Dimensions of Diplomacy

Understanding Private Information in International Relations Using the WikiLeaks Cable Disclosure*

Michael Gill† Arthur Spirling‡

Abstract

We consider the role of private information in international relations—a key concept in 'rationalist' models of interstate relations—from a unique empirical perspective. Noting that little systematic observational data exists regarding the contemporary private information available to state actors, we analyze 163,958 United States diplomatic cables for the period 2005 to 2010—as disclosed by the Wikileaks organization—in order to speak to several aspects of contemporary international relations theory. In this analysis, we show that diplomatic secrecy consists of two distinct 'dimensions': substantive and procedural. The former deals with secrets *per se* as they relate to specific political issue areas, the publication of which would actively damage U.S. interests, especially in terms of revealing the resolve or capabilities of the state. Procedural secrecy, meanwhile, deals with the diplomatic norm of confidentiality in meetings—regardless of the substantive content of any single cable. We relate these two dimensions of diplomacy to different concepts of secrecy in the theoretical IR literature, and demonstrate that both play an important role in dictating the classification decisions of the U.S. State Department. In uncovering these substantive points, our paper presents novel methodological tools of general interest to scholars in the field.

*First version: February 1, 2014. This version: March 6, 2015. Erin Baggott, Amber Boydstun, Andy Hall, Robert Schub, and Anne Sartori provided very helpful comments on an earlier draft. We thank audiences at the Harvard Political Economy Workshop, MIT, Princeton University, the New Directions in Text Analysis conference, and the Midwest Political Science Association meeting for feedback.

†Department of Government, Harvard University. mzgill@fas.harvard.edu

‡Department of Government, Harvard University. aspirling@gov.harvard.edu

1 Introduction

State secrets and the private information possessed by leaders are nearly impossible to observe in practice. Yet such quantities—and the general difficulty of communicating that information in the form of a state's capabilities and resolve—are at the core of the game-theoretic approach to the study of conflict (Frieden and Lake, 2005; Lake, 2010). In that paradigm, private information is a crucial ingredient of understanding why disputes may arise (see, e.g., Powell, 1999, 2002; Tarar and Leventoglu, 2009) and escalate (e.g., Fearon, 1994*a*), how they might end (e.g., Goemens, 2000), their duration (e.g., Slantchev, 2004) and whether they might be avoided in the first place (see, e.g., Fey and Ramsay, 2011). This central place for information, secrecy, and beliefs also holds in theories of international relations that do not invoke the tenants of bargaining directly, including those that rely on ideational motivations (e.g., Wendt, 1999) or the pursuit of material resources and state security (e.g., Waltz, 1979), more broadly. Perhaps unsurprisingly given the importance of communication and information control, recent theoretical research has turned particularly to the use and practice of diplomacy in the international system (see, e.g., Sartori, 2002; Smith and Stam, 2004; Sartori, 2005; Kurizaki, 2007; Trager, 2010; Ramsay, 2011). There it joins a now well-established empirical literature that either explores the plausibility of the fundamental tenants of the rationalist approach (e.g., Fearon, 1994*b*; Partell and Palmer, 1999; Werner, 1999; Reed, 2003; Ramsay, 2008; Reiter, 2009; Potter and Baum, 2014), or assuming that those assumptions are correct, gives methodological advice on how to fit statistical models consistent with them (see, e.g., Signorino, 1999).

Despite the obvious progress made on creating and examining models that do a better job of describing the reality of crises as observed by international relations researchers, we still understand relatively little about the dynamics of information transmission and protection

in *non-crisis* situations—yet these surely constitute the great majority of all international interactions. Accordingly, outside of some specific policy areas relating to conflict (on military operations, see, e.g., Keohane and Nye, 1977) we have a dearth of knowledge regarding the more general diplomatic behavior of leaders and bureaucrats. For example, we know very little, in an empirical sense, of how 'military capabilities' are actually conceived of by important agents, how such information is withheld or protected relative to other international political issues, and how that conception affects what is promoted, concealed or communicated to foreign and domestic actors. This situation is unsurprising, but unsatisfactory. It is unsurprising—given limited resources, including researcher time—the literature has focused on crises and war, which have very clear welfare consequences. Assuming we did want to study more general interactive practices, data limitations are prohibitive in any case: secret information is, by definition, closely guarded by states and even when their files are declassified and actors willing to give interviews, they do so in an obviously selective way (see,e.g, Shapiro and Siegel, 2010, for discussion). Furthermore, were such data possible to obtain, it is not in an obviously usable form. In particular, quantitative scholars have tended to hone their techniques for observational data in which each 'row' represents the incidence of a particular phenomenon of interest (e.g., Ghosn, Palmer and Bremer, 2004), whereas information pertaining to diplomacy is mostly in terms of documents (primarily cables) sent between embassies and bureaucratic departments and ministries. In such a world, what constitutes an individual observation is quite unclear.

This state of affairs is unsatisfactory for more obvious reasons: put very crudely, as a discipline, we do not how secrecy 'works' in contemporary diplomacy, despite the fact that it is a vital part of our theories. Our ignorance in international relations regarding such a core component of our theories may be compared with a much more favorable situation in related areas of political science which have new and fine-grained data to assess the plausibility of

their models, and to expand their understanding of the processes therein. As an example, consider the now voluminous literature on psychology in international relations (starting at least with Jervis, 1976), in which scholars have used new measures of biological responses to assess the effects of emotion on political opinions (e.g., Renshon, Lee and Tingley, Forthcoming).

This paper introduces new data and methods to get precisely at these fundamental issues for United States foreign policy: that is, we characterize diplomacy in terms of what is kept secret, and provide explanations as to why. Our data are the more than one hundred thousand diplomatic cables released to the WikiLeaks organization, dealing with the period between 2005 and 2010—an era in which coverage is relatively dense, and during which the United States had several ongoing military operations in the Middle East. The cables—which are essentially secure emails sent between the Department of State and its embassies and missions in hundreds of cities around the world—include large numbers of documents never intended for public consumption. As such, they include thousands of (officially) secret and confidential missives, and thus allow us extremely rare access to the world of private information in IR unfettered by state censorship. Our primary theoretical contribution is to argue that private information is as much about 'procedure' as it is about 'substance'. Put differently, though the United States certainly does not wish other state actors to discover certain facts about its material capabilities (especially regarding matters of 'high politics' in the sense of Keohane and Nye, 1977), it also wishes to obfuscate the way that it allocates intelligence gathering resources. That is, the United States seeks to hide the information it gathers, from whom it garners it, and what it chooses to disseminate. We call this latter dimension of diplomacy 'procedural' and contend that it refers not to actions regarding specific objects (such as arms, services, plans or strategies) but rather to a method of behavior in general, regardless of political issue areas. In keeping with a pre-existing literature

on diplomacy (particularly Sartori, 2002), we provide corroborative evidence for the notion that developing and maintaining a reputation for confidentiality matters to diplomats and their staffs. We go beyond current accounts, however, in demonstrating that patterns of information protection hold somewhat independently of the policy areas discussed.

While the purpose of this paper is not to 'test the assumptions' of bargaining models *per se*, our hope is that our inspection will aid researchers interested in the empirical implications of such theoretical models. In particular, studying the way that information is shielded from global public view on an everyday basis provides a resource for those interested in 'audience costs' (e.g., Fearon, 1994*a*; Weeks, 2008)—which apply to crises bargaining situations, specifically. Furthermore, since we show that both procedure and substance matters for secrecy, we believe our empirical efforts provide an impetus for theory development in contemporary international relations research. This is in part because our sample is much broader in substantive terms than has been available for previous studies, for which we have theoretical work already. But it is also partly because the cables we use have 'official' designations in terms of their classification status and the justification for that status. That is, our work relies on the officially structured indexing of cables and their topical designations, for which state actors have made conscious decisions. Aside from data cleaning, we do relatively little to restructure the standardized form of the metadata observed on each diplomatic cable, as such a procedure could induce error in the variables we chose to measure. As a result, our inferences are less dependent on somewhat idiosyncratic or arbitrary research rules, and thus sharper than previously possible; we believe this makes them especially ripe for theorizing (for further testing thereafter), as our topical policy codings are *identical* to the internal indexing structure used for information management at the State Department.

In undertaking the study, our paper contributes methodologically and suggests new ways

of working with 'texts-as-data' (Grimmer and Stewart, 2013). In particular, we use machine learning techniques—such as random forests (Breiman, 2001) and the 'lasso' (Tibshirani, 1994)—in tandem with matched sampling designs to identify how 'important' terms discriminate between restricted and unrestricted documents. We provide novel ways of comparing texts, based on matching on the metadata of each document, such that political scientists may think sensibly about the (marginal) influence of secrecy on a document's content.

This paper proceeds as follows. In the next section we describe some literature, and orient the reader to the debates at hand. Then, we introduce the data and provide some summaries of it. We then discuss the methodological contributions. First, we consider an uncovering of the substantive dimension of diplomacy. This boils down to a regression-based approach with internal State Department subject 'tags' as independent variables predicting the classification status of a document (with multi-way fixed effects between location of origin and time). This is followed by an investigation of procedural secrecy, which sees us exactly match on official State Department subject tags, and using popular machine learning techniques to estimate the difference between restricted and unrestricted access cables (i.e., estimating the textual differences between more restricted and less restricted cables, on a sample of cables that are comparable given their meta data). We then provide some discussion of our results and conclude.

2 Theories of Private Information Disclosure

To get much of their purchase on the world, the analytic conclusions drawn from rationalist studies of international dynamics often hinge on differences between private and public information, and the degree of information overlap or 'common knowledge' shared by actors in

a non-cooperative environment (see Powell, 2002, for an overview). Somewhat surprisingly, researchers modeling in this tradition rarely define exactly what constitutes information of either type. If they do, they do not study the difference in detail. Primarily this is a data issue: understanding what states know, what they do not, and what they are keeping secret cannot be determined deductively insofar as official refusal to answer queries about particular issues does not allow scholars to become much more informed about the true state of the world. While it is correct that documents pertaining to international decision making are routinely declassified, this process tends to be slow, somewhat haphazard, and obviously case selective (see Allen and Connelly, 2015, for an overview of U.S. protocols).

Given the obvious data limitations for recent historical periods, scholars interested in examining the plausibility of theories that rely on information—for example, those that utilize 'audience costs'—tend to pursue one of three avenues. First, they use conflict (or survey) observational data in a regression context (e.g., Weeks, 2008). This has the obvious benefit of being straightforwardly replicable, but arguably lacks the kind of internal validity that would be convincing for skeptics (see, e.g., Trachtenberg, 2012, for recent discussion). Second, they enter the archives of governments and produce historically rich case studies (e.g., Schultz, 2001; Snyder and Borghard, 2011), which are necessarily limited to specific times and places. Third, they undertake field or survey experiments on political actors or on less representative samples (e.g., Tomz, 2007). In all cases, it is obvious that if researchers had broader access to cases or incidences of censored versus public materials, they might be able to draw sharper conclusions—better in terms of both internal and external validity—than current approaches allow.

Whatever the research approach, a key insight of the audience cost literature is that increasing the publicity of some political matters may be useful, particularly if the issue at

hand is one of signaling a commitment credibly—and we might expect variation in a nation's willingness to publicize issues given the strategic importance of the issues at hand (e.g., Fearon, 1994*a*). Of course, strategically making information public is helpful beyond the signaling case: citizens sometimes need to be warned to stay away from certain areas, just as investors may be encouraged by the public announcement of tax breaks or shifting interest rates. Along these lines it is not hard to think of 'substantive' facts, especially connected to capabilities, that states may want to conceal—for example, where they locate their nuclear warheads, or how many they have and how easy they are to launch. It would directly hurt a state were any of this information known, since it compromises defense plans, decreases combat effectiveness, and broadly provides potential enemies with additional bargaining power.

Inasmuch as nations have incentives to disclose some information publicly while keeping other information private, one can imagine scenarios in which governments prefer to conceal what they are 'trying to know' or how they acquire and share information independent of any substantive issue at hand. In brief, there may be incentives to practice discretion, on average, independent of a particular policy issue being discussed diplomatically. Suppose, for example, the United States had a rule such that it publicized its diplomatic communications on some policy areas but not on others; competing nations (or public actors) could reasonably update their beliefs on what is likely being negotiated or communicated in private given the absence of that policy in being revealed in public summaries. In such a world, one can see how a practice of selective information disclosure could in fact advantage international competitors relative to a more general norm of secrecy and discretion. In an extension of this logic, nations may have incentives to protect or conceal information concerning the meetings of leaders or public officials irrespective of the topics or issues discussed. Taking this further, suppose it was common knowledge that a group of international leaders had scheduled a

meeting on a certain day to discuss an unknown international issue; if an institutional rule existed such that the topic of the meeting would be made public if and only if it concerned a certain set of political issue areas (e.g., the topic of the meeting is made public if about environmental politics or maritime disputes, but kept private if about military capabilities, terrorism, or territorial disputes), then the absence of a public summary of the meeting would nevertheless reveal some information to the public about the likely issues raised in that private communication—a fact that could make the institutional rule of selective disclosure, at least on the margins, strategically self-defeating. It would be more difficult for the public to update its beliefs on the likely content of a meeting if no such (disclosure) rule were to exist, by contrast.

Although the discussion above concerns the practice of disclosure in international diplomacy, the intuition that selective (i.e., topic-specific) information disclosure may be welfare reducing is at the core of a much broader class of political, bureaucratic, and economic phenomena. There is a vast microeconomic literature on the public dissemination of private information in macroeconomic environments with strategic complementarities—an area of economic research which tends to focus on the conditions under which information disclosure may be welfare improving or influence a system's volatility. Common settings for such formal reasoning are studies of central bank communication and global financial systems, where decision makers may wish on the one hand to disseminate some information to public audiences but avoid being so transparent as to induce financial crises or speculative attacks on the other. In the determination of interest rates and deliberation of monetary policy, central bankers face decision problems that closely resemble problems faced by international diplomats, although this bureaucratic comparison is not specifically discussed in this literature. When making a public statement, social planners may be concerned both about the precision of a public signal (i.e., how closely a message will map to real world outcomes) in addition

to the degree of its publicity (i.e., how many individuals observe a signal). Cornand and Heinemann (2008) provide a careful discussion as to how precision and publicity may both independently and jointly shape the decision to disclose private information. The authors argue, in a refrain largely in step with much of the scholarly theories related to this problem, "The optimal degree of publicity depends on the precision of announcements" (718)—more specifically, that if the precision of a public signal is not guaranteed to be sufficiently high (e.g., if the event or issue described in a communication is not sufficiently likely to occur in reality), it may be dominant for bankers to avoid private information dissemination to the public sphere altogether. Arguments more favorable to the prospect of transparency's welfare improvement can be found in Cukierman (2001) and Angeletos and Pavan (2004), under the requirement that the quality of a public signals is sufficiently high. The model presented in Woodford (2005) provides a sufficient condition for when increased transparency of central banking should lead to a welfare reduction in expectation.

2.1 Implications of Motivating Theories

In international relations, meetings and communications between diplomats and state officials may be in service of multiple ends. Diplomats may wish to signal their intentions about U.S. foreign policy to other leaders, share sensitive information freely and openly to trusted confederates, and—in perhaps an ideal case—update the quality of their private information about political issues that might otherwise be difficult to observe through other means. With that said, there is no perfect way to observe all private information available to diplomatic actors by simply observing records of their communications and their associated handling statuses. Even if our data contained records of all top-secret communications between relevant officials, our sample would still be unable to sensibly measure all officials' private information *per se*.

What can be approximated, however, is the degree to which particular policy issues and language-based features map to higher levels of political protection within our sample. Common sense dictates that the relationship between political issues and handling status ought to be strongly correlated with political actors' preferences over the sensitivity of a political issue area, or more broadly, the incentives to shift information disclosure to the public on a specific political issue. If we are to observe that some political issues are systemically more predictive of document secrecy than others, *a fortiori* this provides evidence of diplomatic preferences over the sensitivity of political issues in the international system.

Formal theories of strategic information disclosure suggest diplomats would have incentives to withhold information as a function of the presumed cooperativeness of a decision environment. All else equal, as incentives between nations to coordinate on policy issues decrease (e.g., on national security discussions or information on nations' relative capabilities), we would expect communications on such issues to be more protected in our sample on average than policy areas with more of a cooperative or 'common-pool' character (such as environmental concerns). In terms of their rank ordering, a natural prediction concerning the 'substantive' secrecy of our sample would suggest a monotonic decrease in the in-sample estimates of cable classification as one moves from traditionally 'non-cooperative' games on one end to 'cooperative' issue areas on the other.

To assess the plausibility procedural dynamics as a driver of information restriction, our aim is to extract features of language that are predictable of cable secrecy after adjusting for the topical focus of diplomatic communications. If formal theories on information dissemination focus on minimizing enemies' abilities to anticipate foreign decisions, one might expect particular textual features of diplomatic cables to be predictive of secrecy conditional

on subject matter. In particular, if meetings between diplomats, ambassadors, or leading political officials serve an information-gathering purpose, one might expect references to such individuals, all else equal, to be positively associated with a document's probability of restriction.

3 Data

The WikiLeaks cables disclosed in the Manning leak of 2010 are the 251,237 diplomatic secure messages sent by the U.S. State Department to U.S. embassies and missions. The date range for the original data is from 1966 to 2010, and in Figure 1 we plot the total number of cables per month from that time period. In our work here, we focus on all cables written and sent between January 1, 2005 and the end of the data, the first few months of 2010. We do this for two reasons: first, because as the figure suggests coverage prior to the year 2000 is somewhat sparse and inconsistent. Second, because we were concerned about changes to security procedure (particularly regarding requirements to copy in embassies and missions on particular types of messages) after the terrorist attacks of September 11, 2001. All told, we are left with around 163,958 documents from which to draw inferences. Additional details on the likely representativeness of this sample relative to all communications between embassies over this period can be found in Gill and Spirling (Forthcoming).

Technically speaking, cables may be classified into one of three categories, depending on the degree of damage to national security that "the unauthorized disclosure of which reasonably could be expected to cause."[1] Furthermore, any classified document must pertain to at least one of a series of topics which *inter alia* include military plans, intelligence, foreign relations of the United States, nuclear programs, weapons of mass destruction and vulnerabilities in

[1] As described in Executive Order 13526, 2009.

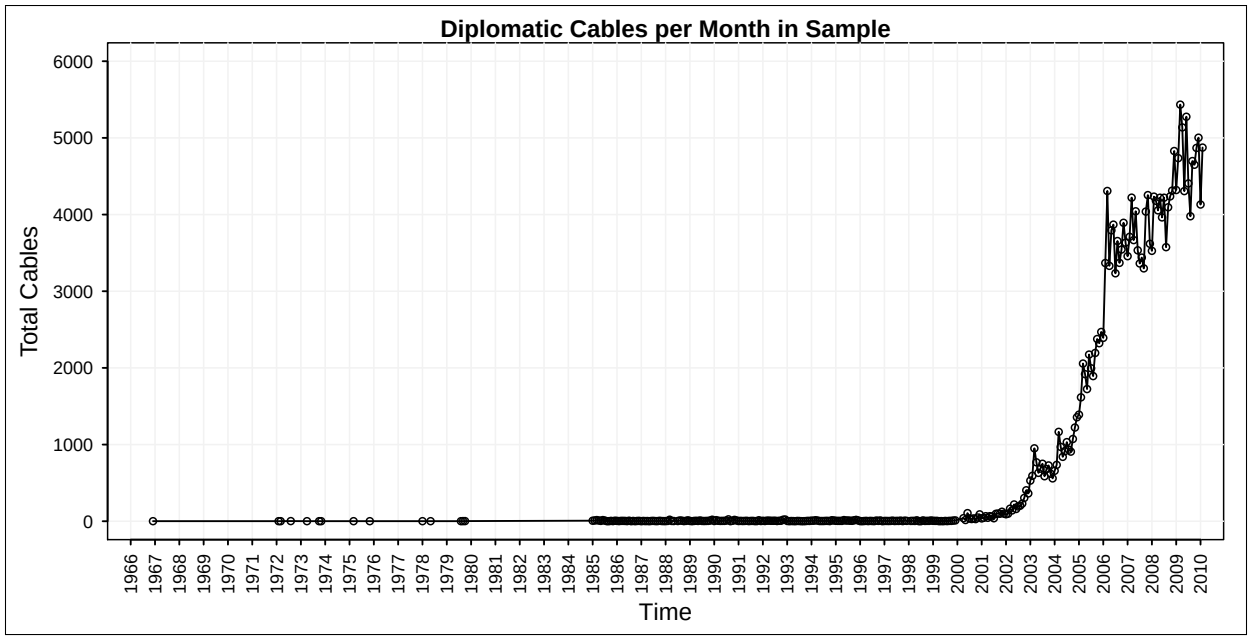

Figure 1: Number of cables per month, 1966–2010. Note that post-2001 period has much more dense coverage.

national security. In descending order of the purported balefulness of unauthorized release, these categories are 'Top Secret', 'Secret' and 'Confidential'. If a cable does not meet the criteria for such restricted access, it is deemed 'Unclassified'. In our particular data, we have the following distribution: zero Top Secret, 10,195 Secret cables, 87,270 Confidential cables, and 66,493 Unclassified. There are, in addition, some extra classifications that appear less frequently in the data, such as 'Confidential and Not For Foreign Distribution' , 'Unclassified for official use only' , and 'Secret and Not for Foreign Distribution'; we ignore these categories for our current efforts.

For our purposes below, we divide the categories into 'restricted' (R), which includes Secret and Confidential communications, and 'unrestricted' (U), which includes the unclassified documents only. The central idea here is to separate documents into more 'private' and more 'public' information, respectively. This measure is somewhat crude, but given that theories in International Relations use similarly binary demarcations we think this is reason-

able. To be clear, the fact that a cable is unrestricted does not mean that it is automatically made public: it is still a government document rather than a press release. But unclassified documents—so long as they are not 'For Official Use Only'—do make their way into the public domain, and are eligible for release under Freedom of Information Act requests. Put otherwise, our unrestricted case covers documents that the public (anyone without specific security clearances) could access; our restricted cables are those that are not released or releasable to the public.

Any given document has a series of subject matter 'tags' assigned to it by its authors, with guidelines for this process provided by the State Department.[2] From our perspective, these tags contribute meta-data that communicates the topic of the content therein, and are assigned to each document following its completion. After cables are written and subject tags are assigned, each cable is given its overall classification status.[3] Examples of subject tags in our data include '`ADCO`' which refers to 'Diplomatic Courier Operations', '`PTER`' which refers to 'Terrorists and Terrorism', '`SMIG`' which pertains to 'Migration' and so on. There a total of 97 tags in our data, though their use varies widely in relative frequency terms. The full list can be seen in Appendix C. The variety in tag number per document can be seen in Figure 2; inspection suggests that the modal number of tags is two or three, though there are 14,451 unique combinations of subject tags (ignoring each cable's location of origin) that appear in our post-2005 sample at least once.

In Figure 3 we report the structure of the data in terms of the way that tags co-occur across

[2]These are literally geopolitical 'TAGS', an acronym for 'Traffic Analysis by Geography and Subject', implemented for diplomatic communication in its modern form by an executive order (number 11652) in June 1974. Their justification was to "[p]ermit more rapid and discriminating distribution of messages", and to "[p]rovide statistics to both offices and posts on what is being communicated in the Department-field system"; they were to "[s]erve as headings for clustering the terms used by professional indexers to identify the content of substantive messages."

[3]Additional details on the origins and formal procedures of cable creation can be found in Appendix A.

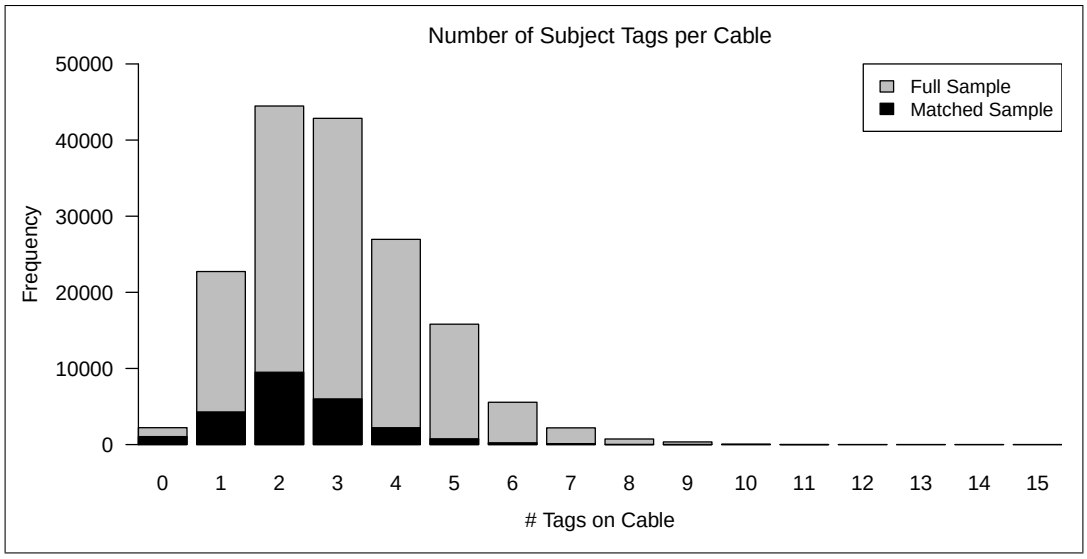

Figure 2: Distribution of tag numbers by document, matched and unmatched sample (defined below)

cables. Areas of darkness in that plot are places where tags coincide. Our main observation is that tags in section 'P' (which denote 'Political' issues) and, to a lesser extent, tags in section 'E' (denoting 'Economic' matters) tend to coexist heavily with other subject indicators, suggesting that these issues play an important organizing role in the U.S. diplomatic service. Machine readable versions of the documents themselves are available at various websites for download, though some pre-processing is then required prior to any analysis. In particular, the tag information must be captured and removed, and some other cleaning performed. Much of what follows involves operations on the 'document-term matrix' (DTM) of the texts, which was 'stemmed' (meaning that words were pruned back to their 'roots' where possible, using the Porter (1980) algorithm), 'stopped' (meaning that function words which are thought to contain little discriminating information were removed), and subject to a 'sparsity' condition of 99-percent (i.e., only words that occur in more than one-percent of all documents but in no more than 99-percent of all documents are included). Such data cleaning is common in text analysis in the social sciences (e.g., Grimmer and Stewart, 2013). The resulting DTM for analysis is matrix with dimensions $163,958 \times 3,755$.

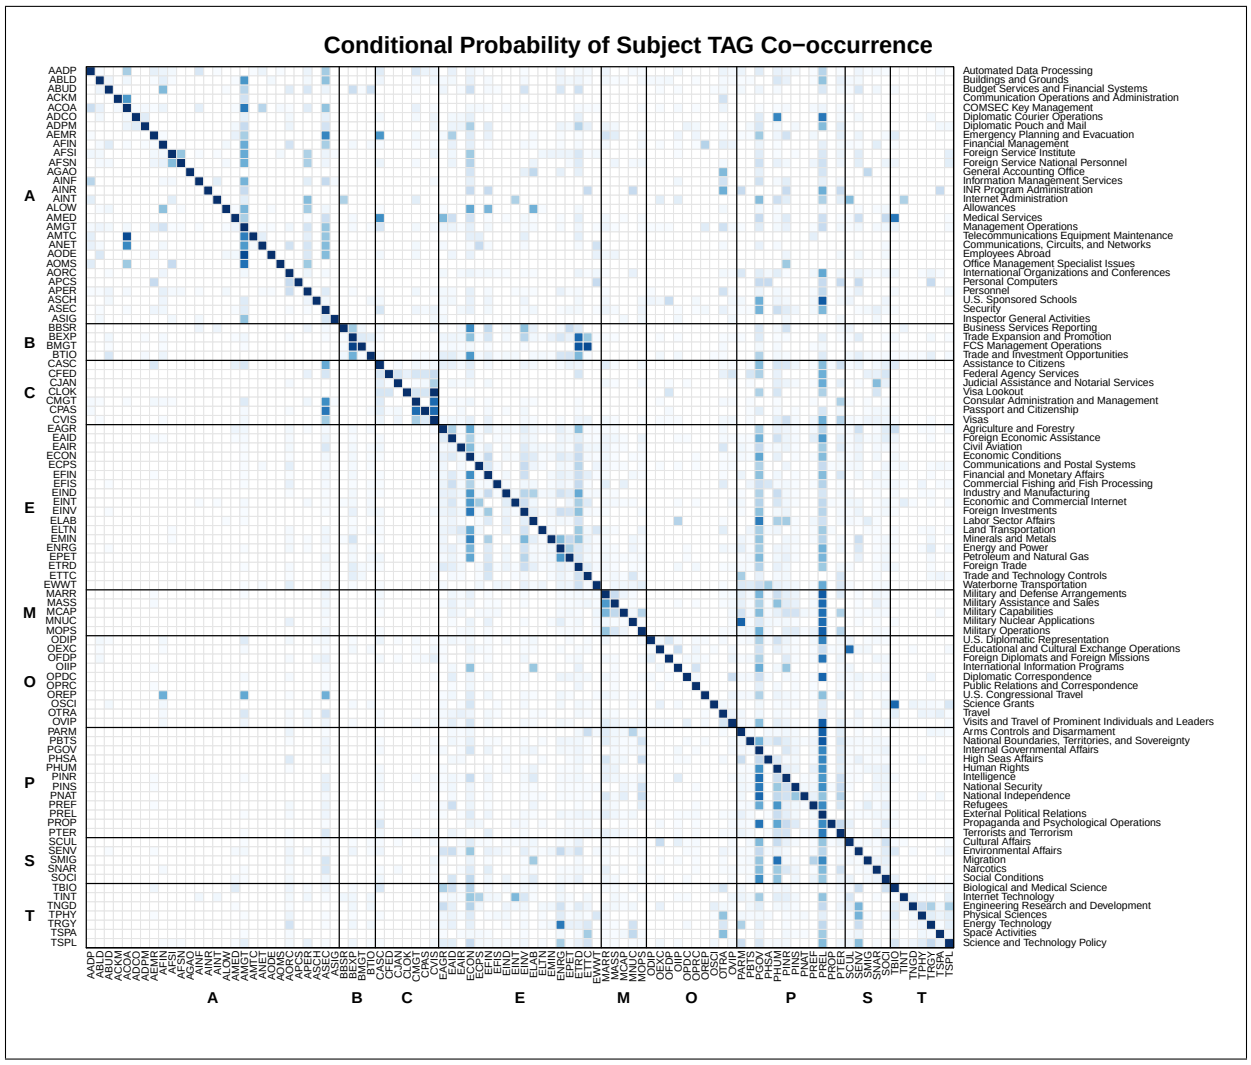

Figure 3: This figure shows the empirical conditional probability of U.S. State Department subject TAG co-occurrences in the post-2005 sample ($n = 163,958$). Subject tags are are presented in alphabetical order with their official U.S. State Department meanings listed in the righthand column. Each cell in the figure represents the conditional probability that a *column subject* will be tagged given that a *row subject* has already been tagged. Darker shaded cells indicate a higher conditional probability.

4 Methods

Our claim above is that the secrecy endemic to diplomacy comes in at least two separable varieties: substantive secrecy—the notion that certain information about a policy area is to be kept confidential because it would be *per se* damaging to security were it released—and procedural secrecy, which is concerned with the notion that secrecy protects foreign or domestic agents from outside consequences of their actions. To assess the evidence for these separate ideas, some care is required in terms of methods. Here we explain our approaches.

4.1 Substantive Secrecy

We first examine the question of *substantive secrecy*—i.e., how a State Department topic or substance of a diplomatic communication, all else equal, influences its probability of restriction. The objective is to quantify both the magnitude and direction of how the presence of official U.S. State Department communication subject tags influence cable secrecy. In suit, we regress each cable's observed restriction status on its subject tags and location of origin. This fixed-effects least squares equation can be written as follows:

$$R_i = \alpha + \sum \beta_t Tag_{it} + \gamma_j + \varepsilon_{ij} \tag{1}$$

where R_i is a dummy variable for cable i that takes the value of 1 if the cable is restricted and 0 if unrestricted, Tag_{it} is a subject tag dummy variable for cable i for each tag t, γ_j is the fixed effect for embassy j, while α and ε_{ij} are the constant and error terms, respectively. Given that each covariate in this regression is binary, each regression coefficient $\hat{\beta}_t$ is a sample estimate of the difference between two conditional expectations: the conditional probability a document will be restricted given the presence of a subject tag minus the conditional

probability of restriction without that subject tag present.[4]

4.2 Procedural Secrecy

Recall that procedural secrecy concerns the diplomatic norms of confidentiality in meetings. If it exists as a quantity that can be identified in our data, then it should emerge as a key discriminator between restricted and unrestricted cables. However, if there is indeed a subject tag imbalance between restricted and unrestricted cables (as suggested above), this implies that a simple comparison of word frequencies between restricted and unrestricted documents is unlikely to isolate how text varies on the margins a function of secrecy status alone, since observed differences are likely to arise directly from ex ante differences in subject matter.

Thus, the question we ask in this section is: having adjusted for cable subject matter (given an observed sequence of subject tags on a document) and locations of origin, all else equal, can restricted diplomatic communications be distinguished from unrestricted communications? This question may be thought of as estimating the *marginal effect* of secrecy on the content of a restricted communication. More precisely, given two documents indexed with identical subject tags and originating from the same source, are there specific textual features that systematically distinguish restricted cables from unrestricted cables? If such textual features exist, is there anything substantively unifying about these features? In particular, does whatever differentiates these communications be considered 'procedural' in nature?

[4]That is, $\hat{\beta}_t = Pr(R_i = \widehat{1|Tag_t} = 1, \mathbf{X}) - Pr(R_i = \widehat{1|Tag_t} = 0, \mathbf{X})$. If standard assumptions hold—e.g., unconfoundedness, overlap, or "selection on observables" (Rosenbaum and Rubin, 1983; Heckman and Robb, 1985)—$\hat{\beta}_t$ tells us on average how much document restriction varies by each subject tag in our sample. Although the outcome of interest is binary, OLS is appropriate when the conditional expectation function (CEF) of each regressor with respect to the outcome exhibits is linear (see, e.g., Angrist and Pischke, 2009, Chapter 3). OLS suitably estimates whereby sample average effect of restriction on each subject tag in the context of our data, as each regression coefficient represents a conditional mean.

4.2.1 Exact Matching on Subject Tags and Origin

To assess whether secrecy, on the margin, is associated with differences in document composition, we restrict our sample to *exactly matched* subsets of cables within each embassy in our sample. More precisely, for each embassy (i.e., each cable's location of origin), we implement the algorithm outlined in Figure 11 in Appendix F to construct datasets of cable pairs that are exactly matched on official U.S. State Department subject tags and their embassies of origin, but differ on their restriction level. The objective of this matching procedure is to restrict the full sample such that there is perfect subject overlap on cables in our study. As a result of the matching procedure, within each embassy, for each restricted cable there will exist an unrestricted cable that has an identical subject tag pattern. We rely only on the State Department's official subject tags for this procedure. If two or more unrestricted matches are found for a single restricted cable, we select the match that is written most closely in time to the restricted cable's date of authorship. For the results presented in this study, matching is performed without replacement, and datasets are stored and analyzed at the embassy level (although pooled analyses are also appropriate with the resulting data).

Since we wish to make inferences about textual differences between restricted and unrestricted cables on the margin—i.e., once cable subject tags have been accounted for—the within-embassy matched sampling design has intuitive appeal. The sampling design allows for a meaningful examination of procedural secrecy. Adjusting the sample directly for differences in subject matter and controlling for embassy-level effects, the design allows us to isolate differences in textual composition that are likely to arise from a document's handling status alone. Intuitively, the aim of our exactly-matched sampling design is to "control" for substantive differences in cables that may be present in the unmatched sample—differences that may arise from hypothetical variation in reporting rules, document disclosure standards, authorship style, or political priorities at the embassy level. If systematic textual differences

remain between restricted and unrestricted cables after subject and location have been accounted for, these differences are likely to arise from procedural rules that are separate from subject-specific handling rules.

The formal appeal of exact matching is that it is nonparametric and approximates the act of "blocking" in randomized experiments (Cox, 1958; Imai, King and Stuart, 2008).[5] Exact matching is often untenable in applied research, however, since in many cases the sampling procedure can dramatically reduce a researcher's final sample size, and the procedure tends to rely on initially large sample sizes. Unsurprisingly, this was a concern for our modeling attempts, along with the possible danger that many documents dealing with sensitive substantive areas would be jettisoned from the final analysis because no match could be found for them. Further, we were concerned that certain 'important' embassies would be, relative to the original dataset, heavily under-represented.

Neither of these concerns appear to be true of the matched sample. In Figure 4 we report the reduction in subject tag imbalance of the exactly matched sample, in addition to information on which subject tags remain present. In the exactly-matched sample, we see both embassy-level and aggregate level subject imbalances have been eliminated. The upper-right plot shows the within-embassy subject proportions are identical between unrestricted

[5]In the context of parametric adjustment, Rosenbaum and Rubin (1983, 1984) show that matching on a correctly specified propensity score (i.e., a unit's conditional probability of being assuaged to treatment, given its covariates) is sufficient to allow for the unbiased of the average effect of a treatment for a given population of interest, i.e., eliminate confounding. But in settings with observational data, a researcher rarely knows whether an appropriate functional relation has been specified in model-based matching procedures (Rosenbaum, 2002). The appeal of an exact covariate matching procedure is that if the appropriate set of conditioning measures has been identified, the unobserved functional relation between between covariates and the assignment to treatment is ignorable due to *perfect balance* on conditioning variables. Under general conditions, exact matching procedures are both *equal percent bias reducing* (Rubin, 1976) and *monotone imbalance bounding* (Iacus, King and Porro, 2011). These traits are not generally true for most distance-based or model-based (parametric adjustment) matching methods, which has led several scholars to conclude that exact matching is close to an "ideal" matching procedure in observational settings (e.g., Stuart, 2010; Imai, King and Stuart, 2008).

and restricted cables, which is why the colored lines appear purple (due to perfect overlap between the red and blue lines). The exact subject balance is also true for the exactly matched sample average. The lower-right plot demonstrates this point in an extreme form: there exist no embassy-level or aggregate level subject imbalances in the exactly-matched sample. The analysis gives us confidence that inferences following from the matched sample will be appropriate to a broad class of diplomatic communications. The distribution of subject tags in the matched sample map to substantively meaningful political issue areas. The majority of cables in the matched samples have to do with Economic Affairs, Military and Defense Affairs, and Political Affairs—each topic within foreign policy that are closely related to formal theories of rational diplomacy. On the other hand, the exact-sampling design is less capable of making credible counterfactual statements about Administrative Affairs and Outreach.[6]

In Figure 5 we report the embassies, and their relative prevalence, in our matched data. Importantly, we note that 'larger' embassies—including the U.S. State Department itself—are most represented; in particular, Ankara, Baghdad, Paris, Cairo and Moscow (all centers of activity in the original data) appear at higher rates in the matched sample. Taken alongside the results of Figure 4, this presents strong evidence that the matched sampling procedure does not leave the general patterns of the whole sample too far behind, and is due to the fact that there are sufficiently high within-embassy subject tag correlations. The diplomatic locations contained in the study sample are represented in a manner proportionate to their

[6]The reason subject tag overlap is important is because exact matching will allow us to inspect textual differences akin to *treatment effects on the treated*. Treatment effects on the treated are not the same as the average effect of treatment unconditionally, nor are they the average effect of treatment in the sample. More precisely, they concern how much potential outcomes would differ for a set of treated units in the sample if they were instead to become untreated. In the present study, therefore, with our exactly-matched sample, the design allows us to estimate answers questions like the following: If a set of treated restricted documents like those in our sample were instead to become unrestricted, on what textual dimensions would we expect those collections of documents to vary?

overall representativeness in the full sample.

4.2.2 Supervised Learning and Penalized Regression

For each of the matched samples described in Section 4.2.1, we implement a set of supervised learning models to identify which words are most *important* to (i.e., predictive of) cable secrecy. The matrix of words used in this classification setting is taken from the full post-2005 document-term matrix described before, but now only includes rows that satisfy the within-embassy, exactly-matched sampling design. On the 'left hand side' we have the (binary) restriction status of a given document which we intend to predict with the words within that document. Quantitatively, we observe how within-sample classification error rates vary as a function of which words are included in the model; qualitatively, we wish to make statements about how a document's restriction status would likely change if particular words within these documents were to vary. Two supervised learning methods are applied to these data: the "random forest" (hereafter RF) algorithm (Breiman, 2001), and the "lasso" (Tibshirani, 1994). Results from both procedures are used alongside the topic model estimates described below to make statements both at the world-level and topic-level about how secrecy, on the margin, influences the content of diplomatic communications. More details on the RF and lasso procedures can be found in Appendix B.

With both RF and the lasso, we obtain embassy-level estimates of word-level dependencies to document restriction. In the context of RF, each exactly-matched dataset for embassy j has a corresponding vector of *word importances*, where importance is defined as an estimate of each variable's in-sample average marginal error reduction. In the context of the lasso, each embassy has a corresponding vector of penalized partial regression coefficients. For both the RF and the lasso procedures, we refer to this collection of embassy importance vectors as the embassy *importance matrix*. Each row in this matrix represents a given embassy, and each

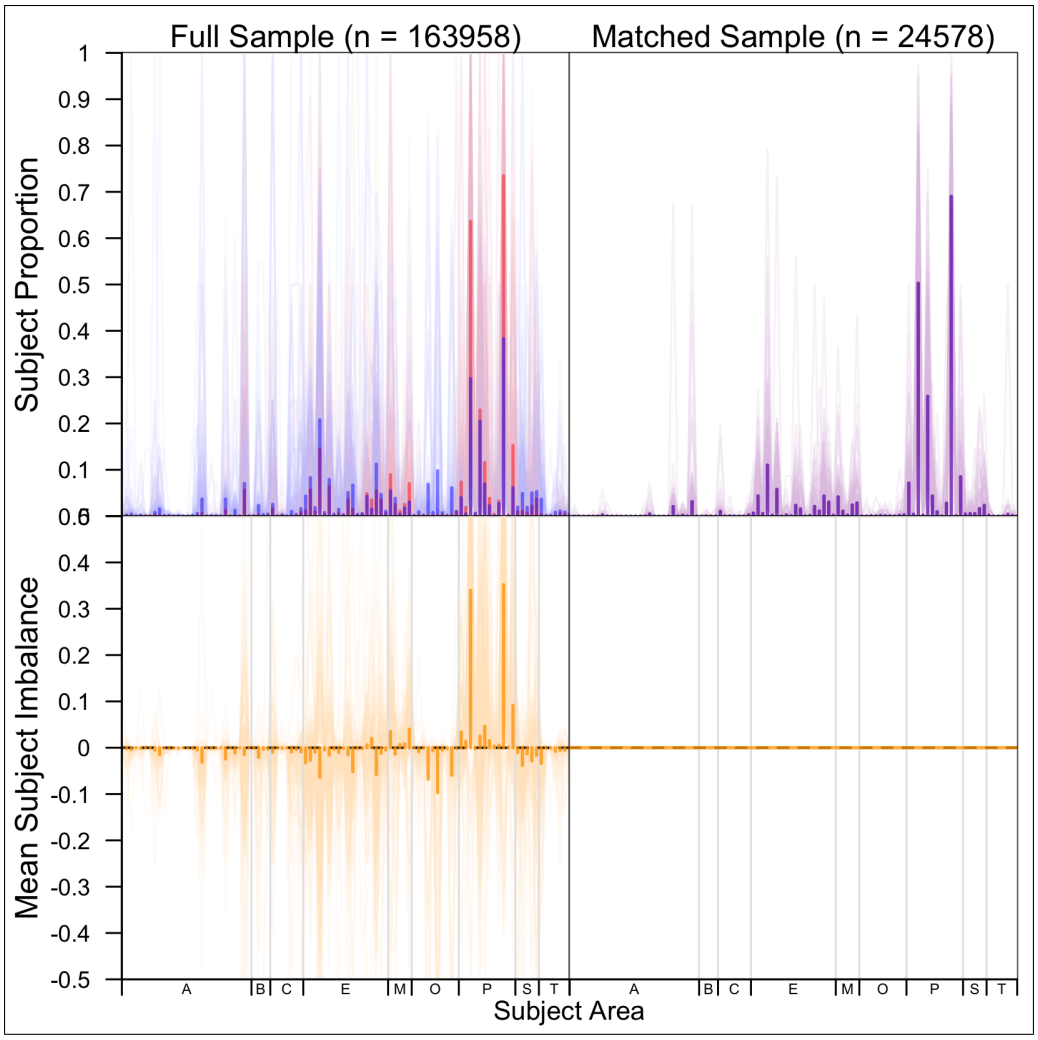

Figure 4: In the upper-left subplot, the thin-transparent lines in red (restricted) and blue (unrestricted) correspond to embassy-level averages of individual tag frequencies in unrestricted and restricted cables. The thicker vertical colored lines in the foreground denote sample averages. The lower-left plot provides much of the same information but in slightly different form: background lines correspond to embassy-level imbalances (subject tag differences in means between restricted and unrestricted cables within embassy) whereas the thicker bar plot in the foreground is the sample level difference in means. These two subplots demonstrate there is subject imbalance between unrestricted and restricted cables both on aggregate levels and, generally speaking, at individual embassy levels. In the post-2005 sample, "A – Administrative Affairs" tend to be more public, "B – Business Services" tend to be more public, "C – Consular Affairs" tend to be more public, "E – Economic Affairs" tend to be more public, "M – Military and Defense Affairs" tend to be more private, "O – Outreach" tends to be more public , "P – Political Affairs" tend to be more private, "S – Social Affairs" tend to be more public, and "T – Technology and Science" issues are slightly more public on average. The overall subject distribution in the matched sample tends to be weighted more towards "Political Affairs," "Economic Affairs," and "Military Affairs."

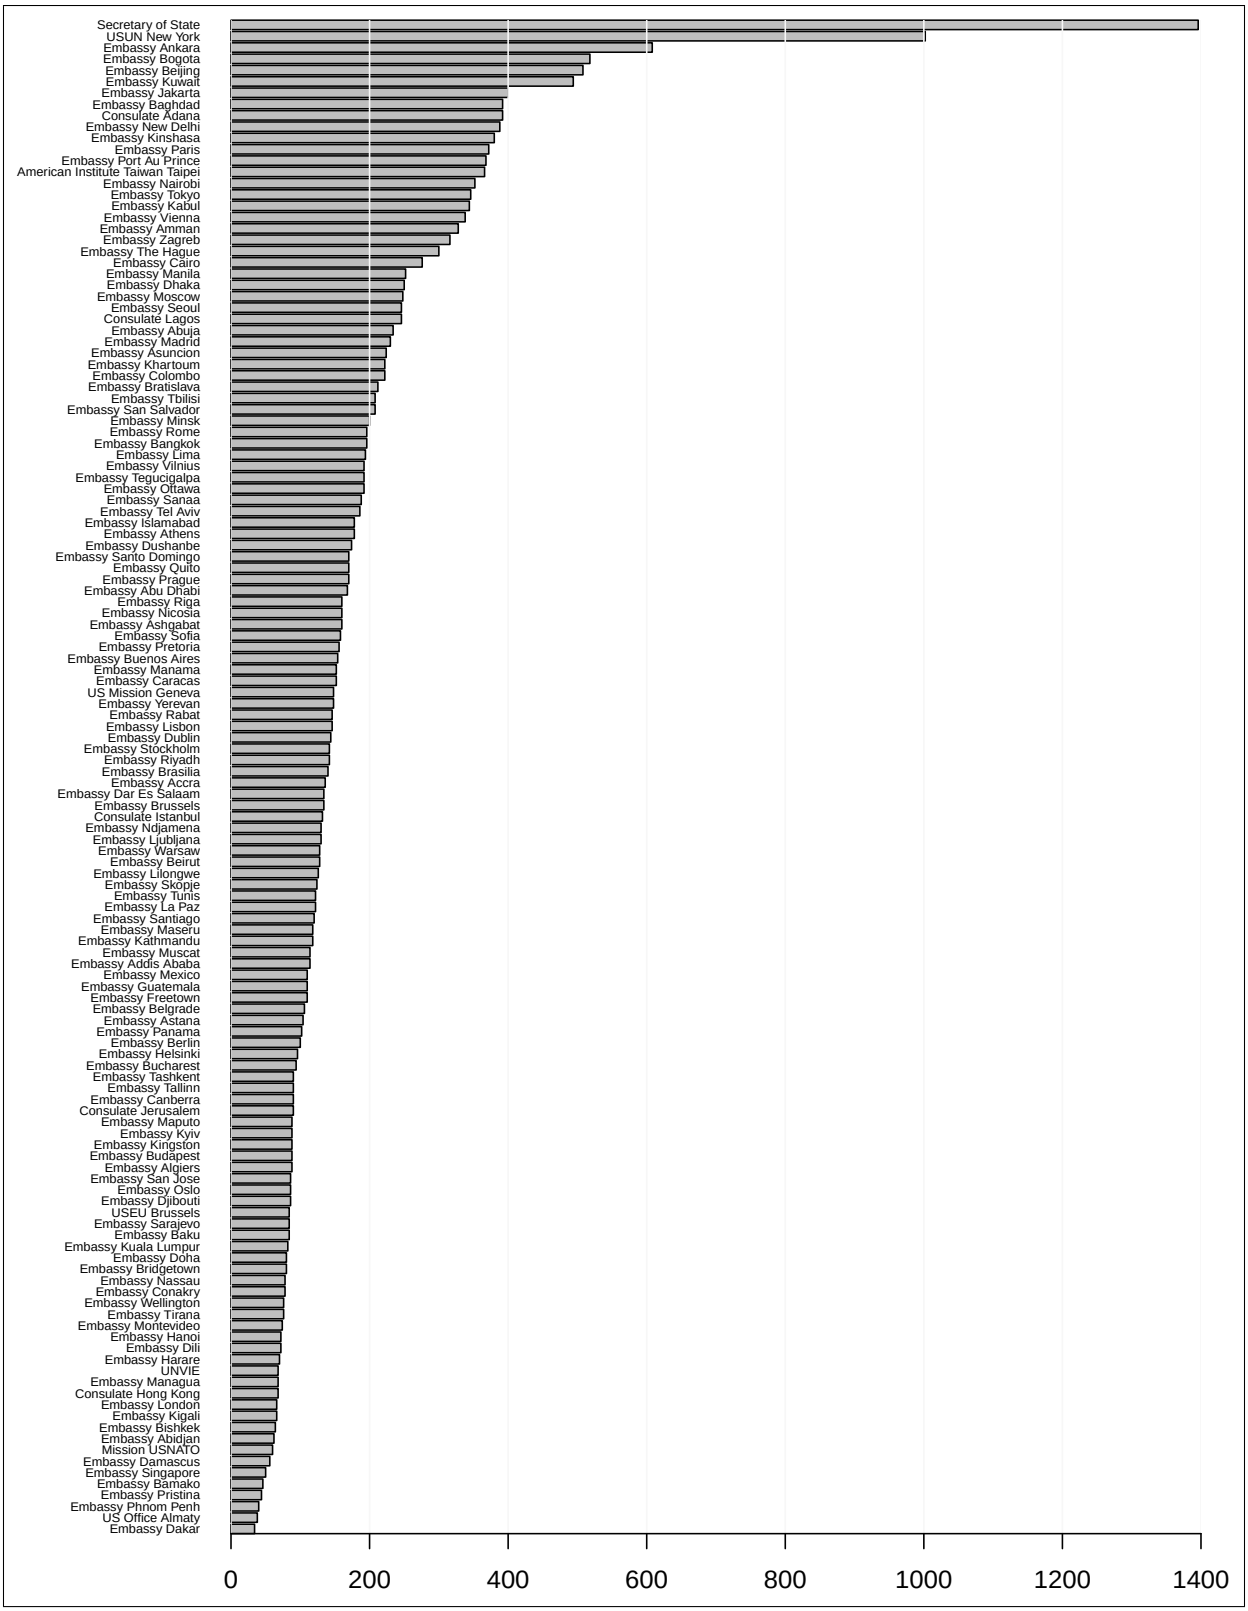

Figure 5: Counts of cables by embassy in the matched sample.

column is a measure of a word's relative importance to prediction accuracy in the embassy's matched sample. Each cell entry is then the RF importance measure for that term in that embassy. To obtain sample-average estimates of word-level importances to prediction, we weight weight the results of each embassy-level importance vector by its relative share of all cables in the exactly matched sample. The prevalence of any given embassy in the matched sample, therefore, proportionately weights the importance terms associated with that embassy (thus, for example, we will up-weight the importance terms associated with the State Department itself and other embassies near the top of Figure 5). Using the sample-weighted results of the RF within-embassy, exactly-matched classification procedure, we then took the top 30 of these terms (recall that they are all positively signed, regardless of their actual signed effect on classification), and recorded their corresponding coefficients from the lasso regressions at the embassy level. The lasso regression coefficients are similarly weighted as sample averages in proportion to each embassy's representation in the matched sample.

4.2.3 Supplementary Analysis: Topics

Some supplementary analyses are performed to address what differentiates more restricted documents from less restricted documents on the margins. In particular, we *topic model* our sample of $n = 163,958$ cables, using the most common probabilistic topic model in contemporary text analysis research, Latent Dirichlet Allocation, henceforth referred to as LDA (Blei et al, 2003). Information on our topic modeling procedure is outlined in Appendix E. Results of the topic modeling procedure are used as an illustrative aid to categorize the words we find to be predictive of document restriction.

5 Results

We first interpret our tag regressions in terms of the nature of the substantive secrecy they reveal, before considering the evidence for our procedural secrecy hypothesis above.

5.1 Substantive Secrecy: High vs. Low Politics

Recall that testing for substantive secrecy boils down to testing whether or not the probability a diplomatic cable is withheld from the public is measurably predicted by the *subject* of the cable communication, adjusting for the cable's location of origin and other factors. Figure 6 presents this analysis, where each point corresponds to an estimate of the sample average effect of a subject TAG on the probability of the cable's restriction. Around each estimate is the 95-percent confidence interval. In terms of coefficient direction, note that the broken line in the center of the plot denotes a point estimate of zero 'effect': tags to the right of this line are generally associated with restricted documents (on average); the presence of tags to the left, generally predict an unrestricted status for the cables. Tags highlighted in red indicate coefficients that are statistically differentiable from zero. Our first observation is that there are a large number of statistically significant predictors: almost every subject matter tag is associated with increasing or decreasing the probability that a particular cable is restricted. Second, we note that the direction of the effects are somewhat in line with our priors. Thus we see that cables concerning "Terrorists and Terrorism", "Military Capabilities", "Intelligence," and "National Independence," for example, are more likely to be kept private than cables concerning "Migration" "Narcotics," "Personnel," or "Environmental Affairs." In particular, we see that dispatches dealing with 'core' state secrets, especially pertaining to information, capabilities and threats are restricted. We note that such subject matter accords with notions of 'high politics'—specifically, state security and survival—as described by Keohane and Nye (1977). On the other hand, cables that discuss more 'public

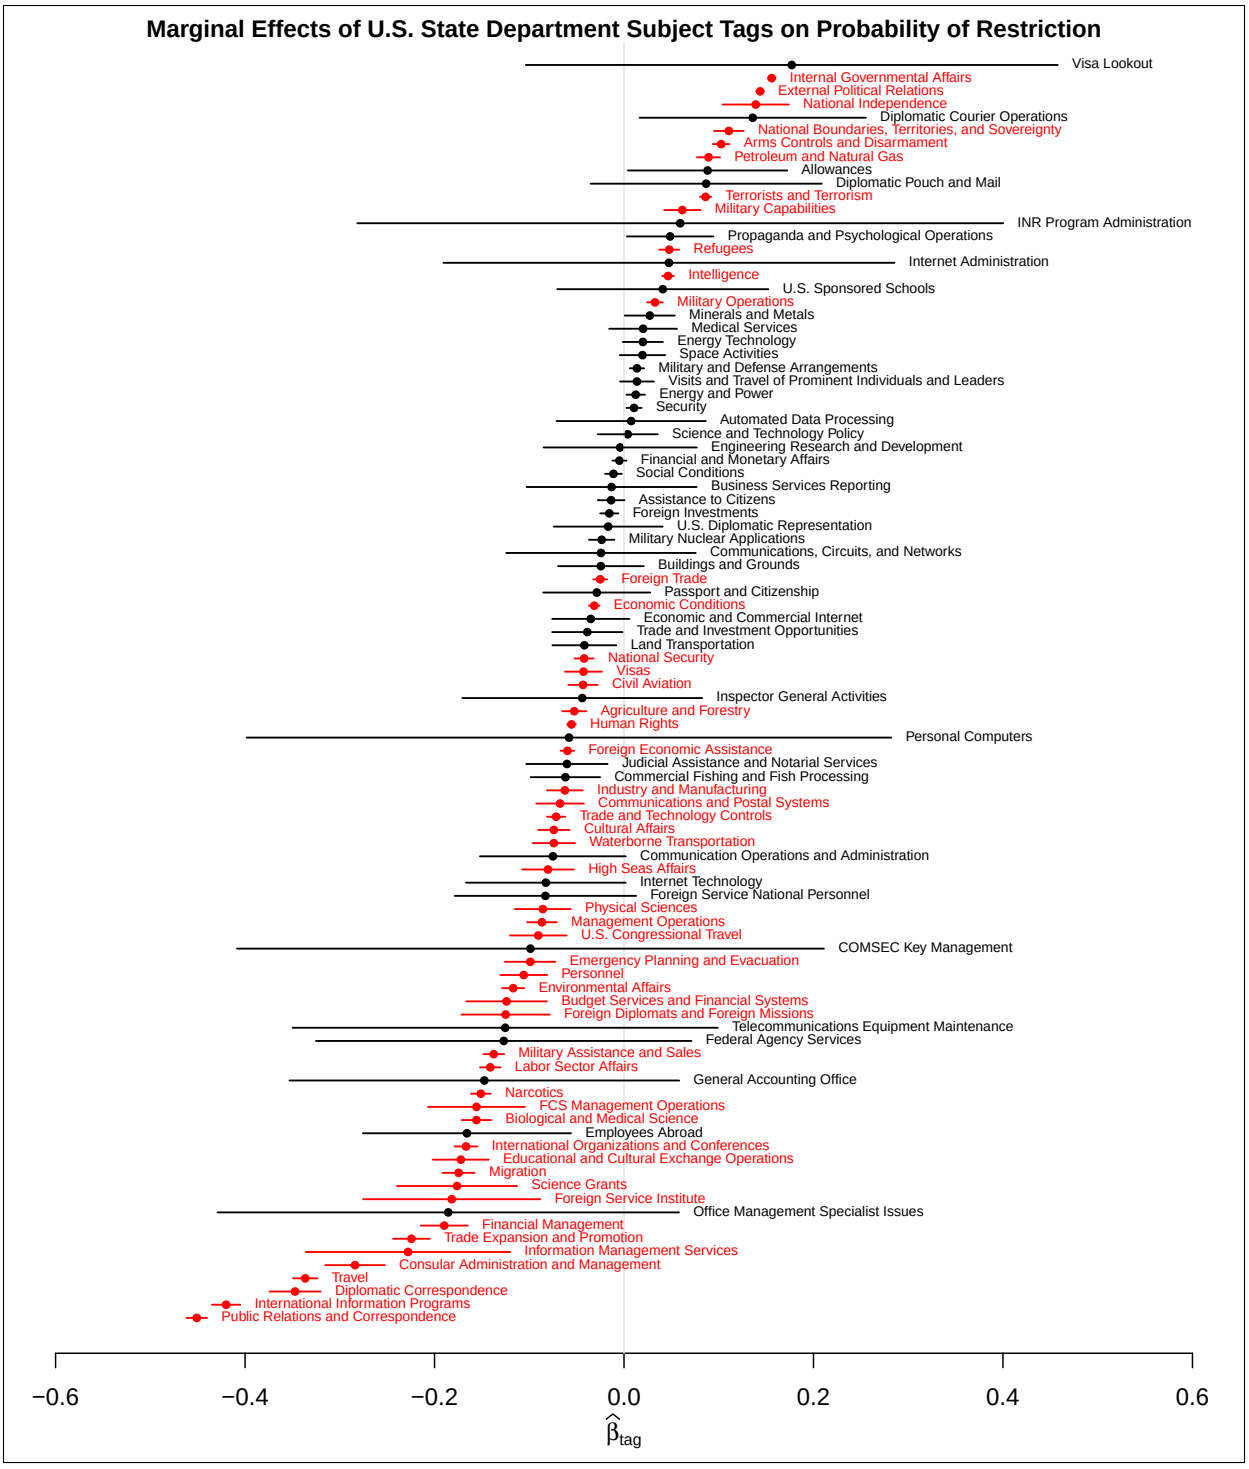

Figure 6: Substantive content as a predictor of secrecy status: estimates in red are statistically distinguishable from zero following the "Holm-Bonferroni" correction for multiple comparisons (Holm, 1979). 95% CIs around each estimate. All cables written in and after 2005 used in estimates, along with cable-origin fixed effects. The central line corresponds to a $\hat{\beta}_t$ of zero.

good' orientated matters—wherein we can imagine that sharing information may not be damaging, and may in fact be optimal—tend to be unrestricted. In this latter category are tags that seem to require or be synonymous with publicity and the dissemination of information: "International Information Programs", "Public Relations and Correspondence", "International Organizations and Conferences", "Educational and Cultural Exchange Operations" and so on. With respect to the work of Keohane and Nye (1977), we might see such matters as 'low politics': issues of more domestic or economic concern.

The fact that cable substance drives at least some part of diplomatic secrecy should not come as a surprise to theorists of rational diplomacy. As noted above, most contemporary theoretical treatments of crisis diplomacy concern agents' incentives to misrepresent their resolve, capabilities, or information in bargaining settings: our results here suggest the United States. acts in a way compatible with that logic.

5.2 Matched Sample Results: Procedural Secrecy

In terms of procedural secrecy, an overview of our main results may be found in Figure 7. Recall that we used the RF algorithm to identify the thirty 'most important' tokens in terms of their ability to discriminate between the unrestricted and restricted cables status of a document. In the second column of the plot, these are clearly seen and include words such as 'said', 'told', 'ambassador', 'want', 'note', 'meet', 'want', 'ask', 'discuss', 'concern', 'state', 'agre[e]' 'support', 'however', 'thank', 'request', 'possibl[e]', 'like' and so on. Our immediate observation is that in stark contrast to our tag regressions, these words do not connote substantive state secrets *per se*; rather, they refer to the holding of meetings and the general protocols of diplomatic exchange with foreign nationals. Related to this idea, note the presence of terms such as 'poloff' (the Embassy's Political Officer), 'usg' (United States Government) and 'minist' (minister): actors who we expect to be involved in daily embassy

interactions. On the left of the figure, we report the lasso (point) estimate associated with the terms. When these points are to the right of the vertical line, the use of that word (on average) increases the probability that a document is restricted; when the points are to the left, this suggests that the word is associated (on average) with a decrease in probability that a document is restricted. Examining this part of our results, we note that terms such as 'said' and 'told', 'request', 'like' are used disproportionately more in restricted cables. To us, this is evidence that once one controls for substantive area, secrecy is mostly about keeping meetings private and confidential, regardless of whether anything intrinsically 'secret' is being discussed.

To evaluate this intuition, we recorded the modal topic—i.e., for each word, the topic with the highest posterior probability from the topic model described earlier—in which our most influential words appeared. If we are correct that secrecy is partly about a norm of discretion rather than content, we would expect to see most of the terms mapping to a single (or perhaps a few) 'administrative' topic(s), rather than topics pertaining to matters of substantive import. On the right-hand side of the plot, we see this is almost entirely the case. There, the solid lines lead from each word to the topic it most likely belongs; the dashed lines are from each word to second most likely topic. We see first that with a few exceptions, all of the words 'belong' to the first, second, or third topics. Inspecting those more closely, we note that those topics generally consist of administrative nouns and verbs, rather than subjects of interest: thus, we find "said" in the first, second, and third topic as a leading word, while 'will' appears in the fourth topic. Importantly, the words that we have identified as discriminating between unrestricted and restricted cables do *not* appear alongside obviously substantive subject matter such as pertains to the Middle East (topic six or topic seven), the Pacific rim (topic 8), nuclear proliferation (topic 16) or Russian aggression (topic 14). Of course, we do see that some terms are likely to appear within certain substantive

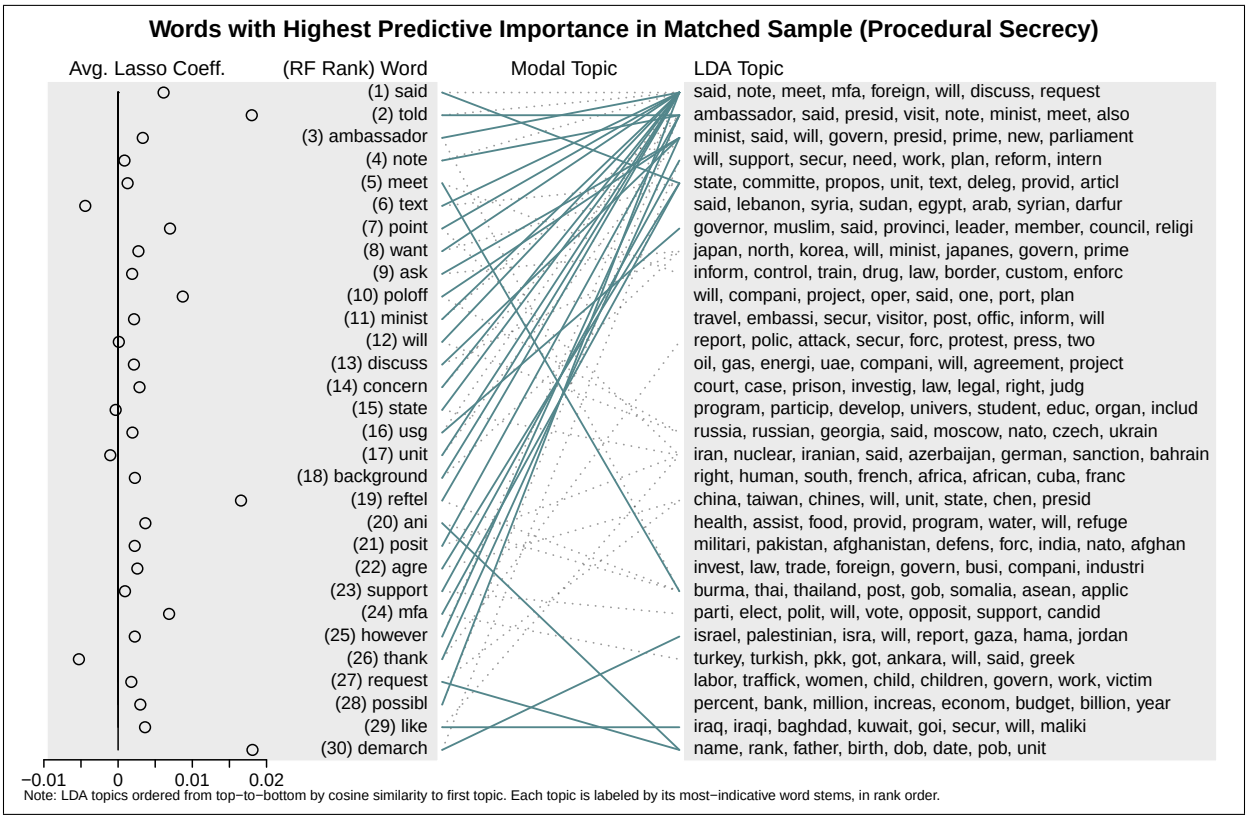

Figure 7: Procedural secrecy: matched sample results. For each of the words listed on the lefthand side of the plot, a solid line maps that word to its most likely topic (given estimates from the LDA model described in Appendix E). A dotted line maps each word to its second most likely topic. The topics listed on the righthand side of the graph are ordered in a specific manner: the uppermost topic is the mode of the modal topic assignments (i.e., the topic that is most frequently the modal topic assignment for the top RF terms), while subsequent topics are presented in descending order according to their similarity to the first topic. Topical similarity determined by the cosine similarity between topic vectors. The plot reveals remarkable concordance on the following: words that are most predictive of secrecy tend to be used in similar topics, and those topics tend to concern the official business of foreign leaders, their meetings, and words relating to information exchange. In general, these words also have positive positive lasso regression coefficient estimates, which implies their use on average is positively associated with cable restriction.

topics (such as 'meet', which appears in a 'Burma' topic and 'demarch[e]' which appears in an Israel topic towards the bottom of the plot). Such occurrences are not the norm, however.

In terms of the theories presented earlier, our finding here seems most closely compatible with the work of Sartori (2002) and Kurizaki (2007) insofar as privacy seems to be intrinsically valued by diplomats, rather than because it allows *per se* information exchange.

5.3 Share of Secrecy: Substance vs. Procedure

Above, we made the claim that while some of observed diplomatic censorship is a consequence of the need to protect state secrets, at least part of it results from the need to keep meetings confidential as a procedural requirement, regardless of what is to be discussed. In our final set of results, we attempt to estimate the relative contribution that these two separate elements make to the practice of restricting information from public view. In Figure 8 we report a comparison of models with this in mind. Here, 'Tags' refers to the tag covariates we noted earlier, 'Embassy' are simply embassy fixed effects, and 'Words' are the top 30 words selected by the Random Forest procedure above. In all cases, the numbers to the right of the bars refer to the percent correctly predicted (unrestricted and restricted) by a given (logit) model in the entire sample of 163,958 documents.

Unsurprisingly, we see that a model with tags, the word information, and the embassy fixed effects does best in terms of the proportion of documents it can classify correctly, at around 92%. The null model, the sample proportion of restricted cables is 59%, and clearly the statistical model improves substantially upon this. More interesting from our perspective is a comparison of the second and third bar ('Words + Embassy' and 'Tags + Embassy'), and the fifth and sixth ('Words' and 'Tags') since the performance of the models using the RF words and tags are so similar. That is, it seems that whether we use the substantive topics

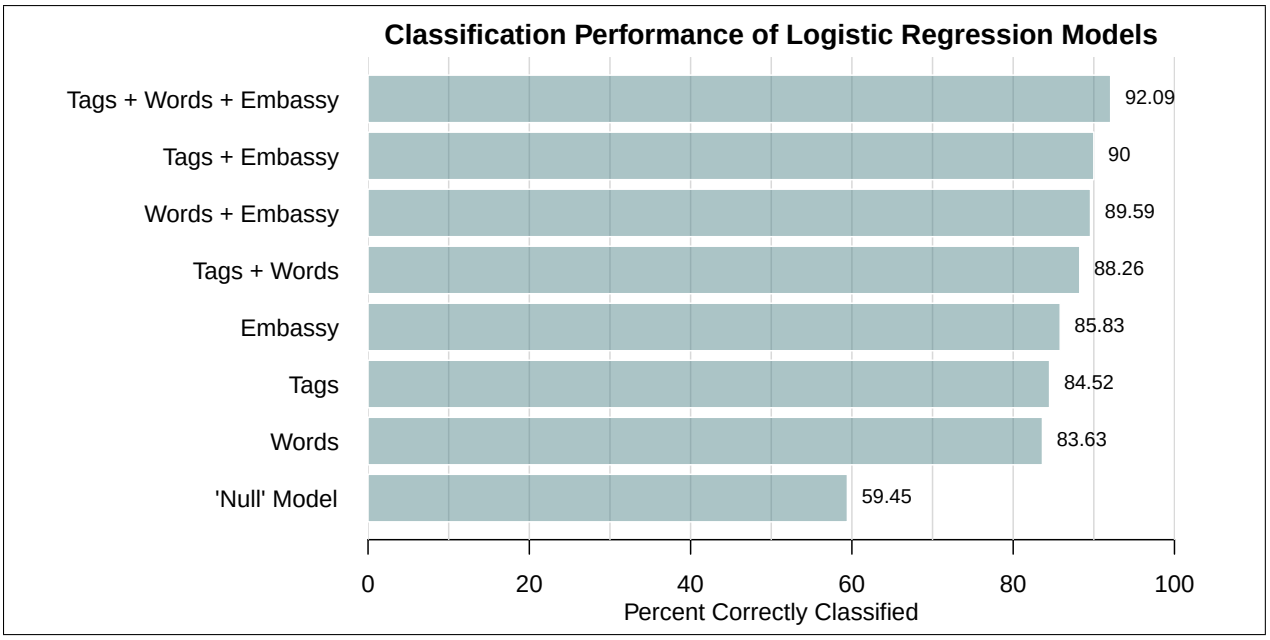

Figure 8: Comparison of classification performance for logistic regression models with 10-fold cross validation on the full post-2005 sample ($n = 163,958$). 'Tags' denote presence of subject tags, 'Words' the top 30 RF words, and 'Embassy' denotes inclusion of embassy IDs. The value of the 'null model' represents the proportion of the modal restriction status in the data set (i.e., restricted).

alone, or the words that we identified as connoting secret meetings rather than substance, our model performs similarly. This suggests, at the very least, that both substantive and procedural secrecy matter for diplomatic communication, and that both the audience cost theories and more recent work on communication have some support in the data.

6 Discussion

Conflict and bargaining have always been at the core of international relations and its study (see, e.g., Thucydides, 1910/431BCE; von Clausewitz, 1832/1989). In recent times, the discipline has amassed an impressive array of theoretical models that make use of, or provide findings for, 'information' and its dissemination between actors. This paper opened by noting that, despite this voluminous literature, there is little systematic statistical work on the subject, and that this is hardly surprising given that secrets—by definition—are difficult to research. In this paper, we made use of the WikiLeaks 'Cablegate' disclosure of diplomatic communications, a new and contemporary dataset that has an unusually large amount of 'uncensored' content (that was not systematically edited), to examine the empirical support for various conceptions of secrecy and communication. We argued that diplomatic confidentiality, i.e. information actively kept from the public, is used in at least two scenarios or 'dimensions': first, in a way pertaining to *substance* and second, pertaining to *procedure*. In the former case, documents that deal with issues that could damage U.S. capabilities were they available to others, are disproportionately kept secret. Meanwhile, in cases where publicity is helpful to the U.S. government are made available either for direct public consumption or for distribution to those who will have the opportunity to influence opinion. More speculatively, and untested here, information may be released because it creates a useful 'audience cost' and encourages commitment to a costly path of action. In the second case, that of the procedural dimension, diplomats ensure that the circumstances and process

of meetings in general—regardless of their actual subject content—are not disclosed. To be clear, we found evidence of both dimensions in our data, and were able to characterize their content and nature. In this way, both the recent literature that emphasizes the importance of diplomacy (e.g., Sartori, 2002; Kurizaki, 2007; Trager, 2010), and the earlier 'rationalist' literature that has it playing little role, finds some support here.

Apart from the preliminary analysis our paper provided, it also contributed methodologically to a growing area of political science: that of text analysis. In particular, we were faced with a situation in which documents had to be compared within particular subject areas, such that their discriminatory terms could be uncovered. We used an exact matching algorithm to get at our textual quantities of interest. In our case, the subject matter was determined by the U.S. State department (via the TAGS system), but the problem is obviously more general than this. For example, one might be interested in the success (or otherwise) of different bills in Congress or the public opinion reception of speeches from primary candidates. Clearly, the subject matter between documents differs and needs to be 'controlled' for in some sense. We provided one way of proceeding in such situations.

Of course, analytically, we have only scratched the surface here. Though we have documented the nature and structure of secrecy and the cables themselves, there is much more to do. First, while we argue that the 'more secret' topics in the TAGS system seem to deal more fully with capabilities than the 'least secret', we are necessarily vague on the details. We would like to know more about why exactly some subjects are kept from public view, and whether such decisions accord with IR theory in the area: for example, is topic secrecy actually dictated by a desire to avoid revealing capabilities on a particular subject, or is it more connected to notions of resolve, or even just the information dispensing machinery itself? Second, although we do not directly engage the plausibility of the audience cost lit-

erature and its critics (see Slantchev, 2012, for a review), our findings are at least minimally consistent with both sides of that debate, insofar as we find some evidence that the U.S. attempts to make more public its views (and thus possibly create such 'audience costs') where helpful, but not always. That is, it seems to preserve 'room for maneuver' in some areas. Subsequent analysis might weigh in more helpfully on this debate by considering the constraints that U.S. officials face in the various areas of international relations with which it deals: for example, we might be interested to know whether, in fact, issues that the United States is seemingly 'open' about with the public are simply those where it cannot be otherwise given commonly held knowledge about the U.S. position (or it weaknesses) in the wider world. This is ultimately a call to incorporate more topic-specific covariates and circumstances in the analysis. Finally, while we have emphasized the importance of private diplomatic meetings as part of the arsenal of U.S. international relations practice, we have done little to explain how or why they are used. That is, we are not much the wiser as to which of the various theories (Sartori, 2002; Kurizaki, 2007; Trager, 2010, e.g.) of diplomatic exchange is correct, if any. A continued digging into the WikiLeaks corpus might allow for more direct tests of these models. We leave such questions for future work.

A How Cables are Written and Classified

U.S. State Department cables are official communications between Foreign Service officials, embassies, consulates, and international organizations (such as the United Nations Headquarters in New York City, The Hague) around the world. These communications serve at least two purposes. First, they exist to share information regarding the daily proceedings of an embassy or nation with other institutions around the world, on topics broadly related to the overall interests of the United States. Second, they exist to create official records (i.e., a database) of political information relevant to particular foreign outposts over time. An early articulation of these twinned objectives can be found in Department of State (1974, http://aad.archives.gov/aad/content/aad_docs/rg59_state_dept_tags_74.pdf). When written, diplomatic cables take a standardized form. For example, the State Department provides glossaries of suggested language for officials to use in discussions of specific political issues in official communications. The "Termdex" chapter of 5 FAH-3 TAGS Terms Handbook serves this role, which "is an alphabetic list of words and phrases frequently found in Departmental communications"(1), requests officials use the terms "Border Incident" instead of "Border Violation" when discussing territorial disputes in official communications (5 FAH-3 H-810, pg. 9), or to use "Liberation Front" instead of "Liberation Movement" (5 FAH-3 H-810, pg. 40). So too, the Termdex recommends that particular phrases, if used, be used in tandem with specific subject TAGS—e.g., cables that discuss "Collective Bargaining" be tagged with the political tag "`ELAB`: Labor Sector Affairs", and discussions of "Collective Security" be tagged with "`MARR`: Military and Defense Arrangements" (13).

After a cable is written and political subject TAGS are marked in metadata, cables are assigned an overall restriction status. Generally speaking, for any single cable in our sample it is impossible to the full set of individuals who may be involved in a single communication's

restriction status, just as it is impossible to know if esoteric operational practices exist across U.S. embassies in the sample. That said, there is sufficient reason to believe a cable's restriction status is determined after a cable's text is complete, and generally this determination is made by a person of authority at the cable's location of origin (e.g., an Ambassador, or a Deputy Assistant Secretary, as articulated in "Original Classification Authority", 5 FAH-3 H-714.1, pg. 2). The Foreign Affairs Manual recommends a communication's "overall classification level is determined by the highest classification level of any of the portions" of its text. In other words, after a communication is written, an embassy's classification authority reviews all portions of a communication to check for the sensitivity of all portions of a message, where a "portion is ordinarily defined as a paragraph but also includes subject lines, titles, subheadings, tables, maps, photographs, graphs, and any other inserts within text" ("Classification Level", 5 FAH-3 H-713, pg. 2). This acknowledges that while not all portions of a cable may be equally sensitive, a cable's overall restriction status is set once all details of its contents have been reviewed. It is each cable's overall restriction status that we use for analysis in this study.

B Additional Details on Lasso and Random Forest

The RF and lasso procedures require brief explanation as they are not widely used in political research, although inevitably many technical details will be left for readers consult in the works cited. Both are widely used in "small n, large p" settings: cases in which there may be there may be a greater number of possible parameters than observations in the sample.

The RF algorithm is a decision tree and resampling-based classification procedure which relies on repeatedly dividing the observed sample of data into random bootstrapped training datasets and fitting decision trees to each random training set, then aggregating the classification results over all independent training sets. In the 'statistical learning' literature, this procedure is commonly referred to as bootstrapped aggregation (i.e., "bagging"), and can be widely applied to improve the classification precision of various models, regression included. A RF algorithm procedure deviates from bagging alone by also randomly sampling the parameter space included in each iteration of this bagging procedure (e.g., Ho, 1998). One result of procedures like RF is it allows researchers to think about the relative *variable importance* of predictors in a classification setting. Due to the fact that at each bagged iteration of the procedure there are random subsets of the feature space included in the decision-trees, not all predicting variables (i.e., "words" in our context) are likely to be included as predictors at each stage of the algorithm. Overall, a predictor's variable importance can be thought of as a result of this process: an estimate of the marginal reduction in classification error that results from a single word's inclusion to the classification procedure overall, given the random inclusion of other predictor variables from the sample.

The lasso is a form of penalized regression, similar to ridge regression, whereby regression coefficients are weighted by "shrinkage factors" such that regression coefficients are

weighted towards zero (Tibshirani, 1994; Hastie, Tibshirani and Friedman, 2009). The lasso is commonly used for feature selection in high-dimensional learning problems to decrease the variance of a particular classifier. In our context, the procedure is similar to an ordinary least squares regression procedure in which the best-model is determined by that which minimizes the in-sample sum of squared residuals, except regression coefficients are penalized according to prior rules (i.e., the shrinkage factor and tuning factor) on the minimum coefficient size a variable is allowed to have to be included in the final classification model. The estimates presented in Figure 7 are obtained from taking the average lasso coefficient for each word over a representative range of shrinkage factors, where each model is estimated at the embassy level using its exactly matched data subset.

C Cable Tags

	Meaning		Meaning
1) AADP	Automated Data Processing	51) ELTN	Land Transportation
2) ABLD	Buildings and Grounds	52) EMIN	Minerals and Metals
3) ABUD	Budget Services and Financial Systems	53) ENRG	Energy and Power
4) ACOA	Communication Operations and Administration	54) EPET	Petroleum and Natural Gas
5) ACKM	COMSEC Key Management	55) ETRD	Foreign Trade
6) ADCO	Diplomatic Courier Operations	56) ETTC	Trade and Technology Controls
7) ADPM	Diplomatic Pouch and Mail	57) EWWT	Waterborne Transportation
8) AEMR	Emergency Planning and Evacuation	58) MARR	Military and Defense Arrangements
9) AFIN	Financial Management	59) MASS	Military Assistance and Sales
10) AFSI	Foreign Service Institute	60) MCAP	Military Capabilities
11) AFSN	Foreign Service National Personnel	61) MNUC	Military Nuclear Applications
12) AGAO	General Accounting Office	62) MOPS	Military Operations
13) AINF	Information Management Services	63) ODIP	U.S. Diplomatic Representation
14) AINR	INR Program Administration	64) OEXC	Educational and Cultural Exchange Operations
15) AINT	Internet Administration	65) OFDP	Foreign Diplomats and Foreign Missions
16) ALOW	Allowances	66) OIIP	International Information Programs
17) AMED	Medical Services	67) OPDC	Diplomatic Correspondence
18) AMGT	Management Operations	68) OPRC	Public Relations and Correspondence
19) AMTC	Telecommunications Equipment Maintenance	69) OREP	U.S. Congressional Travel
20) ANET	Communications, Circuits, and Networks	70) OSCI	Science Grants
21) AODE	Employees Abroad	71) OTRA	Travel
22) AOMS	Office Management Specialist Issues	72) OVIP	Visits and Travel of Prominent Individuals and Leaders
23) AORC	International Organizations and Conferences	73) PARM	Arms Controls and Disarmament
24) APCS	Personal Computers	74) PBTS	National Boundaries, Territories, and Sovereignty
25) APER	Personnel	75) PGOV	Internal Governmental Affairs
26) ASCH	U.S. Sponsored Schools	76) PHSA	High Seas Affairs
27) ASEC	Security	77) PHUM	Human Rights
28) ASIG	Inspector General Activities	78) PINR	Intelligence
29) BBSR	Business Services Reporting	79) PINS	National Security
30) BEXP	Trade Expansion and Promotion	80) PNAT	National Independence
31) BMGT	FCS Management Operations	81) PREF	Refugees
32) BTIO	Trade and Investment Opportunities	82) PREL	External Political Relations
33) CASC	Assistance to Citizens	83) PROP	Propaganda and Psychological Operations
34) CFED	Federal Agency Services	84) PTER	Terrorists and Terrorism
35) CJAN	Judicial Assistance and Notarial Services	85) SCUL	Cultural Affairs
36) CLOK	Visa Lookout	86) SENV	Environmental Affairs
37) CMGT	Consular Administration and Management	87) SMIG	Migration
38) CPAS	Passport and Citizenship	88) SNAR	Narcotics
39) CVIS	Visas	89) SOCI	Social Conditions
40) EAGR	Agriculture and Forestry	90) TBIO	Biological and Medical Science
41) EAID	Foreign Economic Assistance	91) TINT	Internet Technology
42) EAIR	Civil Aviation	92) TNGD	Engineering Research and Development
43) ECON	Economic Conditions	93) TPHY	Physical Sciences
44) ECPS	Communications and Postal Systems	94) TRGY	Energy Technology
45) EFIN	Financial and Monetary Affairs	95) TSPA	Space Activities
46) EFIS	Commercial Fishing and Fish Processing	96) TSPL	Science and Technology Policy
47) EIND	Industry and Manufacturing		
48) EINT	Economic and Commercial Internet		
49) EINV	Foreign Investments		
50) ELAB	Labor Sector Affairs		

Figure 9: List of State Department "Subject TAGS" observed in sample and their meanings. Meanings taken from " Subject TAGS" in the *U.S. Department of State Foreign Affairs Manual* Volume 5, Handbook 3.

D Average Cable Restrictiveness by Embassy

Cable Totals and Restriction Frequency by Place of Origin: post−2005

Origin	# U	# R	% R	Origin	# U	# R	% R	Origin	# U	# R	% R
1) Embassy Baghdad	4970	823	0.14	90) Embassy Wellington	219	410	0.65	179) Mission USOSCE	108	16	0.13
2) Secretary of State	1772	3456	0.66	91) Embassy Maputo	183	442	0.71	180) REO Kirkuk	116	5	0.04
3) Embassy Tokyo	1845	3151	0.63	92) Embassy Brussels	386	224	0.37	181) Embassy Bujumbura	54	56	0.51
4) Embassy Ankara	2689	1787	0.40	93) Embassy Guatemala	186	392	0.68	182) Embassy Luxembourg	62	48	0.44
5) American Institute Taiwan, Taipei	1431	1502	0.51	94) Embassy Budapest	440	135	0.23	183) Consulate Frankfurt	17	88	0.84
6) Embassy Paris	1685	1248	0.43	95) Embassy Stockholm	300	274	0.48	184) Consulate Shenyang	97	7	0.07
7) Embassy Moscow	2172	449	0.17	96) Embassy Bucharest	434	137	0.24	185) Consulate Ho Chi Minh City	68	32	0.32
8) Embassy Tel Aviv	1245	1280	0.51	97) Embassy Lisbon	243	322	0.57	186) Embassy Lome	18	81	0.82
9) Embassy Beijing	1737	787	0.31	98) Embassy Panama	345	217	0.39	187) Embassy Grenada	0	97	1
10) Embassy Madrid	747	1705	0.70	99) Embassy Oslo	403	134	0.25	188) Consulate Kolkata	13	78	0.86
11) USUN New York	1068	1364	0.56	100) Embassy Conakry	394	136	0.26	189) Embassy Windhoek	49	39	0.44
12) Embassy Bangkok	1380	985	0.42	101) Embassy Djibouti	290	237	0.45	190) Consulate Rio De Janeiro	7	79	0.92
13) Embassy New Delhi	1546	773	0.33	102) Embassy San Jose	138	385	0.74	191) Consulate Vladivostok	0	84	1
14) Embassy Jakarta	1498	615	0.29	103) USEU Brussels	337	183	0.35	192) Consulate Monterrey	46	36	0.44
15) Embassy Kuwait	1037	1058	0.51	104) Consulate Sao Paulo	30	489	0.94	193) Consulate Karachi	75	4	0.05
16) Embassy Cairo	1574	413	0.21	105) Embassy Hanoi	186	332	0.64	194) Consulate Nogales	0	78	1
17) Embassy Beirut	1746	144	0.08	106) Embassy Riga	251	258	0.51	195) Embassy Mbabane	12	60	0.83
18) Embassy Kabul	1072	703	0.40	107) Embassy Doha	403	102	0.20	196) Consulate Halifax	0	71	1
19) Embassy Amman	1113	642	0.37	108) Embassy Maseru	107	389	0.78	197) Embassy Valletta	31	39	0.56
20) Consulate Jerusalem	1232	519	0.30	109) Embassy Dar Es Salaam	223	270	0.55	198) Consulate Lahore	60	7	0.10
21) Embassy Caracas	1563	162	0.09	110) Embassy London	309	172	0.36	199) Mission UNESCO	20	47	0.70
22) Embassy Seoul	1066	640	0.38	111) US Mission Geneva	224	256	0.53	200) Consulate Chennai	16	47	0.75
23) Embassy Dhaka	778	913	0.54	112) Mission USNATO	363	109	0.23	201) Embassy Port Moresby	39	23	0.37
24) Embassy Bogota	1032	636	0.38	113) Embassy Pristina	324	136	0.30	202) Consulate Mumbai	22	38	0.63
25) Embassy The Hague	702	942	0.57	114) Embassy Accra	211	247	0.54	203) Consulate Cape Town	14	45	0.76
26) Embassy Islamabad	1089	514	0.32	115) US Interests Section Havana	411	45	0.10	204) Consulate Munich	23	35	0.60
27) Embassy Mexico	308	1281	0.81	116) Embassy Asmara	406	49	0.11	205) Embassy Brazzaville	3	51	0.94
28) Embassy Colombo	1081	453	0.30	117) Embassy Ndjamena	241	212	0.47	206) Consulate Dusseldorf	0	51	1
29) Embassy Buenos Aires	636	881	0.58	118) Embassy Abidjan	318	133	0.29	207) REO Mosul	49	2	0.04
30) Embassy Khartoum	1093	411	0.27	119) Consulate Adana	0	450	1	208) Consulate Hamburg	8	41	0.84
31) Embassy Abuja	903	431	0.32	120) Embassy Nouakchott	391	55	0.12	209) Consulate St Petersburg	0	45	1
32) Embassy Ashgabat	1023	305	0.23	121) Embassy Singapore	267	156	0.37	210) Consulate Surabaya	0	42	1
33) Embassy Baku	1177	145	0.11	122) Consulate Istanbul	287	128	0.31	211) Embassy Praia	5	35	0.88
34) Embassy Kathmandu	1062	253	0.19	123) Embassy Tripoli	352	62	0.15	212) Consulate Toronto	9	29	0.76
35) Embassy Vienna	503	812	0.62	124) Embassy Helsinki	224	189	0.46	213) US Delegation, Secretary	30	6	0.17
36) Embassy Nairobi	673	627	0.48	125) Embassy Kigali	274	128	0.32	214) Consulate Johannesburg	2	33	0.94
37) Embassy Berlin	936	331	0.26	126) Embassy Freetown	209	192	0.48	215) Embassy Podgorica	0	35	1
38) Embassy Damascus	810	415	0.34	127) Embassy Belgrade	144	254	0.64	216) Consulate Montreal	8	23	0.74
39) Embassy Manila	776	432	0.36	128) Mission Geneva	354	43	0.11	217) Embassy Kolonia	20	11	0.35
40) Embassy Kinshasa	702	497	0.41	129) Embassy Dakar	255	138	0.35	218) Consulate Quebec	14	16	0.53
41) Embassy Rome	807	359	0.31	130) Embassy Tallinn	173	209	0.55	219) Consulate Guadalajara	2	27	0.93
42) Embassy Rangoon	1059	82	0.07	131) UNVIE	271	103	0.28	220) Consulate Naha	25	4	0.14
43) Embassy Manama	967	171	0.15	132) Embassy Suva	280	67	0.19	221) Embassy Bangui	2	24	0.92
44) Embassy Santiago	274	861	0.76	133) Consulate Jeddah	298	47	0.14	222) Consulate Milan	18	6	0.25
45) Embassy Muscat	585	526	0.47	134) Embassy Canberra	184	161	0.47	223) Consulate Thessaloniki	0	24	1
46) Embassy Abu Dhabi	700	394	0.36	135) Embassy Lilongwe	108	231	0.68	224) Consulate Guayaquil	18	4	0.18
47) Embassy Bridgetown	271	778	0.74	136) Embassy Niamey	82	252	0.75	225) Consulate Strasbourg	16	6	0.27
48) Embassy Pretoria	426	608	0.59	137) Embassy Skopje	191	143	0.43	226) Consulate Vancouver	5	17	0.77
49) Embassy Tashkent	853	169	0.17	138) Embassy Vatican	271	63	0.19	227) UN Rome	0	21	1
50) Consulate Lagos	707	313	0.31	139) Embassy Bamako	216	115	0.35	228) Consulate Yekaterinburg	0	20	1
51) Embassy Harare	750	268	0.26	140) Embassy Paramaribo	60	268	0.82	229) Consulate Durban	11	7	0.39
52) Embassy Tegucigalpa	616	399	0.39	141) Embassy Montevideo	169	152	0.47	230) Consulate Dhahran	12	5	0.29
53) Embassy Asuncion	391	623	0.61	142) Embassy Phnom Penh	127	194	0.60	231) US Mission CD Geneva	16	1	0.06
54) Embassy Yerevan	775	230	0.23	143) REO Basrah	288	23	0.07	232) Consulate Curacao	3	13	0.81
55) Embassy Brasilia	374	626	0.63	144) Embassy Nassau	154	145	0.48	233) Consulate Tijuana	4	10	0.71
56) Embassy Santo Domingo	229	756	0.77	145) Embassy Georgetown	92	196	0.68	234) Consulate Calgary	0	13	1
57) Embassy Tbilisi	722	255	0.26	146) Embassy Tirana	153	125	0.45	235) Consulate Naples	8	5	0.38
58) Embassy Athens	662	310	0.32	147) Consulate Shanghai	247	26	0.10	236) Consulate Barcelona	2	10	0.83
59) Embassy La Paz	770	181	0.19	148) Embassy Dili	127	128	0.50	237) Consulate Auckland	0	11	1
60) Embassy San Salvador	354	578	0.62	149) Embassy Antananarivo	155	96	0.38	238) Consulate Ciudad Juarez	0	9	1
61) Embassy Port Au Prince	475	433	0.48	150) Embassy Chisinau	212	33	0.13	239) Consulate Florence	2	7	0.78
62) Embassy Riyadh	714	192	0.21	151) REO Hillah	203	39	0.16	240) Consulate Recife	1	8	0.89
63) Embassy Managua	581	313	0.35	152) US Office Almaty	121	117	0.49	241) Embassy Majuro	2	7	0.78
64) Embassy Sanaa	593	298	0.33	153) Embassy Yaounde	141	87	0.38	242) US Delegation FEST TWO	5	4	0.44
65) Embassy Addis Ababa	607	282	0.32	154) Embassy Port Of Spain	80	140	0.64	243) Consulate Fukuoka	1	7	0.88
66) Embassy Kyiv	761	122	0.14	155) Embassy Cotonou	30	189	0.86	244) Consulate Sydney	4	4	0.50
67) Embassy Ottawa	473	404	0.46	156) Embassy Kampala	118	96	0.45	245) Consulate Leipzig	0	7	1
68) Embassy Rabat	590	282	0.32	157) Consulate Chiang Mai	114	98	0.46	246) Consulate Hermosillo	0	6	1
69) Embassy Kingston	257	599	0.70	158) Embassy Vientiane	127	72	0.36	247) Consulate Melbourne	5	1	0.17
70) Embassy Quito	598	256	0.30	159) Embassy Libreville	141	57	0.29	248) Consulate Perth	4	2	0.33
71) Embassy Zagreb	363	480	0.57	160) Embassy Lusaka	114	82	0.42	249) Embassy Koror	0	6	1
72) Embassy Prague	566	268	0.32	161) Embassy Banjul	163	26	0.14	250) Consulate Belfast	4	1	0.20
73) Embassy Minsk	509	312	0.38	162) Embassy Gaborone	74	112	0.60	251) Consulate Hamilton	0	5	1
74) Embassy Bishkek	637	154	0.19	163) Embassy Luanda	119	63	0.35	252) Consulate Marseille	1	4	0.80
75) Embassy Lima	376	381	0.50	164) Embassy Ulaanbaatar	91	91	0.50	253) Consulate Sapporo	1	4	0.80
76) Embassy Warsaw	586	170	0.22	165) Consulate Dubai	149	27	0.15	254) Embassy Apia	0	5	1
77) Embassy Bratislava	482	266	0.36	166) Consulate Guangzhou	92	84	0.48	255) Consulate Matamoros	0	4	1
78) Embassy Kuala Lumpur	517	209	0.29	167) Embassy Bern	144	29	0.17	256) Embassy Malabo	0	4	1
79) Embassy Vilnius	426	295	0.41	168) Embassy Copenhagen	132	39	0.23	257) Consulate Nuevo Laredo	0	3	1
80) Embassy Sofia	465	240	0.34	169) Iran RPO Dubai	163	2	0.01	258) Consulate Osaka Kobe	2	1	0.33
81) Embassy Tunis	520	182	0.26	170) Consulate Peshawar	158	3	0.02	259) Consulate Amsterdam	0	2	1
82) Embassy Sarajevo	542	138	0.20	171) Embassy Belmopan	41	119	0.74	260) Consulate Merida	0	2	1
83) Embassy Algiers	560	118	0.17	172) Consulate Chengdu	144	8	0.05	261) ** Dhahran	1	0	0
84) Consulate Hong Kong	308	361	0.54	173) Embassy Reykjavik	78	73	0.48	262) American Consulate Hyderabad	0	1	1
85) Embassy Nicosia	459	210	0.31	174) Embassy Port Louis	74	74	0.50	263) Consulate Krakow	0	1	1
86) Embassy Astana	291	374	0.56	175) Embassy Monrovia	73	74	0.50	264) Department of State	0	1	1
87) Embassy Dushanbe	350	307	0.47	176) Embassy Bandar Seri Begawan	91	48	0.35	265) DIR FSINFATC	0	1	1
88) Embassy Ljubljana	420	216	0.34	177) Consulate Casablanca	70	66	0.49	266) US Office FSC Charleston	0	1	1
89) Embassy Dublin	348	286	0.45	178) Embassy Ouagadougou	72	63	0.47				

Note: locations in red indicate cable 'senders' with at least 100 restricted and 100 unrestricted cables in sample.

Figure 10: Frequency of restricted versus unrestricted cables by place of origin, for all cables post-2005.

E Probabilistic Topic Model

The field of quantitative text analysis has grown substantially in recent years. In this literature, applied researches extensively use Latent Dirichlet Allocation (Blei, Ng and Jordan, 2003) as a generative model to extract "themes" or "topics" from a collection of documents. See, for example, Blei (2012) for an overview, and Quinn et al. (2010) for recent political science applications of topic models. The model assumes documents are composed of latent topics that are chosen with probabilities following a Dirichlet distribution, and multinomial choice probabilities for word choice conditional on a topic. More precisely, the framework from Blei, Ng and Jordan (2003) has the number of words N in a document be Poisson(ξ), the latent topic probabilities θ be Dirichlet(α), the topics z_n be Multinomial(θ), and the words w_n be Multinomial(β), conditional on z_n. Then, with M documents, they have that $p(\mathcal{C}|\alpha,\beta) = \prod_{d=1}^{M} \int p(\theta_d|\alpha) \left(\prod_{n=1}^{N_d} \sum_{z_{dn}} p(z_{dn}|\theta_d) p(w_{dn}|z_{dn},\beta) \right) d\theta_d$. Computational difficulties arise in this setting, but there are ways to deal with them (e.g., Hoffman et al., 2013).

When a researcher estimates an LDA model, the topics returned are characterized by the multinomial probabilities for all words within each topic, as well as the posterior distribution of topics conditional on a certain word. In practice, the researcher selects the number of topics a priori, although recent efforts have been made to assess how an approximate number of topics may be present in a sample of data (see, e.g., Hoffman et al., 2013). We used these posterior estimates to generate the topical ordering in Figure 7: for each topic, we took the cosine similarity (e.g., Manning, Raghavan and Schütze, 2008) between all pairs of topics, where similarity between topic vectors is determined by the posterior weights placed on each word in each topic.

F Exact Matching Algorithm

1. Let N_j be the set of cables from embassy j that occur during or after the year 2005 in the sample, where $|N_j|$ is the number of cables originating from location j.
2. For each of the $|N_j|$ documents in the sample, record the subject tags present on each diplomatic cable.
3. For all restricted cables in N_j, find all unrestricted cables in N_j that exactly match on subject tags and year of creation.
4. From the subset of restricted cables in N_j with at least one unrestricted exact match,
   (a) Randomly draw a restricted cable and find the unrestricted, exact-matched cable that is written most closely in time (i.e., the cable that minimizes the absolute value between the difference in release days). Each cable may be matched with or without replacement.
   (b) Continue this process until there are no-more restricted and unrestricted cables to pair together.
5. Record the list of exactly-matched pairs of cables, if applicable.

Figure 11: Outline of Exact Matching Algorithm

References

Allen, David and Matthew Connelly. 2015. Diplomatic History After the Big Bang: Using Computational Methods to Explore the Infinite Archive. In *Explaining the History of American Foreign Relations*, ed. Frank Costigliola and Michael Hogan. 3rd ed. New York: Cambridge University Press.

Angeletos, George-Marios and Alessandro Pavan. 2004. "Transparency of Information and Coordination in Economies with Investment Complementarities." *The American Economic Review* 94(2):91–98.

Angrist, Joshua D. and Jörn-Steffan Pischke. 2009. *Mostly Harmless Econometrics: An Empiricists Companion*. Princeton University Press.

Blei, David M. 2012. "Probabilistic Topic Models." *Communications of the ACM* 55(4):77–84.

Blei, David M., Andrew Y. Ng and Michael I Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 2(4-5):993—1022.

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

Cornand, Camille and Frank Heinemann. 2008. "Optimal Degree of Public Information Dissemination." *The Economic Journal* 118:718–742.

Cox, David R. 1958. *Planning of Experiments*. Wiley Publications in Statistics.

Cukierman, Alex. 2001. Accountability, credibility, transparency, and stabilization policy in the Eurosystem. In *The Impact of the EMU on Europe and Developing Countries*, ed. C. Wyplosz. Oxford University Press.

Fearon, James. 1994*a*. "Domestic Political Audiences and the Escalation of International Disputes." *American Political Science Review* 88:577–92.

Fearon, James. 1994*b*. "Signaling Versus the Balance of Power and Interests: An Empirical Test of a Crisis Bargaining Model." *Journal of Conflict Resolution* 38(2):236–69.

Fey, Mark and Kristopher Ramsay. 2011. "Uncertainty and Incentives in Crisis Bargaining: Game-Free Analysis of International Conflict." *American Journal of Political Science* 55(1):149–169.

Frieden, Jeffry and David Lake. 2005. "International Relations as a Social Science: Rigor and Relevance." *Annals of the American Academy of Political and Social Science* 600(1):136–156.

Ghosn, Faten, Glenn Palmer and Stuart Bremer. 2004. "The MID3 Data Set, 1993–2001: Procedures, Coding Rules, and Description." *Conflict Management and Peace Science* 21(2):133–154.

Gill, Michael and Arthur Spirling. Forthcoming. "Estimating the Severity of the Wikileaks United States Diplomatic Cables Disclosure." *Political Analysis* .

Goemens, Hein. 2000. *War and Punishment: The Causes of War Termination and the First World War*. Princeton, NJ: Princeton University Press.

Grimmer, Justin and Brandon Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3):267–27.

Hastie, Trevor, Robert Tibshirani and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. Springer.

Heckman, James J. and Richard Robb. 1985. "Alternative Methods for Evaluating the Impact of Interventions." *Journal of Econometrics* 30(1):239–267.

Ho, Tin. 1998. "The Random Subspace Method for Constructing Decision Forests." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(8):832–844.

Hoffman, Matthew D., David M. Blei, Chong Wang and John Paisley. 2013. "Stochastic Variational Inference." *Journal of Machine Learning Research* 14:1303–1347.

Holm, Sture. 1979. "A simple sequentially rejective multiple test procedure." *Scandinavian Journal of Statistics* 6(2):65–70.

Iacus, Stefano M., Gary King and Giuseppe Porro. 2011. "Multivariate Matching Methods That Are Monotonic Imbalance Bounding." *Journal of the American Statistical Association* 106(493):345–361.

Imai, Kosuke, Gary King and Elizabeth A. Stuart. 2008. "Misunderstandings between experimentalists and observationalists about causal inference." *Journal of the Royal Statistical Society, Series A* 171(2):481–502.

Jervis, Robert. 1976. *Perception and Misperception in International Politics*. Princeton, N.J.: Princeton University Press.

Keohane, Robert and Joseph Nye. 1977. *Power and Interdependence: World Politics in Transition*. Boston: Little, Brown.

Kurizaki, Shuhei. 2007. "Efficient Secrecy: Public versus Private Threats in Crisis Diplomacy." *American Political Science Review* 101(3).

Lake, David. 2010. "Two Cheers for Bargaining Theory: Assessing Rationalist Explanations of the Iraq War." *International Security* 35(3):7–52.

Manning, Christopher D., Prabhakar Raghavan and Hinrich Schütze. 2008. *An Introduction to Information Retrieval*. Cambridge University Press.

Partell, Peter and Glenn Palmer. 1999. "Audience Costs and Interstate Crises: An Empirical Assessment of Fearon's Model of Dispute Outcomes." *International Studies Quarterly* 43(2):389–405.

Porter, Martin F. 1980. "An Algorithm for Suffix Stripping." *Program* 14(3):130–137.

Potter, Philip and Matthew Baum. 2014. "Looking for Audience Costs in all the Wrong Places: Electoral Institutions, Media Access, and Democratic Constraint." *The Journal of Politics* 76:167–181.

Powell, Robert. 1999. *In the Shadow of Power,*. Princeton, NJ: Princeton University Press.

Powell, Robert. 2002. "Bargaining Theory and International Conflict." *Annual Review of Political Science* 5:1–30.

Quinn, Kevin, Burt Monroe, Michael Colaresi, Michael H. Crespin and Dragomir Radev. 2010. "How to Analyze Political Attention with Minimal Assumptions and Costs." *American Journal of Political Science* 54:209–228.

Ramsay, Kristopher. 2008. "Settling It on the Field: Battlefield Events and War Termination." *Journal of Conflict Resolution* 52(6):850–879.

Ramsay, Kristopher. 2011. "Cheap Talk Diplomacy, Voluntary Negotiations, and Variable Bargaining Power." *International Studies Quarterly* 55:1003–1023.

Reed, William. 2003. "Information and Economic Interdependence." *Journal of Conflict Resolution* 47(1):54–71.

Reiter, Dan. 2009. *How Wars End*. Princeton: Princeton University Press.

Renshon, Jonathan, Jooa Julia Lee and Dustin Tingley. Forthcoming. "Physiological Arousal and Political Beliefs." *Political Psychology* .

Rosenbaum, Paul R. 2002. *Observational Studies*. 2nd ed. Springer.

Rosenbaum, Paul R. and Donald B. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70(1):41–55.

Rosenbaum, Paul R. and Donald B. Rubin. 1984. "Reducing Bias in Observational Studies Using Subclassification on the Propensity Score." *Journal of the American Statistical Association* 79(387):516–524.

Rubin, Donald B. 1976. "Multivariate Matching Methods That Are Equal Percent Bias Reducing, I: Some Examples." *Biometrics* 32(1):109–120.

Sartori, Anne. 2002. "The Might of the Pen: A Reputational Theory of Communication in International Disputes." *International Organization* 56(1):121–149.

Sartori, Anne. 2005. *Deterrence by Diplomacy*. Princeton: Princeton University Press.

Schultz, Kenneth. 2001. "Looking for Audience Costs." *Journal of Conflict Resolution* 45:32–60.

Shapiro, Jacob and David Siegel. 2010. "Is this Paper Dangerous? Balancing Secrecy and Openness in Counterterrorism." *Security Studies* 19(1):66–98.

Signorino, Curtis. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review* 93(2):279–297.

Slantchev, Branislav. 2004. "How Initiators End Their Wars: The Duration of Warfare and the Terms of Peace." *American Journal of Political Science* 48(4):813–829.

Slantchev, Branislav. 2012. "Audience Cost Theory and Its Critics." *Security Studies* 21:376–382.

Smith, Alistair and Allan Stam. 2004. "Bargaining and the Nature of War." *Journal of Conflict Resolution* 48(6).

Snyder, Jack and Erica D. Borghard. 2011. "The Cost of Empty Threats: A Penny, Not a Pound." *American Political Science Review* 105(3).

Stuart, Elizabeth A. 2010. "Matching methods for causal inference: A review and a look forward." *Statistical Science* 25(1):1–21.

Tarar, Ahmer and Behar Leventoglu. 2009. "Public Commitment in Crisis Bargaining." *International Studies Quarterly* 53:817–839.

Thucydides. 1910/431BCE. *The Peloponnesian War*. London: J. M. Dent.

Tibshirani, Robert. 1994. "Regression Shrinkage and Selection Via the Lasso." *Journal of the Royal Statistical Society, Series B* 58:267–288.

Tomz, Michael. 2007. "Domestic Audience Costs in International Relations: An Experimental Approach." *International Organization* 61:821–40.

Trachtenberg, Marc. 2012. "Audience Costs: An Historical Analysis." *Security Studies* 21(1):3–42.

Trager, Robert. 2010. "Diplomatic Calculus in Anarchy: How Communication Matters." *American Political Science Review* 104(2).

von Clausewitz, Carl. 1832/1989. *On War*. Princeton: Princeton University Press.

Waltz, Kenneth. 1979. *Theory of International Politics*. New York: McGraw-Hill.

Weeks, Jessica. 2008. "Autocratic Audience Costs: Regime Type and Signaling Resolve." *International Organization* 62(1).

Wendt, Alexander. 1999. *Social Theory of International Politics*. New York: Cambridge University Press.

Werner, Suzanne. 1999. "The Precarious Nature of Peace: Resolving the Issues, Enforcing the Settlement, and Renegotiating the Terms." *American Journal of Political Science* 43(3):912–934.

Woodford, Michael. 2005. "Central Bank Communication and Policy Effectiveness." *NBER Working Paper No. 11898* .