about the content of the Council's second reading proposals. Nonetheless, it largely captures the dynamics of bargaining under asymmetric information—even when both sides can make proposals—and, because of its tractability, serves as a useful platform for building a more complicated model that considers how the Commission mediates codecision bargaining.

2.3.2 Mediated Codecision Bargaining

We are now in a position to examine the Commission's role in codecision bargaining. The game is almost identical to sequential codecision bargaining. First, Nature chooses $\underline{C}$'s type, $k_{\underline{C}}$, again drawing it randomly according to equation 2.3. But now, before the Parliament officially tables its first reading proposal, Nature sends the Commission a private signal, $s \in \{S, W\}$, indicating whether or not the Council is in a strong bargaining position, and thus is likely to reject the Parliament's opening bid, or in a weak position, and therefore likely to accept the Parliament's offer. Specifically, because it is involved in the internal discussions of both the Council and the Parliament prior to Parliament's vote on a proposal at first reading, the Commission has an idea of both the content of the Parliament's initial offer, and the likelihood that the Council will accept that offer. More specifically, building

on the sequential bargaining model, the Commission's signal is such that $s = S$ when the Commission believes that $k_{\underline{C}} < \omega_2^*$, as defined in equation 2.14, and, therefore, believes that the Council will reject the Parliament's initial offer, should the Parliament make that offer based purely on its prior knowledge—as described by equation 2.3—about the Council's bargaining strength. On the other hand, $s = W$ when the Commission believes that the Council will accept Parliament's naive first offer, or when $k_{\underline{C}} \geq \omega_2^*$.

Note that, while the Commission has access to information that the Parliament does not, that information is relatively imprecise. This aspect of the model makes empirical sense. Indeed, while the Commission's special access to intra-Council deliberations affords it private knowledge about the Council's baseline bargaining costs, the salience it places on the file, and intra-Council divisions, it is unreasonable to assume that the Commission perfectly observes the Council's bargaining strength. On the other hand, it is not unreasonable to assume that the Commission would be able to ascertain the Council's appetite for Parliament's most likely proposal. In the model, therefore, the Commission simply knows what the Parliament will eventually discover at second reading before the Parliament has tabled its first reading proposal.[27]

The Commission now sends a message, $m \in \{S, W\}$ to the Parliament, potentially relaying its private information to the lower house. There is no practical way for the Commission to send such a message to over 700 MEPs in private, so m is observed by all actors.[28] This aspect of the model also seems like a plausible representation of the state of affairs that characterizes actual codecision lawmaking. In practice, as I mentioned in section 2.2, the Commission communicates official opinions—with accompanying arguments—on all of the potential amendments that the Parliament tables to the Commission's initial proposal at

[27] An interesting generalization would allow for error in the Commission's private information. For example, s could be correct with probability $1 - \epsilon$. For now I assume that the signal is accurate, if imprecise. I leave the more general model to future research.

[28] Note that every MEP must have access to the same information for the median MEP to maintain its role as pivotal actor within the Parliament. Therefore, the Commission cannot simply relay its private information to the median MEP if it wishes to influence Parliament's proposals.

first reading. While these opinions serve a variety of purposes, the Commission can, in principle, use these public statements to transmit its private information about the Council's willingness to compromise to the Parliament. Therefore, although the model abstracts away from the manner in which the Commission sends messages to the Parliament, the Commission has an institutionalized method for transmitting information to the Parliament in the real world.

After the Commission sends its message the sequential codecision bargaining game plays out, proceeding exactly as it did in the previous section. Initially, the parliament makes a first reading offer, $x_{P,1}$, which the Council accepts or rejects. At second reading, Parliament makes a follow-up offer, $x_{P,2}$, and the Council again decides whether or not to accept the proposal. Should the Council reject the second reading offer, the conciliation committee convenes, ending the game with a lottery.

The Commission—like the Council and Parliament—has preferences over outcomes. The Treaties envision the Commission as an independent, unbiased, and largely apolitical institution. Indeed, Commissioners are formally expected to operate independently, and to make decisions without regard to the preferences of their member states, or national parties, and Commissioners swear an oath to this effect. Of course, academic work often finds evidence for bias within the Commission (see e.g. Thomson 2008). To highlight the role that Commission (in)dependence may play in modulating lawmaking in the EU, I model the Commission's utility in terms of both its own costs to delay, k_M, and the degree to which it is biased in favor of either of the two institutions' most preferred outcomes. Specifically,

$$u_M(x,t) = \beta x - k_M(t-1), \tag{2.16}$$

where $\beta \in \mathbb{R}$. The parameter β is a measure of Commission bias that all players observe. When $\beta = 0$ the Commission is unbiased and maximizes its utility simply by arranging for a speedy conclusion to deliberations. This is the situation envisioned by treaty; the

Commission stands above the disagreements of the political actors and faces an incentive structure that compels it to guide the co-legislators quickly towards compromise. On the other hand, when $\beta > 0$, the Commission is biased—to a greater or lesser extent—towards the Parliament median's point of view. Similarly, when $\beta < 0$, the Commission favors the Council's position. Furthermore, like the Council and Parliament, the Commission pays a fixed cost to delay, which I assume all actors observe. And, like the Parliament, these costs are a function of the resources the Commission must spend when codecision extends into subsequent readings and the degree to which focusing on the file at hand restricts its ability to deploy resources to other, perhaps more salient, codecision files. Yet, perhaps unlike the Parliament, because the Commission plays such an important role in all aspects of codecision—drafting legislation, advising the Council and Parliament's committees during their internal deliberations, rendering opinions on proposals, answering questions in plenary, helping to mediate trilogues, and providing informational and technical support in conciliation—its costs to delay are likely to be substantial.

The analysis of this game proceeds largely along the lines of the baseline sequential bargaining model. Clearly, the mediated codecision bargaining model two pooling Bayesian perfect equilibria where the Commission sends the same message–either $m = W$ or $m = S$—to the Parliament, regardless of the content of the signal that it receives at the beginning of the game. Under such circumstances, the Parliament has no information at its disposal that can influence it to alter its approach to codecision bargaining. Indeed, the sequential codecision bargaining model, that I develop above, describes the Parliament and Council's equilibrium behavior when the Parliament disregards the Commission's message and acts purely on the basis of its prior knowledge. Below, I examine the circumstances under which the mediated game can support a fully separating equilibrium. This answers the question: when does the Commission have an incentive to relay its signal to the Parliament without embellishment, or equivalently, when can the Parliament believe the Commission's signal? I then use the nature of the separating equilibria to examine how the Commission's ability

to honestly transmit information affects both the outcome of codecision bargaining, and the speed with which bargaining is concluded. In the previous section I argued that, because the Parliament can spread its workload out across so many members, MEPs will find costs to delay negligible for any given codecision file. Therefore, and to reduce notational clutter, in this section I assume that $k_P \leq \frac{3}{7}\omega$.[29]

To begin the analysis of the separating equilibria, assume first that $m = S$ and that the Parliament believes that the Commission has truthfully relayed its private information such that, in the Parliament's opinion, $m = s$. Now, at first reading, the Parliament must initially believe that the Council's potential costs are such that

$$k_{\underline{C}} \sim U[0, \omega_2^*]. \tag{2.17}$$

Clearly, from this point, the game is exactly the same as the sequential codecision bargaining game; the only difference is that the potential type space is a subset of the type space that confronted the Parliament prior to receiving the Commission's signal. Therefore, following the logic of the previous section, if $9\omega > 41k_P$,[30] in any separating equilibrium, if the Commission sends the signal $m = S$, the equilibrium offers and cutpoint are

$$x_{P,1,S}^* = \frac{5.4\omega - 14.6k_P}{10}, \quad x_{P,2,S}^* = \frac{1.8\omega - 8.2k_P}{10}, \quad \omega_{2,S}^* = \frac{1.8\omega - 3.2k_P}{5}. \tag{2.18}$$

Similarly, when $9\omega \leq 41k_P$, the equilibrium offers and cutoff are

$$x_{P,1,S}^* = \frac{3\omega - 7k_P}{10}, \quad x_{P,2,S}^* = 0, \quad \omega_{2,S}^* = x_{P,1,S}^*. \tag{2.19}$$

[29]This assumption rules out the corner case where the Parliament, given no information beyond its prior, proposes the Conciliation outcome with certainty at second reading and, perhaps, at first reading. As the equilibrium offers that I derived in the previous section should make clear, the corner case exhibits a flavor that is similar to the situation where the Parliament's second reading offer is unconstrained by its own costs. Therefore, this simplification does little, if anything, to change the theoretical conclusions that one can draw from the model.

[30]Here, the Parliament's delay cost must be almost five times smaller than the Council's highest a priori (i.e. before the Commission's signal) cost, to obtain the unconstrained solution.

Figure 2.3: The Likelihood of a second reading given a strong Council and truthful Commission.

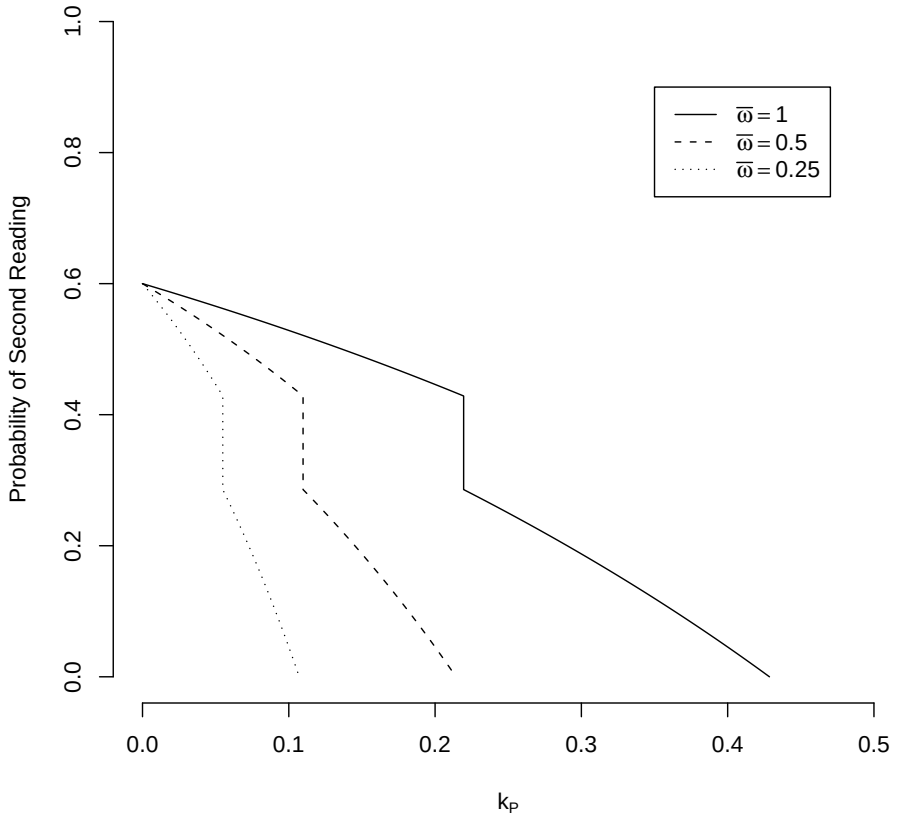

Comparing equation 2.18 to equations 2.12, 2.13, and 2.14, we find an unsurprising result. Specifically, upon initially learning that the Council is relatively strong, the Parliament tones down the series of offers that it presents to the Council, adopting a substantially more accommodating approach. For example, when the Parliament's own costs to delay approach zero, the Parliament's first offer, based on its own prior information, is one and two thirds the size of the first reading offer it tables after receiving, and believing, the message $m = S$.

This finding has important implications for how bureaucratic mediation can influence lawmaking efficiency within the Union. Given a strong Council, bargaining will conclude more quickly when the Parliament can trust the Commission's message. When the Parlia-

ment acts purely on its priors a second reading is guaranteed. On the other hand, when the Commission can obviate the Parliament's need to float the initial screening offer described by equation 2.12, the probability that the Council rejects the first offer is

$$P(k_{\underline{C}} < \omega^*_{2,S} | k_{\underline{C}} < \omega^*_2) = \frac{\omega^*_{2,S}}{\omega^*_2}. \tag{2.20}$$

Figure 2.3 depicts this relationship, as a function of k_P, for three possible values of ω. Clearly, for a fixed cost k_P, the likelihood of delay increases with Parliament's uncertainty about the Council's bargaining strength. Furthermore, while the probability of a second reading is about sixty per cent when Parliament faces no delay costs, this probability quickly drops off as k_P increases. The relationship exhibits a discontinuity at $k_p = \frac{9}{41}\omega$—the point at which Parliament's second round offer is guaranteed to be completely accommodating—and then continues its downward trend, reaching zero when k_P hits its assumed upper bound of $\frac{3}{7}\omega$. A similar dynamic holds at second reading; Parliament's second offer after receiving and believing $m = S$ is always lower than it would be when approaches the game with only its prior information in hand. Therefore, while delay is possible at second reading in both pooling and separating equilibria, a file is less likely to make it all the way to conciliation when the Parliament and Commission can communicate effectively. Overall, when the Parliament trusts it, the Commission can use its mediating influence to reduce bargaining inefficiency, sometimes substantially.

Now, assume that $m = W$ and that the Parliament believes the Commission's message. The perfect Bayesian equilibrium concept requires the Parliament to hold beliefs that are consistent with players' strategies at all points along the equilibrium path. The Parliament always reaches first reading in this game; therefore, in any separating equilibrium, it must initially believe

$$k_{\underline{C}} \sim U[\omega^*_2, \omega] \tag{2.21}$$

after observing $m = W$. Thus, at first reading, it knows that all Council types will accept

the offer $x_{P,1,W} = \omega_2^*$, and thus should never table a more accommodating first proposal.

Unfortunately, the equilibrium concept provides little guidance on how to model the Parliament's beliefs off of the equilibrium path. This is an issue here because—in contrast to the case where $s = m = S$—when players follow their strategies, the game need not reach second reading with positive probability in a separating equilibrium after Parliament observes the message $m = W$. Indeed, lemmas 2 and 3 in appendix A prove that the second reading is never reached when players follow their strategies in any separating equilibrium of the mediated codecision bargaining game, when $s = W$. Therefore, the Parliament can adopt a variety of beliefs about the Council's strength should it—unexpectedly—find itself at second reading. For example, at second reading, Parliament might conclude that the Commission made a mistake and decide that the Council must be strong after all. Alternatively, it could reason that the Commission's message was correct, but that a misunderstanding or other idiosyncratic event within the Council had caused the ministers to err in their response to the Parliament's offer. In situations like the former case, where the Parliament assumes the Council is strong should second reading occur, the Parliament has an incentive to play things safe and to make a relatively accommodating first offer. On the other hand, should the Parliament maintain its beliefs across both readings—it will make an exceedingly aggressive offer, proposing its own ideal point under many circumstances, and no less than $2\omega_2^*$.

Moreover, the range of off-path second reading beliefs that can support a separating equilibrium is distressingly broad. In fact, for certain values of β and k_M, there are separating equilibria where the Parliament holds any belief—so long as the Parliament believes that $k_{\underline{C}}$ lies somewhere on the interval $[0, \omega]$—at second reading, after observing $m = W$. Consider, for example, an extremely "pessimistic" Parliament that, after observing $m = W$ and unexpectedly seeing its first proposal rejected, assumes that the Council must be exceedingly strong and adopts the belief that $k_{\underline{C}} = 0$. Clearly, in this case, Parliament's optimal second reading proposal is $x_{P,2,W}^* = 0$.[31] It is easy to verify that, given this second reading belief,

[31] Technically, when $k_P = 0$, the Parliament is indifferent and may propose any $x_{P,2,W} \in [0, 1]$ but, if we

and corresponding second offer, Parliament's optimal first offer is $x^*_{P,1,W} = \omega^*_2$. Of course, keeping equation 2.21 in mind, it is clear that any weak Council will accept this first round offer, keeping the second round off of the equilibrium path. On the other hand, as I mentioned in the last paragraph, if the Parliament holds "optimistic" beliefs when it finds itself off the equilibrium path, there are equilibria in which $x^*_{P,1,W} = 1$. Therefore, it is possible to make only relatively vague general claims about how the Parliament will behave after observing $m = W$ and believing the Commission's message. First, in such cases, bargaining will end at first reading, again placing the Commission in a position to improve bargaining efficiency through mediation. Second, Parliament will make a first round proposal the lies on the range

$$x^*_{P,1,W} \in [\omega^*_2, 1]. \tag{2.22}$$

So, when will the Commission have an incentive to maintain its separating strategy $s = m$ in equilibrium, given the other players' optimal responses to truthful messages? First, say the Commission observes the signal $s = W$, learning that the Council is weak. Comparing equations 2.18 and 2.19 to equation 2.22, one can see that, should the Parliament believe the Commission, it will always propose a more accommodating first reading offer when $m = S$ than when $m = W$, even if the Parliament holds particularly "pessimistic" beliefs off of the equilibrium path. Additionally, given that $s = W \implies k_{\underline{C}} \geq \omega^*_2$, the weak Council will pounce upon the misled Parliament's over-generous offer at first reading, ending bargaining at the outset, meaning that the Commission can mislead the Parliament in these circumstances without risking delay costs.[32] Therefore, when $\beta < 0$, the Commission always gains from misleading the Parliament into believing that a weak Council is strong and will optimally deviate from its separating strategy; thus, there can be no separating equilibria

assume the Parliament finds delay even minutely costly, it will always propose $x_{P,2,W} = 0$ at second reading if it believes $k_{\underline{C}} = 0$.

[32] A weak Council always at least weakly prefers accepting a first offer that is less than or equal to ω^*_2 to delay, even if it expects the Parliament to fully capitulate and propose an allocation of 0 at the second reading.

when $\beta < 0$. When the Council and Commission have similar preferences, cheap talk from the Commission is no more valuable to the Parliament than signals sent directly from the Council. In such circumstances, the Commission's mediating capabilities are completely undercut.

On the other hand, when $\beta \geq 0$, the Commission has no incentive to deceive the Parliament after observing $s = W$. First, should the Commission maintain its separating strategy, lemmas 2 and 3 show that bargaining will end at first reading. Therefore, the Commission cannot reduce delay costs by tricking the Parliament into adopting a more accommodating stance than it should, as the Parliament would do if it believed the Council strong. Furthermore, when $\beta \geq 0$, the Commission weakly prefers outcomes closer to the Parliament's most preferred bargain to lower allocations. Of course, equations 2.18, 2.19, and, 2.22 show that the Parliament will always make tougher proposals when it believes that the Council is weak than when it thinks that the Council is strong, providing the Commission with no incentive to deviate from its separating strategy when $s = W$ and $\beta \geq 0$.

The final piece in the puzzle hinges on whether or not it is optimal for the Commission to adopt a separating strategy when $\beta \geq 0$ and $s = S$. That is, will a Commission that is weakly biased towards the Parliament's ideal bargain ever prefer to tell a trusting Parliament that the Council is weak when it is actually strong? If the Commission relays the true signal to the Parliament, and the Parliament believes the message, bargaining may end at any of the three readings, and, weighing the relative likelihood of each possible outcome, the Commission can expect to obtain the payoff

$$\beta x^*_{P,1,S}\left(1 - \frac{\omega^*_{2,S}}{\omega^*_2}\right) + \left[\beta x^*_{P,2,S}\left(1 - \frac{x^*_{P,2,S}}{\omega^*_{2,S}}\right) - k_M \cdot \frac{x^*_{P,2,S}}{\omega^*_{2,S}} - k_M\right] \cdot \frac{\omega^*_{2,S}}{\omega^*_2} \tag{2.23}$$

if it follows the separating strategy. On the other hand, if the Commission misleads the Parliament—given that $s = S \implies k_{\underline{C}} < \omega^*_2$—equation 2.22 indicates the Council will certainly reject the Parliament's initial offer. Therefore, in this case, the Commission's

expected utility is

$$\beta x^*_{P,2,W}\left(1-\frac{x^*_{P,2,W}}{\omega^*_2}\right)-k_M\left(\frac{x^*_{P,2,W}}{\omega^*_2}\right)-k_M \tag{2.24}$$

if the Parliament falls for the deception. The Commission will maintain its separating strategy when its expected utility from telling the truth exceeds the payoff it expects to receive upon misleading the Parliament, or when the value of equation 2.23 exceeds the value of equation 2.24. When $\beta = 0$ this condition simplifies to the requirement that $x^*_{P,2,S} - x^*_{P,2,W} \geq \omega^*_2 - \omega^*_{2,S}$, an inequality that will always hold. Therefore, a wide variety of separating equilibria are possible when the Commission is unbiased. When $\beta > 0$ the condition does not simplify so nicely and the resulting inequality is not particularly easy to interpret in general terms. Whether or not a separating equilibria is possible hinges crucially on the degree of the Commission's bias, β, the Commission's cost to delay, k_M, and the Parliament's delay cost, k_P, and off-path beliefs at second reading after observing $m = W$ in a given potential equilibrium. Furthermore, the interplay between these factors is quite complicated. Nonetheless, many parameter configurations and off-path beliefs support separating equilibria when $\beta > 0$.

For example, if—contingent on observing $m = W$—the Parliament adopts the off-the-path second reading belief that $k_{\underline{C}} \sim [0, \omega^*_2]$, then the mediated codecision bargaining game has separating equilibria whenever $\beta \geq 0$. Furthermore, this is an attractive equilibrium from a theoretical point of view because it makes sense for MEPs to update their beliefs in this manner upon reaching second reading after receiving the message $m = W$ and, following their equilibrium strategy, tabling an initial proposal that any weak Council should accept. The Commission is bound to make mistakes from time to time and misinterpret the Council's strength. When this happens we would expect that MEPs might conclude that the Council was strong after all, but would not predict that they adopt particularly specific beliefs about the exact level of Council strength. Thus the vague off-path belief that "the Council is strong" seems quite plausible in this context. Moreover, as long as the Commission is weakly

biased towards the Parliament's bargaining position, such beliefs will induce the Commission to act as an honest mediator, coordinating a speedy compromise between the two chambers. Of course, their are many other belief configurations that can support separating equilibria when $\beta \geq 0$.

On the other hand there are off-the-path Parliamentary beliefs—and resulting second-reading strategies—that cannot support a separating equilibrium when $\beta > 0$. This state of affairs can occur when the Parliament is sufficiently "pessimistic" off the equilibrium path to make a second reading offer that a strong Council has some probability of accepting, but that generate a second offer that is tough enough to yield an expected outcome to the right of the bargain, when played against a strong Council, that the players would expect to strike should the Parliament play its optimal strategy. Given such off-path beliefs, there can be situations where the preferences of the Commission and Parliament diverge, in the sense that the Commission is more willing to stomach risk in the hopes of attaining a larger share of the bargaining spoils. Specifically, when β is sufficiently large, k_M is sufficiently small, and k_P is comparably large, the Commission has an incentive to convince MEPs to take a risk that is not in their best interests. Therefore, one cannot unequivocally say that the Commission will truthfully relay information about the Council's bargaining position to the Parliament when it exhibits a bias towards the Parliament's most preferred outcome. Indeed, depending on how the Parliament reacts to unexpected rejections, the Council can be too biased towards the Parliament's point of view to be trusted, especially if it suffers little costs from delay.

In sum, the analysis that I present here demonstrates that effective communication between the Commission and Parliament is possible only when their preferences are sufficiently similar. More generally, the model indicates that bureaucrats will be most effective at influencing policy-making through mediation when the goals of the bureaucracy are similar to those of those policy-making bodies that are at the greatest informational disadvantage vis-a-vis their peers. Of course, depending on the way in which MEPs form beliefs when

confronted with unexpected behavior by the Council, even pro-Parliament Commissions—and, in particular, exceedingly pro-Parliament Commissions—may not be able to truthfully communicate to the Parliament. Yet such situations are unlikely if, as one would expect, the Commission finds delay reasonably costly. Moreover, pro-Parliament bias always makes persuasive mediation possible for especially plausible off-equilibrium-path beliefs. Finally, the model indicates that the Commission can always play an effective mediating role when unbiased. Nonetheless, one should be cautious when considering the implications that unbiasedness has for successful bureaucratic mediation. In the game that I present here the Commission receives an accurate, if vague, signal of the Council's bargaining strength. In real-world politics, few bureaucrats will have access to such valid measures. Furthermore, error in observation has the potential to affect the decision calculus of the Commission considerably. An unbiased mediator is concerned only with ending negotiations quickly and will accept any outcome as a means to that end. Therefore, when there is a chance that the mediator is wrong about the informed player's bargaining strength, the mediator will wish to err on the side of caution and attempt to convince the proposing actor to compromise. Thus, as Kydd (2003) has argued in the context of conflict mediation, unbiased mediators may actually be unable to convincingly relay information about one negotiator to another.

2.3.3 Empirical Implications

Taken together, the models that I present in this section provide a number of empirical predictions about lawmaking in the European Union, and the ways in which bureaucratic mediators may influence policy-making more generally.

The sequential codecision bargaining model generates implications that mirror standard accounts of bargaining under uncertainty (Sobel & Takahashi 1983, Fudenberg & Tirole 1983, Grossman & Perry 1986, Ausubel, Cramton & Deneckere 2001). Indeed, it shows that negotiations will take longer, and codecision will span more readings, when legislation is more salient, when players face lower fixed bargaining costs, and when the Parliament's

uncertainty about the Council's bargaining constraints grows in comparison to Parliament's own delay costs. Furthermore, Parliament's second reading proposals should always be more accommodating than the first reading proposals that precede them. Yet, while others have previously leveraged bargaining models with asymmetric information to understand the broad contours of how lawmaking gets done in bicameral legislatures (Tsebelis & Money 1997), adding a mediator to the picture yields novel insights.

Because the Parliament can only trust the Commission when preferences align across institutions, Commission arguments—such as opinions on Parliament's amendments, and recommendations to rapporteurs, committees, and speeches in plenary—should only influence decision-making within the Parliament when the median MEP and the Commission share policy goals. This does not mean that MEPs will make decisions based purely on their own ideological predilections. Rather, when the median MEP and the Commission agree, the majority of MEPs can be persuaded to behave against their most preferred outcomes and can be convinced to support second-best compromise positions, because they know that this will maximize their utility in the long term.

When the Parliament and the Commission share preferences, lawmaking should proceed more efficiently than when they are divided on an issue, holding the level of disagreement between the Parliament and Council constant. When effective communication between the Commission and Parliament is possible, the Commission will help the Parliament to avoid the lengthy and wasteful process of floating a relatively aggressive first proposal to test the Council's resolve, shortening negotiations when the Council turns out to be strong. More generally, the preferences of bureaucrats may interact with the information environment to help explain variations in legislative efficiency across political contexts that are, institutionally, similar. Furthermore, the Commission's information can help the Parliament to obtain policy outcomes that it prefers, on average, when compared to bargaining with no help from a mediator. When the Council turns out to be weak, the Commission can relay this information to the Parliament, allowing it to go for the jugular and extract as much bargaining

surplus as possible.[33] Even when the Council turns out to be relatively unimpaired by delay costs, because it can tailor its offers to the Council's strength, the Parliament will make out better, on average, when it can make use of the Commission's messages when planning its course of action.

2.4 Conclusion

Acting only in its capacity as a mediator, the Commission can exert impressive influence on the dynamics of lawmaking. Its control of policy outcomes, on the other hand, is relatively limited—it cannot force either the Council or the Parliament to accept an outcome that it cannot abide—but it can direct compromise towards its preferred policy, as long as its preferences and the Parliament's are similar. This result is consistent with the impression, commonly voiced among practitioners of EU politics, that the Commission and the Parliament often work together to extract concessions from the Council (Rasmussen 2003, pp. 5). Overall, the Commission's mediative capacity may help to explain the disconnect between the Commission's formal powers and many observers' beliefs about its strength under codecision. Acting as a mediator, the Commission will often encourage the Parliament and Council to settle on compromise positions in the early stages of codecision. In so doing, it will make recommendations that both parties follow and will appear to exert substantial influence over the other EU institutions' decisions. And, in a way, it does. The Commission does not simply recommend that the Council and Parliament take the course of action that they would in its absence. Rather, when the appropriate conditions hold, it guides the two chambers to speedy resolutions that they could not have found on their own. Thus, while it lacks power over policy, it is surprisingly persuasive.

The story of Commission influence through mediation has important implications for

[33] Of course, the Parliament's tendency to do so depends on off-equilibrium-path beliefs, and the game also sports equilibria in which the Parliament is more conciliatory when it learns that the Council is weak than it would be if the Commission was not in the picture.

the way in which political scientists conceptualize policy making in governments and international organizations. While bureaucracies in national government derive much of their power from their discretion over the implementation of policy, they nonetheless also enjoy significant advantages based on their multiple points of contact with other governing institutions. How relatively well-connected bureaucracies are—and how their institutional preferences dovetail with other actors—can help explain variation in lawmaking efficiency across polities. It can also shed light on how perceptions about bureaucratic strength form. In systems where bureaucratic mediation plays an important role, civil servants will appear powerful, when perhaps they are only influential. That is, while they are constrained in their ability to change policy outcomes, their interventions have serious consequences for the dynamics of the policy-making process. The present research allows us to draw similar conclusions in the context of IGOs. Furthermore, because IGO staff typically have less control over implementation than their counterparts in national bureaucracies, supporting actors in IGOs may wield a greater proportion of their influence through mediation than via other channels. Thus, the present work dovetails with institutionalist arguments that hold that international actors can modulate interstate bargaining through information provision (Keohane 1984, Koremenos, Lipson & Snidal 2001) and contrasts with the intergovernmentalist contentions about the powerlessness of non-state actors in international interactions (Moravcsik 1998, Moravcsik 1999).

Chapter 3

The Power of Suggestion: Commission Influence on Voting in the European Parliament

No legislators go about the business of lawmaking in a bubble. A variety of sources pepper lawmakers with a constant barrage of information intended to influence their voting behavior. For example, donors indicate support for particular policies, government bureaucrats and lobbyists educate politicians about complicated issues and make recommendations about policy priorities, and party whips cajole lawmakers to follow the party line. When making voting decisions, legislators must weigh the information they receive from external sources against their own private beliefs about the likely outcomes of particular policies. Likewise, they must balance their personal policy priorities with the wishes of actors who may influence their future careers, such as party leaders, campaign donors, and constituents. Furthermore, the way in which politicians balance their ideological predilections and previously held beliefs with new information and external leverage is fundamental to our understanding of democratic lawmaking. Indeed, it would be nearly impossible to exhaustively catalogue the range of academic studies of legislating that deal with this class of questions.

Disentangling preference-congruence from influence is a daunting problem for observational studies of lawmaking in real legislatures, especially in studies of voting behavior, one that often appears futile (Hall 1992). For example, when one observes Democrats and Republicans in the US House voting largely along party lines, should one conclude that parties influence voting behavior or simply that co-partisans have similar preferences (see e.g. Krehbiel 1993)? Similarly, when a lawmaker votes in a manner that pleases a campaign donor, how can one tell whether she would have voted differently had the donor not been in the picture? Sometimes a natural experiment presents itself that helps to overcome

this problem.[1] When one does not, researchers often use an external measure of ideology—perhaps from a survey[2] of legislators—as a control in a model predicting voting behavior. Unfortunately, this approach is often flawed and rarely convenient. The ideological component of a voting decision is driven by the relative utility that the voter derives from the alternatives up for vote, not simply the voter's position in the ideological space. Therefore, because of variation in the location of alternatives across votes, casting a vote against one's ideological preference can be significantly more painful for a legislator on certain votes than it is on others. In other words, sometimes switching one's vote represents a substantial concession to outside interests because the legislator greatly prefers one voting alternative to the other. On the other hand, sometimes both voting alternatives are so close, or so far, from a legislator's preferred policy that pleasing a donor or following the whip is no hardship for the lawmaker, even if she would have voted in the opposite direction in the absence of an outside force. And of course, when alternatives are such that the legislator and outside interest have common preferences, voting in line with external pressure requires no sacrifice from the lawmaker at all. Furthermore, it is typically difficult, if not impossible, to obtain direct measures of legislators' relative utility from voting alternatives. This means that many studies of external influences on voting behavior run the risk of attributing too much, or too little, importance to outside influences because they treat all voting decisions as equal. And finally, in many cases, it is difficult to get reliable external estimates of legislator ideology in the first place, leaving the researcher with little leverage over the question of interest.

Nonetheless, when external pressure is applied selectively across observed votes—for example, if parties whip only particular divisions or a donor expresses interest in only a subset of issues on which lawmakers consider legislation—the voting record itself provides researchers with an exploitable natural experiment. Specifically, because the ideological

[1] See Bronars & John R. Lott (1997) and Stratmann (2002) for some examples from the campaign finance literature.

[2] Of course, it is not uncommon for researchers to use behavioral measures such as NOMINATE (Poole & Rosenthal 1985) scores in this context despite the potential endogeneity problem.

component of voting behavior should remain constant across treated (those votes upon which the outside influence has communicated a preference to the voter) and untreated (those votes that are of little interest to the potentially influential actor) votes, one can adapt existing statistical tools to identify the relative importance that ideology and a given form of external influence play in legislators' voting decisions, taking cross-vote variation in the trade-off between ideological prerogatives and external pressure into account.[3]

In this paper I take advantage of the extensibility of the Bayesian statistical voting model (Clinton, Jackman & Rivers 2004) to identify the extent to which members of the European Parliament (MEPs) follow the recommendations of the European Commission when voting on amendments to proposals considered under the codecision procedure,[4] net of the ideological congruence between the Commission recommendations and individual MEPs' prefrences. While previous work (Tsebelis, Jensen, Kalandrakis & Kreppel 2001, Rasmussen 2003, Kasack 2004) has pointed out the strong correlation between Commission opinion and the success of amendments to codecision bills, my findings demonstrate that this correlation—at the least the portion driven by voting in the EP—is not simply an artifact of preferences shared between actors. By taking advantage of the fact that MEPs vote frequently on questions—such as resolutions and own initiative reports–on which the Commission provides no opinion, and over which the Council holds no veto, I show that many MEPs alter their voting behavior systematically when the Commission presents arguments. Specifically, MEPs hailing from parties in national government—and who are not too ideologically distant from the pivotal players in the Council—make compromises in light of Commission opinions that are likely to facilitate the speedy conclusion of lawmaking negotiations between EU institutions. The Commission's recommendations reflect the bicameral bargaining game between the Parliament and Council that is fundamental

[3]As with all "natural experiments" the viability of this approach is limited by the extent to which treated and untreated votes are truly comparable in all other respects.

[4]Codecision, now also known as the ordinary legislative procedure, is currently the primary procedure under which legislation is developed in the European Union (EU). See European Commission (2010) for an exhaustive explanation of the current version of the codecision process.

to codecision and MEPs use the Commission's opinions to tune their vote choices to the necessary compromises of bicameral lawmaking. However, the Commission does not simply channel the Council when making arguments. It plays a mediating role between the two institutions and provides valuable information to MEPs in its own right. Firstly, because the Commission controls a substantial bureaucracy, it has access to better information about the likely effects of policy choices than do rank-and-file MEPs. MEPs may therefore look to the Commission for guidance on votes concerning issue areas in which they have little personal expertise. Secondly, the Commission has better access to the Council than do EP representatives, and therefore more clearly understands the preferences of Council members than does the average MEP; Commission opinions can serve to transmit information about these preferences to MEPs, facilitating efficiency and compromise in lawmaking.

To obtain these results I build on the standard spatial voting model by incorporating external influence components into MEPs' utility functions, explicitly modeling the relative contributions of internal ideology and external, Commission-driven, influence to MEPs' voting behavior. Therefore, unlike standard approaches, this estimation method is firmly grounded in an explicit model of legislator utility maximization. Linking theory and empirical analysis in this way generates a variety of advantages. First of all, this method takes into account not only MEP ideology but also MEPs' ideological distances from voting alternatives when generating external influence estimates, overcoming the problem I discussed above. Moreover, the method provides researchers with a tool that greatly enhances their ability to interpret the statistical results. Specifically, the model produces estimates of Commission influence that lie on the same scale as MEP ideal points, allowing for comparisons of the relative import of ideology and external pressure in determining MEPs' voting decisions and providing a tool for quantifying external influence that generates substantively meaningful results. Finally, the tools I discuss in this paper also improve our ability to generate NOMINATE-style (Poole & Rosenthal 1985) maps of legislator ideal points and estimates of bill locations, cut-points, or other parameters of interest. Because standard spatial scaling

techniques treat votes as a function of purely ideological voting, they tend to mis-interpret the voting record, allowing external influences to contaminate inferences about ideal points and bill locations. By explicitly modeling the contribution of outside influences on voters' utility functions, one improves the model's other outputs, generating ideal points and bill parameters that are, at least, *less* contaminated by non-ideological considerations.

In what follows, I first discuss codecision and the role that intercameral bargaining plays in EU lawmaking and MEPs' voting calculus, arguing that certain lawmakers will strategically modify their voting behavior under codecision in order to avoid transaction costs and to facilitate compromise between the EU's policy-making bodies. Next, I describe a theory of information transmission between the Commission and Parliament in which the Commission plays a mediating role between the Council and EP, and discuss the implications this theory has for voting patterns in the Parliament. I then describe the methodology and demonstrate how to disentangle ideology from influence on votes over amendments to codecision legislation and introduce the data that I use to test the theory. I explore the model's results, evaluate the consistency of these findings with my theory of information transmission between Commission and Parliament, and examine how augmenting standard ideal point estimation technology with outside influences affects bliss point estimates and other model parameters. Finally, I conclude and describe a variety of other possible applications for the statistical techniques that I demonstrate here.

3.1 Bicameral Lawmaking and Voting in the EP

Lawmakers in bicameral systems experience an interesting trade-off when voting on legislation. On the one hand, politicians in one house of a bicameral legislature face all the standard incentives that occupy legislators; they have personal policy preferences, their party leaderships have expectations about their voting behavior, and they have constituencies—in some cases voters, in others selectorates within their parties—whom they wish to please in order

to win re-election. On the other hand, they must balance their immediate voting preferences with the realities of lawmaking in a bifurcated institution. While they might support a given proposal or amendment for a variety of reasons, they must consider whether or not voting for—or against—a measure will serve them well in negotiations with their partners in lawmaking in the other legislative chamber. This state of affairs is not so different from that facing politicians in unitary systems, in the sense that coalition-building generally requires compromises and trade-offs that sometimes force legislators to temper their voting behavior. Often, lawmakers must support second-best outcomes in order to obtain some movement on policy when their most preferred proposals lack sufficient support or would add aspects to a bill that would fracture coalitions constructed to support the legislation as a whole. Nonetheless, legislators in bicameral systems are likely to face transaction costs that dwarf those weathered by unitary politicians.

For example, under the codecision procedure in the EU, Parliament trades proposals with the Council until both houses agree on the form of legislation. The process takes up to three readings. Initially, the Commission makes a proposal to change policy in some way. The Parliament may then amend the proposal before forwarding it on to the Council. If the Council agrees to the proposal the process ends, but it may further amend the legislation and send it back to the Parliament. The second reading is similar, with the Parliament accepting, rejecting, or amending the Council's counter-offer, while the Council can accept or reject Parliament's second proposal should things come to that. Finally, if the Council rejects Parliament's second reading proposal, the conciliation committee convenes, bringing together representatives from both chambers to hammer out a deal. Should they agree on a compromise, they send the joint text back to both houses for final approval or ultimate rejection. This is a complicated, costly, and lengthy process. Indeed, codecision files can drag on for years. While mistakenly scuttling a compromise in a unicameral legislature can have serious consequences for the legislators involved, it is likely to be easier to work around the impasse—for example, by proposing a new amendment to break the deadlock—than in

a bicameral system. Should the Parliament forward a proposal to the Council that asks for more than the Council is willing to give, a subsequent reading is likely to result. At first reading, this means another round of committee meetings, another debate in plenary, and another round of voting in the EP. A second reading rejection of Parliament's proposal forces the convocation of the conciliation committee which adds significant weight to the affected MEPs' already substantial work-loads.

In both cases, the staffs of EP party groups and national party delegations spend precious resources examining the report, weighing competing interests within the group, and generating voting recommendations for their members on all the amendments that they—and other groups—might table. For groups—and parties, and MEPs—with an interest in passing legislation, this is time better spent working on new bills. Therefore, many MEPs are likely to face strong pressure, both from their party groups and their national party leaderships, to line up behind compromises on codecision legislation, even when this means sacrificing their ideological purity to some extent. This should be especially true for MEPs who hail from mainstream parties, particularly parties that are part of national governments and, therefore, have representation on the Council. Even for parties that are ideologically out-of-step with the pivotal voters in the Council, the weight of responsibility may loom large. These parties are expected to get things done—both at home and in Europe—and they should lean on their MEPs to back policies that have a chance to win Council support and to shy away from unsustainable proposals, at least when compromise is not too ideologically painful. Of course, when a party or group is too at odds with the Council, compromise will be untenable, so this pressure should hold only for MEPs hailing from parties that can stomach policies that are also acceptable to the pivotal Council voter.

Additionally, members of the parties that do not participate in government, especially those on the political fringe, will have little reason to modulate their behavior. Small parties are unlikely to be blamed for the EU's inability to pass legislation and, therefore, are largely free to ignore practical considerations when casting votes. This freedom should be most

notable among extreme, and eurosceptic, parties. Their ideological leanings, and those of their supporters, are likely to be too out of line with the Council's to make compromise profitable. These MEPs have nothing to gain from moderating their votes to help bills through the process and they should adopt voting strategies based on position-taking incentives, holding tightly to their ideological positions. Indeed, if they do break from their preferences it should be to advertise their distinctness from the European mainstream, voting against compromises that obtain substantial majorities, even when supporting a given proposal would not cause them undue ideological aggravation.

Legislators in bicameral legislatures also face an information problem that can tend to exacerbate the tendency to incur bargaining costs in such settings. Of course, politicians who sit together in a unitary legislature will often have trouble observing the bargaining strength—in terms of patience, internal agreement, and the relative salience placed on the issue at hand—of opposition players. Negotiators generally prefer to appear strong—or at least unable to budge—and will not easily relay their willingness to compromise to their opponents. Nonetheless, unicameral legislators interact with each other on a daily basis, participate directly in debates with one another, and have many opportunities to catch glimpses of their opponents' negotiating weak points. This is not true in bicameral legislatures. While party memberships can provide linkages across chambers (Høyland 2006), individual legislators lack the points of contact with one another that they enjoy in unitary systems. Therefore, they may be less able to gauge the bargaining positions—and strengths—of their counterparts in the other chamber. Lacking information, politicians in one house run the risk of proposing legislation that cannot capture the pivotal members of their counterpart institution, precipitating delay. This problem is perhaps especially notable for the EP. While the Parliament holds its sessions in public, records many of its votes, and generally behaves in a transparent manner, the same cannot be said for the Council. Indeed, the Council is a largely closed institution that rarely opens up its internal workings to outside observers. Thus, the potential to misgauge the Council's bargaining resilience is a real

one for MEPs. Nonetheless, they have somewhere to turn to help narrow the information gap: the Commission.

3.1.1 Information, Commission Mediation, and Influence

After initially proposing legislation on a particular issue, the Commission plays a largely supporting role in codecision. Commission representatives take part in informal communications—called trialogues—between the representatives of the Parliament and the Council and the responsible Commissioner generally attends Parliamentary committee meetings and plenary debates on codecision legislation. Furthermore, as I mentioned above, the Commission lodges opinions on Parliament's amendments at both readings.[5] Additionally, should the process proceed all the way to conciliation, Commission delegates serve as facilitators in drafting compromise legislation that is palatable to both the Parliament and the Council and generally work to assure that the process ends with adoption. Yet, the Commission's only formal institutional role in codecision is its ability to set the Council's voting rule by taking negative opinions on Parliament's amendments, forcing the Council to vote unanimously—as opposed to by qualified majority—to adopt positions that are at odds with the Commission's official opinions. Furthermore, because the Commission has no veto power in conciliation, formal spatial models of law-making in the EU categorically find that the Commission has no say over policy outcomes on codecision legislation (see e.g. Crombez 1997).

Nonetheless, these results are based on full-information models of lawmaking and do not consider the Commission's potential role as a broker of information during the course of the legislative process. First of all, while the highest echelons of the Commission are populated by political appointees, each ministry within the Commission—called directorates-general or DGs—employs hundreds, and in some cases thousands, of experienced, full-time, civil service staff (Nugent 2001). These career civil servants provide the Commission with a deep

[5]The Commission also provides a detailed opinion of the common position after it is approved by the Council, should the proposal reach that stage.

reservoir of expertise over all aspects of European policy-making and potentially provide the Commission with both a better understanding of the likely outcomes of particular policies and a more comprehensive view of the policy options available to the Union than that available to the Council, and especially to the average MEP.[6] Indeed, when practitioners were surveyed about the relative legislative advantages afforded to the three EU lawmaking institutions, respondents highlighted the Commission's "great expertise on the policy areas affected by its proposals" (Thomson & Hosli 2006, pp. 398) as one of the its key resources. Furthermore, in order to initially draft codecision legislation, the Commission is forced to invest in substantial proposal-specific information at the start of the legislative process. Therefore, both Council and Parliament members have incentives to conserve their own resources by taking advantage of the Commission's informational investment when they believe that they can trust, or at least discern the truth from, Commission signals. Therefore, we should expect MEPs to turn to the Commission for guidance when they lack information about a particular vote, perhaps when a policy under consideration falls outside a given MEP's area of expertise.

The Commission's informational advantages also extend to knowledge about the inner workings of the Council, providing MEPs with the potential to overcome their knowledge deficit vis-a-vis the Council and reduce the likelihood of extended bicameral bargaining. Indeed, the Commission has direct access to every level of decision-making within the Council, with representatives sitting in on working group meetings, COREPER deliberations, and convocations of ministers (Cini 1996, Nugent 2001). Therefore, when it is in its interests, the Commission may relay information about the Council's bargaining position to the Parliament. In the context of codecision opinions, the Commission can let the Parliament know when a particular proposal pushes too far, or when a tabled amendment under-sells the Parliament's point of view, conceding too much to a willing compromiser. Nonetheless, the

[6]Councilors are cabinet members at the national level and therefore have access to extensive staffs of their own. MEPs, on the other hand, typically only have a handful of full-time staff at their disposal, and these staff are often shared between MEPs of national delegations (Corbett, Jacobs & Schackleton 2003).

Commission's opinions are, largely, a form of cheap talk. The general result of models of cheap talk (Crawford & Sobel 1982, Austen-Smith & Banks 2000) is that a message sender can credibly pass information to a receiver only when the preferences of the two parties dovetail. With respect to the Commission and EP, this means that Commission opinions will carry the most weight with MEPs when the pivotal (median) MEP and the Commission have preferences that are more similar than those of the Commission and Council. When Commission biases turn towards the pivotal MEP it will be in its interests to provide accurate information about the Council's bargaining strength, allowing the Parliament to craft a proposal that is as aggressive as possible without forcing a rebuttal by the Council, and bargaining delay. On the other hand, when the Commission sides with the Council it will wish to make the Parliament believe that the Council is strong, encouraging undue compromise from MEPs. Measuring preference congruence across institutions is a difficult proposition. Nonetheless, researchers commonly argue that the Commission and Parliament, because they are European institutions containing politicians with purely European offices, are more supportive of proposals that expand European power than is the Council, which is composed of ministers from national governments (Tsebelis & Garrett 2000). Therefore, we should expect the Commission to most effectively transmit information to the Parliament when discussing integration-focused proposals, with less communication between the institutions on measures dealing with traditional, left-right, cleavages.

The final reason that Commission opinions may carry weight with MEPs is not a function of information transmission, but rather institutional power. As I noted above, Commission opinions technically set the voting rule in the Council. Most Council decisions are made by qualified majority, but the Council must vote unanimously when adopting positions that differ from the Commission's amended proposal at first reading,[7] or when accepting Parliamentary amendments that the Commission rejects at second reading. Therefore, negative Commission opinions may carry special weight. Because of their influence over the vot-

[7]The amended proposal incorporates all of the Commission's opinions on Parliament's amendments.

ing rule, the Commission can make it difficult for the Parliament and Council to include certain provisions in final legislation. MEPs with a vested interested in actual lawmaking, therefore, may choose to vote against provisions that they support ideologically when the Commission rejects them in its opinion, because they expect to lose the argument in the end anyway. As Rasmussen (2003, pp. 7) quotes one Council official: "...usually, we examine the amendments supported by the Commission, for example if there are 100 amendments and 60 are supported by the Commission, it is useless to examine the others, but it happens sometimes that we do." Therefore, we should expect negative Commission opinions to carry more weight with MEPs than positive ones.

3.2 Statistical Models of Roll Call Voting

As I argued in the introduction, it is difficult to say anything meaningful about how external actors influence voting in legislatures without effectively controlling for the ideological determinants of vote choice. In this section I describe a series of statistical models designed to examine how Commission opinions on EP amendments influence MEPs' voting decisions, taking ideological motivations into account.[8] These models are all generalizations of, or variations on, the canonical statistical roll call voting model developed by Clinton, Jackman & Rivers (2004). Therefore, to develop the logic of the various Commission-influence models introduced here, I begin by describing the canonical model, to which I add a hierarchical prior specification, simplifying model identification, enhancing our ability to interpret the model's results, and improving efficiency.

I next present two extensions to the canonical model that allow the utility that voters (e.g. MEPs) derive from voting yes or no on a given measure to vary as a function of an outside actor's opinion (e.g. the Commission's opinion on a Codecision amendment) while taking standard spatial voting considerations (i.e. the voters' ideal points and votes' proposal

[8]Some of the models that I develop here are applicable to a wide array of political setting beyond lawmaking in the EU. I describe a few of these possible applications in the conclusion of this article.

and status quo points) into account. These models allow one to see whether or not MEPs' voting behavior co-varies with Commission opinions, net of the ideological similarity between individual MEPs and Commissioners. That is, these models make it possible to see whether or not the tendency of MEPs to cast votes that are in line with the Commission's stated opinions is simply a function of ideological congruence between the bulk of the MEPs and members of the Commission, or whether extra-spatial mechanisms are at work. Additionally, because these models nest the canonical model as a special case, it is possible to use standard model selection tools to evaluate the explanatory power of the opinion-augmented approach vis-a-vis the standard model. Furthermore, these models make it possible to examine the conditions under which Commission opinions and MEP votes dovetail, and those situations when the Commission's opinion has little relationship to MEP voting behavior.

3.2.1 The Canonical Statistical Roll Call Voting Model

Following Clinton, Jackman & Rivers (2004) I assume that we can represent each roll call vote in terms of two points in a D-dimensional policy space. Specifically, MEPs must choose between the "Yea" position $\boldsymbol{\zeta}_\mathbf{j}$ and the "Nay" outcome $\boldsymbol{\psi}_\mathbf{j}$ on each of $j \in 1 \ldots m$ votes. Similarly, I assume that each MEP has a quadratic utility function such that legislator $i \in 1 \ldots n$ with ideal point $\mathbf{x_i}$ derives utility

$$U_i(\boldsymbol{\zeta}_\mathbf{j}) = -||\mathbf{x_i} - \boldsymbol{\zeta}_\mathbf{j}||^2 + \eta_{ij} \tag{3.1}$$

from voting for passage, and

$$U_i(\boldsymbol{\psi}_\mathbf{j}) = -||\mathbf{x_i} - \boldsymbol{\psi}_\mathbf{j}||^2 + \nu_{ij} \tag{3.2}$$

for voting for rejection. Again following Clinton, Jackman & Rivers (2004), I assume that the stochastic parts of the utility function, η_{ij} and ν_{ij}, are independent with respect to

both MEPs and votes and normally and jointly distributed with mean $E(\eta_{ij}) = E(\nu_{ij})$ and variance $\text{var}(\eta_{ij} - \nu_{ij}) = \sigma^2$.[9]

Given the $n \times m$ roll call matrix $\mathbf{Y}$—where $Y_{ij} = 1$ when MEP i votes yea on vote j and $Y_{ij} = 0$ when the same legislator votes nay[10] on the vote in question—the probability that MEP i votes in the affirmative on vote j is

$$P_{\text{can}}(Y_{ij} = 1) = \Phi\left(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x}_{\mathbf{i}} - \boldsymbol{\kappa}_{\mathbf{j}})\right) \tag{3.3}$$

where $\boldsymbol{\kappa}_{\mathbf{j}} = \frac{\psi_{\mathbf{j}}+\zeta_{\mathbf{j}}}{2}$ is the cut-point dividing MEPs who support measure j from the who do not, $\boldsymbol{\beta}_{\mathbf{j}} = \frac{2(\zeta_{\mathbf{j}}-\psi_{\mathbf{j}})}{\sigma_j}$ describes the extent to which vote j discriminates between voters,[11] and $\Phi(\cdot)$ is the standard normal distribution function. Multiplying across MEPs and votes yields the likelihood function

$$L(\boldsymbol{\theta}_{\mathbf{can}}|\mathbf{Y}) = \prod_{i=1}^{n}\prod_{j=1}^{m} \Phi\left(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x}_{\mathbf{i}} - \boldsymbol{\kappa}_{\mathbf{j}})\right)^{Y_{ij}} \times \left[1 - \Phi\left(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x}_{\mathbf{i}} - \boldsymbol{\kappa}_{\mathbf{j}})\right)\right]^{1-Y_{ij}} \tag{3.4}$$

where $\boldsymbol{\theta}_{\mathbf{can}} = \{\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}\}$.

At this point it would be traditional to adopt uninformative conjugate prior distributions for all parameters in the model and to identify the model by specifying spike priors for particular ideal points and/or bill parameters. This can be a complicated task, especially when dealing with a multi-dimensional policy space (Jackman 2001). Therefore, I use a different approach, recommended by Bafumi et al. (2004), that takes advantage of the wealth of available prior information on MEPs' ideological predilections. Specifically, I use the prior

[9] It is possible to make the less restrictive assumption $\text{var}(\eta_{ij} - \nu_{ij}) = \sigma_j^2$ here but a strict common error variance assumption greatly facilitates estimation of the influence models that I develop below.

[10] I treat abstentions as missing values in this analysis.

[11] This parameterization of the statistical roll call voting model differs from the standard approach in Clinton, Jackman & Rivers (2004) and instead follows Bafumi, Gelman, Park & Kaplan (2004). I use this parameterization because it makes it easy to work with vote cut-points, greatly simplifying certain aspects of the analysis and exposition.

specification

$$\begin{aligned} p_{\text{can}}(\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}) &= p(\mathbf{x})p(\boldsymbol{\beta})p(\boldsymbol{\kappa}) \\ &= \prod_{i}^{n} \mathcal{N}_D(\boldsymbol{\mu}_{\mathbf{x_i}}, \boldsymbol{\Sigma}_{\mathbf{x}}) \prod_{j}^{m} \mathcal{N}_D(0, \sigma^2_{\boldsymbol{\beta}_\mathbf{j}} \mathbf{I}_D) \prod_{j}^{m} \mathcal{N}_D(0, \sigma^2_{\boldsymbol{\kappa}_\mathbf{j}} \mathbf{I}_D) \end{aligned} \tag{3.5}$$

where I assume that $\sigma^2_{\boldsymbol{\beta}_\mathbf{j}}$ and $\sigma^2_{\boldsymbol{\kappa}_\mathbf{j}}$ are known a priori but

$$\boldsymbol{\mu}'_{\mathbf{x_i}} = \boldsymbol{\Lambda}\mathbf{c_i} \tag{3.6}$$

and

$$\boldsymbol{\Sigma}_{\mathbf{x}} \sim \mathcal{W}_K^{-1}\left(\upsilon, \sigma_x^2 \mathbf{I}_K\right) \tag{3.7}$$

where $\mathbf{c_i}$ is a $k \times 1$ vector of covariates describing MEP i and each $\boldsymbol{\Lambda}$ is a $D \times k$ matrix of coefficients mapping MEP characteristics into D-dimensional ideal point space. To complete the specification, I use a simple conjugate hyper-prior for the hierarchical coefficient matrix. Specifically, I assume

$$\boldsymbol{\Lambda} \sim \mathcal{N}(0, \sigma^2_{\boldsymbol{\Lambda}} \mathbf{I}_k), \tag{3.8}$$

where $\sigma^2_{\boldsymbol{\Lambda}}$ is known a priori.

Thus, I adopt traditional uninformative conjugate priors for the bill parameters. On the other hand, I use an informative prior for the ideal point vectors. Specifically, I bring additional data to the table, modeling each MEP's prior ideal point distribution in terms of observable covariates. We have a lot of information about MEPs beyond their voting behavior that can help tell us identify their ideological preferences. Notably, Hooghe, Bakker, Brigevich, de Vries, Edwards, Marks, Rovny & Steenbergen (2008) conducted a survey that asked experts on European party systems to provide quantitative ratings of the ideological positions of national parties across a variety of ideological dimensions in 2006, right in the middle of the sixth EP term. While we lack expert judgements about the ideologies of

individual MEPs, we should expect that MEPs' own ideologies will covary closely with those of the national parties as a whole. Furthermore, because each national party delegation is reasonably small, these rating vary substantially across the population of MEPs. Therefore, each $\mathbf{c_i}$ in equation 3.6 holds the vector of ideology scores describing MEP i's party in Hooghe et al.'s (2008) expert survey. Specifically, experts provided ratings of national party viewpoints on the role of government in the *economy*, their ideological stances on *social* issues and civil liberties, and their general positions on European *integration*.[12]

Building expert ideology ratings into the prior specification of the statistical roll call voting model serves a variety of purposes. First of all, it provides a data-driven approach to identifying the model. Traditional identification techniques require the modeler to make strong assumptions about the ideal points of individual MEPs or the manner in which individual roll call votes cut across the ideological spectrum. Especially when dealing with multi-dimensional issue spaces, it can be difficult to select spike priors that identify the model in the first place. Furthermore, even when identification seems to have been achieved, poor choices by the analyst can skew results considerably. In contrast, the data-driven approach effectively identifies the space without undue interference by the researcher.

Moreover, the hierarchical modeling approach that I adopt here enhances our understanding of ideological picture painted by the model. We commonly ascribe meaningful

[12]Raters placed each party on a scale from 0 to 10 on the first two dimensions. The *economy* ratings ranged from "extreme left" to "extreme right" where "... [p]arties on the economic left want government to play an active role in the economy. Parties on the economic right emphasize a reduced economic role for government: privatization, lower taxes, less regulation, less government spending, and a leaner welfare state." The *social* issue scores again range from "extreme left" to "extreme right" where parties on the left "... favor expanded personal freedoms, for example, access to abortion, active euthanasia, same-sex marriage, or greater democratic participation," while right parties "... often reject these ideas; they value order, tradition, and stability, and believe that the government should be a firm moral authority on social and cultural issues." Finally, experts scored each party's position on European *integration* from 1 to 7 where parties with a score of 1 were "strongly opposed" and parties given a 7 were "strong in favor" European integration in 2006 (Hooghe et al. 2008, pp. 9–10). The survey also asked experts to score the party ideology on much finer-grained dimensions, coding attitudes towards the internal European market, redistribution, multiculturalism, and EU foreign policy, to name only a few issues. Using these scores generates results that are quite similar to those produced by a model incorporating only the broad ideological ratings, but high multicollinearity between the fine-grained scores makes interpretation difficult. Therefore, I focus my attention on the broad issue categories.

descriptions to the issue dimensions that statistical roll call voting models spit out. For example, in the EP literature, researchers often argue that the dominant dimension produced by various scaling techniques captures the traditional left-right political divide, and maintain that the second-most-dominant dimension describes MEPs' positions on the breadth and depth of European integration (Hix, Noury & Roland 2006, Hix, Noury & Roland 2007). Of course, researchers make these judgements by extrapolating from the pattern of MEP ideal points implied by the model, using their own knowledge of MEPs' ideological orientations to tell a post-hoc story about the model's output. Using prior expert ratings of MEPs' general ideological leanings allows for a more scientific approach to ideal point interpretation. Indeed, after estimation, $\mathbf{\Lambda}$ describes the relationship between each ideological dimension rated by the experts in Hooghe et al.'s (2008) survey and the dimensions produced by the model. In other words, we can use $\mathbf{\Lambda}$ to see whether MEPs' positions on the dominant dimension produced by the statistical voting model are consistent with their parties' stances on the traditional left-right divide. Similarly, the prior specification allows us to examine the relationship between MEP's parties' viewpoints on European integration and their voting behavior. Furthermore, the estimated variance-covariance matrix for the hierarchical prior, $\mathbf{\Sigma_x}$, describes the extent to which the theoretical concepts measured by each $\mathbf{c_i}$ account for variation in MEPs' estimated ideal points. Therefore, as I will demonstrate below, we can use the hierarchical approach to link the atheoretical output of the statistical roll call voting model to explicit ideological concepts.

Finally, incorporating additional information about MEPs into the prior specification improves the efficiency of the statistical model, reducing the error in estimated ideal points. Additionally, the model demonstrates how to combine two forms of information about MEP ideology—subjective expert ratings and objective information about voting behavior—into a single measure of legislator ideology.

Figure 3.1 presents the results of fitting the canonical model to EP voting data[13] from

[13]I describe the data in more detail, below.

Figure 3.1: MEP ideal points, as estimated by the canonical model.

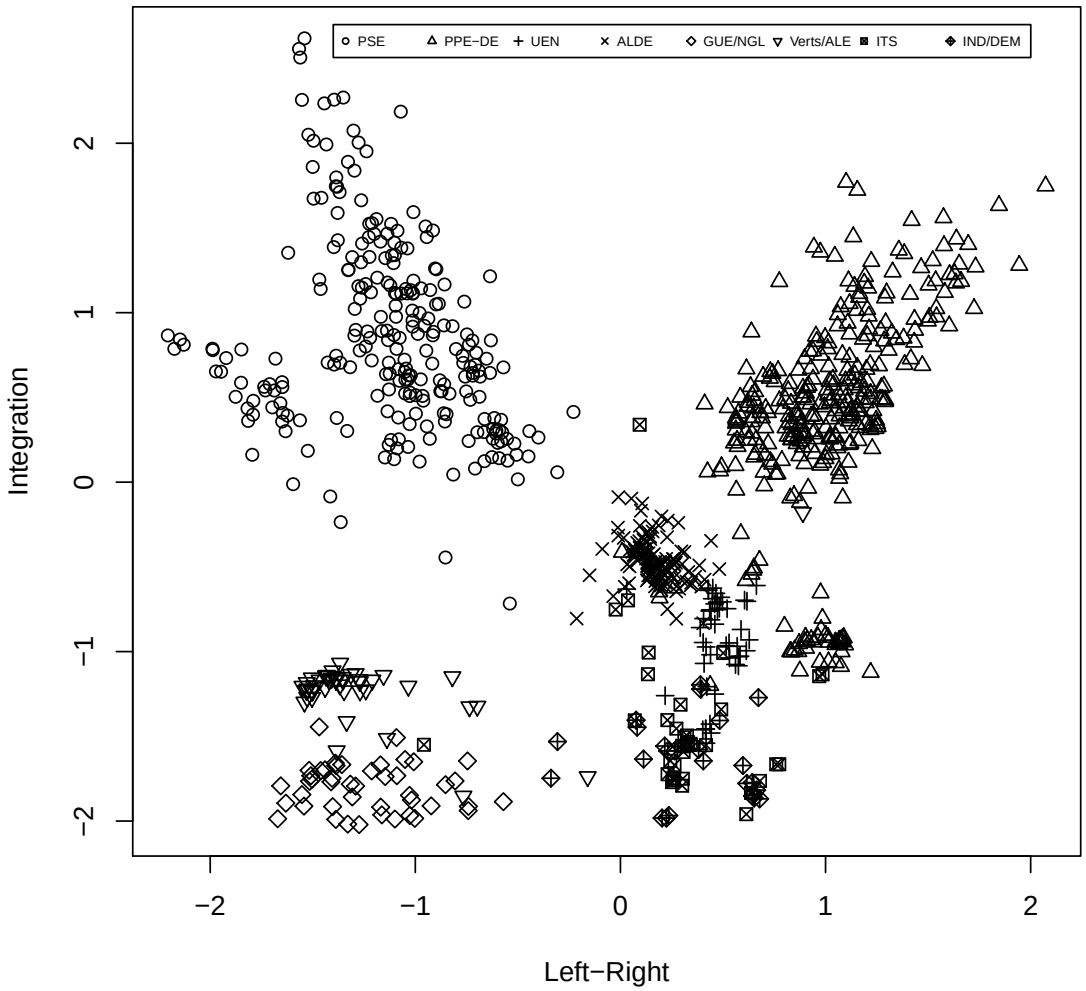

Table 3.1: Explaining the dimensions of EP issue space.

	Left-Right	Integration
Intercept	-1.37 (0.08)*	1.48 (0.10)*
Economy	0.26 (0.01)*	0.04 (0.01)*
Social	0.13 (0.01)*	0.01 (0.01)
Integration	0.12 (0.01)*	0.34 (0.02)*

Parenthetical values are posterior standard deviations and the stars indicate 95% HPD intervals exclude zero.

the 6th term, displaying MEP ideal point estimates for a 2-dimensional model, and painting a picture of EP ideology that is largely in tune with previous research. I have labeled the horizontal dimension "left-right" and the vertical dimension "integration," in keeping with standard custom in EU studies (Hix, Noury & Roland 2007), but the results of the hierarchical model indicate that these terms only partially capture the true meaning of these dimensions. Table 3.1 displays the $\mathbf{\Lambda}$ estimates generated by the fitted model, showing how expert ratings of MEPs' national party ideologies map into the estimated issue space. The results show that the left-right dimension does, indeed, tap strongly into left-right preferences and both the coefficients for attitudes towards the *economy* and *social* policy are statistically significant. Note, furthermore, that economic considerations play a stronger role in this dimension than do social values—the difference between the *economy* and *social* coefficients is statistically significant. This makes sense; the EP has less influence over European social policy than it does economic policy and recorded votes are likely to reflect this aspect of the agenda. On the other hand, *integration* ideology is also significantly related to this dimension and one cannot distinguish between the magnitude of this relationship and the degree of covariance between expert ratings of ideology on social issues and ideal points on this axis. Therefore, because this primary dimension of conflict in the EP corresponds to aspects of both left-right ideology and preferences over European integration, we should take care when interpreting its meaning.

Turning to the second dimension, table 3.1 shows that it is largely related to parties' attitudes towards European integration. While the coefficient for *economy* is also statistically significant, the *integration* coefficient is 8.5 times its size. Thus, we can be quite confident when treating this dimension as a measure of MEP's attitudes towards European integration, assuming we trust the expert ratings upon which this conclusion is based. Interestingly, the model estimates $\mathbf{\Sigma_x} = \begin{bmatrix} 0.39 & -0.05 \\ -0.05 & 0.60 \end{bmatrix}$, indicating that expert ratings explain more of the variance on the first dimension than they do the second, and this difference is statistically significant. Therefore, we may wish to turn to other sources of prior information to help us understand what the "integration" dimension really captures.

3.2.2 A Baseline Commission Influence Model

I next extend this model by allowing Commission opinions to influence MEPs' preferences through a modification of their utility functions. Specifically, I extend equations 3.1 and 3.2 by assuming that legislator i obtains utility

$$U_i(\boldsymbol{\zeta}_\mathbf{j}) = -||\mathbf{x_i} - \boldsymbol{\zeta}_\mathbf{j}||^2 + a_j \cdot \delta_i^{aa} - r_j \cdot \delta_i^{ra} + \eta_{ij} \tag{3.9}$$

from voting yea on j, and

$$U_i(\boldsymbol{\psi}_\mathbf{j}) = -||\mathbf{x_i} - \boldsymbol{\psi}_\mathbf{j}||^2 - a_j \cdot \delta_i^{ar} + r_j \cdot \delta_i^{rr} + \nu_{ij} \tag{3.10}$$

from voting to reject the tabled measure. Here, a_j (r_j) is an indicator variable that equals one when the Commission accepts (rejects) the "Yea" position on vote j and equals zero otherwise,[14] δ_i^{aa} (δ_i^{ra}) is the utility that MEP i stands to gain (lose) from voting yes on

[14]Note that $a_j = r_j = 0$ on votes, such as votes on amendments to resolutions or own-initiative reports, where the Commission does not lodge an opinion. I currently assume that the Commission cannot simultaneously accept and reject an amendment, although this coding might make sense for partially accepted motions. For the moment, I leave this extension to future research.

a measure that the Commission supports (rejects), and δ_i^{ar} (δ_i^{rr}) represents the utility loss (gain) to MEP i from voting against when the Commission counseled acceptance (rejection). This modified utility function leads directly to the following equation for the probability that MEP i casts a yea vote on measure j:

$$P_{\text{bin}}(Y_{ij} = 1) = \Phi\left(\boldsymbol{\beta_j}(\mathbf{x_i} - \boldsymbol{\kappa_j}) + a_j \cdot \delta_i^a - r_j \cdot \delta_i^r\right) \tag{3.11}$$

where $\boldsymbol{\beta_j} = \frac{2(\boldsymbol{\zeta_j} - \boldsymbol{\psi_j})}{\sigma_j}$, $\boldsymbol{\kappa_j} = \frac{\boldsymbol{\psi_j} + \boldsymbol{\zeta_j}}{2}$, $\delta_i^a = \frac{\delta_i^{aa} + \delta_i^{ar}}{\sigma}$, and $\delta_i^r = \frac{\delta_i^{ra} + \delta_i^{rr}}{\sigma}$.

In equation 3.11, $\boldsymbol{\delta^a}$ ($\boldsymbol{\delta^r}$) is a vector of MEP-specific fixed effects, capturing each legislator's propensity to follow the Commission's recommendation to vote for (against) an amendment, above and beyond her ideological affinity for the measure.[15] Therefore, a positive (negative) value for δ_i^a (δ_i^r) indicates that MEP i derives positive utility from voting for an amendment that the Commission accepts (rejects) after taking MEP i's proximity to the "Yea" and "Nay" positions for vote j into account. Note that this functional form allows MEPs to be responsive to either positive opinions, negative opinions, or both.[16] The fixed effects are identified by the fact that the Commission provides no opinion on a large proportion of votes. Thus, even when susceptibility to influence and ideology are correlated, the influence-free votes allow the model to separate the two effects.[17]

[15]Note that I am using "ideological affinity" as convenient short-hand here and do not claim that ideal points estimated by this model are "pure" measures of ideology. A host of other factors, including national party and parliamentary party group pressure will influence MEPs' estimated ideal points.

[16]One might reasonably simplify the model by adopting the restriction $\boldsymbol{\delta^a} = -\boldsymbol{\delta^r}$, yielding a single Commission-influence term for each MEP. Nonetheless, there is reason to believe that MEPs might react differently to positive and negative Commission opinions. For one thing, negative opinions by the Commission limit the ability of the Council to accept EP amendments under qualified majority—as opposed to unanimity rule—generating a clear qualitative difference in the institutional impact of the two types of opinions.

[17]I ran a number of simulations to check the model's ability to identify parameters effectively. Specifically, I first simulated roll call matrices from the assumed data generating process described by equation 3.11, drawing parameters from their prior distributions, and attributing influential opinions to half of the simulated votes, and verified the model's ability to recover parameter values. Next, I replicated the process, but attributed an ideal point to the opinion-giver, and simulated opinions, based on that ideal point, from equation 3.3. Furthermore, I held the true fixed influence effects of the simulated legislators to zero. Thus, in this second simulated dataset, the correlation between vote choices and opinions was purely a function of ideological congruence between the voters and opinion-giver. The model did a good job of recovering parameters when applied to this dataset and did not erroneously attribute influence to the opinion-giver,

Comparing equation 3.11 to 3.3 we can see that the canonical model is nested within the baseline influence model. That is, the canonical model is a special case of the baseline influence model where $\boldsymbol{\delta}^{\mathbf{a}} = \boldsymbol{\delta}^{\mathbf{r}} = 0$. Therefore, this model provides for a straightforward test of the proposition that MEP votes and Commission opinions covary for reasons beyond simple ideological congruence: statistically significant estimates of values in the $\boldsymbol{\delta}^{\mathbf{a}}$ and $\boldsymbol{\delta}^{\mathbf{r}}$ vectors would indicate that the relationship between MEP voting behavior and Commission opinion hinges on factors not considered by the canonical spatial voting model. Furthermore, one can use these estimates to see who behaves in such an extra-spatial manner. Note finally, that while many MEPs will tend to vote with Commission to a greater extent than purely spatial factors would predict, other MEPs will tend to deviate from the Commission's opinion in a similar manner, and this model provides a window into both behaviors.

Multiplying across MEPs and votes yields the likelihood function for the baseline influence model,

$$\begin{aligned} L(\boldsymbol{\theta}_{\text{bin}}|\mathbf{Y}, \mathbf{a}, \mathbf{r}) = & \prod_{i=1}^{n}\prod_{j=1}^{m} \Phi\left(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x}_{\mathbf{i}} - \boldsymbol{\kappa}_{\mathbf{j}}) + a_j \cdot \delta_i^a - r_j \cdot \delta_i^r\right)^{Y_{ij}} \\ & \times \left[1 - \Phi\left(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x}_{\mathbf{i}} - \boldsymbol{\kappa}_{\mathbf{j}}) + a_j \cdot \delta_i^a - r_j \cdot \delta_i^r\right)\right]^{1-Y_{ij}}, \end{aligned} \tag{3.12}$$

where $\boldsymbol{\theta}_{\text{bin}} = \{\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}\}$. To complete the model, I again specify hierarchical priors for the ideal points and independent, conjugate prior distributions for all other model parameters:

$$\begin{aligned} p_{\text{bin}}(\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}) &= p_{\text{can}}(\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}) p(\boldsymbol{\delta}^{\mathbf{a}}) p(\boldsymbol{\delta}^{\mathbf{r}}) \\ &= p_{\text{can}}(\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}) \prod_{i}^{n} \mathcal{N}(\mu_{\delta_i^a}, \sigma^2_{\delta_i^a}) \prod_{i}^{n} \mathcal{N}(\mu_{\delta_i^r}, \sigma^2_{\delta_i^r}). \end{aligned} \tag{3.13}$$

consistently generating $\boldsymbol{\delta}^{\mathbf{a}}$ and $\boldsymbol{\delta}^{\mathbf{r}}$ estimates with confidence intervals spanning zero.

3.2.3 A Conditional Commission Influence Model

The baseline influence model provides a window into the extra-spatial voting behavior of MEPs and allows one to identify the extent to which individual MEPs vote for or against the Commission's wishes in an extra-spatial manner. While useful, this model is a poor tool for examining broad trends in MEPs' non-spatial tendencies towards voting with (or against) the Commission or for summarizing the circumstances under which extra-spatial (in)congruence occurs. The baseline influence model can only tell us what each MEP's average extra-spatial tendency is, shedding little light on what drives MEPs' behavior. To better understand what modulates the relationship between Commission opinions and MEP votes I extend the baseline model, developing a conditional influence model that lets one examine the role that arbitrary covariates play in this process. While the baseline influence model provided only for idiosyncratic Commission opinion-following tendencies among MEPs, the conditional approach allows one to model how MEPs with shared characteristics, or legislators facing similar circumstances, behave similarly to one another with respect to Commission opinions. Specifically, I further extend the baseline model by assuming that legislator i obtains utility

$$U_i(\boldsymbol{\zeta_j}) = -||\mathbf{x_i} - \boldsymbol{\zeta_j}||^2 + a_j(\delta_i^{aa} + \boldsymbol{\gamma}^{\mathbf{aa}}\mathbf{z_{ij}}) - r_j(\delta_i^{ra} + \boldsymbol{\gamma}^{\mathbf{ra}}\mathbf{z_{ij}}) + \eta_{ij} \tag{3.14}$$

when voting for measure j and expects utility

$$U_i(\boldsymbol{\psi_j}) = -||\mathbf{x_i} - \boldsymbol{\psi_j}||^2 - a_j(\delta_i^{ar} + \boldsymbol{\gamma}^{\mathbf{ar}}\mathbf{z_{ij}}) + r_j(\delta_i^{rr} + \boldsymbol{\gamma}^{\mathbf{rr}}\mathbf{z_{ij}}) + \nu_{ij} \tag{3.15}$$

when voting against j. This utility specification implies the following equation for the probability that MEP i voters in favor of measure j:

$$P_{\text{cin}}(Y_{ij} = 1) = \Phi\left(\boldsymbol{\beta_j}(\mathbf{x_i} - \boldsymbol{\kappa_j}) + a_j(\delta_i^{a} + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{ij}}) - r_j(\delta_i^{r} + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{ij}})\right) \tag{3.16}$$

where $\boldsymbol{\beta_j} = \frac{2(\boldsymbol{\zeta_j} - \boldsymbol{\psi_j})}{\sigma}$, $\boldsymbol{\kappa_j} = \frac{\boldsymbol{\psi_j} + \boldsymbol{\zeta_j}}{2}$, $\delta_i^a = \frac{\delta_i^{aa} + \delta_i^{ar}}{\sigma}$, $\delta_i^r = \frac{\delta_i^{ra} + \delta_i^{rr}}{\sigma}$, $\boldsymbol{\gamma}^{\mathbf{a}} = \frac{\boldsymbol{\gamma}^{\mathbf{aa}} + \boldsymbol{\gamma}^{\mathbf{ar}}}{\sigma}$, and $\boldsymbol{\gamma}^{\mathbf{r}} = \frac{\boldsymbol{\gamma}^{\mathbf{ra}} + \boldsymbol{\gamma}^{\mathbf{rr}}}{\sigma}$.

Thus, as in the baseline influence model, the probability that a MEP votes yes on a given vote is a function of spatial considerations and the Commission's opinion, as indicated by a_j and r_j. But while the baseline model allowed only for an idiosyncratic sensitivity to the Commission's opinion, the conditional influence model allows MEPs to share a general responsiveness to Commission recommendations, as a function of a vector of l covariates, $\mathbf{z_{ij}}$, in addition to their individual, idiosyncratic, responses to opinion, δ_i^a and δ_i^r. This shared responsiveness is captured by the coefficient vectors $\boldsymbol{\gamma}^{\mathbf{a}}$ and $\boldsymbol{\gamma}^{\mathbf{r}}$.

For example, one might hypothesize that MEPs are more likely to vote in line with opinions lodged by commissioners who hail from their own national parties than they are to support Commission opinions in issue areas under the purview of commissioners from other parties, or countries. In this case, each z_{ij} would be a simple scalar dummy variable, coded one if MEP i and the commissioner responsible for the bill considered during vote j belong to the same party, and zero otherwise. In turn, a positive (negative) estimate for the single coefficient γ^a (γ^r) would be consistent with the above hypothesis, while a zero or negative (positive) estimate would not be.

As in the baseline model, I allow for differential responses to positive and negative Commission opinions. Thus, when $a_j = 1$ (this implies $r_j = 0$, of course), the probability that MEP i votes in the affirmative is $\Phi(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}) + \delta_i^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{ij}})$. On the other hand, when the Commission voices a negative opinion on an amendment, the probability that the MEP votes yes is $\Phi(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}) - \delta_i^r - \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{ij}})$. And, of course, when the Commission provides no opinion ($a_j = r_j = 0$), MEP i's behavior is purely a function of spatial considerations and she votes yes with probability $\Phi(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}))$.

Again, multiplying across MEPs and votes generates the likelihood function for the con-

ditional influence model,

$$
\begin{aligned}
L(\boldsymbol{\theta}_{\text{cin}}|\mathbf{Y},\mathbf{a},\mathbf{r},\mathbf{Z}) = \prod_{i=1}^{n}\prod_{j=1}^{m} \Phi\left(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i}-\boldsymbol{\kappa}_{\mathbf{j}}) + a_j(\delta_i^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{ij}}) - r_j(\delta_i^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{ij}})\right)^{Y_{ij}} \\
\times \left[1 - \Phi\left(\boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i}-\boldsymbol{\kappa}_{\mathbf{j}}) + a_j(\delta_i^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{ij}}) - r_j(\delta_i^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{ij}})\right)\right]^{1-Y_{ij}},
\end{aligned}
\tag{3.17}
$$

where $\boldsymbol{\theta}_{\text{cin}} = \{\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}, \boldsymbol{\gamma}^{\mathbf{a}}, \boldsymbol{\gamma}^{\mathbf{r}}\}$. Finally, I adopt a prior specification that resembles that of the baseline commission influence model with the addition of improper priors on the regression parameters.[18] Therefore,

$$
p_{\text{cin}}(\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}, \boldsymbol{\gamma}^{\mathbf{a}}, \boldsymbol{\gamma}^{\mathbf{r}}) \propto p_{\text{bin}}(\mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}). \tag{3.18}
$$

I estimated all of the models described here using Gibbs sampling, simulating draws from the conditional posterior distributions of model parameters. Appendix B provides estimation details.

The applicability of these models—and their ability to identify the effects of external pressure—hinges on the availability of votes that are not colored by external influence that interests the analyst. The EP provides an excellent arena in which to apply such tools because it considers both legislative measures—which are subject to the pressures of bicameral lawmaking, and on which the Commission renders explicit verdicts—and non-legislative resolutions and initiatives that have no binding legal ramifications. Indeed, the Parliament still spends the bulk of its voting time considering such non-legislative questions. These votes have no purpose beyond position-taking and allow MEPs to wave their ideological flags without worrying about the practical constraints of lawmaking in the EU. Therefore, they make an excellent set of control votes, and can help researchers to disentangle MEPs' ideological

[18] I have little a priori information about the regression parameters in the context in which I apply this model. Therefore, improper priors provide a reasonable way to incorporate my lack of prior information into the model. Another approach that would yield similar results would be to adopt vague conjugate normal priors for the regression parameters. Given more specific prior information, adopting conjugate priors, or even a hierarchical prior structure could prove useful.

voting motivations from strategic their ones. Of course, political parties value consistency in their members and may apply pressure on both sorts of votes. Thus, I do not claim that MEPs vote in a manner on non-legislative proposals that perfectly reveals their ideological proclivities. I argue only that, because the Commission renders no explicit verdicts on these votes, and because MEPs do not need to worry about potential bicameral bargaining costs when voting on such measures, they represent a useful tool for gaining leverage on questions of external influence over legislative voting in the EP.[19]

3.3 Data

Using the Parliament's online archive (European Parliament 2009*a*), I collected vote data from the 6th EP, covering a period from the beginning of the term in July 2004 through May 2008. The Parliament voted 18,493 times over this period but only recorded 4086 of these votes. I included only votes on codecision amendments and votes regarding own-initiative reports and EP resolutions—both roll calls on amendments and final votes—in the dataset.[20] The Commission lodges opinions only on amendments to Union legislation. Therefore, the codecision amendments are "treated" observations, where Commission opinions have the potential to influence MEP voting, while the votes on the initiatives and non-legislative resolutions serve as a "control" group where Commission opinions can play no role in MEPs' voting decisions.

I augmented the core voting data by collecting Commission opinions on individual code-

[19]Of course, all observational studies are imperfect and one can think of ways in which the natural experiment that I describe here could be compromised. For instance, MEPs may use non-legislative resolutions to signal their intentions on legislative issues, potentially contaminating the results presented here. MEPs may also know how the Commission feels about particular measures, even when they do not officially lodge an opinion. Nonetheless, the strategic situation facing MEPs differs drastically across these two types of votes implying that they should provide significant leverage over my question of interest.

[20]I general, I dropped final codecision votes from the dataset because the Commission's position towards such votes—while not officially specified—is a function of their positions on related amendments. Nonetheless, I included final votes when the Parliament voted on an unaltered Commission proposal as the Commission implicitly supported the text in question. I did not include votes on legislation considered under any other procedure, such as consultation, in the dataset.

cision amendments from multiple sources. In many cases, the Commission's opinions on amendments are listed at the end of the transcripts of EP plenary debates available from the online archive; where possible I transcribed Commission opinions from this source. Additionally, I consulted PreLex (European Commission 2009), the EU's legislative database, and extracted Commission opinions on amendments from the documents describing the Commission's first and second-reading positions on EP's proposals. In many cases, I was able to obtain opinions from both sources. While discrepancies were extremely rare, I used the debate transcripts when the two sources disagreed, because debates clearly reflect the Commission's opinion prior to the Parliament's vote. When the debate transcript did not clearly indicate the Commission's attitude towards all of the Parliament's amendments, I relied on the positions published in PreLex.[21] I dropped codecision votes for which I could not find opinion information from the dataset, leaving 572 codecision amendments and single votes available for analysis, along with 2879 initiative and resolutions in the "control group." Where possible, I also collected information on Council positions from PreLex. Specifically, when the Parliament amended the Commission's proposal on first readings, I recorded which of the Parliament's amendments the Council included in the Common position. That is, while there is no way to observe Council opinions on tabled amendments that the majority of the Parliament rejects, one may nonetheless record the Council's opinion on passing amendments. I recorded Council positions on 133 of the 547 codecision votes.[22]

I gathered information about individual MEPs—their EP group and national party affiliations and parliamentary committee memberships—from the EP's MEP database (European Parliament 2009*c*), and retrieved data on MEPs' national parties—notably their partici-

[21]The College of Commissioners is responsible for deciding Commission opinions, and vote internally when making their decisions. Individual Commissioners do not have the ability to alter the Commission's opinions on their own. Thus, while it is technically possible for the Commission to "change its mind" on an amendment after the Parliament's vote, such switches are rare in practice, because of the bureaucratic difficulties involved.

[22]Admittedly, this is a very restrictive window on Council opinions towards Parliamentary amendments. Clearly, selecting only passing amendments has the potential to bias conclusions drawn from these positions. Nonetheless, they can be helpful when one wishes to separate Commission influence over Parliamentary voting from Parliament's expectations about the Council's response to its proposals.

pation in national government—from the *European Journal of Political Research*'s yearly country reports. Furthermore, I collected bill information from the Parliament's Legislative Observatory (European Parliament 2009*b*), including the bill's procedure, the committee responsible for the bill, and the identity of the bill's rapporteur.

3.4 Influence or Agreement?

So, do MEPs alter their votes in a manner that varies systematically with Commission opinions, or is the correlation between the Commission's recommendations and voting outcomes simply an artifact of ideological congruence? To answer this question, I fit the baseline commission influence model to a dataset containing 540 of the 547 votes for which commission opinions are available and a random sample of 540 of the 2879 roll calls on EP resolutions and own initiative reports.[23]

The model indicates that a substantial number of MEPs alter their voting behavior systematically when a Commission opinion is available. Table 3.2 displays the number of MEPs that experience statistically significant changes in voting behavior, on average, when considering measures on which the Commission has rendered an opinion. As the table shows, around 39 per cent of MEPs tend to vote in a manner that is inconsistent with the way in which they vote on non-legislative issues when the Commission indicates support for a proposal. In other words, the 95 per cent highest posterior density (HPD) intervals around the estimated $\boldsymbol{\delta}^{\mathbf{a}}$ parameters exclude zero for 344 MEPs. Of these 344, 267—or about 78 per cent—are more likely to vote for proposals graced with Commission approval than their ideal point estimates would suggest. On the other hand, 77 MEPs are more likely to vote

[23]I restricted analysis to votes that were at least somewhat contested, in the sense that at least 25 MEPs voted for each alternative, losing 7 codecision votes in the process. Furthermore, I dropped MEPs who participated in less than 100 total votes from the analysis, leaving 893 of the 905 MEPs for whom at least one vote was recorded. Finally, I replicated the analysis with the entire set of 2879 non-codecision votes, yielding similar results. I use the results from the sample here for comparability with later results. Fitting these models takes a lot of time, especially as the number of votes grows, and it was impractical to use the entire dataset to fit all the variety of model specifications that I explored when conducting this research.

Table 3.2: How MEPs react to Commission opinions.

		Commission Opinion	
		Accept	Reject
MEP Bias	+	267	144
	−	77	237
		Total MEPs: 893	

against a measure that has been approved by the Commission. Similarly, the estimates for $\boldsymbol{\delta}^{\mathbf{r}}$ show that around 43 per cent of MEPs alter their behavior when deciding whether or not to support a measure that the Commission has rejected. Again, more MEPs are biased in the direction of the Commission recommendation than against it, yet a substantial number of MEPs—about 38 per cent of those displaying statistically significant responses to negative Commission opinions—actually tend to support proposals that they would not otherwise favor in the face of Commission opposition. Interestingly, of the 344 MEPs that change their behavior in a statistically significant manner when voting on measures towards which the Commission is positive, and the 381 MEPs that likewise alter their behavior on proposals that the Commission dislikes, only 144 overlap. Therefore, there are 437 MEPs that are strongly moved only by one type of opinion, and a total of 581 MEPs—or about 65 per cent—change their behavior in some way when voting in the shadow of a Commission verdict. Moreover, this lack of overlap in behavior across opinion types implies that different processes may drive MEP responses across positive and negative recommendations.

These findings provide telling evidence for an extra-ideological relationship between Commission recommendations and voting behavior in the EP. That is, one cannot simply explain the correlations between MEP voting choices and Commission recommendations in terms of ideological congruence. The relationship between MEP behavior and Commission opinions is not epiphenomenal; if MEPs voted the same way on codecision amendments that they do on non-legislative issues then the estimated $\boldsymbol{\delta}$ parameters would not significantly differ from zero. Furthermore, the process that drives Commission opinion formation is systematically

related to strategic voting by MEPs on codecision amendments. Nonetheless, these findings do not tell us much about what drives the relationship between Commission opinions and MEP voting, nor do they establish that the Commission is influential, in and of itself. For example, the Commission's opinions might just reflect what the Parliament already knows about the Council, rather than providing new information, or otherwise directly influencing MEP votes. In what follows I delve deeper into these questions, examining who follows—and who rejects—Commission recommendations, and under what circumstances they do so.

3.5 Who Listens to the Commission?

Table 3.2 shows that a substantial number of MEPs vote differently when the Commission voices an opinion than when it does not. Raw counts tell us that Commission opinions matter, but they do not tell us much else. In section 3.1 I argued that MEPs from larger groups, and especially those in government would be more likely to modulate their voting behavior in the face of bicameral pressures than small, non-governing parties, especially those with fringe or eurosceptic policy stances. Indeed, I noted that eurosceptic partisans might even reject the mainstream positions advocated by the Commission just to say they did. Figures 3.2 and 3.3 display 95% HPD intervals for estimated δ^a and δ^r parameters for MEPs hailing from Germany (DE), France (FR), the United Kingdom (GB), and Spain (ES). Colors indicate national party membership and the scales are normalized across figures so that positive values indicate a bias towards voting in favor in both figures.[24]

Figures 3.2 and 3.3 provide mixed evidence with respect to my hypotheses; MEPs behave largely as expected on rejections but their responses to positive opinions are somewhat surprising. In Germany, the CDU/CSU and SPD sat in a grand coalition—with the CDU holding the Chancellor's office—for the bulk of the EP's 6th term, yet MEP reactions to commission opinions differed across governing parties. On acceptances the CDU and CSU

[24]That is, positive values show that a MEP tends to vote with the commission in figure 3.2, and against it in figure 3.3.

Figure 3.2: MEP δ^a values for four countries, 95% HPD intervals.

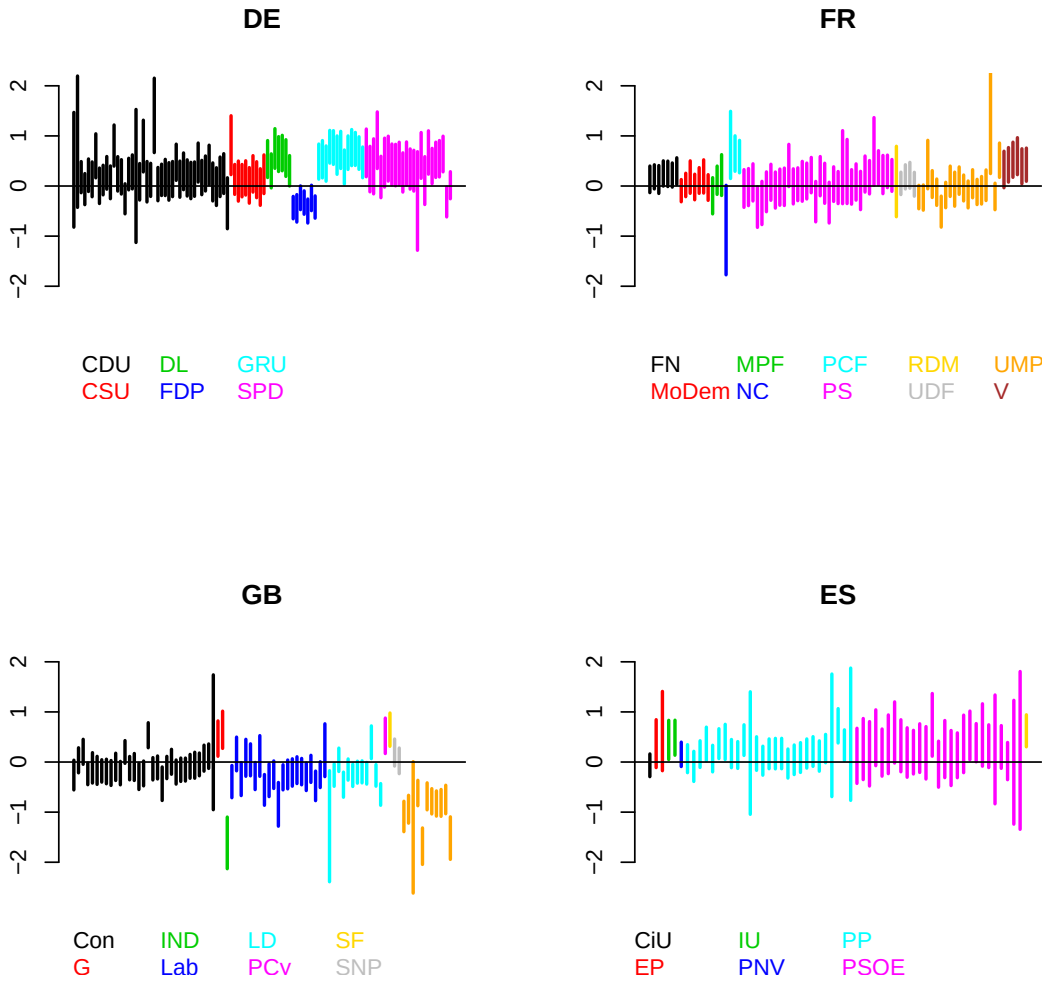

Figure 3.3: MEP δ^r values for four countries, 95% HPD intervals.

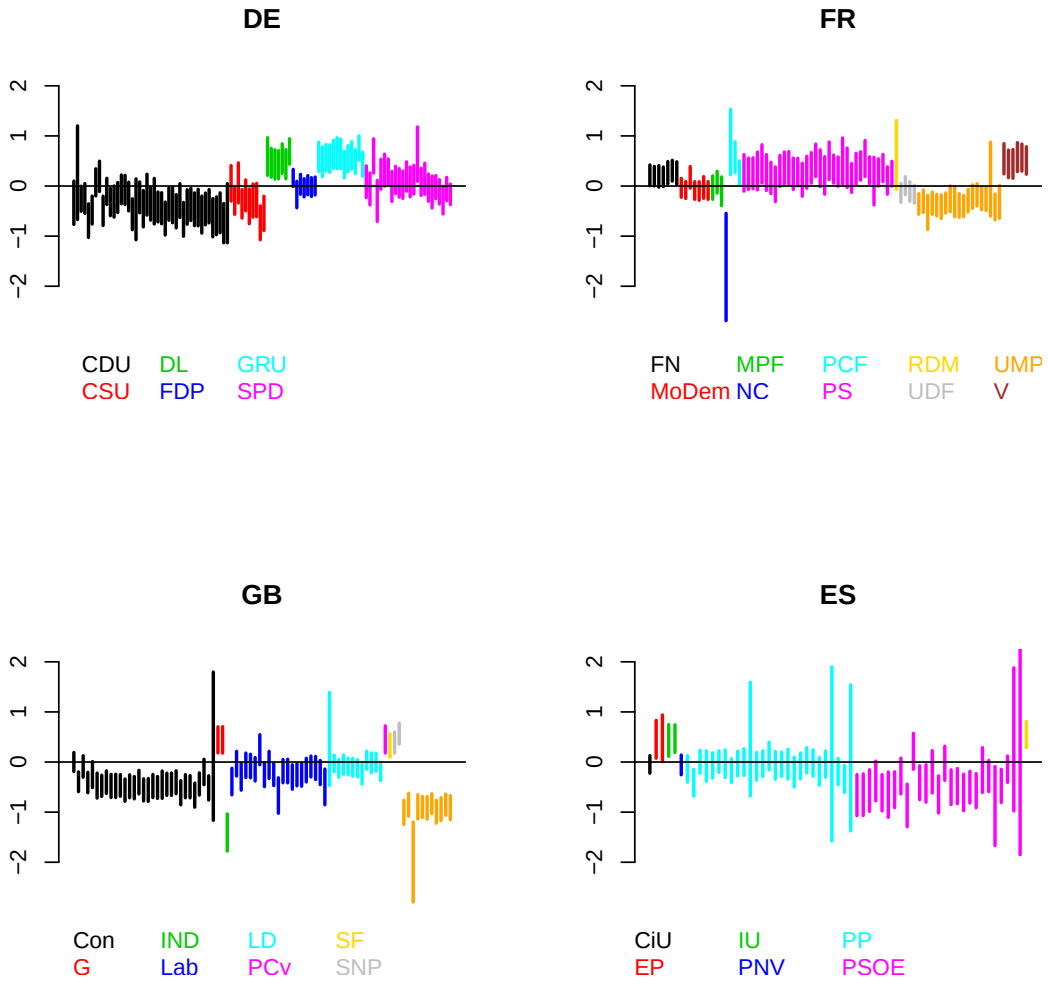

sat largely unmoved (their HPD intervals span zero) while the SPD MEPs reacted with favor to Commission endorsements. Surprisingly, members of Germany's left-leaning parties, the Greens and Die Linke, also tended to vote for amendments supported by the Commission, while the liberal FDP generally opposed accepted amendments at a higher rate than their non-codecision voting tendencies would predict. On the other hand, when the Commission rejected amendments, the CDU and CSU MEPs were swayed by their arguments, the SPD and FDP were largely unmoved, and members of the smaller fringe parties actually supported such amendments at greater than expected rates. During this period the Commission, Council, and Parliament were generally dominated by center-right majorities. Therefore, German reactions to negative recommendations are largely in line with my expectations. The in-government CDU and CSU appear willing to compromise when instructed by the Commission. The SPD, in government, but not at the head of the government nor particularly ideologically disposed towards compromise with the Council, is ambivalent, and smaller fringe parties with little stake in compromise tend to vote against the Commission's recommendation, engaging in "anti-government" voting even when their ideological convictions might guide them in the other direction. On the other hand, when the Commission accepted an amendment the pattern is less clear.

In France, only the Greens (V) show statistically significant tendencies to react to Commission endorsements; again it seems that fringe parties support amendments, to a greater extent than they otherwise would, when the Commission does. As was the case in Germany, French reactions to Commission rejections are in line with my expectations. The UMP controlled the government for the entire observation period and it is the only party with members that go out of their way to follow the Commission's recommendation to reject. Fringe parties—the National Front (FN), the Communists (PCF), and the Greens (V)—all appear to be fond of protest votes. Interestingly, the out-of-government, but sizeable Socialist party (PS), also tends to go against the grain on negative opinions, although many Socialist MEPs' opinion effects are statistically indistinguishable from zero. As we might

expect, the large, governing and center-right UMP moderates its position in a manner that is consistent with the avoidance of protracted inter-institutional bargaining. On the other hand, the Socialists, while large, are firmly in opposition and appear moderately inclined towards protest voting.

In the United Kingdom, parties are again largely impervious to positive opinions, although the MEPs from especially small parties do exhibit positive biases. Strangely, UK Independence Party MEPs (UKIP) seem to strongly oppose Commission-supported amendments, but as figure 3.3 shows, they also dislike amendments that the Commission rejects.[25] British MEPs' reactions to negative opinions are reasonably well-aligned with expectations, although government-opposition dynamics do not appear to play a key role in this case. Indeed, while Labor held Downing Street for the entire observation period, the Conservatives were more likely to moderate their voting behavior in the Commission's preferred direction than Labour MEPs, who, while slightly tending to follow the Commission line, largely walked to the beat of their own drum. Indeed, the fact that the out-of-government Conservatives clearly join their ideologically right-leaning peers in the CDU and UMP in following rejections, while governing, but center-left Labour, does not, may indicate that the ability to stomach compromise trumps the drive to get legislation done in MEPs' voting calculus.

On the other hand, Spanish MEPs behave in a manner that supports the hypothesis that governing parties will seek compromise because they are responsible for results. Indeed, while few Spanish MEPs change their voting behavior in the light of positive Commission opinions, it is the governing Socialist Workers' Party (PSOE) that follows Commission recommendations to reject, while a few fringe MEPs from the United Left (UL) and Europa de los Pueblos (EP)—a coalition of left-wing regional parties—engage in protest votes. Perhaps the differences between Labour and PSOE reactions to Commission opinions are explained by variation in attitudes towards Europe in Britain and Spain, with eurosceptic Brits less

[25]As far as I can tell, the UKIP just likes to vote no on codecision roll calls, regardless of the content of the proposal under consideration.

Table 3.3: Who follows the Commission?

	δ^a	δ^r
Intercept	0.186 (0.044)*	-0.016 (0.034)
$\sqrt{\text{Votes in Govt.}}$	0.001 (0.002)	-0.004 (0.002)*
LR Ideal	-0.037 (0.028)	-0.161 (0.036)*
Integration Ideal	0.007 (0.039)	-0.100 (0.026)*
R^2	0.011	0.298

Parenthetical values are posterior standard deviations and the stars indicate 95% HPD intervals exclude zero. While the coefficients incorporate estimation uncertainty using the method of composition, the R^2 values are based on simple regressions using posterior means.

concerned about their government's role in the speedy production of European legislation than their Spanish counterparts.[26]

These four countries paint a picture that is, largely, representative of behavior of MEPs from across the Union. Table 3.3 summarizes these patterns for all MEPs and provides a more formal test of a number of the arguments that I made in section 3.1. Specifically, table 3.3 presents the results of two linear models in which I regress MEPs' estimated extra-spatial responses to Commission opinions (their δs) on an indicator of governing status—the square root[27] of the number of codecision votes during which the given MEP's party held a portfolio in national government—and MEPs' estimated ideal points.[28] I argued that belonging to a party in government would compel MEPs to take intercameral bargaining costs into account when making codecision voting decisions. Conversely, opposition parties hold little responsibility for passing legislation quickly and thus are more free to vote their preferences. Therefore, MEPs from governing parties should be more willing to vote against their ideological tendencies on codecision measures than opposition legislators. Nonetheless, ideology

[26]Unfortunately, my data do not speak to this question and many other explanations exist. Nonetheless, this difference is intriguing and perhaps worth further examination.

[27]I use the square root transformation because these are counts. Results are robust to transformation.

[28]Both the response variable and some of the explanatory variables in these regressions are estimated parameters themselves. I use the method of composition (see e.g. Tanner 1993, pp. 30) to propagate measurement error in these parameter estimates to the regression estimates.

should temper this dynamic because, while MEPs may be willing to make compromises to make sure legislation passes smoothly, there will be a limit to how far legislators will be willing to go. In the context of the 6th EP, given the Council's largely center-right composition, this means that left-wing MEPs should be less willing to modulate their voting than legislators on the right. And, indeed, the results in table 3.3 are consistent with these arguments, at least with respect to negative Commission opinions. As the second column of the table shows, government MEPs—and MEPs on the right—vote in concert with the Commission, even at the cost of their own preferences, to a degree that is statistically significant.

Interestingly, integration ideology also predicts MEPs' tendencies to vote with the Commission, and MEPs that are more supportive of European integration reject amendments that the Commission does. This finding is somewhat surprising because MEPs tend to support integration to a greater degree than ministers in the Council, even when they hail from the same parties. For instance, as figure 3.1 shows, MEPs from the large Christian and Social Democratic groups stand strongly on the side of integration, something that may not be true of their counterparts at home. We would not expect compromises with the Council to favor integration-oriented MEPs over eurosceptics. Therefore, this result is not particularly consistent with the compromise hypothesis. On the other hand, the finding is weakly consistent with the argument that the Commission communicates best with the Parliament in situations where their preferences coincide. Being a European institution, the Commission is often thought to back integration to a greater degree than the Council and pro-integration MEPs may, therefore, be more willing to follow Commission arguments than other legislators.

Table 3.3 also conclusively demonstrates that the relationship between government status, ideology, and MEPs' willingness to pay intercameral bargaining costs cannot explain legislators' behavior when the Commission supports amendments. Indeed, none of the factors considered in the regressions correlate significantly with MEPs' δ^a estimates. Clearly, the mechanisms underlying MEP responses to positive and negative acceptances differ, as

figures 3.2 and 3.3 also imply. Indeed, the R^2 values in table 3.3 indicate that the three explanatory variables explain almost a third of the variance in MEP behavior on amendments the Commission dislikes but virtually none of the variance in reactions to Commission approval. One possible explanation is that Commission acceptances are better conceptualized as non-rejections, rather than as actual endorsements. The Commission's institutional power in codecision stems from the fact that its rejections force the Council to act unanimously. Passing an amendment that the Commission dislikes, the Parliament runs a real risk of bargaining delay because the Council may find it difficult to construct unanimous support for a version of the bill that incorporates the tainted measure. Therefore, the results strongly support the idea that the Commission wields influence in codecision—perhaps substantial influence—simply because of its control over voting procedures in the Council. Importantly, this implies that the content of Commission opinions matters and that the relationship between recommendations and MEPs' votes is not driven entirely by what legislators already know about the Council's likelihood of accepting particular amendments.

3.6 Influence or Reflection?

This last result provides some evidence that the Commission is influential, in and of itself, but we can do better. Specifically, the conditional Commission influence model provides a tool with which to control for the possibility that Commission opinions influence the Parliament's behavior only because they mirror what the Parliament already knows about the Council. To this end, I fit the conditional model to a reduced dataset, containing the same non-legislative votes as the data that I used above, but holding only those codecision votes for which I was able to collect Council opinions. While this sample suffers from obvious selection issues, it provides a useful testbed for examining how dependent Commission influence is on Council preferences. When fitting the conditional model to these data I include a single covariate in $\mathbf{Z}$, indicating the Council's eventual decision—positive or negative—on the amendment in

Table 3.4: Variance in Commission-Council verdicts.

		Commission Opinion	
		Accept	Reject
Council Decision	Adopt	42	16
	Reject	23	52

question. Table 3.4 summarizes the variation in these data; note that we observe cases in all four cells, although Commission-Council agreement is clearly the norm.

The conditional model provides evidence that the Commission is, indeed, influential. While the coefficients that the model estimates ($\bar{\gamma^a} = 1.16$ and $\bar{\gamma^r} = -0.99$ and both are statistically significant) indicate that eventual Council decisions explain quite a bit of variation in the tendency of MEPs to react to Commission opinions, a substantial proportion of MEP behavior is left unaccounted for. Specifically, the model estimates statistically significant δ^a and/or δ^r parameters for 45 per cent of MEPs. Therefore, MEPs change their voting behavior in the face of Commission arguments to an extent that Council preferences cannot explain alone. It is difficult to tease out exactly what is driving this result. Nonetheless, the finding implies—if somewhat tangentially—that the Commission provides MEPs with information that they do not already have, as I have hypothesized. While the Commission's institutional powers may account for a significant degree of its influence, they operate through the behavior of the Council. If MEPs know ahead of time how the Council will behave, they can use this information to decide when the Commission's negative opinions will bite. Therefore, the fact that Council opinions do not explain large components of the relationship between Commission opinion and MEP vote behavior is consistent with information transmission by the Commission.

The argument above shows that MEP behavior is consistent with information transmission by the Council, but it is far from conclusive. One useful aspect of the models that I introduce here is that they allow one to see how allowing for external influence alters estimates in comparison to the canonical ideal point model. Interestingly, ideal point estimates

Figure 3.4: Mean cutpoint estimates across models.

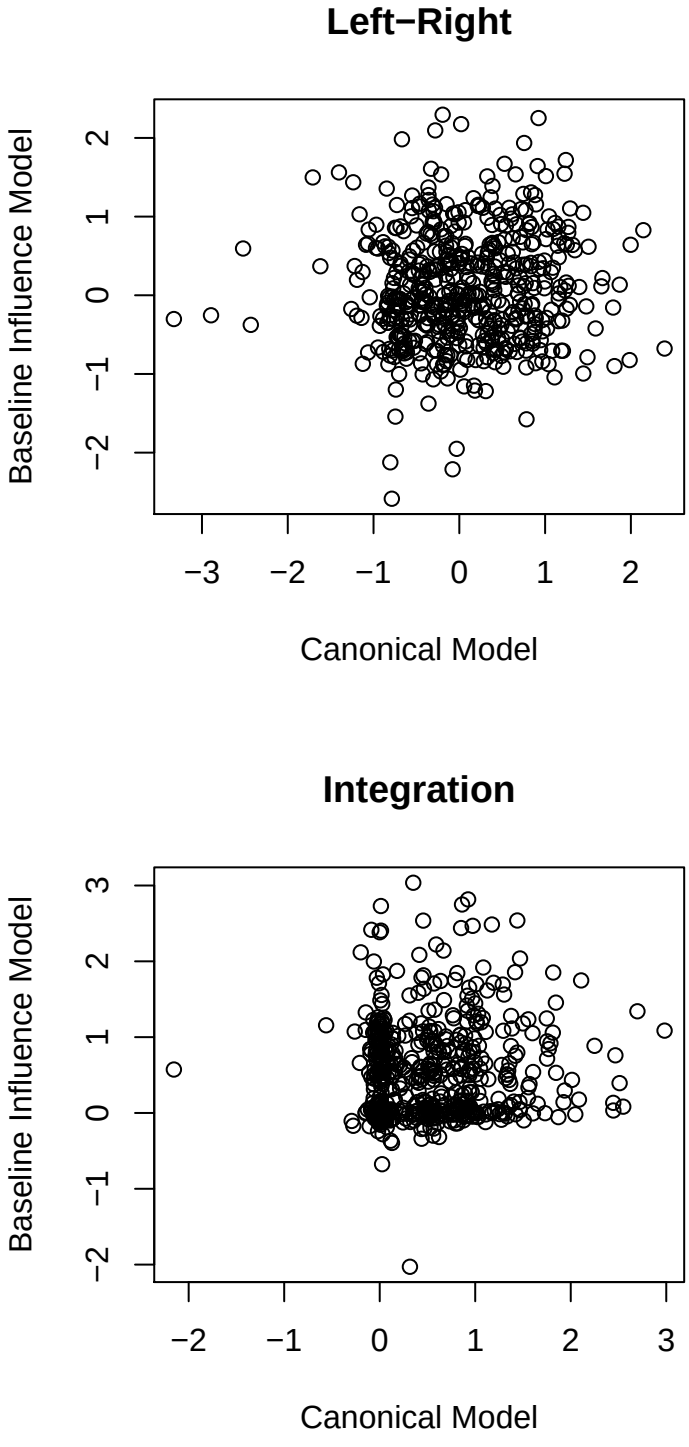

do not change substantially when one fits the baseline influence model to the data. Indeed, the map of ideal points produced by the expanded model looks just like figure 3.1 and ideal points correlate across models almost perfectly. This is not surprising because the influence model relies on non-legislative votes to identify ideal points, attributing differences in voting behavior across procedures, not to ideology, but strategic response to outside influence. On the other hand, the models generate substantially differing pictures of the characteristics of the codecision votes.[29] Figure 3.4 compares mean point estimates of the cutpoints—the $\boldsymbol{\kappa}$ matrices—on codecision votes produced by both models. Strikingly, the correlation in point estimates of cutpoints across models is around 0.06. On the left-right dimension there is little pattern in the cross-model relationship. Cutpoints estimates are quite uncertain, and the lack of agreement across specifications may stem simply from the model's inability to nail down cutpoints on this dimension. On the integration dimension, on the other hand, there is a clear pattern to the disagreement. Specifically, the baseline influence model moves cutpoints that the canonical model sees as extreme towards the center of the issue space. What look like rah-rah votes to the canonical model often appear to the influence model as votes on controversial proposals where influence swayed MEPs to vote against their ideological priorities.

Therefore, when influence does occur, it appears to happen with respect to the integration dimension. This is consistent with mediation by the Commission, as I argued in section 3.1, because mediation is most likely to occur on issues where the Commission and the average MEP share priorities. Because the Commission's opinions are mere cheap talk, the MEPs should trust the Commission's recommendations when they expect the Commission to have the pivotal MEP's interests at heart. While suggestive, this evidence for information transmission by the Commission is, also, less than conclusive. The Parliament median and the Council are likely to disagree with one another on integration votes; perhaps this finding

[29]Bill parameters for non-legislative votes are almost completely identical across the two model fits, as one would expect.

is simply a result of MEPs adjusting their voting behavior only when holding to their own preferences becomes potentially costly. In either case, the result is interesting. And, together with the finding that Commission opinions influence MEPs even taking the Council's preferences into account, the evidence in figure 3.4 lends plausibility to the argument that the Commission acts as a strategic mediator under codecision. Nonetheless, nailing down the role of this information transmission mechanism will require further research.

3.7 Conclusion

This paper provides a window into the strategic behavior of MEPs, helping us to better understand how their institutional context—specifically the pressures of bicameral lawmaking—influences their voting behavior. I demonstrate that MEPs from governing parties, and who are ideologically open to compromise with the Council, tailor their votes in a manner that is consistent with a concern for intercameral bargaining outcomes. Specifically, these legislators vote against proposals that they might otherwise support when such proposals are likely to cause bargaining delay and impede legislative efficiency. The European Commission plays a key role in this process, relaying information about Council bargaining positions to MEPs, and using their wealth of expertise to guide MEPs in their voting choices. While it is difficult to prove that the Commission selectively reveals Council bargaining strength to the Parliament, it does appear to influence Parliamentary voting, and its opinions do no not simply mirror Council preferences. Furthermore, its one institutional power during the codecision process—beyond its prerogative to introduce legislation—appears to pack a punch. MEPs who should be concerned with getting laws passed quickly react strongly to negative Commission recommendations and are loath to support amendments that could force unanimity within the Council.

The findings demonstrate the importance of considering the wider inter-institutional bargaining environment when studying the voting behavior of legislators. MEPs approach

legislation that requires compromise with the Council differently from intra-parliamentary resolutions and initiatives. Thus, the mechanisms that drive voting behavior differ depending on the institutional context. Standard techniques for modeling legislative voting, such as common ideal point estimation models, treat all parliaments equally. This will often provide us with an inaccurate picture of what drives legislative behavior; we can improve our comparative understanding of lawmaking by taking key underlying determinants of vote choice—such as the need to strike intercameral bargains—into account. Furthermore, this paper highlights the potential that informational advantages give the Commission to influence the outcome of codecision legislation, notwithstanding the Commission's complete lack of veto-power over other actors' amendments to the Commission's proposals. Thus, as others have argued (Rasmussen 2003), the Commission maintains an important role in a procedure in which the formal rules render it "irrelevant" (Crombez 2001, pp. 101). More generally, the Commission's ability to leverage its access to information highlights an important advantage—information garnered through the control of ministries and through multiple points of contact with legislating institutions—available to most executive branches of government. Indeed, because the Commission looks a lot like a parliamentary government, but has a composition that is not a function of the seats in the Parliament, and has no recourse to such institutional devices as votes of confidence,[30] the EU provides an excellent laboratory within which to examine the role that informational advantages play in allowing governments to influence legislators' voting decisions.

Finally, the techniques that I introduce in this paper may travel to a variety of other contexts. For example, they may help us to explain when Presidential veto threats influence Congressional voting in the US, or provide a new way to model the role that pressure groups play in swaying lawmakers' votes. Nonetheless, the models can be tricky to deploy effectively. First, the analyst must have access to a set of roll call votes that are not contaminated by the

[30]Although the Parliament invests the Commission and can remove it from office with a supermajority vote.

influence process that she wishes to model, something that will not occur in all circumstances in which outside influences weigh on legislators. Additionally, researchers using these models must be able to effectively measure the preferences of the influential actor and directly link these preferences to individual votes. This can be quite a challenge, and the EU represents a rare instance where institutional rules require an influential party to officially register its preferences on legislative votes. Nonetheless, when these barriers are surmountable, the influence models that I present here represent powerful tools for separating strategy from preference when analyzing voting behavior.

Chapter 4

Predicting and Explaining Roll Call Sponsorship in the European Parliament with Debate Speech

Recorded votes in parliaments, or roll call votes (RCVs), occupy a key place in the analysis of legislatures. In democracies, voting is at the heart of lawmaking and patterns of votes can potentially tell us a great deal about the size and shape of the policy space (see e.g. Poole & Rosenthal 1985), the roles played by parliamentary parties (see e.g. Rohde 1991, Cox & McCubbins 1993, Krehbiel 1998, Desposato 2006, Hix, Noury & Roland 2007), how accountable politicians are to voters and interest groups (see e.g. Carey 2008), and a multitude of other topics. Furthermore, the use of recorded votes spans the discipline; while roll calls have long featured in work on the U.S. Congress, researchers have since applied roll call analysis to legislatures ranging from the Chilean Senate (Londregan 2000) to the European Parliament (Hix, Noury & Roland 2006) and the United Nations' General Assembly (Voeten 2000), and roll call analyses even play key roles in large-scale cross-national studies (Carey 2008).

Nonetheless, scholars have long cautioned against an over-reliance on recorded votes (Greenstein & Jackson 1963, Arnold 1990) and recent work has highlighted the potential for selection bias inherent in RCV analysis (Carrubba et al. 2006, Hug 2010). Indeed, only a subset of votes in many parliaments are recorded (Saalfeld 1995, Hug 2010), and, even in the U.S., where roll call votes are often treated as ubiquitous, less than 15% of public statutes are subjected to roll call of any kind, in either the House or Senate (Clinton & Lapinski 2008). The selection mechanism driving vote recording is unlikely to be random, and factors such as the issue area examined by the legislation, the salience of the bill, and parliamentary procedure correlate with RCV requests in both the European Parliament and U.S Congress (Carrubba et al. 2006, Clinton & Lapinski 2008). Furthermore, the decision

to call roll is often endogenous—in many parliaments, individual legislators, groups of members meeting some minimum threshold, or political parties have the right to request recorded votes (Saalfeld 1995, Hug 2010)—and the selection processes underlying roll call votes are, therefore, likely to result from strategic political considerations. For example, political parties may call roll to discipline their members, embarrass opposing political parties, or signal their policy positions to a variety of audiences (Saalfeld 1995, Thiem 2006, Carrubba, Gabel & Hug 2008). Thus, there are likely to be interactions between the way legislators vote on RCVs and the reasons that the votes that we observe were recorded in the first place. This means that we may draw biased inferences from RCVs not only because they represent a non-random sample of all votes, but because legislators may change the way they vote depending on whether the tally is public or not. In fact, Hug (2010) takes advantage of the Swiss parliament's decision to keep track of un-publicized, but technically recorded, votes to demonstrate that Swiss legislators' voting behavior does indeed change depending on whether or not they expect their votes to be observed.

Yet, in most parliaments, unrecorded votes are genuinely unrecorded and we lack tools to evaluate the extent to which legislators' voting behavior differs across recorded and hidden votes. One route to solving this problem is theoretical. Because we typically cannot observe how legislators vote on secret votes, nor how their voting behavior would have changed as a function of vote publicity, we should carefully consider the likely mechanisms underlying the interplay between legislative voting behavior and the publication of individuals' vote choices. Indeed, modeling the strategic situation that faces parties and legislators with respect to RCVs provides insight into when we should trust the inferences we draw from roll call data, and when we should suspect pernicious selection effects. Unfortunately, theoretical accounts of roll calling are thin on the ground, and this research area is still in infancy (Carrubba & Gabel 1999, Carrubba, Gabel & Hug 2008, Thiem 2006). Furthermore, it is difficult to test theories of strategic roll calling for the very same reasons that we need them in the first place: the key primitive in theories of roll call selection—expected voting

behavior—is hidden from analysts. While clever research design can overcome this problem to some extent, we need new tools to explore the similarities and differences between RCVs and unrecorded votes. Only with new techniques will we be able to test theories of RCV selection, and use these theories to improve the use of roll call data in empirical analyses. This challenge poses a difficult problem because the things we typically observe about both RCVs and unrecorded votes—the identity of the legislator or party that proposed the bill or amendment, the procedure under which the vote was taken, the issue area with which the legislation is concerned, and so forth—provide relatively little information about the factors, such as the ideological content of the alternatives in question, that should drive roll call selection. For example, Hug (2010) evaluates the ability of a simple Heckman selection model to correct the bias in cohesion scores generated from Swiss roll-calls and finds the results somewhat wanting because of the lack of information about the selection process carried by observable variables.[1]

To help remedy these problems, I examine the potential that the automated analysis of legislative text—such as speeches in floor debates on legislation, legislative reports attached to particular bills, and the text of legislation itself—has to shed light on the circumstances under which politicians request RCVs. Specifically, I use the text of floor debates to predict RCV requests on final votes on proposals considered by the European Parliament (EP), a parliament where recorded votes are not the norm, but where representatives of the party groups or any group of 40 or more members can request a roll call on any vote. Effective predictive models of roll calling are a first step on the path towards statistically correcting for RCV bias, potentially yielding "propensity-to-be-roll-called" estimates for votes that could underpin the application of Heckman-style selection models or matching techniques to RCV bias, at least in certain domains. I show that, when naively applied, state-of-the-art machine learning tools for text classification (Joachims 1998) do a poor job predicting RCV requests

[1]The peculiarities of the Swiss case—notably, the fact that the Parliament tallies all votes electronically and saves roll call records for "unrecorded" votes—allow him to observe true cohesion levels.

from floor debates. Yet, by using these same tools to model speaker sentiment towards proposals, and by using information about the distribution of sentiment in a given debate to model the probability of a recorded vote request, I demonstrate that one can substantially improve RCV prediction accuracy.

Furthermore, I use the model of speaker sentiment that underpins the RCV prediction technique to examine the circumstances under which political parties request roll call votes. The sentiment model predicts the vote choices of speakers in plenary and, therefore, provides a window into how legislators would cast their votes on a given measure if the voting record were made public. Thus, the sentiment measure that I derive from the debates provides a clear proxy for one of the key unobservable quantities in roll-call analysis. Indeed, while we can observe actual voting behavior only on votes that are subject to roll call, the sentiment model generates a measure of that behavior that is observable across all final votes, public and secret. This variable captures patterns in expected public voting behavior across legislators who speak in plenary and provides researchers with a fine-grained tool with which to examine RCV requests, greatly improving our inferential potential when compared to blunt measures of likely preference patterns among legislators, such as bill issue area, procedural features, and authorship.[2] I use this new measure to explore how political parties react to expected patterns in public voting behavior when deciding whether or not to sponsor roll calls on particular votes. The analysis demonstrates that European party groups call roll primarily on salient and divisive issues, and especially to demonstrate their opposition to policies that they oppose, and for which they are not responsible. As one Danish staff member highlighted in an interview, party groups often call roll to embarrass their opponents when they take the "wrong" side of an issue.[3] Groups also use roll call votes to show support for their own members' reports and rarely seek to signal their preferences through RCVs when they

[2]Co-sponsorship, which may be useful in this regard in some legislatures, provides little leverage in the EP. Non-legislative resolutions are technically introduced by entire party groups and legislative proposals originate in the European Commission.

[3]I conducted a series of open-ended interviews with MEPs, EP functionaries, and national delegation staff members in the fall and winter of 2007.

strongly support another groups' measure. Therefore, the analysis provides strong support for a story of roll call requests within the EP based on taking credit and assigning blame. The results are also weakly consistent with theories of calling roll to maintain party discipline, although the present study has less to say on this account. Finally, the findings imply that researchers should be wary when using roll call votes to draw conclusions about within and across-group cohesion in the EP, because cohesion co-varies with RCV selection.

4.1 Why Parties Call Roll

Theories of roll call vote requests fall broadly into two categories: position-taking and discipline-enhancing. First of all, political parties may request open votes in order to publicize their positions—or the positions of other players—on a given issue (Saalfeld 1995). Through RCVs, party members may signal their preferences to third parties such as constituents, interest groups, other governing institutions, and even coalition members or other allies within the parliament itself. In the context of the EP, where elections are mainly won and lost based on the domestic fortunes of national parties, interest groups and other EU institutions—such as the European Commission and Council—may represent the dominant audience for such signals. As Thiem (2006) argues, party group leaders can use roll calls not only to demonstrate the preferences of their members but to demonstrate their ability to act cohesively as a group. Such cohesion can be an attractive signal to potential coalition partners and like-minded interest groups and a warning to opposition parties and members of other institutions that engage in intercameral bargaining with the legislature, such as the Council in the context of the EP. If EP party group leaders use RCV requests for such purposes, therefore, we would expect them to focus their roll call sponsorships on circumstances where they expect their rank-and-file members to maintain the party line publicly.

On the other hand, parties may use RCVs to engage in position-taking more to claim credit for particular policies—and to distance themselves and lay blame for others—than to

demonstrate the value of the party group as a cohesive unit. Within the EP, party groups are intrinsically tied to particular proposals by the group-membership of the rapporteur, the MEP responsible for guiding a bill through the legislative process. Thus, observers will generally be able to attribute credit—or blame—to the party group of the rapporteur. Yet, it may be unclear where groups stand on proposals that they do not report; in these cases the leadership may request a roll call to publicly register a group's position. Furthermore, such circumstances can provide a window into the motivation behind groups' roll call requests. If groups sponsor RCVs purely for informational purposes—simply signaling their position to third parties—then they should request RCVs on other groups' reports with similar frequency regardless of their disposition towards the proposal. On the other hand, if RCV requests serve mostly as a way for group leaders to apportion credit and blame, they should target their public vote requests on other groups' reports to those measures that they disagree with.

While position-taking accounts explain some of the reasons why parties request public votes, they may also use RCVs to monitor their own members in order to maintain discipline within the party and obtain the policy outcomes that they most prefer (Fennel 1974, Jenkins & Stewart 2003). This account of RCV sponsorship is somewhat controversial with respect to the EP. Indeed, Kreppel (2002) and Thiem (2006) argue that party groups lack the ability to effectively sanction their members for defection in open votes, primarily because party group leaderships are composed of the leaders of national party delegations. These delegation leaders, those authors argue, have little interest in sanctioning their own members. Nonetheless, many authors attribute significant sanctioning ability to the EP party groups, maintaining that angering one's group can have negative consequences for an MEP's European career, even when the actions that draw the group's ire are tacitly condoned by one's national delegation (Hix 2002, Faas 2003, Hix 2004, Hix, Noury & Roland 2007, Meserve, Pemstein & Bernhard 2009). Clearly, EP groups may not have the sanctioning powers of parties in many national legislatures, but it seems reasonable that they might, nonetheless, swing around considerable weight. Furthermore, Carrubba, Gabel & Hug (2008) develop a

Table 4.1: Leveraging speaker sentiment.

	RCV	**Secret**
Vote|RCV	Observed	Unobserved*
Vote|Secret	Unobserved	Unobserved

model of discipline-based RCV sponsorship that they argue is directly applicable to the EP (Carrubba, Gabel & Hug 2009).

Position-taking accounts of RCV sponsorship provide predictions in terms of what group leaderships expect voting behavior within their own groups to look like should the vote take place in public and do not engage in a discussion of how individual MEPs might change their votes as a function of their publication. While they need not hold any particular assumption on the matter, theories of RCV requests that argue that sponsorships are solely a function of informational factors implicitly assume that voting behavior remains constant across public and secret votes. On the other hand, discipline-based stories argue that legislators may change their votes if they must make them in the cold hard light. Specifically, when MEPs have personal policy preferences that differ sufficiently from those of their party leadership, they may defect from the group line on private votes; yet when party leaders have sufficient sanctioning ability to offset the gains of defection, these same legislators will change their votes if forced to reveal them.

Table 4.1 highlights the distinctions between these two broad categories of RCV selection theories, depicting legislative votes in terms of what researchers observe on a given type of vote and the counterfactuals that are relevant to both types of theory. The columns in table 4.1 represent vote types while the rows describe (counter)factuals. Position-taking theories of RCV selection focus on the top two cells in the table. The patterns of selection that they predict deal with the differences in expected voting patterns across how legislators would vote on a given vote—RCV or secret ballot—if it were held publicly. To evaluate such theories, and to diagnose the extent of the selection problem that plagues a given set

of roll call votes, we need only the observable roll call votes and a good measure of how legislators would have behaved on secret votes, had they cast those votes in public. On the other hand, effective tests of discipline-based theories—and effective evaluations of selection effects driven by discipline-based RCV sponsorship—require measures of all four cells in table 4.1. Such theories ask questions like "How would a given MEP have voted on a given RCV if it were held in secret?" and "If that secret vote had been an RCV, how would legislators have altered their behavior?"

The text-based analysis that I present in this paper provides a measure of the unobserved votes described by the top right cell in table 4.1. Specifically, I fit a model that uses legislative speech to predict voting behavior in recorded RCVs. This model, in turn, provides a predictive measure of how MEPs who spoke in debate would have voted on a secret vote, had it been taken publicly. Thus, I restrict my examination of theoretical accounts of RCV selection primarily to position-taking. Nonetheless, it is possible to make some conjectures that relate to discipline-based theory. Most notably, because third parties are largely able to attribute credit or blame to party groups on bills that their own members report, position-taking cannot explain why party groups would request RCVs on final votes in such circumstances. On the other hand, group leaders might request such votes to maximize the level of support for their parties' measures if they believe that public voting will influence their members to toe the party line.

In what follows I first describe the dataset and explain how to use debate speech to predict voting behavior in the EP. Next I demonstrate that patterns in expected voting behavior explain substantial variation in RCV sponsorship, in contrast to other predictors, and demonstrate a method for predicting RCV requests from floor speeches. Finally, I use the new measure of expected voting behavior to explain the circumstances under which party group leaders request public votes.

4.2 Data

I collected floor debates, voting data, and roll call vote tallies from the European Parliament's online archive (European Parliament 2009*a*). The dataset spans the majority of the 6th term of the European parliament, from July 2004 through May 2008. During this time-span the Parliament considered over 1,800 proposals and held debates on almost 800. Though the MEPs voted 18,493 times during this period, the vast majority of these votes dealt with amendments and only 1,706 votes were "final" votes on whole proposals.[4] The dataset contains all proposals that were actively debated by the Parliament, and on which the Parliament held a final vote, during the observation period, yielding 769 debated final votes.[5]

Each debate consists of a series of speech segments generated by members of the EP (MEPs) and representatives of other European institutions, such as the Commission and the Council. Each speaker in the debate is identified by name and role and I coded whether speakers were MEPs, representatives of the Commission, Council, or other bureaucracy, or the (acting) President of the Parliament.[6] Furthermore, some MEPs formally spoke on behalf of their parliamentary party groups, while others spoke for themselves, and I recorded this information for each MEP segment. I cross-referenced MEP speech segments with the EP's MEP database (European Parliament 2009*c*) which provides information on each speaker, including age, nationality, party group and national party affiliations, and EP committee,

[4]I code a vote as final if it was a single vote on a proposal, a vote for resolution on a non-legislative proposal, or a vote for legislative resolution on a legislative proposal. I exclude votes on amended proposals because, while they do consider the bill in question as a whole, they are generally followed by a vote for legislative resolution on the same bill. These two votes are linked and should not be treated as separate observations.

[5]The EP voted on every proposal it debated, but did not debate every piece of legislation on which it held a vote. The members of the EP are unlikely to have selected proposals to debate in a random fashion. Therefore, debated bills may systematically vary from proposals that the EP voted on without discussion. Clearly, the results I present here have little to say about the determinants of RCV requests for bills without debates. Nonetheless, these proposals are associated with a variety of textual information, including legislative reports and the texts of the bills themselves, and it may eventually be possible to extend the methods I describe here to these data.

[6]The EP President handles the formal aspects of the debate, such as introducing the topic at hand, limiting speakers to time, and ensuring that speeches are germane. Therefore I excluded the President's speech segments from analysis.

bureau, and delegations memberships. I also recorded a variety of information about the bill under discussion from the Parliament's Legislative Observatory (European Parliament 2009*b*), including the parliamentary procedure under which the Parliament considered the bill,[7] the issue area covered by the proposal,[8] and the identity of the bill's rapporteur, the MEP responsible for guiding the legislation through Parliament. Finally, I cross-referenced the debates with the vote and RCV data available in the online archives, associating each debate with a final vote. I recorded the outcome of each vote, the type of vote—single, non-legislative resolution, or legislative resolution—and whether or not the vote was roll-called. For RCVs, I also recorded the overall vote tallies and the vote decision—yea, nay, abstain, or missing—of every MEP who spoke in the debate on the bill.

4.3 Roll Call Prediction

Predicting RCV requests from debate text is, at its heart, a classification problem. Specifically, we require a classifier that can take debates on proposals on the floor and, using only the text of the debates and related meta-data, group those debates into two categories: debates on bills that political actors will choose to roll-call, and debates on proposals that the floor will vote on sans roll. Text classification is a well-developed field in machine learning (Mitchell 1997, Manning & Schütze 1999) and researchers have applied a variety of methods, including naive Bayes classifiers (Lewis 1998), maximum entropy classification (Berger, Pietra & Pietra 1996), and support vector machines (SVMs) (Joachims 1998), to text categorization tasks. I use SVMs[9] in this paper because they often outperform other classifiers at both the traditional problem of categorizing texts by topic (Joachims 1998) and the higher-

[7]The EP uses a variety of legislative procedures including Assent, Codecision, Cooperation, and Consultation, and numerous non-legislative procedures such as own-initiative reports, decisions, and resolutions.

[8]The EP bureaucracy provides a four-level issue classification for every proposal. I coded only the first level in the classification scheme, which groups bills into eight issue areas: citizens' rights, internal market, agricultural fisheries and economies, economic and social cohesion, economic and monetary system, common foreign and security policy, justice, and the state and evolution of the Union.

[9]I also experimented with random forests (Breiman 2001) with similar results.

order problem of textual sentiment classification (Pang, Lee & Vaithyanathan 2002), which features prominently in my approach.

SVMs are based on rigorous statistical learning theory (Vapnik 1995) and are exceptionally flexible learners that are capable of finding highly non-linear relationships in data. Moreover, they tend to strike a nice balance between pattern matching and data over-fitting and, in practice, often outperform similar techniques at out-of-sample classification. At their core, SVMs are a two-type classification algorithm. They take a series of training examples, each consisting of a vector of predictor variables and a binary outcome variable, and attempt to learn a function mapping predictors to outcomes that is reasonably robust to idiosyncrasies in the training data. Once trained, an SVM should be able to take predictors from previously unobserved data points and predict their outcomes with high accuracy.

The SVM algorithm represents each observation as an n dimensional point in space, with one dimension for each predictor variable. The most basic SVM is a linear classifier that finds the $n-1$ dimensional hyperplane in the predictor-space that simultaneously divides the training data in a way that minimizes outcome classification error and maximizes the "margin," or the distance between the separating hyperplane and the nearest observations from both classes. Thus, the SVM algorithm works by finding the hyperplane that both separates the training data in one class (e.g. RCV requests) from the other (e.g. unrecorded votes), and maximizes the degree of separation between classes. Figure 4.1 illustrates this concept for a simple case, with two predictor dimensions, x_1 and x_2.[10] Each point in the figure represents a training case; filled in circles correspond to one possible outcome, while empty circles represent the other. The three lines—H_1, H_2, and H_3—depict three possible separating hyperplanes. H_3 does not cleanly separate the training examples, but both H_1 and H_2 do. The SVM algorithm would select H_2 because it not only separates the training data cleanly, but also maximizes the margin between the two classes.

[10]This figure is a public-domain image housed on the Wikipedia, available at `http://en.wikipedia.org/wiki/File:Svm_separating_hyperplanes.png`. I retrieved this copy on August, 22, 2010.

Figure 4.1: Separating hyperplanes.

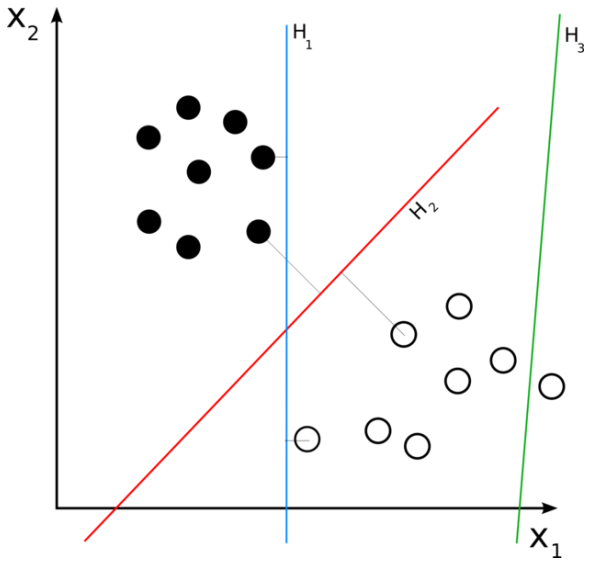

Of course, real-world applications do not typically sport observations that are subject to clear linear separation. Thus, researchers have augmented basic SVMs in two ways. First, Cortes & Vapnik (1995) extended the basic algorithm to allow for non-separability through "soft margins," using techniques that penalize the classifier for mis-classified training instances while still maximizing the margin between the correctly classified instances. This means that modern SVM software can deal with training data for which no perfectly separating hyperplane exists. Second, researchers have also extended SVMs to learn non-linear relationships between predictors and outcomes by first using kernel functions to map non-linear input spaces into linear feature spaces and then applying the standard SVM methods to the transformed space (Boser, Guyon & Vapnik 1992). Using these extensions, one can use SVM techniques developed for linearly separable data on error-prone datasets that are characterized by highly non-linear relationships between predictors and outcomes. Note, finally, that once fitted, an SVM can take a given observation—described by a vector of predictor variables—and calculate its distance in feature space to the maximal-margin separating hyperplane—also known as the decision plane—that the SVM learned from the training

data. Thus, the SVM chooses a likely classification for new instances simply by observing the sign of the distance calculated for the given case; observations that sit on one side of the hyperplane are likely to fall into one outcome class, while instances residing on the other side should represent examples of the other class. Furthermore, an observation's distance from the fitted hyperplane represents how "certain" the SVM is that the given observation comes from a given class because examples that sit close to the hyperplane are more likely to be misclassified than those that are further from the dividing line. These hyperplane distance measures play an important role in my analysis. As we shall see, I use them as a tool to measure, using only what MEPs said in plenary, how likely a given MEP is to support a particular piece of legislation.

A more detailed discussion of the mathematics underlying SVMs is beyond the scope of this article and Burges (1998) provides a detailed introduction to both theory and practice that interested readers may find useful. I used the e1071 R package (Dimitriadou, Hornik, Leisch, Meyer & Weingessel 2009) to fit all the SVMs presented here, and used default parameter settings in each case.[11]

Before proceeding, I randomly divided the data into training, development, and testing sets containing 70%, 10%, and 20% of the observations, respectively, following standard machine learning conventions. I fit models to the training data, using the development set to evaluate out-of-sample accuracy while tweaking model parameters and choosing which text features and meta-data to include when fitting the models. The test set provides an out-of-sample accuracy benchmark for the final fitted models. Each observation corresponds to a single debate/final vote and Table 4.2 provides a breakdown of the number of debates, speech segments, and segments per debate across the three data subsets.

[11]Specifically, I use the Gaussian radial basis kernel with $\gamma = 1/k$, where k is the number of predictor variables.

Table 4.2: Debate statistics.

	Total	Train	Test	Development
Debates	769	538	152	79
Speech Segments	15017	10517	2869	1631
Avg. Segments per Debate	19.5	19.5	18.9	20.6

4.3.1 Easily Observed Covariates

Previous research has shown that the decision to call roll in the EP co-varies with both procedural factors and issue area (Carrubba et al. 2006). If such easily observably characteristics are capable of explaining a large part of the variation in recorded vote requests, they could provide a low-cost set of variables with which to build selection models of roll call voting behavior. Therefore, as a baseline, I looked at how effective easily observable characteristics of bills are at predicting whether or not the Parliament will subject final votes to roll call. Specifically, I fit an SVM to the training data, using easy-to-code aspects of the legislation recorded in the vote data and the EP's online Legislative Observatory—such as the type of final vote (single, resolution, or legislative resolution), issue area of the bill, and bill procedure—to predict RCV occurrence. Using the development set to monitor out-of-sample accuracy I settled on a specification that included final vote type, issue area, and four procedural dummy variables coding Codecision bills, Consultation legislation, own-initiative reports, and EP resolutions, as predictor variables.

4.3.2 Bill Salience

Another useful predictor of variance in vote recording is bill importance; at least in the U.S. Congress it appears that RCVs are more frequent on especially salient legislation (Clinton & Lapinski 2008). Politicians wishing to take public positions on popular or controversial legislation and party leaders calling roll either to embarrass the opposition or monitor the votes of their own rank and file should all focus their roll call requests on salient legislation. It

makes little sense to take positions on bills that nobody cares about; similarly, the opposition is unlikely to be embarrassed about their position on trivial legislation, however far it is from the mainstream, and it is wasteful for parties to expend time and political capital on monitoring and disciplining members for their behavior on low-priority votes. While there exists no direct measure of bill importance for the EP, the debate data provides information that may proxy for bill salience: number of speakers. Debates with more speakers are likely to deal with topics that are important to a wider audience than the bills discussed by only a few MEPs. For example, the 6th term debate with the most speakers focused on EP bill A6-0070/2004, titled "Parliament's opinion concerning the draft treaty establishing a constitution for Europe," an issue of utmost importance to MEPs, and a topic ripe for position-taking.[12] 119 MEPs spoke in the debate, representing the largest turnout in any debate in the dataset.

I use the number of MEPs that spoke in each debate to capture the relationship between proposal salience and RCV requests. The number of speakers ranges between one and 119, with a mean just under 20, and a standard deviation around 14. Thus, there is wide variability in the number of speakers across debates. I used these counts to fit a second SVM and also include counts of MEPs speaking formally on behalf of party-groups, and speaker counts by party-group, in the salience model.

4.3.3 Naive Text-Based Classification

I next evaluated the effectiveness of using entire debate transcripts, and a bag-of-words approach, to predict whether or not MEPs will choose to record their final vote on a given proposal. To do so, I ignored the speech segments in the debates, and collapsed each observation into a single debate-level document. Next, I used the RSNL R library (Fader, King, Pemstein & Quinn 2009) to tokenize the debates, converting each document into a series of individual, lower-case, words. I also removed punctuation and garbled tokens, stemmed each

[12]The final vote on the bill was, indeed, recorded.

word to reduce it to its root (e.g. running and runner were converted to run), and replaced all numbers with the word "NUMBER." Finally, I converted the debates in each data subset into a term-frequency matrix, containing debates along the rows and normalized[13] word frequencies along the columns. Thus, I represented each debate as a vector of 7108 word frequencies. Therefore, for example, the term-frequency matrix for the training dataset contained 538 rows—one for each debate—and 7108 columns, where each column represented the relative frequency with which a given word—say "unacceptable"—appeared in a given debate. I fit an SVM to the training data, using only these vectors of word frequencies to predict roll call vote occurrence.[14]

4.3.4 Speaker Sentiment

Bag-of-words classifiers have shown themselves to be highly accurate in topic categorization (Joachims 1998). In the context of the EP, for example, a 200 word speech with 30 occurrences of the word "fish" in it is likely to deal with fisheries legislation. But the relationship between word frequency and higher order classification tasks, like the one considered in this paper, is less clear-cut. Therefore, classifiers trained on debates represented as simple bags of words may predict roll call requests poorly. I take advantage of the structure of the debate data and political science theory to build a more nuanced RCV request classifier. Specifically, I make use of the fact that the pattern of roll call requests in legislatures should reflect not only the salience of the proposal under consideration, as discussed above, but the pattern of support for the bill in the parliament. Fundamentally, recorded votes should occur only when there is reasonable variation in support for the proposal under debate. This variation

[13] I used a term frequency by inverse document frequency (tf.idf) weighting scheme (Manning & Schütze 1999, pp. 543). This approach first takes the number of times a given token occurs in an observation and then normalizes by the number of observations in which the given token occurs. Thus, common words like "the" and "and" receive relatively low weights, while unusual words are given higher priority.

[14] I experimented with a variety of schemes for weighting the term-document matrix, including raw term frequency and presence-of-feature (rather than frequency-of-feature) approaches. Similarly, I experimented with filtering out especially (un)common tokens from the term-frequency vectors. The various approaches all yielded similar development set prediction accuracy.

may come in multiple forms, depending on the reasons underlying the RCV request. First of all, when politicians call roll to publicize their differences in opinion with their opponents there should be significant inter-party variation in proposal support. On the other hand, if a party requests a RCV to help monitor and discipline its members, or to expose cracks in the opposition, there should be localized intra-party or intra-coalition differences in bill support. In cases of pure position-taking, patterns of support may be more haphazard, especially in parliaments with low party cohesion, but we should, nonetheless, be more likely to observe roll calls on bills with significant numbers in support and opposition, than on proposals that are uniformly regarded by the legislature.[15]

Taking advantage of the likely relationships between bill support and recorded vote requests, I built a two-stage classifier that first uses basic sentiment analysis (Pang & Lee 2008) to predict individual speakers' support for the proposal under debate and then predicts RCV requests based on variables summarizing the pattern of speaker support for the bill, as estimated by the first stage classifier. To train the first stage classifier, I created a subset of the training data containing only debated final votes on which the Parliament held a roll call vote. I dropped speakers who could not vote—bureaucrats, visitors, etc—or who spoke in a purely formal capacity, from the data subset. Furthermore, to focus on debates with reasonable variation in sentiment, I limited the data subset to cases where at least one speaker voted against the proposal in the RCV on the bill. These restrictions yielded a sample of 216 debated final votes containing 5291 speech segments. In contrast to the previous section, I kept observations at the speaker level. I then followed the same procedure for converting textual observations into term vectors that I used for the naive text-based classifier, except that, following conventions in sentiment analysis (Thomas, Pang & Lee 2006), I represented each speech in terms of a term-presence vector rather than a term-frequency array. That is, I coded only whether or not a given term appeared in a particular speech, rather than how

[15]One exception may be proposals on highly salient issues like decisions on military action. In these cases, politicians may wish to publicize their positions even if the entire legislature is of one mind.

many times it occurred. Furthermore, I removed especially common and rare terms that occurred in more than 99 per cent or less than 1 per cent of the observed speech segments, reducing the length of the term-presence vectors to 1965 elements each. I then separated the RCV subset into training, testing, and development sets, just like I did with the whole dataset. The training, testing, and developments sets held 3758, 963, and 570 speech segments, respectively. I fit an SVM to the training set, using the term-presence patterns in individual speeches to predict whether or not the speaker who made each speech voted in support of the bill in question on the final recorded vote for that proposal.[16] Note that the naive RCV classifier that I described in the previous section operated at the level of individual debated final votes, represented each debate as a vector of 7108 term frequencies, and used those vectors to predict whether or not a particular vote was held publicly. In contrast, the first stage of the sentiment classifier focused on the speech level, represented each speech segment as a vector of 1965 term-presence indicators, and used those vectors to predict whether or not the speaker who generated the given segment voted yes on the associated final vote.

I then used the results of the fitted sentiment model to generate inputs for a second-stage RCV classifier. Specifically, I took the trained sentiment classifier and used it to predict speaker support for all of the bills in the original training, development, and testing sets. That is, I used a speech-level classifier trained on bills with recorded votes to predict the voting behavior of speakers in every debate, both debates on bills with recorded final votes, and debates on proposals that faced no roll call. Therefore, my approach tacitly assumes that the patterns of speech that predict positive votes by individual MEPs in debated bills which were subjected to roll call will also predict positive votes by MEPs on votes that were not subsequently put to public vote. This assumption seems warranted because positive

[16] On a given vote attending MEPs may support or oppose a bill or formally abstain. Furthermore, MEPs can simply skip votes. For the purposes of this paper, I attempted to predict whether or not MEPs would vote yes on a given bill, treating no votes, abstentions, and missing votes equivalently. In reality, whether an MEP votes no, abstains, or simply misses a vote provides information about strength of a MEP's support for a bill. I plan to exploit this information in future research.

and negative speech are rather general concepts, and it seems unlikely that MEPs would substantially alter the ways in which they express their support or opposition to legislation across types of votes, even if they could anticipate public votes.[17] Of course, many words that will predict speaker sentiment in a given debate will tell us little about how speakers feel in the context of another piece of legislation. For example, the term "sovereignty" could covary highly with negative votes in a debate over a piece of legislation seeking to expand European powers into areas—like tax policy—that are the traditional domain of member states, but could predict support for legislation devolving monitoring of states' compliance from the Commission to local organs. Therefore, the sentiment classifiers' success rests on its ability to find patterns in speech that are generally related to speaker sentiment. Such patterns could include the use of positive and negative adjectives, for example, although substantially more complex work presence patterns are likely to play a role. SVMs are perhaps the best known tool for striking the balance between uncovering complex patterns in data and avoiding over-fitting that is necessary to perform successfully on such a task; nonetheless, the success of the sentiment classifier rests crucially on the variance provided by the training data.

The sentiment classifier works at the speech segment level, mapping verbiage to voting behavior. To use the sentiment classifier's estimates to predict RCV requests, it is necessary to aggregate the sentiment classifier's predictions upwards to the debate—or, equivalently, final vote—level. Let d_{ij} be the distance between the term-presence vector representing speaker i in debate j and the first stage sentiment classifier's decision plane. Each d_{ij} represents how supportive or antagonistic speaker i's speech on bill j appears to the sentiment-based SVM. I calculated the mean distance and overall standard deviation of speaker distances for each debate j, as well as the within-group standard deviations for each of the seven major party groups within the EP.[18] These measures capture the predicted average level of support among

[17]That is not to say that we should not expect selection effects in floor speech that may covary with RCV requests. Nonetheless, while selection issues are likely to weaken the predictive accuracy of the speaker sentiment model, they do not fundamentally undermine its utility.

[18]I set group standard deviations to their average values for debates in which no representative of a given

Table 4.3: Sentiment classifier accuracy (in percent).

	Development Set	Test Set
Modal Category	75	74
Sentiment Classifier	82	81

speakers for the proposal under debate, as well as the variance in support across speakers, both as a whole, and within groups. Therefore, they describe the level of disagreement on the floor during the debate and should help predict RCV requests. Of course, not every MEP speaks in every debate and these summaries of estimated speaker sentiment will miss patterns in opinion within the chamber that MEPs do not voice in plenary. Nonetheless, the variation in speaker sentiment provides substantial information about prevailing attitudes among MEPs and the level of division in the chamber prior to a given vote.

Finally, I used these debate-level sentiment summary scores to train a second-stage RCV classifier.[19] That is, using debate-level summaries of patterns in the estimated level of support expressed by speakers in a given debate, I trained an SVM to predict whether or not any party group would request an RCV on the final vote attached to the debate in question.

4.3.5 Roll Call Prediction Results

Before proceeding to a comparison of the various RCV classifiers, it is worth noting the performance of the first-stage classifier in the sentiment model. Table 4.3 compares the predictive accuracy of the sentiment classifier to a simple pick-the-modal-category approach to prediction. Clearly, participants in debates are quite likely to vote in support of the bills they discuss, and simply guessing that the speaker will vote for the proposal is accurate about

group spoke.

[19] I experimented with a variety of other sentiment summaries, including measures of the overall inter and intra-group variance contributions, but the classifier trained on the average distances, distance standard deviations, and group-specific standard deviations provided the best predictive accuracy on the development set.

three quarters of the time. This makes good sense; final votes represent the culmination of extensive legislative wrangling and can reflect the outcomes of a multitude of previous votes on amendments or portions of the bill. Amendments allow MEPs to better tailor legislation to preferences on the floor and final votes may not reflect all of the disagreements voiced by MEPs when the bill was initially debated. Nonetheless, speakers' utterances during debates do help to predict their voting behavior on whole bills and the term-presence SVM modestly outperforms modal guessing, correctly predicting 81% of test-set outcomes. Sentiment classification is a difficult problem in natural language processing and previous research has shown that one may substantially improve the accuracy of sentiment classifiers by making use of information that goes beyond simple term-presence vectors (Pang & Lee 2008). Therefore, in future work, it may be useful to take advantage of indicators of agreement between speakers, such as the words legislators use when referring to other speakers in the debate (Thomas, Pang & Lee 2006), or speaker characteristics, such as group membership and nationality, to improve the performance of the stage one classifier.

Figure 4.2 presents the development and test set accuracies for modal guessing, the six RCV predicting classifiers, and a seventh classifier, sentiment+, that combined the variables from the easy-to-observe, importance, and sentiment classifiers. The EP recorded about 40% of final votes and picking the modal category (i.e no recorded vote) generated accurate predictions 61% of the time in the development set and 58% in the test set. The easy observables, importance-based, and naive text-based classifiers all performed at around modal category accuracy. Of these three fitted SVMs, only the naive text-based classifier outperformed modal guessing in both development and test sets, and the easily-observed and importance models actually did worse than modal guessing on the development data. Therefore, while they correlate with roll call voting, neither easily observed predictors like bill issue area and legislative procedure, nor indicators of the salience of a bill explain much of the variance in RCV requests, at least when held to the high standard of out-of-sample prediction. This is an important result because it makes one question the likely utility of

Figure 4.2: RCV classifier accuracy (in percent).

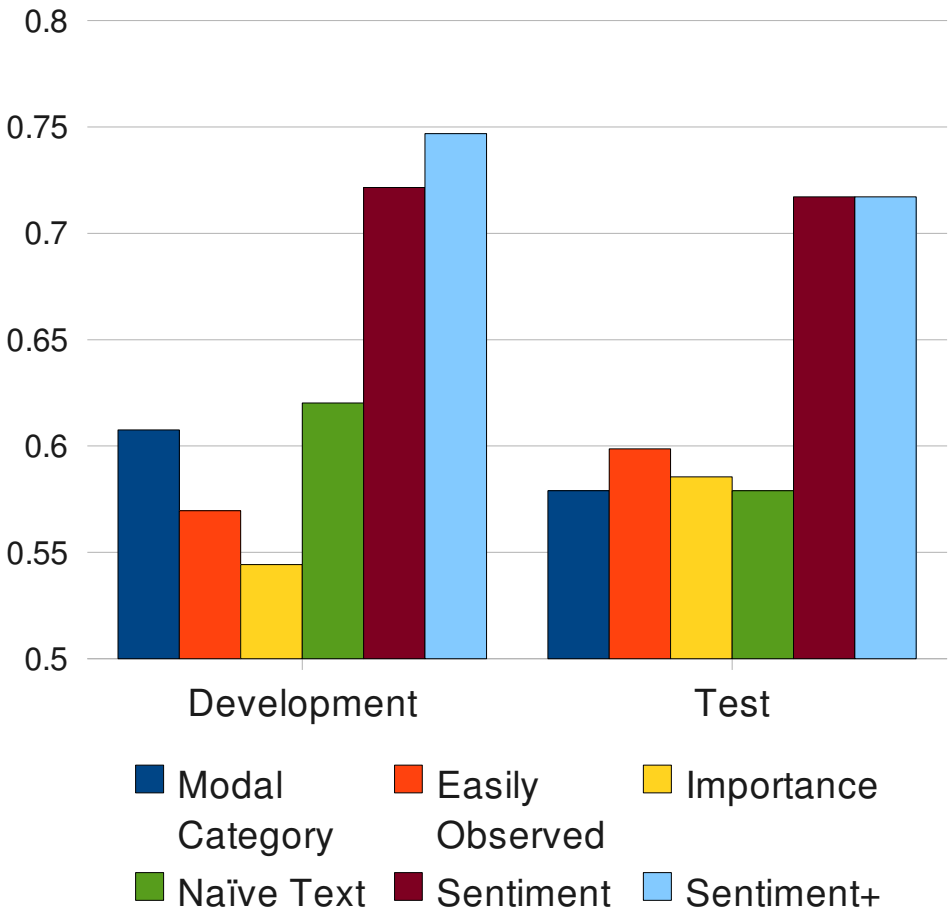

selection models of roll call voting based purely on simple bill characteristics like issue area and procedure. Political parties use roll call votes to reveal their preferences or those of their opponents to the public, or to monitor the behavior of their members. Thus, underlying preference patterns are likely to be the true determinants of roll call vote requests. While preferences may tend to align in a manner that makes an RCV more likely under certain broad sets of circumstances than others, such gross instruments will always predict RCV

requests quite poorly. Furthermore, one cannot hope to effectively correct for selection in RCVs using predictors that explain so little variance in RCV requests.

Similarly, a naive bag-of-words representation of floor debates does a poor job of predicting when MEPs will request recorded votes. In contrast, an identical bag-of-words approach is quite accurate at identifying the issue area of a bill under debate and predicted 57% of the test-set cases' 8-level issue area codes correctly, compared to a modal category predictive accuracy of 22%. Therefore, while the text of the debates contain a great deal of information about the bills under discussion, a structure-free (i.e. naive text) approach to text classification appears unable to effectively tease out relationships between what politicians say when debating a bill and their propensity to request RCVs when voting on the bill as a whole.

On the other hand, figure 4.2 shows that the sentiment-based classifier significantly outperforms the modal category baseline, and the other classifiers, at RCV prediction. Representing debates in terms of the distribution of sentiment that the speakers in the debate have towards the bill under discussion generates 72% accuracy in both the development and test sets, a substantial improvement over modal guessing. By taking advantage of things we can observe—speeches on the floor and voting in roll calls—we can generate useful summaries of the level of disagreement surrounding a particular piece of legislation. These summaries, in turn, account for a substantial amount of the variation in RCV requests and, when coupled with standard machine learning techniques, are useful for predicting RCV requests in previously unobserved data. Furthermore, and somewhat surprisingly, incorporating other information into the classifier—such as measures of bill importance, issue area, and legislative procedure—add little predictive accuracy to the model, as the results for the sentiment+ model in figure 4.2 reveal. While the sentiment+ model outperformed the basic sentiment-based classifier on the development set, there was no difference in accuracy between the two classifiers on the test data.

The sentiment-based approach to RCV prediction outperforms the competing models not just because it draws on a rich source of information—indeed, the naive bag-of-words

classifier uses exactly the same dataset—but because it is better wedded to theoretical accounts of roll call voting than the other approaches. While we can construct stories for why the sorts of preference distributions that encourage RCV requests should occur more often under particular institutional contexts, on especially salient legislation, or when actors consider certain issues, these factors are—at best—weak proxies for the determinants of roll call vote selection. Similarly, while the speeches that MEPs give during legislative debates contain information that can predict roll call requests, one needs to take advantage of theory to best leverage this information. The predominant theories of roll call requests all rely on patterns in legislators' voting intentions to predict roll call votes; the sentiment model works by using floor debates to explicitly model these voting intentions and, in turn, predicts roll call votes based on patters in legislators' expected voting behavior.

4.4 Speaker Sentiment and Roll Call Requests

The sentiment-based model is interesting, not only because of its predictive accuracy, but because of what it can tell us about the circumstances under which politicians endogenously request RCVS. The first-stage sentiment model generates measures of the likelihood that a given speaker will choose to vote in favor of the proposal under debate. While the model is based on theory, it also provides a novel tool with which to test hypotheses about when roll call votes are most likely to occur. Specifically, the sentiment measures produced by the first stage of the classifier allow one to examine how the expected pattern of support for a bill influences politicians' decisions to request recorded votes.

4.4.1 Broad Patterns

Figure 4.3 shows predicted RCV probabilities for debates as a function of average predicted speaker support for the bill (left panel) and the standard deviation of support for the proposal

Figure 4.3: Predicted probability of RCV by 1st-stage sentiment classifier decision value debate mean and standard deviation.

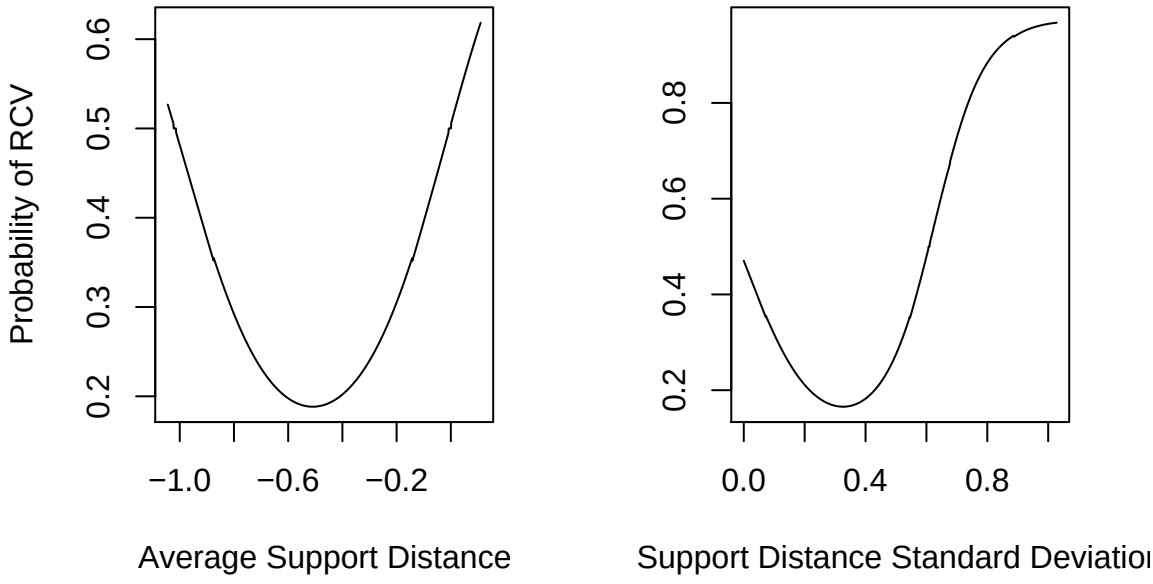

under debate (right panel).[20] Remember that each speaker's sentiment is captured by her speech's distance[21] from the decision plane of the 1st-stage sentiment classifier. To generate these graphs I held all other inputs to the sentiment-based classifier at their sample means, while varying the quantity of interest over its observed range. Looking first at the left panel of figure 4.3, note that RCVs are most likely at the tails of the average support distribution. Especially negative distances from the stage one classifier's decision plane indicate likely support for the bill under debate while especially positive distances imply that the speaker

[20]SVMs are a non-parametric machine learning tool and are not built upon an explicit probability model. Therefore, the SVM does not generate predicted probabilities directly. I generated the predicted probabilities in figures 4.3 and 4.4 by fitting a logistic distribution to the support plane distances calculated for each observation by the fitted SVM, using maximum likelihood. More specifically, after fitting the SVM to the training set, I regressed RCV requests on the fitted distances from the decision plane for each case in the training set, using logistic regression. To generate out-of-sample predicted probabilities one takes the out-of-sample observation and feeds it into the fitted SVM, generating a decision-plane distance. Next, one uses the fitted logistic model to predict the probability of a RCV request from the given decision plane distance.

[21]More formally, by the vector of term-presence dummies representing the speech.

will fail to vote for the proposal. Note that the observed range of average support distances ranges only from -1.0 to around 0.1, in contrast to the speaker-level distances, which range between -1.8 and 1.5. Therefore, the left panel of figure 4.3 indicates that RCVs are most likely for votes characterized by especially high levels of support, and by votes where the floor is divided. The second panel in figure 4.3 further nails home this point but adds some refinement to the story. First of all, RCVs are most likely on bills characterized by substantial disagreement among debate participants. But, RCVs are also more likely—although, to a lesser degree—when everyone taking part in the debate agrees—as captured by variation in the speakers' estimated sentiments—about the proposal under discussion. For the most part, MEPs call roll to publicize differences in opinion. This is consistent with standard stories of roll call voting, which predict recorded votes when politicians see an opportunity to differentiate themselves from their competitors on controversial issues (Saalfeld 1995). On the other hand, MEPs are also somewhat likely to request public votes when they are highly unified in support of a particular bill. This seems likely to represent position-taking on universally popular proposals, or perhaps may indicate shows of solidarity on issues that pit the EP against other institutional actors such as the European Commission or Council.

Figure 4.4 is analogous to the right panel of figure 4.3 and displays predicted RCV probabilities as a function of within-group support standard deviations for the seven of the EP's parliamentary party groups.[22] I plotted each curve in figure 4.4 by varying each within-group support standard deviation over its empirical range, while holding all other model inputs at their means.[23] RCVs become more likely as within-group disagreement grows, although the strength of the relationship varies across groups. The largest groups—

[22]These include the Group of the European People's Party (Christian Democrats) and European Democrats (PPE-DE), the Socialist Group in the EP (PSE), the Group of the Alliance of Liberals and Democrats for Europe (ALDE), the Union for Europe of the Nations Group (UEN), the Group of the Greens/European Free Alliance (Verts/ALE), the Confederal Group of the European Left-Nordic Green Left (GUE/NGL), and the Independence/Democracy Group (IND/DEM). I exclude the Identity, Tradition, and Sovereignty Group (ITS) because it was very short-lived.

[23]The within-group support standard deviations are surprisingly uncorrelated across groups, with the PPE-DE and IND/DEM groups topping the scale at $r = 0.25$. The curves in figure 4.4 are different lengths because their observed supports vary.

Figure 4.4: Predicted probability of RCV by 1st-stage sentiment classifier standard deviation, across EP party groups.

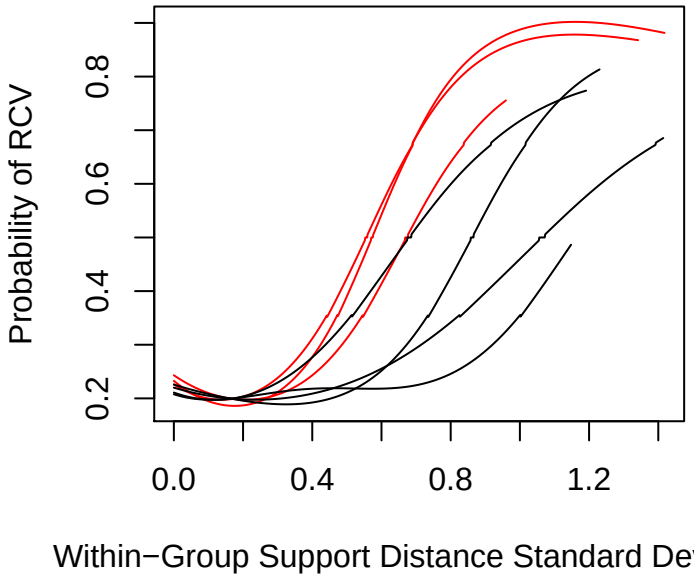

the largely Christian Democratic PPE-DE, the socialist PSE, and the liberal ALDE—are highlighted in red in the graph and display the strongest association between intra-group strife and RCV requests. Therefore, while overall levels of conflict among debate participants are indicative of RCV likelihood, intra-party squabbles also make recorded vote requests more common. This result is potentially consistent with a story of roll calling by parties in order to monitor their own members and maintain cohesion (Carrubba, Gabel & Hug 2008), and also with the idea that parties may request public votes to expose cracks in their opponents' discipline (Saalfeld 1995). It is exciting, and somewhat surprising, that we observe this effect. For one thing, we might expect to observe little open disagreement in debates within parties for exactly the same reasons we might expect them to maintain high discipline in RCVs. Furthermore, disgruntled partisans may say one thing in a debate and then vote the other way when pressured by the whip. The sentiment measures I use here are not

pure measures of support, per se, but rather predictions about likely voting behavior given speaking behavior; insofar as politicians say one thing and do another, they will tend to underestimate differences within party groups. Nonetheless, one must be cautious when interpreting figure 4.4 because it describes the relationship between expected intra-group cohesion and RCV requests across roll calls sponsored by all of the party groups in the EP. Thus, it is difficult to differentiate between parties' use of roll calls to monitor and discipline their own members, and their tendency to request roll to expose divisions within the opposition.

4.4.2 Group-by-Group RCV Requests

To more carefully investigate the mechanisms underlying roll call requests I constructed a panel of potential group roll-calls based on the same set of 769 final votes that I used to train and test the sentiment model. Each observation in this dataset represents a single group's opportunity to request an RCV on a given final vote. The dependent variable in the analysis is a simple dummy indicating whether or not a particular group decided to sponsor an RCV on a given vote.

While previous research has examined such decisions by characterizing them as functions of broad categories of votes, I rely primarily on the outputs of the sentiment classifier—or speech-based measures of expected patterns in public voting behavior by MEPs—to model the circumstances under which groups request roll call votes. Remember that the stage-1 support classifier produces a measure, d_{ij}, representing the distance between the term-presence vector representing speaker i's speech in debate j, and classifier's decision plane. Each d_{ij} is a measure of how likely speaker i is to vote in support of proposal j, as predicted by the classifier. Thus, d_{ij} is a proxy for how speaker i would vote on j if the vote were recorded. In other words, d_{ij} provides a measure of both of the quantities described by the cells in the top row in table 4.1. Of course, if vote j was held publicly, we are likely to

have an even better measure of this concept, namely MEP i's actual vote.[24] But, when no roll call vote is held, the vector $\mathbf{d_j}$, containing decision plane distance measures for every speaker that took part in debate j, provides a novel proxy for the, otherwise unobservable, concept described by the top right cell in table 4.1. Therefore, I use various summaries of $\mathbf{d_j}$ as predictors when modeling roll call vote requests by particular party groups.[25]

Both position-taking and discipline-based accounts of RCV selection posit that RCVs will be more common when legislators disagree on a proposal than when a vote is uncontroversial. I measure the level of agreement within the EP, or *chamber disagreement* on a given vote j, in terms of the standard deviation of $\mathbf{d_j}$. Similarly, to examine how intra-group agreement co-varies with RCV requests, I operationalize *group disagreement* for group k on vote j as the standard deviation of the subset of elements in $\mathbf{d_j}$ representing speeches from members of group k. Political parties that call roll to take positions may either wish to claim credit for a proposal that they support or to distance themselves from, or blame another party for, bills that they disagree with. Therefore, I measure the *group sentiment* of group k towards proposal j in terms of the average decision value of speakers from group k in the debate on bill j. Additionally, I consider a number of non-speech-based predictors of RCV requests as control variables, including whether or not bill j was especially *salient*,[26] if the vote was on *legislation* or a parliamentary resolution, and whether or not the bill under consideration represented a *group report*. This last variable is an indicator that equals one when the rapporteur for bill j—the MEP assigned to research the measure and guide it through the

[24] This is not always the case. For instance, voter i may not have attended plenary on the day that vote j was held.

[25] The process I use to model group-level roll call requests is analogous to what I did when fitting the two-stage sentiment classifier above. Nonetheless, the the two models differ substantially in purpose. In the previous section I sought to determine if patterns in expected voting behavior could explain the variance in overall roll call vote predictions, and my focus was, primarily, classification accuracy. In this section the sentiment scores serve to examine the predictive accuracy of common theories of roll call selection.

[26] As in the previous analysis, I measure salience in terms of the number of speakers present for the debate on bill j. Initial exploratory analyses indicated that the relationship between number of speakers and roll call requests, while generally positive, is non-linear. Furthermore, the distribution of speakers is highly skewed towards the low end and contains quite a few outliers. Therefore, I classified votes as salient if the debates that preceded them featured an especially large number of speakers. Specifically, I coded a vote as salient if the number of MEPs that spoke in the debate was greater than one standard deviation above the mean.

Table 4.4: Sentiment and party group roll call requests.

	Model 1	Model 2	Model 3
Group Sentiment	-0.66 (0.15)*	-0.82 (0.15)*	-0.98 (0.16)*
Group Disagreement	-0.26 (0.22)	-0.23 (0.22)	-0.06 (0.22)
Chamber Disagreement	2.30 (0.30)*	2.06 (0.32)*	1.88 (0.32)*
Salient	0.81 (0.34)*	0.84 (0.34)*	0.81 (0.34)*
Legislation	-0.28 (0.15)	-0.41 (0.15)*	-0.40 (0.15)*
Group Report		1.20 (0.16)*	0.07 (0.46)
Group Sentiment × Report			1.54 (0.55)*
Observations	4376	4376	4376
Log-Likelihood	-1099.10	-1068.17	-1062.75

Standard errors are in parentheses and $* \equiv p < 0.05$.
Party group fixed effects not shown.

parliamentary process—hails from group k, and which equals zero otherwise.

Table 4.4 presents three logistic regression models that predict groups' roll call vote requests using combinations of the covariates described above.[27] Each model includes group fixed effects to account for idiosyncrasies in predilections to request roll call votes across groups, which table 4.4 suppresses. Furthermore, the table reports clustered robust standard errors which reflect possible dependence between observations within votes. Model 1 in the first column of table 4.4 predicts requests solely in terms of patterns in speaker sentiment, bill salience, and legislative procedure, while models 2 and 3 examine how a group's ownership of a given report affects its RCV sponsorship calculus, both individually, and in the context of group sentiment towards the bill.

The best predictor of roll call vote requests is *chamber disagreement*. The coefficient on this variable is positive and statistically significant across all three specifications. Furthermore, the substantive effect of chamber disagreement is substantial. Across all of the observations, the average predicted probability that a given group will sponsor a roll call on

[27]There are seven major groups (ALDE, GUE/NGL, IND/DEM, PPE-DE, PSE, Verts/ALE, and UEN) and 769 debated final votes, yielding 5383 possible observations. I was forced to omit 1007 of these potential observations when no members of a given group spoke in the debate and for which I have no measure of group sentiment or disagreement.

a vote is around 0.08, according to model 1. Yet, for those cases in the sample where *chamber disagreement* is at least one standard deviation above its mean, the average probability that a group will sponsor a roll call doubles to 0.16. Similarly, the probability of sponsorship drops to 0.05 on votes on measures where *chamber disagreement* is more than one standard deviation below the sample mean. Thus, as virtually all theoretical accounts of roll call voting would predict, and as the findings from the sentiment-based classifier depicted in figure 4.3 indicate, roll call vote requests are most common on controversial measures. This finding is not surprising but this paper provides the first empirical evidence for this result that does not rest on a substantial inferential leap. The sentiment scores used here derive directly from the behavior of individual MEPs, and as the previous section demonstrated, they are effective predictors of MEP voting behavior. Thus, while still a proxy for legislators' voting intentions, they do a much better job of measuring theoretical primitives than lumping votes into categories based on their issue areas, procedural status, or even bill authorship. Furthermore, they vary across votes at a much finer level of detail than do previously used proxies for patterns in vote intention.

The advantages of using such fine-grained measures of vote intention become especially clear when we directly address the reasons that individual groups call roll, something that is inherently difficult to do with existing measures. For example, the current analysis demonstrates that EP party groups use roll call votes to stake out their positions against bills with which they disagree. Indeed, the coefficient on *group sentiment* is negative and statistically significant across all three specifications in table 4.4. Again, the substantive influence of this variable is non-negligible; when the value of this variable is less than one standard deviation below the mean, within the sample, the probability that a given group calls roll nearly doubles to just above 0.14, according to model 1. As I argued in section 4.1, EP groups are inherently tied to particular proposals by the identity of the rapporteur—and constituents, national parties, and interest groups can readily guess the position of MEPs who hail from the same party group as the rapporteur. Therefore, they concentrate their RCV requests on

bills on which their positions may be less clear.

Interestingly, the relationship between *group disagreement* and RCV sponsorship is statistically insignificant across all three specifications. It seems that party groups may not take as much care to ensure that the entire group is on the same page before sponsoring a roll call vote as we might expect. One possible explanation for this finding is that groups sometimes call roll to discipline or monitor members but that their sanctioning powers are not sufficient to maintain high levels of party cohesion. Of course, this argument is largely speculative and the models that I present here cannot speak directly to this question. Indeed, because the sentiment measures that I use here only tell us something about how MEPs are likely to vote on recorded votes, they are less well-suited to examining theories of discipline-based roll calling than accounts based on position-taking. A perhaps simpler explanation is that, because EP groups are made up of so many individual national parties, group leaders are forced to put up with some degree of defection by party delegations on any given vote (Kreppel 2002). Indeed, multiple MEPs and party delegation staff members emphasized this point during personal interviews that I conducted in the fall of 2007.

Turning to the controls, one finds, not surprisingly, that all of the models in table 4.4 indicate that RCV requests are more likely on especially salient votes. In addition, groups are less likely to request vote recording on actual legislation than they are on parliamentary resolutions, at least once one takes the group membership of the rapporteur into account. Thiem (2006) argues that roll call votes should be less frequent on legislation in the EP because national party constraints are more binding on votes that can affect national policy than on pure position-taking measures, and finds strong empirical support for her claim in a model that predicts roll call vote requests with easily observable measures. On the other hand, the tendency of party groups to request roll call votes less frequently on legislation may have more to do with how much party groups disagree with one another than their internal cohesion because MEPs put far more effort into hammering out sustainable compromises—both within the Parliament and across EU institutions—on legislation than they do on

resolutions.

Finally, models 2 and 3 investigate how a party group's ownership of a given report affects its RCV sponsorship calculus. Interestingly, while, overall, groups call roll to signal their positions against legislation, they also call roll more often on their own reports than they do on other groups' reports. At first blush, this may seem counter-intuitive, but the interactive effect revealed by model 3 clarifies the result. Specifically, party groups sponsor RCVs on their own reports when the average group member agrees with the rapporteur and request roll calls on other groups' reports when they dislike the contents of the bill. Therefore, even though groups are already clearly linked to reports that belong to them, they nonetheless emphasize their support for bills that they favor when full responsibility for the report can be attributed to the group. On the other hand, they virtually never request roll calls on other groups' reports that are in line with their preferences, and only infrequently sponsor RCVs on their own reports when the majority of the group stands against passage. Therefore, groups appear to engage in RCV sponsorship more to take credit and assign blame than to simply signal their preferences to third parties.

4.5 Conclusion

The results in this paper demonstrate the potential that political speech has to help researchers to better understand roll call votes. Furthermore, they emphasize the important role that patterns in legislators' sentiment towards bills play in decisions to publicly record parliamentary votes. They underscore the serious nature of the selection problem inherent in RCV data while simultaneously developing a foundation upon which we may build tools to overcome this issue. While modest, the predictive accuracy of the sentiment-based model is encouraging, especially in light of the rather esoteric nature of the RCV classification task. This paper provides a solid first step towards predicting RCVs with legislative text, and highlights the importance of capitalizing on our theoretical understanding of politics

to best take advantage of natural language's ability to help predict political behavior. Furthermore, the approach presented here has the potential to provide measures that can help researchers statistically deal with selection bias in RCV-based analyses. Distances between debates and the decision plane of the sentiment-based RCV classifier represent predictive—if noisy—propensity-to-be-roll-called scores and I hope to examine their utility in this context in future research. I am also currently working on improving the accuracy of the stage one sentiment classifier, using agreement modeling techniques introduced by Thomas, Pang & Lee (2006). Hopefully an improved stage one classifier will lead not only to better speaker sentiment scores, but to greater accuracy in second-stage RCV prediction.

Furthermore, the current research provides a new tool for testing theories of roll call vote sponsorship, by generating a new way to measure expected voting patterns within legislatures. By using a non-parametric model to predict vote behavior with legislative speeches, I was able to tease out the circumstances under which party groups in the EP are most likely to request roll. The findings show that groups sponsor RCVs predominantly to distance themselves from controversial opinions with which they disagree, but also request public votes to demonstrate their support for their own rapporteurs. In general, they behave in a manner that is consistent with a process of credit-taking and blame assignment. Interestingly, there is limited empirical support for the argument that parties will call roll only when their own rank-and-file are in step with the general group line.

These findings have important implications for how we use roll call vote data from the EP. For one thing, they imply that measures of cohesion in the Parliament will be skewed and will tend to overestimate the overall level of disagreement within the EP. And, insofar as RCV sponsorship strategies may have changed over time, raw cohesion scores based on RCVs may misrepresent general trends in party group polarization. Furthermore, while the evidence on the relationship between within-group disagreement and RCV sponsorship is mixed, there is some reason to believe that RCVs may under-represent the level of within-group disagreement in the Parliament. This finding behooves researchers to re-examine some

of the most prominent findings in research on the EP. Notably, Hix, Noury & Roland (2007) use RCV data to demonstrate that intra-group cohesion has increased over time. Furthermore, they argue that party group cohesion is a function of a variety of factors, including group size, national fractionalization, and the percentage of national parties within the group that hold governing status at home. If these variables also predict RCV sponsorship then there is the potential that selection bias may influence these findings.[28] Furthermore, the propensity-to-be-roll-called scores generated by the sentiment-based RCV classifier represent an excellent measure with which to model RCV selection in this context. Thus, re-evaluating these findings in the context of selection bias represents another interesting avenue for future research.

The approach that I use here to measure sentiment on the floor has a variety of potential applications and generalizes easily. Specifically, whenever one observes a particular political indicator only infrequently but has access to speeches, legislation, reports, or other forms of political text that are likely to contain content that covaries with the variable of interest, there exists the potential to use standard machine learning tools to generate proxy measures for that indicator across an expanded sample of observations. Indeed, as large collections of digitized political text become increasingly available, these techniques may allow researchers to study a variety of topics that we currently know little about. More generally, as social scientists' access to large swaths of digitized data grows, techniques for data reduction are likely to become invaluable to researchers. Nonetheless, the applicability of these tools is not universal. As the present study demonstrates, simply throwing a bunch of text at a support vector machine is unlikely to generate meaningful results. Rather, one must carefully leverage theory to get the most out of these powerful data reduction techniques.

[28]While the current paper does not engage this issue directly, it is perhaps worth noting that the coefficient on group size is statistically significant when one adds the indicator to the model in table 4.4.

Chapter 5

Conclusion

While the EU has existed in one form or another since the 1950s, its modern incarnation is very young. Indeed, the codecision procedure is less than twenty years old and truly bicameral lawmaking has only recently applied to the majority of Union legislation. Political scientists have built a strong theoretical foundation for explaining inter-institutional bargaining within the Union, but our understanding of the role that information asymmetry, and transmission, plays in European lawmaking remains limited. Each of the above chapters examined some way in which politicians in the EU manipulate or react to information flows in the process of coordinating around policy compromises. I use this conclusion to outline some ideas for new research in this area that build on the work in the preceding chapters.

The formal model in Chapter 2 shows how the European bureaucracy can act as a go-between in bicameral negotiations between the Council and Parliament, overcoming information asymmetries between the two institutions when it is in the Commission's interest to do so. This work demonstrates that avenues of information transmission can help determine the speed of lawmaking in Europe and emphasizes the power that can accrue to political actors—namely the Commission—simply through access to information. But this model represents only a partial explanation of the politics of information sharing between European institutions. Most notably, there are actors beyond the Commission that have privileged access to information within the European legislative process. Perhaps most importantly, other authors have emphasized the advantages that rapporteurs—the MEPs responsible for guiding particular pieces of legislation through the lawmaking process—have in European lawmaking (Corbett, Jacobs & Schackleton 2003, Mamadouh & Raunio 2003, Farrell &

Hèritier 2004, Kaeding 2004, Kaeding 2005, Hausemer 2006).

Rapporteurs are more deeply engaged in the legislation they oversee than are other MEPs and regularly interact with representatives of the other EU institutions in informal negotiations over the bills that they chaperone. Thus, like the Commission, they have two informational advantages over the average MEP. First, they know more about the their own legislation's content than other members, and therefore are likely to have a better grasp of available policy options, and their likely implications. Second, through their multiple contacts with the Commission and Council, rapporteurs may be in a position to observe aspects of the Council's bargaining resolve than other MEPs cannot.[1] Indeed, Høyland (Unpublished) develops a model of codecision lawmaking that implies that bicameral bargaining in the EU should proceed more quickly when the rapporteur assigned to a bill hails from a party represented on the Council than when an opposition party controls the report. He argues that MEPs from governing parties will be well informed about the Council's bargaining position, encouraging speedy compromise.

Combining Commission-based and rapporteur-focused accounts of information transmission during European lawmaking is likely to be a fruitful avenue for future research. Specifically, MEPs may be able to look to two sources—the Commission and the rapporteur—for information about the Council's bargaining strength. Using evidence from two signal-sending parties may provide MEPs with more information about the Council's resolve than they could obtain from a single source, but this may depend on how preferences are distributed across actors. While the Commission's composition varies only infrequently, rapporteurs are drawn from all of the EP's major party groups and have preferences that span the ideological spectrum. Therefore, the Commission and the rapporteur may face differing strategic incentives and, therefore, may send inconsistent signals to other MEPs. In turn, the receivers of these signals may be able to turn such inconsistency to their advantage, extracting value

[1] Although I would argue that the Commission's direct access to internal Council organs provides informational advantages that should outstrip those available to rapporteurs through their engagement in tripartite negotiations.

from multiple messages under circumstances in which a single sender would be strategically uninformative.

Similarly, if we consider the (likely) possibility that actors are incompletely informed about the mapping between policy choices and outcomes, but that both the rapporteur and the Commission know more than most MEPs, then the ideological breakdown across the two informed parties may hold striking implications for when other MEPs may be able to elicit truthful information about the mapping between policies and outcomes. Specifically, the degree of Commission-rapporteur preference divergence will modulate how much information the Parliament can extract from these well-informed actors. This, in turn, may create information-based incentives for party groups to pursue particular reports and to leave ownership of other pieces of legislation to ideologically opposed groups. MEPs take rapporteur assignments very seriously. Indeed, in a series of interviews I conducted in 2007, MEPs repeatedly cited reports as one of their highest priorities, and as their main tool for wielding influence within the Parliament. While studies that evaluate the ways in which party groups distribute reports abound (Mamadouh & Raunio 2003, Kaeding 2004, Kaeding 2005, Hausemer 2006), merging this line of research with a theoretical account based on asymmetric information could pay dividends.

Furthermore, because rapporteurs are more ideologically varied than the Commission, they provide a useful tool for gaining empirical leverage over how MEPs react to the Commission's recommendations to the EP. The analysis in chapter 3 provides strong evidence that MEPs alter their voting behavior in response to bicameral bargaining constraints, and shows that MEP voting varies systematically with Commission recommendations. Nonetheless, there is room to improve on these findings, especially in terms of more accurately identifying the extent to which the Commission provides new information to MEPs that alters their voting behavior. An examination of situations in which rapporteurs and Commissioners differ in their opinions on tabled amendments could prove effective in this regard. Specifically, because both actors should have access to information unavailable to other MEPs,

how MEPs react to divergence in recommendations—and how these reactions interact with Commissioner, MEP, and rapporteur preferences—should help to isolate the independent effects of Commission, and rapporteur, influence.

Finally, there is substantial room for new scholarship linking the inter-institutional bargaining model of chapter 2 to the internal organization of the EP. While chapter 3 explores how Commission signals, transmitted as a part of this over-arching game, drive voting decisions in the EP, it does not directly discuss the role that party groups play in this dynamic. While the key finding of this analysis—namely that MEPs from governing parties are more susceptible to Commission influence than their counterparts—it leaves one wondering where the party groups fit into the picture. And, while chapter 4 demonstrated the role that party groups can play in publicizing legislative voting, potentially changing the costs of compromise for rank-and-file MEPs, it does not directly link this phenomenon to bicameral bargaining.

In their highly influential work Hix, Noury & Roland (2007) argue that the EP has grown increasingly cohesive over time as a result of the internal organization of party groups and the work of group leaderships. This impression extends to the EP itself; Inger Segelström, a MEP in the Socialist group, told me that the Socialists had become more cohesive during the 6th EP term precisely because of the efforts of PSE president Martin Schultz.[2] But others argue that group cohesiveness is attributable not to party organization, but to limited information and satisficing by MEPs. In particular, Ringe (2010) maintains that most MEPs are uninformed about the majority of legislation dealt with in the legislature and simply follow the lead of more informed members, such as rapporteurs and committee members, who share their preferences. Both of these stories have interesting implications for how MEPs coordinate around policies as they anticipate bicameral bargaining. On the one hand, if party groups are indeed strong, they should play a key role in organizing their MEPs into coalitions around viable bargains. On the other, if MEPs simply follow the leader, then the identity of the coordinating influences will vary from bill to bill. Each model of

[2]Interview with MEP Inger Segelström, November 27, 2007.

parliamentary organization is likely to generate different implications for the dynamics of bicameral bargaining between the EP and the Council. Explicitly incorporating theories of EP organization into a story of bicameral bargaining would be a first step in building a model of European policymaking upon full-fledged micro-foundations.

Over the past two decades scholars have increasingly, and fruitfully, applied the tools of rational choice to lawmaking in the EU, illuminating how the Union's ever-changing institutions interact to channel actors' preferences into policy. The spatial bargaining models that are at the heart of this endeavor are likely to remain central to our understanding as the Union continues to evolve. Nonetheless, if this line of research is to remain relevant, future work in this vein must relax the overly simplistic assumptions that have underpinned the majority of our models of European politics. Crucially, we must reduce our reliance on the assumption of complete information and expand the small set of studies that explore how asymmetries in knowledge affect EU politics. Just as importantly, we must continue to draw links between what we know about interactions amongst European institutions and what we have learned about how politicians operate within these same structures. Only by linking the two levels of analysis can we truly understand how European politicians coordinate compromise.

Appendix A

Details of the Formal Model

Lemma 1. *In any perfect Bayesian equilibrium of the sequential codecision bargaining game*

$$x_{P,1} - x_{P,2} \geq \frac{x_{P,1}}{2}.$$

Proof. Say, instead that

$$\frac{x_{P,1}}{2} > x_{P,1} - x_{P,2}, \tag{A.1}$$

implying, from equation 2.9 that

$$\omega_2 = \frac{x_{P,1}}{2}. \tag{A.2}$$

First, consider the case where $\omega_2 > k_P$. Then, plugging the appropriate case from equation 2.8 into equation A.1, we get

$$\frac{x_{P,1}}{2} > x_{P,1} - \left(\frac{\omega_2 - k_P}{2}\right) \implies \omega_2 > x_{P,1} + k_P.$$

Now, plugging equation A.2 into the above result yields

$$\frac{x_{P,1}}{2} > x_{P,1} + k_P \implies k_P < \frac{-x_{P,1}}{2}.$$

But $\frac{-x_{P,1}}{2} \leq 0$ in equilibrium, contradicting the assumption that $k_P \geq 0$.

Now, look at the case where $\omega_2 \leq k_P$. Inserting the relevant case from equation 2.8 into A.1 yields

$$\frac{x_{P,1}}{2} > x_{P,1} \implies x_{P,1} < 0.$$

But we know that P will never table a proposal below zero in equilibrium, again creating a contradiction. □

Lemma 2. *In any separating equilibrium of the mediated codecision game,*

$$x^*_{P,2,W} = \omega^*_2$$

when the second reading is on the equilibrium path after P observes $m = W$.

Proof. When the second reading is on the separating equilibrium path when $m = W$, P must believe $k_{\underline{C}} \sim U[\omega^*_2, \omega_{2,W}]$, where $\omega^*_2 \leq \omega_{2,W} \leq \omega$, upon reaching the second reading. Thus, P's expected utility at second reading is

$$u_P(\omega_{2,W}) = x_{P,2,W} \cdot \left(1 - \frac{x_{P,2,W} - \omega^*_2}{\omega_{2,W} - \omega^*_2}\right) - k_P \cdot \frac{x_{P,2,W} - \omega^*_2}{\omega_{2,W} - \omega^*_2}.$$

Maximizing the above equation with respect to $x_{P,2,W}$ shows that P's optimal second round proposal is

$$x^*_{P,2,W} = \begin{cases} \frac{\omega_{2,W} - k_P}{2} & \text{if} \quad \frac{\omega_{2,W} - k_P}{2} > \omega^*_2, \quad \text{and} \\ \omega^*_2 & \text{otherwise.} \end{cases}$$

Furthermore, $\frac{\omega_{2,W} - k_P}{2} > \omega^*_2 \implies 5\omega_{2,W} > 6\omega + k_P$. This implication contradicts the fact that $\omega_{2,W} \leq \omega$ because I assume $\omega \geq 0$ and $k_P \geq 0$. Therefore $x^*_{P,2,W} = \omega^*_2$ in any separating equilibrium where the second reading is reached with positive probability after P observes $m = W$. □

Lemma 3. *The mediated codecision game never reaches the second reading in any separating equilibrium when $s = W$.*

Proof. First note that $s = W \implies m = W$ in any separating equilibrium. Now assume that there exists a separating equilibrium in which the players reach the second reading with positive probability when $s = W$. Under such circumstances the Council will always accept

P's second reading offer $x^*_{P,2,W} = \omega^*_2$ (see lemma 2) because the Council's payoff, $1 - \omega^*_2$, for accepting the offer is at least as good as what the Council can expect from conciliation $(1 - k_{\underline{C}})$.[1] Therefore, in such an equilibrium, the Council will only reject P's first reading proposal if it expects to obtain greater utility from waiting for the second proposal than from accepting the initial offer. That is, C rejects $x_{P,1,W}$ if and only if

$$k_{\underline{C}} < x_{P,1,W} - x_{P,2,W}. \tag{A.3}$$

Thus, remembering that the Council will accept $x_{P,2,W}$ should the game reach the second reading, and using lemma 2, P's optimal first round proposal, $x_{P,1,W}$, maximizes

$$x_{P,1,W} \cdot \left(1 - \frac{x_{P,1,W} - 2\omega^*_2}{\omega - \omega^*_2}\right) + (\omega^*_2 - k_P) \cdot \frac{x_{P,1,W} - 2\omega^*_2}{\omega - \omega^*_2}.$$

Maximizing the above equation with respect to $x_{P,1,W}$ shows that P's optimal first round proposal is

$$x^*_{P,1,W} = \begin{cases} \frac{11\omega - 9k_P}{10} & \text{if} \quad \omega^*_2 \leq \frac{11\omega - 9k_P}{10} \leq 1, \\ 1 & \text{if} \quad \frac{11\omega - 9k_P}{10} > 1, \quad \text{and} \\ \omega^*_2 & \text{otherwise.} \end{cases}$$

Therefore $x_{P,1,W}$ is at most $\frac{11\omega - 9k_P}{10}$ in the sort of separating equilibrium that we are considering here. Plugging this result and the result of lemma 2 into equation A.3 indicates that there is a chance that the Council will reject the Parliament's first offer in equilibrium if, and only if,

$$k_{\underline{C}} < \frac{11\omega - 9k_P}{10} - \omega^*_2 \implies k_{\underline{C}} < \frac{\omega - k_P}{2} \implies \omega + k_P < 0.$$

But both ω and k_P are greater than or equal to zero by assumption, implying that the Council will always accept the Parliament's first offer. Therefore, the mediated codecision game has no separating equilibria in which the players adopt strategies that would allow

[1] Remember that $k_{\underline{C}} \geq \omega^*_2$ when $s = W$.

them to reach the second reading with positive probability when $s = W$. $\square$

Appendix B

Statistical Estimation

This appendix contains descriptions of the Gibbs sampling algorithms I used to fit the models described in section 3.2 and various other estimation details. Gibbs samplers work by successively drawing values from the posterior distributions of a subset of model parameters, conditional on current draws for the rest of the parameters in the model.[1] The analyst allows the sampler to iterate for $t \in 1 \ldots T$ iterations, generating T draws from the conditional posterior distributions of the model parameters. These draws can then be used to summarize the joint posterior distribution of the given model. Note that the equations below assume that $\mathbf{x}$ is a $D \times n$ matrix with i^{th} column $\mathbf{x_i}$, $\boldsymbol{\beta}$ is an $m \times D$ matrix with j^{th} row $\boldsymbol{\beta_j}$, $\boldsymbol{\kappa}$ is a $D \times m$ matrix with j^{th} column $\boldsymbol{\kappa_j}$, $\mathbf{C}$ is an $k \times n$ matrix with i^{th} column $\mathbf{c_i}$, $\boldsymbol{\Lambda}$ is a $D \times k$ matrix, and $\boldsymbol{\Sigma_x}$ is a $D \times D$ matrix.

B.1 Gibbs Sampler: Canonical Model with Hierarchical Priors

The Gibbs sampler for the canonical model takes advantage of a data augmentation approach, first sampling latent utilities, $\hat{Y}_{ij} = U_i(\boldsymbol{\zeta_j}) - U_i(\boldsymbol{\psi_j})$, from the truncated normal distribution

$$\hat{Y}_{ij} \mid \boldsymbol{\theta}_{\hat{\mathbf{Y}}_{\mathbf{ij}}} \sim \begin{cases} \mathcal{TN}_{(-\infty,0)}(\mu_{\hat{Y}_{ij}}, 1) & \text{if} \quad Y_{ij} = 0 \\ \mathcal{TN}_{[0,\infty)}(\mu_{\hat{Y}_{ij}}, 1) & \text{if} \quad Y_{ij} = 1 \end{cases} \tag{B.1}$$

[1] See Gelman, Carlin, Stern & Rubin (2004) or Gill (2002) for an introduction to Gibbs sampling.

where $\boldsymbol{\theta}_{\hat{\mathbf{Y}}_{\mathbf{ij}}} = \{Y_{ij}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}\}$ and $\mu_{\hat{Y}_{ij}} = \boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}})$. Next, the algorithm samples legislator ideal points from the conditional posterior distribution

$$\mathbf{x_i} \mid \boldsymbol{\theta}_{\mathbf{x_i}} \sim \mathcal{N}_D \left(\left[\boldsymbol{\beta}'\boldsymbol{\beta} + \boldsymbol{\Sigma}_{\mathbf{x}}^{-1} \right]^{-1} \left[\boldsymbol{\beta}'\mathbf{w}_{\mathbf{i}}^{\mathbf{x}\prime} + \boldsymbol{\Sigma}_{\mathbf{x}}^{-1}\boldsymbol{\Lambda}\mathbf{c_i} \right], \left[\boldsymbol{\beta}'\boldsymbol{\beta} + \boldsymbol{\Sigma}_{\mathbf{x}}^{-1} \right]^{-1} \right) \tag{B.2}$$

where $\boldsymbol{\theta}_{\mathbf{x_i}} = \{\mathbf{C}, \hat{\mathbf{Y}}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\Lambda}, \boldsymbol{\Sigma}_{\mathbf{x}}\}$ and $w_{ij}^{\mathbf{x}} = \hat{Y}_{ij} + \boldsymbol{\beta}_{\mathbf{j}}\boldsymbol{\kappa}_{\mathbf{j}}$, and normalizes each $\mathbf{x_i}$ to have zero mean and unit variance on each dimension. The sampler subsequently draws bill cutpoints from the conditional posterior distribution

$$\boldsymbol{\kappa}_{\mathbf{j}} \mid \boldsymbol{\theta}_{\boldsymbol{\kappa}_{\mathbf{j}}} \sim \mathcal{N}_D \left(\left[\mathbf{B_j}'\mathbf{B_j} + \sigma_{\boldsymbol{\kappa}_{\mathbf{j}}}^{-2}\mathbf{I}_D \right]^{-1} \left[\mathbf{B_j}'\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\kappa}} + \sigma_{\boldsymbol{\kappa}_{\mathbf{j}}}^{-2}\mathbf{I}_D\boldsymbol{\mu}_{\mathbf{k_j}} \right], \left[\mathbf{B_j}'\mathbf{B_j} + \sigma_{\boldsymbol{\kappa}_{\mathbf{j}}}^{-2}\mathbf{I}_D \right]^{-1} \right) \tag{B.3}$$

where $\boldsymbol{\theta}_{\boldsymbol{\kappa}_{\mathbf{j}}} = \{\hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\mu}_{\boldsymbol{\kappa}_{\mathbf{j}}}, \sigma_{\boldsymbol{\kappa}_{\mathbf{j}}}\}$, $\mathbf{B_j}$ is the $n \times D$ matrix formed by stacking $\boldsymbol{\beta}_{\mathbf{j}}$ n times, and $\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\kappa}} = \left[\boldsymbol{\beta}_{\mathbf{j}}\mathbf{x_1} - \hat{Y}_{1j},\ \boldsymbol{\beta}_{\mathbf{j}}\mathbf{x_2} - \hat{Y}_{2j},\ \ldots,\ \boldsymbol{\beta}_{\mathbf{j}}\mathbf{x_n} - \hat{Y}_{nj} \right]'$. After drawing cutpoints, the algorithm samples each bill discrimination parameter according to

$$\boldsymbol{\beta}_{\mathbf{j}} \mid \boldsymbol{\theta}_{\boldsymbol{\beta}_{\mathbf{j}}} \sim \mathcal{N}_D \left(\left[\mathbf{A}\mathbf{A}' + \sigma_{\boldsymbol{\beta}_{\mathbf{j}}}^{-2}\mathbf{I}_D \right]^{-1} \left[\mathbf{A}\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\beta}} + \sigma_{\boldsymbol{\beta}_{\mathbf{j}}}^{-2}\mathbf{I}_D\boldsymbol{\mu}_{\boldsymbol{\beta}_{\mathbf{j}}} \right], \left[\mathbf{A}\mathbf{A}' + \sigma_{\boldsymbol{\beta}_{\mathbf{j}}}^{-2}\mathbf{I}_D \right]^{-1} \right) \tag{B.4}$$

where $\boldsymbol{\theta}_{\boldsymbol{\beta}_{\mathbf{j}}} = \{\hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\kappa}, \boldsymbol{\mu}_{\boldsymbol{\beta}_{\mathbf{j}}}, \sigma_{\boldsymbol{\beta}_{\mathbf{j}}}\}$, $\mathbf{A}$ is a $D \times n$ matrix with typical column $\mathbf{A_i} = \mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}$, and, $\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\beta}} = \hat{\mathbf{Y}}_{\mathbf{j}}$. The sampler then moves on to the hierarchical parameters, drawing the $\boldsymbol{\Lambda}$ coefficient matrix from the conditional posterior distribution

$$\mathrm{vec}(\boldsymbol{\Lambda}) \mid \boldsymbol{\theta}_{\boldsymbol{\Lambda}} \sim \mathcal{N}_D \left(\mathrm{vec}\left(\left[\mathbf{C}'\mathbf{C} + \sigma_{\boldsymbol{\Lambda}}^{-2}\mathbf{I}_k \right]^{-1} \mathbf{C}'\mathbf{x}' \right), \boldsymbol{\Sigma}_{\mathbf{x}} \otimes \left[\mathbf{C}'\mathbf{C} + \sigma_{\boldsymbol{\Lambda}}^{-2}\mathbf{I}_k \right]^{-1} \right) \tag{B.5}$$

where $\boldsymbol{\theta}_{\boldsymbol{\Lambda}} = \{\mathbf{C}, \mathbf{x}, \boldsymbol{\Sigma}_{\mathbf{x}}, \sigma_{\boldsymbol{\Lambda}}\}$. Finally, the algorithm samples the hierarchical variance-covariance matrix, $\boldsymbol{\Sigma}_{\mathbf{x}}$, from

$$\boldsymbol{\Sigma}_{\mathbf{x}} \mid \boldsymbol{\theta}_{\boldsymbol{\Sigma}_{\mathbf{x}}} \sim \mathcal{W}^{-1} \left(n + \upsilon, \left[\mathbf{x}' - \mathbf{C}\boldsymbol{\Lambda} \right]' \left[\mathbf{x}' - \mathbf{C}\boldsymbol{\Lambda} \right] + \sigma_x^2\mathbf{I}_D \right) \tag{B.6}$$

where $\boldsymbol{\theta}_{\boldsymbol{\Sigma}_{\mathbf{x}}} = \{\mathbf{C}, n, \upsilon, \mathbf{x}, \boldsymbol{\Lambda}, \sigma_{\mathbf{x}}^2\}$ and $\mathcal{W}^{-1}(\cdot)$ is the inverse Wishart density function.

B.2 Gibbs Sampler: Baseline Influence Model

The Gibbs sampler for the baseline influence model is largely analogous to the sampler for the canonical model. Throughout this section, $\mathbf{s} = \begin{bmatrix} \mathbf{a} & \mathbf{r} \end{bmatrix}$, such that $\mathbf{s_j} = \begin{bmatrix} a_j & r_j \end{bmatrix}$, and $\boldsymbol{\delta_i} = \begin{bmatrix} \delta_i^a \\ -\delta_i^r \end{bmatrix}$. Again, the sampler starts by drawing latent utilities from the conditional posterior distribution described by equation B.1, but where $\boldsymbol{\theta}_{\hat{\mathbf{Y}}_{\mathbf{ij}}} = \{\hat{Y}_{ij}, \mathbf{s}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}\}$ and $\mu_{\hat{Y}_{ij}} = \boldsymbol{\beta_j}(\mathbf{x_i} - \boldsymbol{\kappa_j}) + \mathbf{s_j}\boldsymbol{\delta_i}$. . Next, the algorithm samples legislator ideal points from equation B.2 except, now, $\boldsymbol{\theta}_{\mathbf{x_i}} = \{\mathbf{s}, \mathbf{C}, \hat{\mathbf{Y}}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\Lambda}, \boldsymbol{\Sigma}_{\mathbf{x}}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}\}$ and $w_{ij}^{\mathbf{x}} = \hat{Y}_{ij} + \boldsymbol{\beta_j}\boldsymbol{\kappa_j} - \mathbf{s_j}\boldsymbol{\delta_i}$. Again, the sampler normalizes each ideal point to have zero mean and unit variance on each dimension. The sampler next draws bill cutpoints according to equation B.3 where, now, $\boldsymbol{\theta}_{\boldsymbol{\kappa}_{\mathbf{j}}} = \{\mathbf{s}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\mu}_{\boldsymbol{\kappa}_{\mathbf{j}}}, \sigma_{\boldsymbol{\kappa}_{\mathbf{j}}}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}\}$ and $\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\kappa}} = [\boldsymbol{\beta_j}\mathbf{x_1} - \hat{Y}_{1j} + \mathbf{s_j}\boldsymbol{\delta_1},\ \boldsymbol{\beta_j}\mathbf{x_2} - \hat{Y}_{2j} + \mathbf{s_j}\boldsymbol{\delta_2},\ \ldots,\ \boldsymbol{\beta_j}\mathbf{x_n} - \hat{Y}_{nj} + \mathbf{s_j}\boldsymbol{\delta_n}]'$. The algorithm subsequently draws bill discrimination parameters using equation B.4, but now $\boldsymbol{\theta}_{\boldsymbol{\beta}_{\mathbf{j}}} = \{\mathbf{s}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\kappa}, \boldsymbol{\mu}_{\boldsymbol{\beta}_{\mathbf{j}}}, \sigma_{\boldsymbol{\beta}_{\mathbf{j}}}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}\}$ and $\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\beta}} = \left[\hat{Y}_{1j} - \mathbf{s_j}\boldsymbol{\delta_1},\ \hat{Y}_{2j} - \mathbf{s_j}\boldsymbol{\delta_2},\ \ldots,\ \hat{Y}_{nj} - \mathbf{s_j}\boldsymbol{\delta_n}\right]'$. The sampler must now draw the MEP fixed effects from their conditional posterior distribution

$$\boldsymbol{\delta_i} \mid \boldsymbol{\theta}_{\boldsymbol{\delta_i}} \sim \mathcal{N}\left(\begin{bmatrix} \frac{\phi_{\delta_i^a}}{\sum_{j=1}^m a_j + \sigma_{\delta_i^a}^{-2}} \\ \frac{\phi_{\delta_i^r}}{\sum_{j=1}^m r_j + \sigma_{\delta_i^r}^{-2}} \end{bmatrix}, \begin{bmatrix} \frac{1}{\sum_{j=1}^m a_j + \sigma_{\delta_i^a}^{-2}} & 0 \\ 0 & \frac{1}{\sum_{j=1}^m r_j + \sigma_{\delta_i^r}^{-2}} \end{bmatrix}\right) \tag{B.7}$$

where $\phi_{\delta_i^a} = \sum_{j=1}^m \left(a_j \left[\hat{Y}_{ij} - \boldsymbol{\beta_j}(\mathbf{x_i} - \boldsymbol{\kappa_j})\right]\right) + \mu_{\delta_i^a}\sigma_{\delta_i^a}^{-2}$, $\phi_{\delta_i^r} = \sum_{j=1}^m \left(r_j \left[\hat{Y}_{ij} - \boldsymbol{\beta_j}(\mathbf{x_i} - \boldsymbol{\kappa_j})\right]\right) + \mu_{\delta_i^r}\sigma_{\delta_i^r}^{-2}$, and, $\boldsymbol{\theta}_{\boldsymbol{\delta_i}}{}^{(t)} = \{\mathbf{s}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \mu_{\delta_i^a}, \mu_{\delta_i^r}, \sigma_{\delta_i^a}, \sigma_{\delta_i^r}\}$. Finally, the algorithm updates the hierarchical ideal point parameters exactly as in the baseline model, using equations B.5 and B.6.

B.3 Gibbs Sampler: Conditional Influence Model

First note that $\mathbf{Z}$ is an $n \cdot m \times l$ matrix such that each row $\mathbf{z_{ij}}$ of $\mathbf{Z}$ represents a vector of l covariates for MEP i's vote on question j. Again, the sampler for the conditional influence model closely resembles the samplers for the canonical and baseline influence models. It starts by drawing latent utilities using equation B.1, but now $\boldsymbol{\theta}_{\hat{\mathbf{Y}}_{\mathbf{ij}}} = \{\hat{Y}_{ij}, \mathbf{Z}, \mathbf{a}, \mathbf{r}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}, \boldsymbol{\gamma}^{\mathbf{a}}, \boldsymbol{\gamma}^{\mathbf{r}}\}$ and $\mu_{\hat{Y}_{ij}} = \boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}) + a_j(\delta_i^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{ij}}) - r_j(\delta_i^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{ij}})$. Next, it draws ideal points from their conditional posterior distributions according to equation B.2 where $\boldsymbol{\theta}_{\mathbf{x_i}} = \{\mathbf{Z}, \mathbf{a}, \mathbf{r}, \mathbf{C}, \hat{\mathbf{Y}}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\Lambda}, \boldsymbol{\Sigma}_{\mathbf{x}}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}, \boldsymbol{\gamma}^{\mathbf{a}}, \boldsymbol{\gamma}^{\mathbf{r}}\}$ and $w_{ij}^{\mathbf{x}} = \hat{Y}_{ij} + \boldsymbol{\beta}_{\mathbf{j}}\boldsymbol{\kappa}_{\mathbf{j}} - a_j(\delta_i^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{ij}}) + r_j(\delta_i^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{ij}})$. Again, the sampler normalizes each ideal point to have zero mean and unit variance on each dimension. The sampler next draws bill cutpoints according to equation B.3 where, now, $\boldsymbol{\theta}_{\boldsymbol{\kappa}_{\mathbf{j}}} = \{\mathbf{Z}, \mathbf{a}, \mathbf{r}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\mu}_{\boldsymbol{\kappa}_{\mathbf{j}}}, \sigma_{\boldsymbol{\kappa}_{\mathbf{j}}}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}, \boldsymbol{\gamma}^{\mathbf{a}}, \boldsymbol{\gamma}^{\mathbf{r}}\}$ and $\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\kappa}} = [\boldsymbol{\beta}_{\mathbf{j}}\mathbf{x_1} - \hat{Y}_{1j} + a_j(\delta_1^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{1j}}) - r_j(\delta_1^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{1j}}),\ \boldsymbol{\beta}_{\mathbf{j}}\mathbf{x_2} - \hat{Y}_{2j} + a_j(\delta_2^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{2j}}) - r_j(\delta_2^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{1j}}),\ \ldots,\ \boldsymbol{\beta}_{\mathbf{j}}\mathbf{x_n} - \hat{Y}_{nj} + a_j(\delta_n^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{nj}}) - r_j(\delta_n^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{1j}})]'$. The algorithm subsequently draws bill discrimination parameters using equation B.4, but now $\boldsymbol{\theta}_{\boldsymbol{\beta}_{\mathbf{j}}} = \{\mathbf{Z}, \mathbf{a}, \mathbf{r}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\kappa}, \boldsymbol{\mu}_{\boldsymbol{\beta}_{\mathbf{j}}}, \sigma_{\boldsymbol{\beta}_{\mathbf{j}}}, \boldsymbol{\delta}^{\mathbf{a}}, \boldsymbol{\delta}^{\mathbf{r}}, \boldsymbol{\gamma}^{\mathbf{a}}, \boldsymbol{\gamma}^{\mathbf{r}}\}$ and $\mathbf{w}_{\mathbf{j}}^{\boldsymbol{\beta}} = [\hat{Y}_{1j} - a_j(\delta_1^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{2j}}) + r_j(\delta_1^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{1j}}),\ \hat{Y}_{2j} - a_j(\delta_2^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{2j}}) + r_j(\delta_2^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{1j}}),\ \ldots,\ \hat{Y}_{nj} - a_j(\delta_n^a + \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{2j}}) + r_j(\delta_n^r + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{1j}})]'$. Next, it samples MEP fixed effects from equation B.7 where $\phi_{\delta_i^a} = \sum_{j=1}^m \left(a_j \left[\hat{Y}_{ij} - \boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}) - \boldsymbol{\gamma}^{\mathbf{a}}\mathbf{z_{ij}}\right]\right) + \mu_{\delta_i^a}\sigma_{\delta_i^a}^{-2}$, $\phi_{\delta_i^r} = \sum_{j=1}^m \left(r_j \left[\hat{Y}_{ij} - \boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}) + \boldsymbol{\gamma}^{\mathbf{r}}\mathbf{z_{ij}}\right]\right) + \mu_{\delta_i^r}\sigma_{\delta_i^r}^{-2}$, and, $\boldsymbol{\theta}_{\boldsymbol{\delta}_{\mathbf{i}}}{}^{(t)} = \{\mathbf{Z}, \mathbf{a}, \mathbf{r}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \mu_{\delta_i^a}, \mu_{\delta_i^r}, \sigma_{\delta_i^a}, \sigma_{\delta_i^r}\boldsymbol{\gamma}^{\mathbf{a}}, \boldsymbol{\gamma}^{\mathbf{r}}\}$. In the next step the sampler draws

$$\boldsymbol{\gamma}^{\mathbf{a}} \mid \boldsymbol{\theta}_{\boldsymbol{\gamma}^{\mathbf{a}}} \sim \mathcal{N}\left(\left[\mathbf{Z_a}'\mathbf{Z_a}\right]^{-1}\mathbf{Z_a}'\mathbf{w}_{\boldsymbol{\gamma}_{\mathbf{a}}}, \left[\mathbf{Z_a}'\mathbf{Z_a}\right]^{-1}\right) \tag{B.8}$$

where $\boldsymbol{\theta}_{\boldsymbol{\gamma}^{\mathbf{a}}} = \{\mathbf{Z_a}, \mathbf{a}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\delta}^{\mathbf{a}}\}$ and where $\mathbf{Z_a}$ is a matrix contain only those rows of $\mathbf{Z}$ for which $a_j = 1$. Similarly, the $n \cdot m$ column vector $\mathbf{w}_{\boldsymbol{\gamma}} = \hat{Y}_{ij} - \boldsymbol{\beta}_{\mathbf{j}}(\mathbf{x_i} - \boldsymbol{\kappa}_{\mathbf{j}}) - \mathbf{s_j}\boldsymbol{\delta}_{\mathbf{i}}$ and $\mathbf{w}_{\boldsymbol{\gamma}_{\mathbf{a}}}$ is a vector of length $\sum_{i=1}^n a_i$ containing only those elements of $\mathbf{w}_{\boldsymbol{\gamma}}$ for which $a_j = 1$. Turning

to the other set of regression parameters, the sampler draws

$$\boldsymbol{\gamma}^{\mathbf{r}} \mid \boldsymbol{\theta}_{\boldsymbol{\gamma}^{\mathbf{r}}} \sim \mathcal{N}\left(\left[\mathbf{Z_r}'\mathbf{Z_r}\right]^{-1}\mathbf{Z_r}'\mathbf{w}_{\boldsymbol{\gamma}_{\mathbf{r}}}, \left[\mathbf{Z_r}'\mathbf{Z_r}\right]^{-1}\right) \tag{B.9}$$

where $\boldsymbol{\theta}_{\boldsymbol{\gamma}^{\mathbf{r}}} = \{\mathbf{Z_r}, \mathbf{r}, \hat{\mathbf{Y}}, \mathbf{x}, \boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\delta}^{\mathbf{r}}\}$ and with $\mathbf{Z_r}$ and $\mathbf{w}_{\boldsymbol{\gamma}_{\mathbf{r}}}$ defined analogously to $\mathbf{Z_a}$ and $\mathbf{w}_{\boldsymbol{\gamma}_{\mathbf{a}}}$. Finally, the algorithm updates the hierarchical ideal point parameters exactly as in the baseline model, using equations B.5 and B.6.

B.4 Model Fitting Details

I used the Scythe Statistical Library (Pemstein, Quinn & Martin Forthcoming) to implement software to fit all three models. I ran every sampler for 60,000 iterations, discarding the first 10,000 "burn-in" iterations. Standard Markov chain Monte Carlo diagnostics and visual inspection of sampled parameter chains showed little evidence for non-convergence.

I used the prior values $\sigma^2_{\boldsymbol{\beta}_{\mathbf{j}}} = \sigma^2_{\boldsymbol{\kappa}_{\mathbf{j}}} = 25\ \forall j$, $\sigma^2_x = 1$, $\sigma^2_{\boldsymbol{\Lambda}} = 25$, and $\upsilon = 1$ when fitting the canonical model. For the baseline and conditional influence models, I also set $\mu_{\delta^a_i} = \mu_{\delta^r_i} = 0\ \forall i$ and $\sigma^2_{\delta^a_i} = \sigma^2_{\delta^r_i} = 1\ \forall i$. These priors are uninformative and vague and have little influence over parameter estimates. In the case of the ideal points, of course, the models make use of detailed prior information encoded in the hierarchical priors described in the main text. On the other hand, the vague priors for the bill parameters, fixed effects, and hierarchical ideal point coefficients (i.e. $\boldsymbol{\Lambda}$) reflect my lack of prior information about these parameters.

The canonical model is identified by the priors, by restricting one coefficient in each column of $\boldsymbol{\Lambda}$ to be greater than or equal to zero to fix the orientation of the ideal point space, and by normalizing the ideal points to have zero mean and unit variance on each dimension. The identification of the influence parameters in the extended models rests on the assumptions that MEP ideal points are constant across codecision and non-codecision votes, and that the Commission only exerts influence when it lodges an opinion (i.e. on

codecision votes).

References

Alt, James E. & Robert C. Lowry. 1994. "Divided Government, Fiscal Institutions, and Budget Deficits: Evidence from the States." *American Political Science Review* 88(4):811–828.

Arnold, Douglas R. 1990. *The Logic of Congressional Action*. New Haven: Yale University Press.

Austen-Smith, David & Jeffrey S. Banks. 2000. "Cheap Talk and Burned Money." *Journal of Economic Theory* 91(1):1–16.

Ausubel, Lawrence M., Peter Cramton & Raymond J. Deneckere. 2001. Bargaining with Incomplete Information. In *Handbook of Game Theory, Vol. 3*, ed. Robert J. Aumann & Sergiu Hart. Amsterdam: Elsevier Science B.V.

Bafumi, Joseph, Andrew Gelman, David K. Park & Noah Kaplan. 2004. "Practical Issues in Implementing and Understanding Bayesian Ideal Point Estimation." *Political Analysis* 13(2):171–187.

Berger, Adam L., Stephen A. Della Pietra & Vincent J. Della Pietra. 1996. "A Maximum Entropy Approach to Natural Language Processing." *Computational Linguistics* 22(1):39–71.

Boekhoorn, Vincent, Adrian M.A. Van Deemen & Madeleine O. Hosli. 2006. A Cooperative Approach to Decision-Making in the European Union. In *The European Union Decides*, ed. Robert Thomson, Frans N. Stockman, Christopher H. Achen & Thomas König. Cambridge: Cambridge University Press pp. 178–210.

Boser, Bernhardt E., Isabelle Guyon & Vladimir Vapnik. 1992. A Training Algorithm for Optimal Margin Classifiers. In *Fifth Annual Workshop on Computational Learning Theory*. pp. 144–152.

Brehm, John & Scott Gates. 1997. *Working, Shirking, and Sabotage: Bureaucratic Response to a Democratic Public*. Ann Arbor: The University of Michigan Press.

Breiman, Leo. 2001. "Random Forests." *Machine Learning* 45(1):5–32.

Bronars, Stephen G. & Jr. John R. Lott. 1997. "Do Campaign Donations Alter How a Politician Votes? Or, Do Donors Support Candidates Who Value the Same Things That They Do?" *Journal of Law & Economics* 40:317–350.

Burges, Christopher J.C. 1998. "A Tutorial on Support Vector Machines for Pattern Recognition." *Data Mining and Knowledge Discovery* 2:121–167.

Burns, Charlotte. 2004. "Codecision and the European Commission: A Study of Declining Influence?" *Journal of European Public Policy* 11(1):1–18.

Cameron, Charles M. 2000. *Veto Bargaining: Presidents and the Politics of Negative Power*. Cambridge: Cambridge University Press.

Cameron, Charles M. & Susan Elmes. 1994. "Sequential Veto Bargaining." `http://www.princeton.edu/~ccameron/VetoBargain.pdf`.

Carey, John M. 2008. *Legislative Voting and Accountability*. New York: Cambridge University Press.

Carrubba, Clifford J., Matthew Gabel, Lacey Murrah, Ryan Clough, Elizabeth Montgomery & Rebecca Schambach. 2006. "Off the Record: Unrecorded Legislative Votes, Selection Bias and Roll-Call Vote Analysis." *British Journal of Political Science* 36:691–704.

Carrubba, Clifford & Matthew Gabel. 1999. "Roll-Call Votes and Party Discipline in the European Parliament: Reconsidering MEP Voting Behavior." European Parliament Research Group Working Paper.

Carrubba, Clifford, Matthew Gabel & Simon Hug. 2008. "Legislative Voting Behavior, Seen and Unseen: A Theory of Roll-Call Vote Selection." *Legislative Studies Quarterly* 33:543–572.

Carrubba, Clifford, Matthew Gabel & Simon Hug. 2009. "Voting at the Surface: Roll Call Votes in the European Parliament." Paper prepared for presentation at the Colloque "Le parlement europen: un parlement comme les autres?" IEE-ULB, April 29, 2009.

Cini, Michelle. 1996. *The European Commission: Leadership, Organization, and Culture in the EU Administration*. Manchester: Manchester University Press.

Clinton, Joshua D. & John Lapinski. 2008. "Laws and Roll Calls in the U.S. Congress, 1891–1994." *Legislative Studies Quarterly* 33(4):511–541.

Clinton, Joshua, Simon Jackman & Douglas Rivers. 2004. "The Statistical Analysis of Roll Call Data." *American Political Science Review* 98:355–370.

Corbett, Richard, Francis Jacobs & Michael Schackleton. 2003. *The European Parliament*. 5th ed. London: Harper.

Cortes, Corinna & Vladimir Vapnik. 1995. "Support Vector Networks." *Machine Learning* 20:273–297.

Cox, Gary W. & Mathew D. McCubbins. 1993. *Legislative Leviathan: Party Government in the House*. Berkeley: University of California Press.

Crawford, Vincent P. & Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50(6):1431–1451.

Crombez, Cristophe. 1997. "The Co-Decision Procedure in the European Union." *Legislative Studies Quarterly* 22(1):97–119.

Crombez, Cristophe. 2001. The Treaty of Amsterdam and the Co-decision Procedure. In *The Rules of Integration: Institutionalist Approaches to the Study of Europe*, ed. Gerald Schneider & Mark Aspinwall. Manchester: Manchester University Press pp. 101–122.

Crombez, Cristophe, Bernard Steunenberg & Richard Corbett. 2000. "Understanding the EU Legislative Process: Political Scientists' and Practitioners' Perspectives." *European Union Politics* 1(3):363–381.

Curtin, Deirdre. 1993. "The Constitutional Structure of the Union: A Europe of Bits and Pieces." *Common Market Law Review* 30:17–69.

Desposato, Scott W. 2006. "Parties for Rent? Ambition, Ideology, and Party Switching in Brazil's Chamber of Deputies." *American Journal of Political Science* 50(1):62–80.

Dimitriadou, Evgenia, Kurt Hornik, Friedrich Leisch, David Meyer & Andreas Weingessel. 2009. "e1071: Misc Functions of the Department of Statistics (e1071), TU Wien v1.5-19." `http://cran.r-project.org/web/packages/e1071/index.html` (last accessed March 5, 2009).

Epstein, David & Sharyn O'Halloran. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers*. Cambridge: Cambridge University Press.

European Commission. 2009. "PreLex: Monitoring of the Decision-Making Process Between Institutions." `http://ec.europa.eu/prelex/apcnet.cfm?CL=en` (last accessed December 17, 2009).

European Commission. 2010. "Ordinary Legislative Procedure "Step by Step"." `http://ec.europa.eu/codecision/stepbystep/text/index_en.htm` (last accessed March 18, 2010).

European Parliament. 1992. *Maastricht, the Treaty on the European Union, Position of the European Parliament*. Luxembourg: Office for Official Publications of the European Communities.

European Parliament. 2009*a*. "The Archives of the Activities: Parliamentary Business." `http://www.europarl.europa.eu/activities/archives.do?language=EN` (last accessed April 20, 2009).

European Parliament. 2009*b*. "Legislative Observatory of the European Parliament." `http://www.europarl.europa.eu/oeil/index.jsp` (last accessed April 20, 2009).

European Parliament. 2009*c*. "Your MEPs." `http://www.europarl.europa.eu/members/archive.do` (last accessed April 20, 2009).

Faas, Thorsten. 2003. "To Defect or not to Defect? National, Institutional and Party Group Pressures on MEPs and their Consequences for Party Group Cohesion in the European Parliament." *European Journal of Political Research* 42(841–866).

Fader, Anthony, Gary King, Daniel Pemstein & Kevin Quinn. 2009. "RSNL: R Statistics for Natural Language v0.10." http://gking.harvard.edu/RSNL.

Farrell, Henry & Adrienne Hèritier. 2004. "Interorganizational Negotiation and Intraorganizational Power in Shared Decision Making: Early Agreements Under Codecision and Their Impact on the European Parliament and Council." *Comparative Political Studies* 37(10):1184–1212.

Fennel, Lee. 1974. "Reasons for Roll Calls: An Exploratory Analysis with Argentine Data." *American Journal of Political Science* 18:395–403.

Fudenberg, Drew & Jean Tirole. 1983. "Sequential Bargaining with Incomplete Information." *Review of Economic Studies* 50:221–247.

Fukumoto, Kentaro. 2008. "Why Do Bicameral Chambers Usually, but Not Always, Agree? An Incomplete Information Game Model." Presented at the 66th Annual Meeting of the Midwest Political Science Association.

Galloway, David. 2001. *The Treaty of Nice and Beyond*. Sheffield: Sheffield Academic Press.

Garrett, Geoffrey. 1995. "From the Luxembourg Compromise to Codecision: Decision Making in the European Union." *Electoral Studies* 14(3):289–308.

Garrett, Geoffrey & George Tsebelis. 1996. "An Institutional Critique of Intergovernmentalism." *International Organization* 50:269–299.

Gelman, Andrew, John B. Carlin, Hal S. Stern & Donald B. Rubin. 2004. *Bayesian Data Analysis*. 2nd ed. Boca Raton: Chapman & Hall.

Gibbons, Robert. 1992. *Game Theory for Applied Economists*. Princeton, New Jersey: Princeton University Press.

Gill, Jeff. 2002. *Bayesian Methods: A Social and Behavioral Sciences Approach*. Boca Raton: Chapman & Hall.

Golub, Jonathan. 2002. Institutional Reform and Decision-Making in the European Union. In *Institutional Challenges in the European Union*, ed. Madeleine O. Hosli, Adrian M.A. van Deemen & Mika Widgrén. New York: Routledge pp. 134–154.

Golub, Jonathan. 2007. "Survival Analysis and European Union Decision-Making." *European Union Politics* 8(4):555–566.

Greenstein, Fred & Elton Jackson. 1963. "A Second Look at the Validity of Roll-Call Analysis." *Midwest Journal of Political Science* 7:156–166.

Grossman, Sanford J. & Motty Perry. 1986. "Sequential Bargaining Under Asymmetric Information." *Journal of Economic Theory* 39:120–54.

Haas, Ernst B. 1958. *The Uniting of Europe*. Stanford: Stanford University Press.

Hagemann, Sarah & Bjørn Høyland. 2010. "Bicameral Politics in the European Union." *Journal of Common Market Studies* 48(4).

Hall, Richard L. 1992. "Measuring Legislative Influence." *Legislative Studies Quarterly* 17(2):205–231.

Hammond, Thomas H. & Gary J. Miller. 1987. "The Core of the Constitution." *American Political Science Review* 81(4):1155–1174.

Hausemer, Pierre. 2006. "Participation and Political Competition in Committee Report Allocation: Under What Conditions do MEPs Represent Their Constituents?" *European Union Politics* 7(4):505–530.

Hawkins, Darren G., David A. Lake, Daniel L. Nielson & Michael J. Tierney, eds. 2006. *Delegation and Agency in International Organizations*. Cambridge: Cambridge University Press.

Heller, William B. 1997. "Bicameralism and Budget Deficits: The Effects of Parliamentary Structure on Government Spending." *Legislative Studies Quarterly* 22:485–516.

Heller, William B. 2007. "Divided Politics: Bicameralism, Parties, and Policy in Democratic Legislatures." *Annual Review of Political Science* 10:245–269.

Hix, Simon. 2002. "Parliamentary Behavior with Two Principals: Preferences, Parties, and Voting in the European Parliament." *American Journal of Political Science* 46(3):688–698.

Hix, Simon. 2004. "Electoral Institutions and Legislative Behavior: Explaining Voting Defection in the European Parliament." *World Politics* 56(2):194–223.

Hix, Simon. 2005. *The Political Systemo of the European Union*. New York: Palgrave Macmillan.

Hix, Simon, Abdul Noury & Gerard Roland. 2006. "Dimensions of Politics in the European Parliament." *American Journal of Political Science* 50(2):494–511.

Hix, Simon, Abdul Noury & Gerard Roland. 2007. *Democratic Politics in the European Parliament*. New York: Cambridge University Press.

Hoffman, Stanley. 1966. "Obstinate or Obsolete? The Fate of the Nation-State in Western Europe Today." *Daedalus* 95(4):862–915.

Hooghe, Lisbet, Ryan Bakker, Anna Brigevich, Catherine de Vries, Erica Edwards, Gary Marks, Jan Rovny & Marco Steenbergen. 2008. "Reliability and Validity of Measuring Party Positions: The Chapel Hill Expert Surveys of 2002 and 2006." Unpublished ms.

Hörl, Björn, Andreas Warntjen & Arndt Wonka. 2005. "Built on Quicksand? A Decade of Procedural Spatial Models on EU Legislative Decision-Making." *Journal of European Public Policy* 12(3):592–606.

Høyland, Bjørn. 2006. "Allocation of Codecision Reports in the Fifth European Parliament." *European Union Politics* 7(1):30–50.

Høyland, Bjørn. Unpublished. "Government-Opposition in Bicameral Negotiations: Decision-Making in the European Union's Codecision Procedure." EPRG Working Paper, No. 12.

Huber, John D. & Charles R. Shipan. 2002. *Deliberate Discretion: The Institutional Foundations of Bureaucratic Autonomy*. Cambridge: Cambridge University Press.

Hug, Simon. 2010. "Selection Effects in Roll Call Votes.".

Jackman, Simon. 2001. "Multidimensional Analysis of Roll Call Data vi Bayesian Simulation: Identification, Estimation, Inference and Model Checking." *Political Analysis* 9:227–41.

Jenkins, Jeffery & Charles Stewart. 2003. "Out in the Open: The Emergence of Viva Voce Voting in the House Speakership Elections." *Legislative Studies Quarterly* 28:481–508.

Joachims, Thorsten. 1998. Text Categorization with Support Vector Machines: Learning with Many Relevant Features. In *Proceedings of the European Conference on Machine Learning*.

Kaeding, Michael. 2004. "Rapporteurship Allocation in the European Parliament: Information or Distribution?" *European Union Politics* 5(3):355–373.

Kaeding, Michael. 2005. "The World of Committee Reports: Rapporteurship Assignment in the European Parliament." *Journal of Legislative Studies* 11(1):82–104.

Kasack, Christiane. 2004. "The Legislative Impact of the European Parliament under the Revised Co-Decision Procedure." *European Union Politics* 5(2):241–260.

Keohane, Robert. 1984. *After Hegemony: Cooperation and Discord in the World Political Economy*. Princeton, New Jersey: Princeton University Press.

König, Thomas. 2007. "Convergence or Divergence? From Ever-Growing to Ever-Slowing European Decision-making Process." *European Journal of Political Research* 46:417–444.

König, Thomas. 2008. "Analysing the Process of EU Legislative Decision-Making: To Make a Long Story Short...." *European Union Politics* 9(1):145–165.

König, Thomas & Mirja Pòter. 2001. "Examining the EU Legislative Process: The Relative Importance of Agenda and Veto Power." *European Union Politics* 2(3):329–351.

Koremenos, Barbara, Charles Lipson & Duncan Snidal. 2001. "The Rational Design of International Institutions." *International Organization* 55(4):761–800.

Krehbiel, Keith. 1993. "Where's the Party?" *British Journal of Political Science* 23:235–266.

Krehbiel, Keith. 1998. *Pivotal Politics: A Theory of U.S. Lawmaking*. Chicago: University of Chicago Press.

Kreppel, Amie. 2002. *The European Parliament and the Supranational Party System*. Cambridge: Cambridge University Press.

Kydd, Andrew. 2003. "Which Side Are You On? Bias, Credibility, and Mediation." *American Journal of Political Science* 47(4):597–611.

Lewis, David D. 1998. Maive (Bayes) at Forty: The Independence Assumption in Retrieval. In *Proceedings of the European Conference on Machine Learning*. pp. 4–15.

Lindberg, Leon N. 1963. *The Political Dynamics of Economic Integration*. Oxford: Oxford University Press.

Londregan, John B. 2000. *Legislative Institutions and Ideology in Chile*. New York: Cambridge University Press.

Mamadouh, Virginie & Tapio Raunio. 2003. "The Committee System: Power, Appointments, and Report Allocation." *Journal of Common Market Studies* 41(2):333–351.

Manning, Christopher D. & Hinrich Schütze. 1999. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: The MIT Press.

Meserve, Stephen A., Daniel Pemstein & William T. Bernhard. 2009. "Political Ambition and Legislative Behavior in the European Parliament." *Journal of Politics* 71(2):1015–1032.

Mitchell, Tom M. 1997. *Machine Learning*. New York: McGraw Hill.

Moravcsik, Andrew. 1993. "Preferences and Power in the European Community: A Liberal Intergovernmentalist Approach." *Journal of Common Market Studies* 31(4):473–524.

Moravcsik, Andrew. 1998. *The Choice for Europe*. Ithaca: Cornell University Press.

Moravcsik, Andrew. 1999. "A New Statecraft? Supranational Entrepreneurs and International Cooperation." *International Organization* 53(2):267–306.

Muthoo, Abhinay. 1999. *Bargaining Theory with Applications*. Cambridge: Cambridge University Press.

Napel, Stefan & Mika Widgrén. 2006. "The Inter-Institutional Distribution of Power in EU Codecision." *Social Choice and Welfare* 27:129–154.

Nash, John F. 1950. "The Bargaining Problem." *Econometrica* 18:155–162.

Nugent, Neil. 2001. *The European Commission*. New York: Palgrave.

Pang, Bo & Lillian Lee. 2008. "Opinion Mining and Sentiment Analysis." *Foundations and Trends in Information Retrieval* 2(1–2):1–135.

Pang, Bo, Lillian Lee & Shivakumar Vaithyanathan. 2002. Thumbs Up? Sentiment Classification Using Machine Learning Techniques. In *Proceedings of EMNLP*. pp. 79–86.

Pemstein, Daniel, Kevin M. Quinn & Andrew D. Martin. Forthcoming. "The Scythe Statistical Library: An Open Source C++ Library for Statistical Computation." *Journal of Statistical Software* .

Poole, Keith T. 2005. *Spatial Models of Parliamentary Voting*. Cambridge: Cambridge University Press.

Poole, Keith T. & Howard Rosenthal. 1985. "A Spatial Model for Legislative Roll-Call Analyis." *American Journal of Political Science* 29(2):357–384.

Rasmussen, Anne. 2003. "The Role of the European Commission in Co-decision – A Strategic Facilitator Operating in a Situation of Structural Disadvantage." *European Integration online Papers (ELoP)* 7(10):`http://eiop.or.at/eiop/texte/2003--010a.htm`.

Rauchhaus, Robert W. 2006. "Asymmetric Information, Mediation, and Conflict Management." *World Politics* 58:207–241.

Ringe, Nils. 2010. *Who Decides, and How? Preferences, Uncertainty, and Policy Choice in the European Parliament*. Oxford: Oxford University Press.

Rogers, James R. 2001. "An Informational Rationale for Congruent Bicameralism." *Journal of Theoretical Politic* 13(2):123–151.

Rohde, David. 1991. *Parties and Leaders in the Postreform House*. Chicago: University of Chicago Press.

Rubenstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50:97–109.

Saalfeld, Thomas. 1995. On Dogs and Whips: Recorded Votes. In *Parliaments and Majority Rule in Western Europe*, ed. Herbert D oring. New York: St. Martin's Press pp. 528–565.

Schackleton, Michael. 2000. "The Politics of Codecision." *Journal of Common Market Studies* 38(2):325–342.

Schulz, Heiner & Thomas König. 2000. "Institutional Reform and Decision-Making. Efficiency in the European Union." *American Journal of Political Science* 44:653–666.

Selck, Torsten J. & Bernard Steunenberg. 2004. "Between Power and Luck: The European Parliament in the EU Legislative Process." *European Union Politics* 5(1):25–46.

Sobel, J & I Takahashi. 1983. "A Multi-Stage Model of Bargaining." *Review of Economic Studies* 50:411–426.

Steunenberg, Bernard. 1994. "Decision Making Under Different Institutional Arrangements: Legislation by the European Community." *Journal of Institutional and Theoretical Economics* 150:642–669.

Steunenberg, Bernard. 2000. Constitutional Change in the European Union: Parliament's Impact on the Reform of the Codecision Procedure. In *Government Institutions: Effects, Changes, and Normative Foundations*, ed. H Wahenaar. Dordrecht: Kluwer pp. 89–109.

Steunenberg, Bernard & Antoaneta Dimitrova. 1999. "Interests, legitimacy and constitutional choice: the extension of the codecision procedure in Amsterdam." University of Twente, Unpublished Manuscript.

Steunenberg, Bernard & Torsten J. Selck. 2006. Testing Procedural Models of EU Legislative Decision-Making. In *The European Union Decides*, ed. Robert Thomson, Frans N. Stockman, Christopher H. Achen & Thomas König. Cambridge: Cambridge University Press pp. 54–85.

Stone Sweet, Alec & Wayne Sandholtz, eds. 1998. *Supranational Governance: The Institutionalization of the European Union*. Oxford: Oxford University Press.

Stratmann, Thomas. 2002. "Can Special Interests Buy Congressional Votes? Evidence from Financial Services Legislation." *Journal of Law & Economics* 45:345–373.

Tanner, Martin A. 1993. *Tools for Statistical Inference: Methods for the Exploration of Posterior Distributions and Likelihood Functions*. New York: Springer-Verlag.

Taylor, Paul G. 1982. "Intergovernmentalism in the European Communities in the 1970s: Patterns and Perspectives." *International Organization* 36(4):741–66.

Thiem, Janina. 2006. "Explaining Roll Call Vote Request in the European Parliament." Mannheim, Arbeitspapiere - Mannheimer Zentrum für Europäische Sozialforschung, 90.

Thomas, Matt, Bo Pang & Lillian Lee. 2006. Get Out the Vote: Determining Support or Opposition from Congressional Floor-Debate Transcripts. In *Proceedings of the EMNLP*. pp. 327–335.

Thomson, Robert. 2008. "National Actors in International Organizations: The Case of the European Commission." *Comparative Political Studies* 41(2):169–192.

Thomson, Robert, Frans N. Stockman, Christopher H. Achen & Thomas König, eds. 2006. *The European Union Decides*. Cambridge: Cambridge University Press.

Thomson, Robert & Madeleine Hosli. 2006. "Who Has Power in the EU? The Commission, Council and Parliament in Legislative Decision-making." *Journal of Common Market Studies* 44(2):391–417.

Tsebelis, George. 1994. "The Power of the European Parliament as a Conditional Agenda Setter." *American Political Science Review* 88(1):128–142.

Tsebelis, George. 1997. "Maastricht and the Democratic Defecit." *Aussenwirtschaft* 52:29–56.

Tsebelis, George. 2002. *Veto Players: How Political Institutions Work*. Princeton, New Jersey: Princeton University Press.

Tsebelis, George, Christian B. Jensen, Anastassios Kalandrakis & Amie Kreppel. 2001. "Legislative Procedures in the European Union: An Empirical Analysis." *British Journal of Political Science* 31:573–599.

Tsebelis, George & Geoffrey Garrett. 2000. "Legislative Politics in the European Union." *European Union Politics* 1(1):9–36.

Tsebelis, George & Jeannette Money. 1997. *Bicameralism*. Cambridge: Cambridge University Press.

Vapnik, Vladimir N. 1995. *The Nature of Statistical Learning Theory*. New York: Springer.

Voeten, Erik. 2000. "Clashes in the Assembly." *International Organization* 54(2):185–215.