From diffusion to diffuse-ability:
A text-as-data approach to explaining the global diffusion of Corporate Sustainability Policy

Forthcoming in *International Studies Quarterly*

Adam William Chalmers[1]
Robyn Klingler-Vidra[2]
Onna Malou van den Broek[3]

Abstract
This paper argues that attributes of diffusion objects, in their own right, shape the form and extent of policy diffusion. To date, diffusion scholarship focuses on actor-level attributes (e.g., connections, culture, physical proximity, etc.) to explain what is diffused, and how much. Extending existing theory on the impacts of policies' textual properties on diffusion patterns, we argue that policies that are easier to understand, specific in their applicability, and that do not mandate specific behavior have their text diffused with less adaptation, regardless of the attributes of the authoring organization. We test our argument in the context of the global diffusion of Corporate Sustainability Policy (CSP), analyzing a novel dataset of 1,429 CSPs from 100 countries, 20 international organizations, and 12 regional organizations over a 65-year period. Offering a precise measure of diffusion as the extent to which a source text is copied into an adopter text, we find statistical support for our hypothesis. We contribute to diffusion scholarship by helping to mainstream natural language processing methods and by theorizing how attributes of policy documents affect how much adaptation occurs.

Keywords: diffusion object, corporate sustainability policy, policy diffusion, text-as-data

Acknowledgements
Research for this project would not have been possible without generous funding from the King's College London Faculty Research Fund, the King's Together Multi- and Interdisciplinary Research Grant, and the John Templeton Foundation. We are grateful to Ruben Ruiz Rufino, Christine Reh, Matia Vannoni, Christel Koop and Dimiter Toshkov for their comments on earlier drafts of this article as well as to the participants of the Comparative Politics Research in Progress Workshop at King's College London (28.01.2021), the Hertie School's EU Governance Colloquium (01.03.2021), the 5th International Conference on Public Policy (08.07.2021) and the 37th EGOS (European Group for Organizational Studies) Colloquium (09.07.2021). Finally, we would like to thank our talented and hardworking team of research assistants: Anais Julin, Karissa Shah, Sixtine Gaudry, Julie Eriksen, Lucie Skopkova, and Gaurav Kathuria.

[1] adam.w.chalmers@ed.ac.uk
[2] robyn_klingler.vidra@kcl.ac.uk
[3] O.M.Van-Den-Broek@exeter.ac.uk

Introduction

To what extent does policy diffusion depend on the attributes of the object being diffused? Existing research largely explains diffusion in terms of the actors involved: potential adopters' evaluations of available models (Acharya 2004), characteristics of the intermediaries involved (van den Broek and Klingler-Vidra 2022), and the spatial proximity (Elkins and Simmons 2005; Shipan and Volden 2008; Weyland 2007) and network connections (Simmons, Lloyd, and Stewart 2018; Sommerer and Tallberg 2019) between actors. By contrast, relatively little attention has been paid to how the *nature* of the diffusion object itself affects the diffusion process.

This oversight has not gone unnoticed. Rogers's seminal work ([1962] 2003, 219) on the diffusion of innovation explains that while we know a lot about the characteristics of adopters, little effort has been devoted to analyzing "innovation differences" (i.e., the attributes of the object). Makse and Volden (2011, 108), writing nearly 50 years later, point out that the "literature on diffusion [still] suffers from the exact malaise that Rogers first described." This insufficient attention to policies' attributes reduces the transferability of diffusion theories and ultimately limits the accumulation of knowledge. Today, only a handful of policy diffusion studies pay more attention to the analytical role that attributes, not actors, play in shaping diffusion outcomes (e.g., Karsh et al. 2016; Weyland 2007).

In this paper we address this gap by exploring how much a policy is diffused (i.e., the degrees of textual copying between a source policy and an adopter policy) as a function of the attributes of the diffusion object. Theorizing diffusion object attributes in terms of policies' textual and linguistic characteristics, our focus is on written policy documents rather than general policy ideas, models, or principles (see Radaelli 2005; Weyland 2007). To draw a parallel, we are interested in the diffusion of the text in the Constitution of the United States rather than the wider model of presidential democracy. Accordingly, our analytical framework draws on affordance theory (Gibson 1979) and extends diffusion studies that have focused on policy characteristics (Karsh et al. 2016; Makse and Volden 2011; Mooney and Lee 1995, 2000; Nicholson-Crotty 2009; Weyland 2007).

We conceive of diffusion as "more or less convergence" (Klingler-Vidra and Schleifer 2014), in which degrees of textual copying occur, ranging from loose paraphrasing by changing words and/or their order to direct, word-for-word copy–pasting of text. We then argue that there are three main textual attributes that shape how much text is used in other policy documents: (1) complexity, (2) flexibility, and (3) restrictiveness. Complexity refers to the readability of a policy document, an attribute that creates a higher cognitive burden for the potential adopter and therefore reduces diffusion. Flexibility refers to the use of general and equivocating policy language that allows for policy adaptation, translation, and localization and hence reduces diffusion. Finally, restrictiveness captures the language of compellence, which is defined as language that compels and deters action (sometimes called "hard and soft law") and, insofar as it legally binds actors to specific behaviors, results in less diffusion.

We test our arguments using a novel dataset of 1,429 Corporate Sustainability Policy (CSP)[4] documents from 100 countries, 20 international organizations (IOs), and 12 regional organizations over a 65-year period. While existing research has explored the diffusion of firms' CSP strategies and activities (e.g., Knudsen and Moon, 2017; Pope and Lim 2020), ours is the first project to assess the multilevel diffusion of CSPs as issued by national governments,

[4] CSPs include guidelines and legislation pertaining to "corporate responsibility," "corporate social responsibility," "environment, social, governance," "ESG," "materiality," "non-financial materiality," "shared value," and "social value." They do not include the broader suite of labor-related governance policies, such as "industrial relations," "labor reforms," and "labor regulation." Our understanding is consistent with the approaches used by the databases we aggregated for this study, including Carrots and Sticks, the European Corporate Governance Institute (ECGI), Principles for Responsible Investment (PRI), and the Sustainable Stock Exchange Initiative (SSE), as well as leading CSP scholarship (Knudsen and Moon 2017; Pope and Liam 2020).

IOs, regional organizations, and multi-stakeholder initiatives. CSPs have some unique features as a policy area, given the extent to which these efforts involve multiple stakeholders and operate across levels of governance. However, CSPs are also representative of other policy areas incentivizing firm behavior in particular directions, such as industrial policy.

We analyze CSPs using state-of-the-art natural language processing (NLP) methods to develop a data-driven, quantitative measure of how attributes of policy documents affect how much text diffuses. We operationalize diffusion in terms of textual and linguistic similarity: namely, how much of the text in the source policy document is *copied* into the adopted policy. This allows us to systematically measure the extent to which textual diffusion occurs and, hence, understand why some CSP text diffuses more than others. Leveraging NLP techniques helps us get to the heart of what we mean when we say that "something diffuses." It also helps us to overcome the problem of confusing diffusion with the coincidence of two or more actors having a "similar response to a common stimulus or shock" (Simmons, Lloyd, and Stewart 2018, 255), as we can measure similarity on a word-by-word basis. Measuring diffusion in terms of *degrees of textual copying* helps establish the intentionality of the adopter.

We find strong evidence that policy document attributes – when controlling for key characteristics of the actors involved – shape the extent to which text is diffused. More precisely, there is more diffusion (i.e., textual copying) when the text of policy documents is not overly complex and does not legally bind the adopter into specific actions or behaviors. In contrast to our expectation that textual ambiguity diminishes diffusion, we find instead that it increases it: policies replete with ambiguous language are replicated "as is" by adopters. Our findings are robust across numerous regression model specifications, different methods for measuring text similarity, and alternative indicators measuring policy document attributes.

We make two main empirical and theoretical contributions. First, we advance the analytical tools of policy diffusion by helping to bring theories centering on the attributes of diffusion objects into mainstream diffusion research. We do so in a way that is distinct from existing approaches: by conceptualizing and empirically testing how the inherent attributes of the policy documents shape textual diffusion. We provide empirical evidence that policy document attributes have independent analytical power when explaining how much text diffuses. The key implication of our finding is that diffusion research that studies only the actors and mechanisms involved may be overestimating the analytical power of these explanations.

Second, we extend the conceptual and methodological tools available for studying diffusion by using NLP methods to measure the precise similarity between source and adopter documents. This more fine-grained method of studying diffusion in terms of how much the text is diffused into adopters' texts advances existing approaches that distinguish the extent of diffusion in broad categories (see, for example, Sommer and Tallberg 2019, for their delineation of imitation, adaptation, and inspiration categories of "how much" diffusion).

Our findings also offer policy implications: by knowing which textual characteristics "travel best," policymakers can write policies in a way that will help aid the diffusion, and hence widespread use, of their policies. This is relevant for policies diffusing from international and regional bodies to national contexts as well as the national to global direction of travel. This is important for CSPs as they, arguably, promote environmental and socially responsible behavior, benefitting society. We are, in this sense, taking up Gilardi's (2016, 16ff) challenge to explicitly discuss practical implications as one key avenue for improving diffusion studies.

Diffusion and Adaptation

Diffusion is the spreading of an object – be it an idea, norm, policy, or practice – from one actor to another (Rogers [1962] 2003). For policy diffusion research, the core idea is that what happens in one locale (typically a state) affects the agenda, and ultimately the choice set, elsewhere (Simmons, Lloyd, and Stewart 2018; Solingen 2012). As such, policy diffusion is a form of interdependent decision-making: when the "prior adoption of a trait or practice in a population alters the probability of adoption for remaining non-adopters" (Strang 1991, 325).

Diffusion research has advanced an impressive range of analytical tools for explaining why models diffuse, as well as the extent to which models are adapted as they diffuse. Addressing why objects diffuse, seminal studies focus on the mechanisms of coercion, competition, emulation, and learning to explain patterns in the global diffusion of democracy and liberal market policies (Simmons, Dobbin, and Garrett 2006) and social policy (Gilardi, 2010; Weyland 2007). Actor-centric explanations point to spatial proximity (Elkins and Simmons 2005) and connections among actors (e.g., in supranational organizations; Sommerer and Tallberg 2019) as predominant explanations of where policies diffuse from and to. Cognitive psychology-inspired studies have offered insights into why actors value particular models and, as a result, either copy them or adapt them to "fit" their local context (Acharya 2004; Ansari et al. 2010; Klingler-Vidra 2018).

Whereas mainstream diffusion research has advanced insights for studying how actor-level attributes explain diffusion patterns, there has been relatively less theoretical attention dedicated to the role of the attributes of the diffusion object itself. Our contention is that being precise about the characteristics of the diffusion object is critical to understanding diffusion patterns and outcomes. Hence, our central aim is to advance understandings of the analytical power of the diffusion object's "diffuse-ability" in terms of its textual attributes.

To this end, we develop a set of theoretical expectations about how the textual attributes of a written policy affect the extent to which the text is diffused into an adopter's text. It is important to first delineate what defines the attributes of the diffusion object. We contend that policy document attributes are aspects that *inhere* in the object itself. As such, they do not require reference to other objects, nor are they determined by actors' evaluations. Certainly, actor-centered explanations are critical to explaining policy diffusion. But, as we will show, they are just part of the story. A complete story, we argue, requires isolating the object attributes and assessing the extent to which they explain diffusion.

To explain the different degrees by which text travels, we need greater precision in our conceptualization of document-level attributes. Affordance theory (Gibson 1979) helps in this regard. Affordances are described as endogenous characteristics that can offer opportunities for action but may also place constraints upon action, making it more or less likely that a practice will be adopted (Scarbrough, Robertson, and Swan 2015, 368). Specifically, the diffusion object provides the *possibility* for some action that exists independently of an actor. Just as "the walk-ability of a surface exists whether or not someone walks on it" (Ansari et al. 2010, 82), the *diffuse-ability* of a policy exists whether or not someone adopts it. We seek to isolate these affordances and theorize how they shape the extent to which text is adapted as it diffuses.

Drawing together insights from the nascent policy characteristics literature, as well as affordance theory, we theorize three main object attributes that may determine how much text is diffused: (1) complexity, (2) flexibility, and (3) restrictiveness. Table 1 provides a summary of our theoretical framework.

Table 1. Theoretical Expectations

	Diffusion Object Attributes		
	(1) Complexity	(2) Flexibility	(3) Restrictiveness
Description	Use of jargon; highly technical; difficult to read/parse/understand	Use of equivocating language; lack of focus; does not target specific actors or activities	Use of the language of compellence and deterrence; constraining language
Analytical expectations	Increases cognitive burden and uncertainty; requires technical expertise; may require delegation to a third party	Allows for flexibility in interpretation and understanding; encourages adaptation for the application of adopter policies in different contexts	Binds adopter to long-term obligations; increases conflict among actors; no opting out

Expected impact on diffusion outcomes	Highly complex source texts are adapted more when they diffuse	Equivocating source texts that do not target specific firm activities are adapted more when they diffuse	Highly restrictive source texts are adapted more when they diffuse

Complexity

Complexity, following Rogers ([1962] 2003, 257), refers to the degree to which the diffusion object is "perceived to be relatively difficult to understand or use." While Rogers was analyzing the diffusion of innovations, these insights can be directly applicable to the study of policy diffusion.[5] Indeed, Rogers's definition is used by the likes of Nicholson-Crotty (2009, 195), Makse and Volden (2011, 111), and Karsh et al. (2016, 89) in the study of policy diffusion. In this case, complexity relates to a policy's technicality and readability – namely, the complexity of writing, the use of jargon, the length of the document, and, more generally, how difficult it is to read and understand the text.

Complexity impacts diffusion in several ways. Highly complex policies may require specific expertise to properly link solutions to policy problems. Generalist decision-makers, without the requisite expertise to fully grasp the solutions made available in the diffusion object, may "face considerable uncertainty regarding the policy solution that will be most likely to produce their preferred outcome" (Nicholson-Crotty 2009, 195). Similarly, Ansari et al. (2010, 84) and Karsh et al. (2016, 89) explain how highly complex diffusion objects negatively affect the ability of adopters to assess and predict the potential success of solutions in addressing policy problems. According to Ansari et al. (2010, 84), "a complex practice consists of more components and more uncertainty regarding the links between these components, as well as more uncertainty about the causal links between inputs and outputs." By contrast, "low complexity practices contain only a few parts, and the causal relationships between them are usually fairly well understood by the potential adopters". Thus, complexity increases

[5] "Meta-analyses find [the attributes listed in Rogers's work] to be successful in predicting adoption rates across a wide range of disciplines" (Makse and Volden 2011, 109).

uncertainty (Ansari et al. 2010, 84; Karsh et al. 2016, 89; Nicholson-Crotty 2009, 195), which makes it difficult for actors to anticipate the consequences, and hence the success, of adoption.

Managing complexity can be costly. Abbott and Snidal (2000, 441), theorizing differences between soft and hard law arrangements at the international level, argue that increased complexity "makes it hard to adapt agreements rapidly without some coordinating authority," and may force adopters to rely on third parties. Adopters, following this line of reasoning, either avoid adopting complex models (in which case we would not see diffusion) or adopt the models with little adaptation of the text due to their inability to confidently adjust the text for their local context. Indeed, empirical findings consistently point to a negative relationship between complexity and diffusion (Ansari et al. 2010, 83). These insights lead to our first hypothesis:

> H1. The lower the complexity of the text in a policy document, the more the text will diffuse *without* adaptation.

Flexibility

Flexibility is the degree to which a diffusion object can be applied in different contexts. For written policies, this means the use of less targeted and more equivocating language. As such, the less a diffusion object targets specific activities, behaviors, and actors, the more the text will be adapted as it diffuses. Similarly, Weyland (2007) asserts that the level of specificity of the object diffusing determines the need for interpretation and, as a result, will determine the need for adaptation. Blueprints, which are specific policy documents, are susceptible to relatively less adaptation because they are precise texts, rather than concepts, and so less prone to different interpretations. By contrast, principles, like democracy or social welfare, are general ideas that do not prescribe specific action and are therefore more open to interpretation when being invoked for use in different contexts. Principles, then, are expected to be adapted more than specific blueprints; the more room for interpretation, the more a policy will be adapted.

Diffusion objects' flexibility may also be found in their use of ambiguous language. Affordance theory refers to this as "interpretive viability," meaning that "certain practices have a greater likelihood of *adaptation* because they lend themselves to multiple interpretations and can be adapted to multiple agendas" (Ansari et al. 2010, 83, emphasis in original; see also Scarbrough et al. 2015). Interpretive viability is directly related to ambiguity (see Giroux 2006): ambiguous language is open to definition and redefinition. Empirical research suggests that interpretive viability positively impacts how much text is diffused (Ansari et al. 2010). As such, flexible language increases the scope for "reasonable interpretation" and creates broad areas of discretion (Abbott et al. 2000, 414). Given the different levels of specificity, and therefore the variation in the extent to which interpretation is needed or is possible, principles are expected to be adapted more than blueprints as they diffuse (Weyland 2007). Rather than needing to fit particular "institutional riverbeds" (Radaelli 2005, 934), general models can diffuse in different ways in a greater number of contexts. However, greater ambiguity also means that these objects will be adapted to a greater extent when they diffuse. This leads to our second hypothesis:

> H2. The greater the flexibility of the text in a policy document, the more the text will diffuse *with* adaptation.

Restrictiveness

Restrictiveness inheres in the diffusion object to the degree that it mandates specific actions or places specific (legal) limits on other actions. For a policy document, restrictiveness can take the form of a directive, legal requirement, or mandate. This is often reflected in using the language of compellence and deterrence. For Abbott et al. (2000, 408), the more restrictive a document is, the more it imposes "a particular type of binding obligation on states and other subjects." In her study of the diffusion of environmental policy, Tews (2005) relates this to the degree to which the diffusion object requires adopters to implement shorter-term versus longer-term (and more binding) changes. Tews notes that adopters are less inclined to adopt overly restrictive, hard law polices that place the adopter under specific, long-term obligations. The adopter, in other words, seeks to hold out the "possibility of escaping the policy obligation" (Tews 2005, 18).

Adopting policies that are low in restrictiveness will also require fewer major substantive changes (Linos and Pegram 2016, 596). This allows adopters to "learn about the consequences," such as hidden costs or unforeseen contingencies, and work through problems without being constrained by binding rules (Abbott and Snidal 2000). Finally, less restrictive diffusion objects are also more politically feasible in that they provoke fewer conflicts with powerful interest groups (Kern et al. 2001, 14). Indeed, one of the main enabling factors of soft legalization is allowing state officials to avoid taking sides in polarized domestic conflict with the hope "not to alienate either group of constituents" (Abbott and Snidal 2000, 453). Taken together, these insights lead to our third and last hypothesis.

> H3. The lower the restrictiveness of the text in a policy document, the more the text will diffuse *without* adaptation.

Research Design

Creating the Dataset

Our empirical analysis focuses on the diffusion of CSPs issued by policymakers, including national governments (specifically executives, legislatures, ministries, agencies, and regulators), IOs (e.g., United Nations), regional organizations (e.g., European Union), and international multi-stakeholder initiatives (e.g., the Climate Disclosure Standards Board). Consequently, an important part of our research is our answer to recent calls for a reassessment of the role of the state in sustainability policymaking (e.g., Kourula et al. 2019; Knudsen and Moon 2017; Knudsen, Moon, and Slager 2015).

To test our hypotheses, we created a novel dataset of 1,429 CSP documents – 1,040 at the national level, 332 at the international level, and 57 at the regional level – spanning a 65-year period (1955–2020). This includes 100 countries from all seven major world regions (Africa, Asia, Europe, Middle East, North America, South America, and Oceania), as well as 20 IOs and 12 regional organizations. Further details about our policy dataset are available in the Appendix.

We constructed our dataset in several steps. First, for national-level CSPs, we drew together data from four databases, each of which aggregates CSPs, broadly conceived. These include Carrots and Sticks,[6] ECGI,[7] PRI,[8] and SSE.[9] The data – as of April 2021 – were merged, cleaned, and coded, and duplicates were omitted.

In line with the approach taken in these databases, our dataset is inclusive in terms of what counts as a CSP, reflecting the full range of CSPs, from hard law (i.e., laws, bills, legislation, and binding regulation) to soft law (i.e., recommendations, guidelines, action plans, and frameworks). The policies are authored by various governmental actors, including government ministries, central banks, regulators, public research institutes, stock exchanges, and various associations. In some cases, non-governmental organizations (NGOs), including firms, co-authored a policy. When a policy explicitly names a government as an author, we treated the regulation as a governmental policy; private sustainability policies that are only authored by corporate entities, stock exchanges, or associations are not included. Next, IO-level and regional organization-level CSPs were only partly covered in the mentioned databases; therefore, we conducted desktop research to identify the policies of the pertinent

[6] https://www.carrotsandsticks.net/ (accessed January 3, 2021).
[7] https://ecgi.global/ (accessed January 3, 2021).
[8] https://www.unpri.org/ (accessed January 3, 2021).
[9] https://sseinitiative.org/ (accessed January 3, 2021).

IOs (those that have some CSP remit) and the full population of major regional organizations in addition to the databases. A more detailed overview of the 1,429 policies in our dataset can be found in the Appendix.

To the best of our knowledge, our CSP dataset is the most comprehensive compilation, both in terms of breadth (across the state, IO, and regional levels) and time (given our 65-year time frame). The earliest policies included in the dataset appear in the early post-war period. We see this as a reflection of the fact that the post-war period is largely considered to be the beginning of the international regulatory order as we know it, and the time during which the concept of Corporate Social Responsibility (CSR) was first theorized in Howard Bowen's 1953 book, *Social Responsibilities of the Businessman* (Acquier, Gond, and Pasquero 2011). While CSR and CSP are relatively new concepts, research shows that policies that we consider to fit these categories today have been around for a long time, albeit described using different language (Carroll 2008). This broad timescale helps us to comprehensively assess the diffusion of text *over time*.[10] Put simply: it can take time for policy to diffuse and studying change over a long timeframe captures these longer diffusion trajectories.

In taking this comprehensive policy dataset approach, we note that there are more CSPs than actors because nearly every actor in our database has issued more than one CSP over time. In fact, the average number of CSPs per actor in our dataset is 32. This underscores our expectation that diffusion happens at the (dyadic) policy document level; therefore, studies should go beyond the binary assessment of whether an actor has adopted a policy.

Variable Operationalization

Diffusion

[10] Results of robustness tests assessing a shorter timeframe for policies issued after 1993, the year of the seminal UNCED Earth Summit, can be found in the Appendix.

We measure diffusion in terms of textual similarity: namely, the degree to which the language of a source policy shapes the language of an adopter policy. For our purposes, diffusion encompasses the full spectrum of textual copying, from paraphrasing, or what other scholars have interpreted as *inspiration*, to *adaptation*, to outright *imitation* (e.g., plagiarism; Sommerer and Tallberg 2019). Our textual approach to diffusion not only allows us to systematically measure the *extent* to which diffusion occurs but also to understand why the text of some CSPs diffuse more than others. To this end we measure diffusion using three NLP methods: cosine similarity, substring matching, and a combined approach using cosine similarity of five-grams (strings comprising five words). Each method has its advantages and disadvantages. Our central contention is that only in triangulating the results of the three methods will we obtain a clear picture of the different degrees to which the text of a policy document diffuses. We detail these methods below.

It is important to note here that, for all three methods, our unit of analysis is the policy document dyad. Each dyad has a source policy and an adopter policy. In all cases, we use a one-year lag between the adopter and source policies to allow time for diffusion. This means that we omit all dyads where the source and adopter policies were published in the same year. We allow for a year before a document can diffuse, given the practicalities of it needing time to travel. We also omit all dyads where both source and adopter are the same actor (e.g., two policies from the UK). In short, policy diffusion is observed in degrees of textual copying – by an adopter policy issued by one actor at time t_1 from a source policy issued by a different actor at time $t-n$. The time between the adopter and source policies ranges from one year to 65 years.

Cosine similarity is a "bag of words" approach that measures the similarities between two texts "based on the cosine of the angle between two vectors" of word occurrences within the texts (Grimmer et al. 2022, 402; see also Huang 2008). Values for cosine similarity range from 0, where two texts are completely different, to 1, where two texts have identical textual features. Importantly, cosine similarity is insensitive to the location and sequence of words as well as to document length (Huang 2008, 51).[11] This is useful for capturing textual copying that occurs when words in the adopter text are reshuffled and reordered. While we are not aware of any previous studies using cosine similarity to measure the diffusion of CSPs, this method has been used widely to assess textual similarities in public policy documents. This includes the study of international financial regulation (Pagliari and Wilf 2020), legislation (Linder et al. 2020; Jansa, Hansen, and Gray 2019), political speeches (Hager and Hilbig 2020), and European Central Bank speeches (Moschella and Diodati 2020). CSPs are issued by similar policy actors (namely, elected officials, policymakers, regulators, etc.) and with the intention of shaping business activities. We therefore contend that the use of cosine similarity is appropriate in the context of measuring the diffusion of CSP texts. Our variable, *Cosine similarity*, is measured using the Quanteda package in RStudio (Benoit et al. 2018). Following standard practice, we prepared all documents in our database by first removing stop words, punctuation, and numbers; transforming all letters to lower case; and stemming all remaining words. Given our focus on the textual and linguistic attributes of policies, we exclude all CSPs that were not available in English (n = 201).

The second method we use to capture textual diffusion is substring matching. Unlike cosine similarity, this NLP method accounts for location, sequence, and word order when comparing two texts. Substring matching requires that the texts of two policies match "literally word for word (with no deviations)," typically "for sequences of six or more words" (Alle and Elsig 2019, 606). This approach is similar to that used by plagiarism-detection software, and only catches the most literal forms of copying. Critically, while substring matching is better at measuring outright copying compared to cosine similarity, it cannot capture the nuances of

[11] Given two documents $\vec{t_a}$ and $\vec{t_b}$, their cosine similarity is $Sim_C(\vec{t_a},\vec{t_b}) = \frac{\vec{t_a} \cdot \vec{t_b}}{|\vec{t_a}| \times |\vec{t_b}|}$ where $\vec{t_a}$ and $\vec{t_b}$ are m-dimensional vectors over the term set $T = \{t_1, \ldots, t_m\}$ (Huang 2008, 51).

textual copying, such as adaptation, inspiration, or paraphrasing. While it is less widely used compared to cosine similarity (Huang 2008, 51), substring matching has been recently used by Allee and Elsig (2019) in their study of preferential trade agreements, and by Pagliari and Young (2020) in their study of information exchange among interest groups.[12] Using the textreuse package in RStudio and applying a six-gram string length used by Hinkle (2015) and Alle and Elsig (2019), we created the variable *Text-reuse*, where values range from 0 (no similarities) to 1 (identical texts).

Finally, a recent study by Peacock, Milewicz, and Snidal (2019) of textual similarities between international trade agreements proposes a third approach where cosine similarity scores are measured at the level of five-gram word strings in text dyads (see also Alschner and Skougarevskiy 2016). This approach takes a novel middle path between cosine similarity and substring matching. It introduces some degree of accounting for word location into the cosine similarity method without the strict word sequence requirement used in substring matching. At the same time, in taking a middle path between cosine similarity and substring matching, this approach limits our ability to capture paraphrasing and adaptation while still falling short of detecting actual word-for-word copying (imitation). Using the Quanteda and tidyr packages in RStudio and applying a five-gram string length as used by Peacock et al. (2019), we created the variable *Cosine similarity w/ 5-grams*, where values range from 0 to 1, and where values approaching 1 reflect greater textual copying.

In Figure 1 we use an example to illustrate how textual copying captures policy diffusion.[13] On the left is Tanzania's 1994 "Guidelines on Corporate Governance Practices by Public Listed Companies in Tanzania" (an example *source* policy document); on the right is

[12] For a similar approach in the context of the diffusion of legal texts, see Hinkle (2015), who uses the open source plagiarism detection software WCopyfind (https://plagiarism.bloomfieldmedia.com/; accessed June 24, 2022).

[13] Pagliari and Young (2020) use a similar approach to visualizing textual copying in their study of information sharing among interest groups.

Uganda's 2003 "Capital Markets Corporate Governance Guidelines" (an example *adopter* policy document). These policy documents are clearly issued by different actors and in different years. Cosine similarity between these documents is extremely high (0.96), indicating that a large amount of the source text was diffused into the adopter's policy document. Looking at the highlighted sections of just one page of each document helps illustrate the basis for this high similarity score. We can see a great deal of direct "copying and pasting." Much of the highlighted text was copied word for word from the source Tanzanian text into the adopter Ugandan policy document. Where the text differs is telling of the intentionality of the adopting Ugandan author(s): particularly in the pink highlighted text, "Tanzania" is simply replaced by "Uganda."

Figure 1. Comparison of Tanzania's (1994) "Guidelines on Corporate Governance Practices by Public Listed Companies in Tanzania" and Uganda's (2003) "Capital Markets Corporate Governance Guidelines."

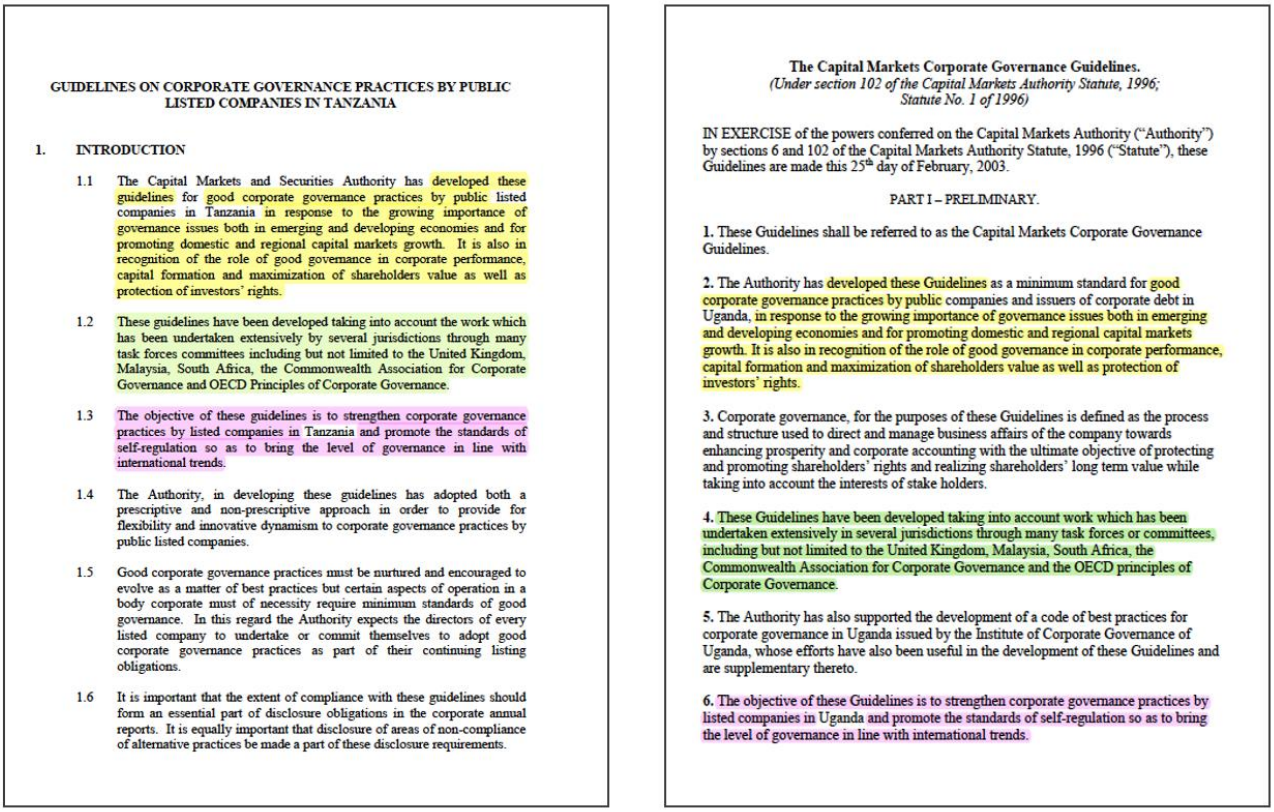

GUIDELINES ON CORPORATE GOVERNANCE PRACTICES BY PUBLIC LISTED COMPANIES IN TANZANIA

1. INTRODUCTION

1.1 The Capital Markets and Securities Authority has developed these guidelines for good corporate governance practices by public listed companies in Tanzania in response to the growing importance of governance issues both in emerging and developing economies and for promoting domestic and regional capital markets growth. It is also in recognition of the role of good governance in corporate performance, capital formation and maximization of shareholders value as well as protection of investors' rights.

1.2 These guidelines have been developed taking into account the work which has been undertaken extensively by several jurisdictions through many task forces committees including but not limited to the United Kingdom, Malaysia, South Africa, the Commonwealth Association for Corporate Governance and OECD Principles of Corporate Governance.

1.3 The objective of these guidelines is to strengthen corporate governance practices by listed companies in Tanzania and promote the standards of self-regulation so as to bring the level of governance in line with international trends.

1.4 The Authority, in developing these guidelines has adopted both a prescriptive and non-prescriptive approach in order to provide for flexibility and innovative dynamism to corporate governance practices by public listed companies.

1.5 Good corporate governance practices must be nurtured and encouraged to evolve as a matter of best practices but certain aspects of operation in a body corporate must of necessity require minimum standards of good governance. In this regard the Authority expects the directors of every listed company to undertake or commit themselves to adopt good corporate governance practices as part of their continuing listing obligations.

1.6 It is important that the extent of compliance with these guidelines should form an essential part of disclosure obligations in the corporate annual reports. It is equally important that disclosure of areas of non-compliance of alternative practices be made a part of these disclosure requirements.

The Capital Markets Corporate Governance Guidelines.
(Under section 102 of the Capital Markets Authority Statute, 1996; Statute No. 1 of 1996)

IN EXERCISE of the powers conferred on the Capital Markets Authority ("Authority") by sections 6 and 102 of the Capital Markets Authority Statute, 1996 ("Statute"), these Guidelines are made this 25th day of February, 2003.

PART I – PRELIMINARY.

1. These Guidelines shall be referred to as the Capital Markets Corporate Governance Guidelines.

2. The Authority has developed these Guidelines as a minimum standard for good corporate governance practices by public companies and issuers of corporate debt in Uganda, in response to the growing importance of governance issues both in emerging and developing economies and for promoting domestic and regional capital markets growth. It is also in recognition of the role of good governance in corporate performance, capital formation and maximization of shareholders value as well as protection of investors' rights.

3. Corporate governance, for the purposes of these Guidelines is defined as the process and structure used to direct and manage business affairs of the company towards enhancing prosperity and corporate accounting with the ultimate objective of protecting and promoting shareholders' rights and realizing shareholders' long term value while taking into account the interests of stake holders.

4. These Guidelines have been developed taking into account work which has been undertaken extensively in several jurisdictions through many task forces or committees, including but not limited to the United Kingdom, Malaysia, South Africa, the Commonwealth Association for Corporate Governance and the OECD principles of Corporate Governance.

5. The Authority has also supported the development of a code of best practices for corporate governance in Uganda issued by the Institute of Corporate Governance of Uganda, whose efforts have also been useful in the development of these Guidelines and are supplementary thereto.

6. The objective of these Guidelines is to strengthen corporate governance practices by listed companies in Uganda and promote the standards of self-regulation so as to bring the level of governance in line with international trends.

Diffusion Object Attributes

Complexity

Complexity is the degree to which a policy document is difficult to read and understand. NLP measures complexity in two ways: (a) textual readability and (b) lexical diversity. Textual readability is understood as the level of education needed to read a certain text or, more concretely, the number of structural obstacles in a text that make reading difficult. These structural obstacles can include the use of jargon, long words, and lengthy sentences within many sub-sentences. We measure *Readability* using the Flesch–Kincaid Grade Level index (Flesch 1948). Employing a formula including the number of sentences, syllables, and words in a document, this index results in values that approximate American educational grade levels, ranging from low values (corresponding to texts that are very easy to read) to high values (corresponding to texts that are very difficult to read).[14]

Lexical diversity, by contrast, measures the "complexities of text as they relate to the actual informational content of the text" (Aizenberg and Muller 2020, 9). For instance, text A is more complex than text B if it contains more information within the same unit of text (either characters, sentences, or words). We measure *Lexical Diversity* as a type–token ratio where type refers to the number of unique words in a document and token is the raw number of words in a document (regardless of repetition).[15] The central idea underpinning lexical diversity is that the use of a more varied vocabulary corresponds to a larger number of unique words and, hence, more complex texts. *Lexical Diversity* ranges from 0 to 1, where values closer to 1 correspond to documents being more complex.

Flexibility

[14] Flesch Kincaid Grade Level index = 0.39 * number of words / number of sentences + 11.8 * number of syllables / number of words – 15.59.

[15] Lexical diversity is measured as a type–token ratio ($ttr = \frac{V}{N}$), where V is the number of types and N is the number of tokens.

Our second policy attribute is flexibility, or the degree to which a CSP lends itself to being applied in different contexts. Like complexity, we measure flexibility at the text level using two indicators: *Ambiguity*, or the use of general and equivocating language, and *Focus*, or the use of language that does not target specific firm activities.

We measure *Ambiguity* using a dictionary of 298 "uncertainty terms" developed by Loughran and McDonald (2011). This includes terms related to ambiguity, approximation, contingency, lack of specificity, and randomness.[16] Research has shown that "sentiment dictionaries," especially those examining "tone" (positive and negative sentiments) are sensitive to "domain specificity" (Grimmer and Stewart 2013, 268f; Breen, Hodson, and Moschella 2019; Loughran and McDonald 2011). The current gold standard in NLP research is to use bespoke dictionaries developed using supervised methods and in specific domains. The dictionary developed by Loughran and McDonald (2011), which includes lists for "uncertainty" terms as well as lists for "constraining" and "litigious" terms, for example, have been shown to be appropriate for text analysis in accounting, business, and finance contexts (Breuer and Ghufran 2020). Specifically, recent research has applied these dictionaries in contexts directly relevant to our analysis of CSP, including NLP analyses of public policy and regulation (Breen, Hodson, and Moschella 2019; Beaulieu-Guay, Tremblay-Faulkner, and Montpetit 2020; Möller and Rechmann 2021), as well as firms' CSR reports (Du and Yu 2021; Mucko 2021).

To create our *Ambiguity* indicator, we identified the frequency of Loughran and McDonald's 298 "uncertainty terms" in each CSP in our corpus, and then calculated a percentage based on the total number of words in the documents. The validity of this measurement was assessed by manually coding a random sample of 30 CSPs from the corpus,

[16] Ambiguity terms can be found at https://sraf.nd.edu/textual-analysis/resources/#Master%20Dictionary (accessed August 10, 2021).

comparing this to our automated text-as-data analysis, and then calculating a Cohen's Kappa statistic (a common method of ascertaining inter-coder reliability). The result was a very high level of agreement (0.95) between the manual coding and automated text-as-data analysis. The small discrepancy between the two approaches is largely the result of small and rare inconsistencies in spelling, spelling mistakes in corpus texts, and hyphens added to words where words wrap around multiple lines in a corpus text.

Next, CSPs vary in terms of the degree to which firms operating in specific industries are targeted. Measuring the *Focus* of each CSP required several steps. First, we used an NLP approach developed in Al-Ubaydli and McLaughlin (2017) to assess how many different economic activities are targeted in each CSP. Economic activities are coded using the North American Industry Classification System (NAICS), a widely used classification scheme comprising 20 broad categories of economic activity. A full list of activities is available in the Appendix. Each NAICS category has a corresponding set of unique n-grams.[17] We used these n-grams to determine the degree to which the text of each CSP targets specific sectors of firms' economic activity. This is expressed as the percentage of each NAICS n-grams per policy document and is measured as the sum of the squared proportions of economic activities for each document. Hence, our final indicator, *Focus*, ranges from values close to 0, which indicate CSPs that target an extremely diverse range of firms' economic activities with little sectoral focus, to values close to 1, which indicate CSPs that target just a few economic sectors.

Restrictiveness

Our third textual attribute is *Restrictiveness*, or the degree to which CSPs compel or mandate specific action. Typically, researchers make a distinction between policies enshrined in legislation, policies that are binding, or "hard law," and all other types of non-binding or "soft law" policies (Abbott and Snidal, 2000). For those focusing on CSP, a central distinction is

[17] See https://www.quantgov.org/ (accessed August 10, 2021).

made between guidelines that recommend and policies that mandate action (Knudsen, Moon, and Slager 2015, 4). A CSP mandate, which is enshrined in legislation, would be highly restrictive in the sense of placing long-term and binding obligations on firm behavior, and hence will have a negative impact on diffusion. We created the indicator *Mandate* by hand-coding all CSP documents in our dataset in terms of their legal document type. All *Mandate* policies are legally binding (which includes, for example, legislation, bills, directives, laws, regulations) and are coded as 1, and all other (soft) policies are coded as 0.

In addition to this rather blunt binary indicator, we also include an approach that draws on NLP techniques. This indicator, *Restrictions*, is based on a dictionary of 183 constraining terms developed by Loughran and McDonald (2011) that includes terms related to commitments, compulsion, dictates, mandates, and obligations.[18] To create our *Restrictions* indicator, we identified the frequency of these restriction terms in each CSP and then calculated a percentage based on the total number of words in the document. As noted above, the Loughran and McDonald dictionaries are sufficiently specified to our domain of study, namely, CSPs. Nevertheless, to assess the robustness of our results, we use two alternative dictionaries that have been developed to assess stringency in the context of public policy: Al-Ubaydli and McLaughlin (2017) and Beaulieu-Guay, Tremblay-Faulkner, and Montpetit (2020). CSPs are, after all, public policies. What is more, using multiple dictionaries as evidence of robustness follows best practice in the literature (Boudt and Thewissen 2015). As with *Ambiguity*, we also carried out a test for the validity of this measure by comparing manual coding and automated text-as-data analysis. In this case, the Cohen's Kappa statistic measuring the validity of this measure indicated very high agreement (0.97).

Control Variables

[18] Restrictions terms can be found at https://sraf.nd.edu/textual-analysis/resources/#Master%20Dictionary (accessed August 10, 2021).

Our analytical aim is to study how policy documents' textual attributes affect the extent to which the text diffuses. Given that the thrust of diffusion scholarship has, to date, focused on actor-level explanations, our analysis includes several control variables derived from the mainstream diffusion literature. The actor-level control variables emanating from diffusion research include geographical distance and membership-based connections between issuers and adopters.

Geographical Distance

Diffusion studies often invoke the geographical distance separating actors as a key determinant of policy convergence (e.g., Elkins and Simmons 2004; Weyland 2007). According to this literature, the greater the geographical distance between two actors, the less diffusion is expected to occur, as diffusion has often been found to occur in spatial clusters. We therefore include the variable *Distance*, which measures the distance, in kilometers, between the source and adopter. For national governments, this is measured as the distance between each country's capital city. For IOs and regional organizations, we use Sommerer and Tallberg's (2019, 417) approach, which takes an IO's or region's headquarters to measure distance. Data for this variable were gathered from the CEPII database.[19]

Membership-Based Connections

CSP diffusion involves interdependent decision-making among the actors authoring the documents. For policies to be diffused, actors need to be aware of existing policies and policy documents available to them. This awareness can stem from the interactions actors have through their shared memberships in various international and regional organizations (Sommerer and Tallberg 2019). Indeed, IOs have also been identified as key venues for actor interaction and socialization (Simmons, Lloyd, and Stewart 2018, 259), as well as relational structures that coerce policy diffusion (Simmons, Dobbin, and Garrett 2006).

---

[19] See http://www.cepii.fr/CEPII/en/cepii/cepii.asp (accessed August 10, 2021).

To control for the extent to which connectivity shapes diffusion outcomes, we measure the connections between actors through international and regional organizations by mapping their shared memberships in IOs and regional bodies. Following Sommerer and Tallberg (2019) and Simmons, Lloyd, and Stewart (2018), we assess connectivity using data from the Correlates of War-Intergovernmental Organizations dataset.[20]

However, we adapt this approach to accommodate our inclusion of state-level, IO, and regional actors. This means accounting for all nine dyadic combinations: IO–IO, IO–state, IO–region, state–IO, state–state, state–region, region–IO, region–state, region–region. For IO–IO, IO–region, region–IO, and region–region dyads, we measure connectivity in terms of the proportion of member states shared by each IO or region dyad for each year of our dataset (e.g., we measure the number of member states shared by the International Labor Organisation each year).

Following the approach in Sommerer and Tallberg (2019, 416), a value of 1 is given if the share of identical member states within a given dyad is greater than 90%. This amounts to nearly 43% of all relevant dyads in our dataset. For state–IO, state–region, IO–state, and region–state dyads, we use a dummy variable for whether the specific state was a member of the IO or region at the time of diffusion, where 1 indicates membership. Finally, for the state–state dyad, we use a measure of the proportion of IOs and regional bodies of which both states in a dyad are members at the time of diffusion. Again, we use the 90% threshold employed by Sommerer and Tallberg (2019). In this case, however, the highest proportion of shared membership was just 34% (i.e., no dyad met the threshold for inclusion), reflecting a large disparity between national governments and IOs or regional organizations. In a final step, we concatenated and recoded all these binary indicators to create the single binary variable *Membership Overlap* (where 1 = overlap).

[20] These data are available online at https://correlatesofwar.org/data-sets/IGOs (accessed June 12, 2021).

Types of Diffusion Actors

In addition to the analytically motivated control variables, we also control for the fact that the CSPs in our dataset have been authored by many different types of actors, including international and regional organizations as well as national governments. This multilevel approach helps us comprehensively study CSP diffusion, but it also presents the challenge that there are different types of actors involved. These actors differ in terms of their ability, and ambition, to author diffuse-able CSPs. We control for issuer types by creating two variables, one for the issuer of the source document (*Source Issuer*) and one for the issuer of the adopted document (*Adopter Issuer*). To do this, we coded each issuer using the following issuer types: (1) government ministries, (2) financial market regulators, (3) industry associations, (4) stock exchanges, (5) NGO and research institutes, (6) central banks, (7) IOs, and (8) regional organizations. These categories were generated inductively from the policies in our dataset.

Time

We also control for the fact that diffusion processes are temporal in nature. Diffusion tends to begin slowly, accelerates, and then declines, reflecting saturation in the population of potential adopters (Rogers [1962] 2003). As such, diffusion has been shown to follow an S-shaped curve (Colyvas and Jonnson 2011, 32). More generally, research suggests that the "factors associated with [diffusion] are likely to shift over time" (Colyvas and Jonnson 2011, 33). To control for the time dimension, we include a final variable, *Time*, in our analysis. *Time* measures the number of years between the publication of the source policy document (at $t-n$) and the publication of the adopter document (at t_1). Since we include a lag of one year between the publication of the source and adopter policies, *Time* ranges from a minimum of 1 year to a maximum of 64 years. As such, *Time* captures how long it takes for an existing policy to be copied in an adopter policy. This treatment of time differs from other diffusion studies that tend to focus on *when* policies are enacted. Shipan and Volden (2006), to take one prominent example, use survival analysis to assess policy diffusion, but in terms of the amount of time between a fixed year and the adoption of a policy (see also Hinkle 2015, 1006). Our study of time is different given that policies can be adopted and updated multiple times by each policy actor; so, we are not testing for a binary yes/no regarding whether or not a policy has been adopted by an actor at some point in time.

Analysis

As noted above, we use three methods to test our hypotheses: cosine similarity, substring matching, and a third approach using cosine similarity on five-gram strings. First, diffusion measured as cosine similarity between policy dyads, with or without five-gram strings, is beta distributed such that values are bounded by 0 and 1 but also exclude the end points of 0 and 1. An appropriate regression method for this type of dependent variable is beta regression (Ferrari and Cribari-Neto 2004), which was used in this study. Second, and by contrast, diffusion measured as substring matching between policy dyads is a proportion where values are bounded by 0 and 1 and can include the end points of 0 and 1. In this instance, a generalized linear model with a logit link (fractional logistic regression), a common method for this type of dependent variable (Papke and Woolridge 1996), was used to estimate the regression models. Given the dyadic nature of our data, we expect some interdependencies among our dyads. As such, we used clustered standard errors for which our cluster variable is the specific source issuer ID. All regression results are presented in Table 2. Further results using a stepwise approach where each independent variable is introduced individually into the regression model can be found in the Appendix.

Table 2. Regression Analysis Explaining the Diffusion of CSPs Using Three Methods

	Cosine sim. (1)	Substring matching (2)	Cosine sim. w/ 5-grams (3)
Policy Attributes			
H1a. Readability	−0.000873***	−0.0105***	−0.00137***

	(0.0000748)	(0.00104)	(0.000119)
H1b. Lexical diversity	−0.608***	−2.685***	0.0718***
	(0.00805)	(0.0622)	(0.0133)
H2a. Focus	0.0206***	0.226***	0.254***
	(0.00440)	(0.0255)	(0.00680)
H2b. Ambiguity	2.685***	2.190	0.340
	(0.114)	(1.252)	(0.174)
H3a. Restrictions	−3.764***	−15.97***	−3.889***
	(0.290)	(1.908)	(0.461)
H3b. Mandate	−0.355***	−0.608***	−0.161***
	(0.00201)	(0.00979)	(0.00307)
Control variables			
Distance	−0.00000718***	−0.0000627***	−0.0000156***
	(0.000000165)	(0.00000143)	(0.000000254)
Membership overlap	−0.188***	−0.758***	−0.260***
	(0.00185)	(0.0166)	(0.00297)
Time	−0.00235***	−0.0175***	−0.00280***
	(0.000113)	(0.000671)	(0.000172)
Source issuer	Yes	Yes	Yes
Adopter issuer	Yes	Yes	Yes
Constant	−0.734***	−6.741***	−7.594***
	(0.00454)	(0.0302)	(0.00703)
N	773,485	718,960	537,080

Notes: Beta regression is used in models 1 and 3; fractional logit is used in model 2. Clustered standard errors are in parentheses; observations are clustered on the Source Issuer ID.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Our results support our main theoretical expectation that the textual attributes of the diffusion object have an important bearing on the extent to which a CSP's text is adapted as it diffuses. However, support for our specific hypotheses is mixed. Starting with our Complexity hypothesis (H1), our results show, as expected, that the extent to which a policy's text diffuses is negatively correlated with both *Readability* and *Lexical Diversity*. In other words, a highly technical, difficult to read, and hard to understand source text is adapted more than the text of a source policy that is easy to read and understand. The results for *Readability* are consistent across all three regression models; however, those for *Lexical Diversity* are only statistically significant in models 1 and 2. We explore this further in our robustness tests below. Taken together, these results partially support our Complexity hypothesis, as existing literature informed expectations that highly complex policies, when they are diffused, are not adapted due to adopters' inability to skillfully reword them.

Results provide mixed support for our Flexibility hypothesis (H2). First, and across all three models, we observe a positive and statistically significant correlation between CSP diffusion and *Focus*. As expected, the text of policies that target just a few specific sectors of economic activity tend to diffuse more than polices that either target a large and diverse range of sectors or remain ambiguous about targeted sectors. Second, we find some evidence that, contrary to our theoretical expectations, greater *Ambiguity* increases the extent to which the text of policy documents diffuses without adaptation. However, these results are only statistically significant in model 1, with models 2 and 3 showing no significant differences for *Ambiguity*. Despite being mixed, these results call into question the assumption that adopters' differing interpretations of policies will result in adaptations of the source model (Weyland 2007). However, we propose that rather than favoring flexibility in a policy's *interpretation*, adopters might favor flexibility in a policy's *application*. Ambiguous language may be copied from the source policy document to the adopted text to retain flexibility in the *application* to their local context. Rather than the additional cognitive work implicit in interpreting how the text of a policy should be adapted for local use, our findings here suggest that ambiguous texts are partly replicated to retain flexibility in policy application.

Finally, we find considerable support for our Restrictiveness hypothesis (H3). CSPs that *Mandate* specific behavior through legal means, or that use more restrictive language (*Restrictions*), tend to have their text adapted more than policies that are softer, less binding, or that are less restrictive. These results are unambiguously consistent across all three models using different methods for measuring textual copying.

To get a better sense of the magnitude of the effects presented in the three models in our regression table, we estimate marginal effects for our six policy document attribute indicators for our results in model 1 (using cosine similarity). These are presented in Figure 2. Additional figures for marginal effects for models 2 and 3 can be found in the Appendix. For many of our covariates, the predicted change in value for cosine similarity is rather small. In part, this reflects the distribution of our dependent variable, where values cluster tightly around the mean (0.267). To aid the interpretation of the marginal effects, we have not standardized the y-axis.

Figure 2. Marginal Effects of Policy Attributes on Cosine Similarity.

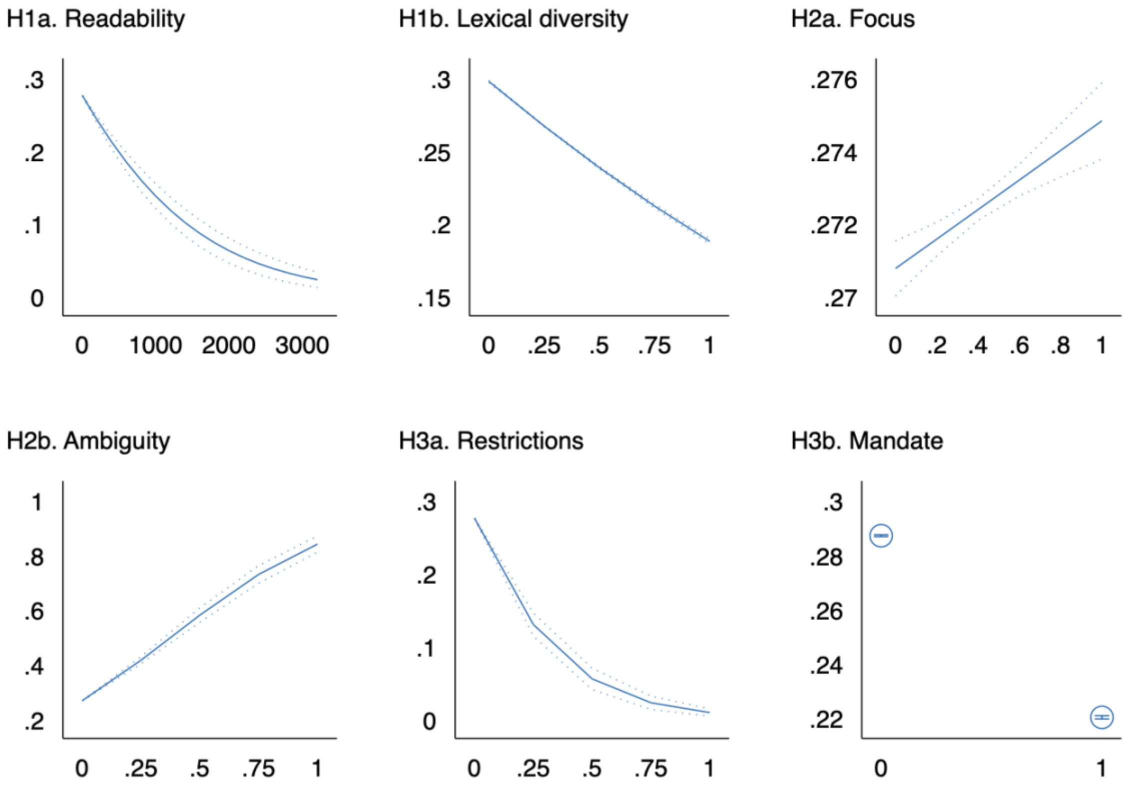

Note: Y-axis values plot the conditional mean of *Cosine similarity*. Estimations are from Table 2, model 1.

Starting with the indicators for Complexity, as we move from the lowest to the highest values of *Readability* and *Lexical Diversity*, which correspond to CSPs that are easy to read and understand, and those that are extremely complex, respectively, we see a significant decrease in the predicted values of cosine similarity. These trends support our expectation that the text of more complex policies will tend to diffuse less than simpler policies. Moving across the full spectrum of *Readability* from very easy to read (equivalent to a US grade level of kindergarten) to incredibly difficult to read policies (equivalent to post graduate levels), we see a change in cosine similarity scores from about 0.3 (well above the mean) to nearly 0 (no similarity between policies at all). For *Lexical diversity*, measuring the complexity of policy documents in terms of the use of unique words relative to sentence length, values for cosine similarity decrease modestly from about 0.3 to about 0.2, as *Lexical diversity* increases from its minimum value of 0 to its maximum value of 1.

Turning next to our marginal effects for our Flexibility indicators, we see very different results. For *Focus*, measuring the extent to which polices target just one sector of economic activity (values close to 1), or target a large and diverse number of sectors (values close to 0), the predicted impact on cosine similarity is very small, moving only slightly from 0.29 to 0.24. Our second indicator for Flexibility – *Ambiguity* – appears to have the largest impact on cosine similarity. As we move from an *Ambiguity* score of nearly 0 (those policies that use little or no equivocating language) to a score of nearly 1 (those policies replete with equivocating language), the cosine similarity between two policies increases from about 0.29 (modest similarity) to 0.9 (a great deal of direct copying from the source document to the adopter text). The magnitude of this effect gives further purchase to the notion that adopters embrace and replicate ambiguous language from the source policy rather than see ambiguity as an opportunity, or need, to adapt. However, we need to recall that these findings are sensitive to the methods used to measure textual copying, with models 2 (using substring matching) and 3 (using cosine similarity on five-grams) showing results for *Ambiguity* that are not statistically significant.

Finally, for Restrictiveness we again see significant effects. For *Restrictions*, the predicted values for cosine similarity decrease from about 0.3 (above the mean) to nearly 0 (no similarities between policies) as we move from CSP documents that use few "restriction" terms to CSP documents that use restriction terms substantially. This is an important effect: the texts of source policy documents replete with restriction terms are rarely if ever copied into adopter policies. However, the texts of source polices with few, or no, restriction terms are copied in an adopter text to a significant degree. Finally, *Mandate*, which differentiates between hard law (bills, binding regulations, laws) and soft law (frameworks, guidelines, recommendations), shows the latter clearly diffusing without adaptation more than the former. Specifically, moving from soft law to hard law equates to a decrease of about 0.05 in cosine similarity. Together, these findings lend support to our third hypothesis, that the text of source policies that use the language of compellence and deterrence (e.g., hard law) are adapted as they diffuse into adopter policies. On a broader scale, this suggests greater convergence around voluntary CSPs (e.g., soft law) that does not necessitate specific action.

The results for the control variables in Table 2 provide important insights into various actor-centric explanations for "how much" text diffuses. First, the results for *Distance* show a negative correlation with cosine similarity. This effect is statistically significant and consistent across all three models. Greater geographical distance impedes diffusion; this finding is consistent with those on the role of spatial proximity in driving diffusion in empirical areas as diverse as democracy (Elkins and Simmons, 2004), social policy (Weyland 2007), and the criminalization of human trafficking (Simmons, Lloyd, and Stewart 2018).[21] Next, our results for *Membership overlap* (a negative and significant relationship) suggest that diffusion is not facilitated by actors interacting through shared membership in international or regional organizations. While our results are specific to CSP diffusion, they do stand in contrast to Sommerer and Tallberg's (2019) findings that connectivity through IOs facilitate the diffusion of participatory governance arrangements among IOs. Finally, our regression results suggest

[21] To be sure, Simmons, Lloyd, and Stewart (2018, 262) assess geographical proximity in terms of the number of roads linking two countries together and the length of passage over water.

that *Time* is an important determinant for diffusion. The more time passes after a policy is issued, the less likely it is that its text will be copied into an adopter policy. This finding lends further support to expectations about how diffusion is not consistent over time (Westphal et al. 1997). Thus, without studying the attributes of the policy itself, our study of the diffusion of 1,429 CSPs would have emphasized spatial proximity and not actor connectivity. In both cases, the study would have overestimated the impact of actors' closeness on diffusion outcomes.

Robustness Tests

We test the robustness of our results using several alternative indicators for policy documents' textual attributes. As noted above, there are many alternative measures for *Readability* as well as different type–token ratios used to measure *Lexical diversity*. We have opted for two popular alternatives, first replacing the Flesch–Kincaid Index with Gunning's Fog Index (Gunning 1969; calculated as the number of three-syllable words in a document divided by the average sentence length in the document) and the type–token ratio with Herdan's C (calculated as the product of the log of the number of "types" over the log of the number of "tokens"). We are also aware of several different approaches to measuring policy restrictiveness: (1) Loughran and McDonald's (2011) lexicon of 904 "litigious terms" (*Litigiousness*)[22]; (2) Beaulieu-Guay, Tremblay-Faulkner, and Montpetit's (2020) Stringency Index comprising 14 restriction n-grams; and (3) Al-Ubaydli and McLaughlin's (2017) five restriction n-grams. These alternative indicators were tested in separate beta regression models using cosine similarity as the dependent variable and including all control variables. Marginal effects are plotted in Figure 3, while a complete regression table is available in the Appendix. Our results using these alternative indictors are again robust, supporting our original regression results.

[22] A list of litigiousness terms can be found at https://sraf.nd.edu/textual-analysis/resources/#Master%20Dictionary (accessed August 10, 2021).

Only the alterative indicator for Lexical diversity – Herdan's C – stands out. The impact on cosine similarity is still negative and statistically significant, but the marginal effects show only a modest predicted change in our dependent variable. This, coupled with inconsistent results for *Lexical diversity* in Table 2, further undermines our expectation about the linkages between how much the text of policy documents diffuses and the policy document's "complexity." If anything, we can say that, while *Readability* is a good predictor of CSP diffusion, policy complexity, measured using a type–token ratio or Herdan's C is not. The three alternative measures for restrictiveness (namely, Litigiousness terms, or restriction terms developed by Al-Ubaydli and McLaughlin 2017 and Beaulieu-Guay, Tremblay-Faulkner, and Montpetit 2020) are remarkably consistent with our findings using our *Restrictions* indicator.

Figure 3. Marginal Effects with Alternative Indicators

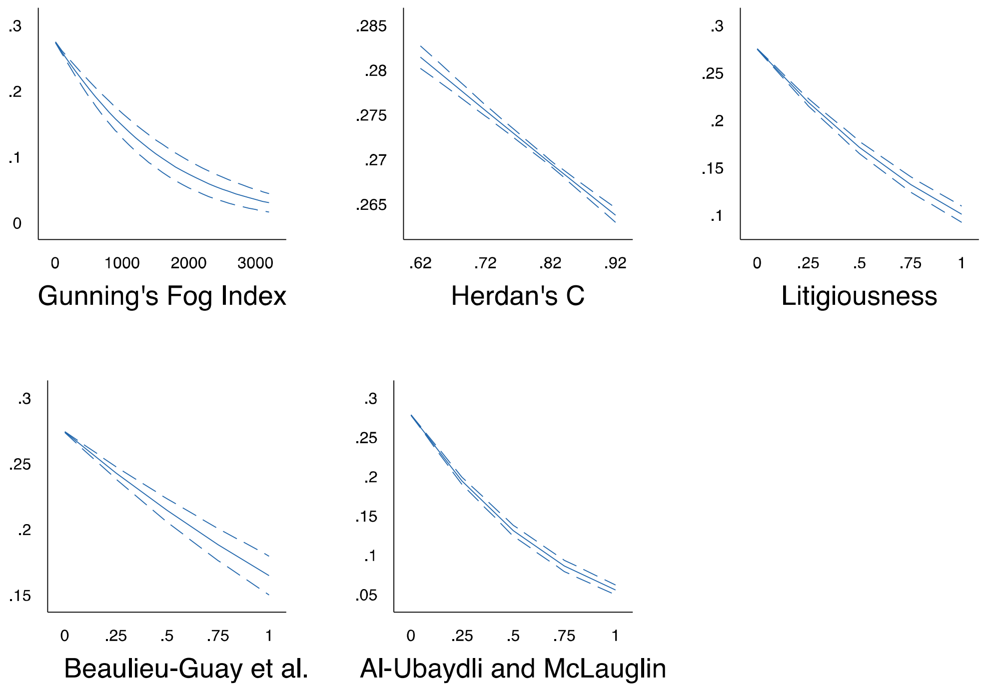

Note: Y-axis values plot the conditional mean of *Cosine similarity*.

Conclusion

We argue that whether diffusion results in "more or less convergence" (Klingler-Vidra and Schleifer 2014) is not solely determined by actor-centric explanations. By extending research on policy attributes and drawing on insights from affordance theory, we show how diffusion outcomes – in terms of the extent to which the diffusion object is adapted in the process, rather than just whether the objected diffuses or not – are also shaped by the textual attributes of the policy document itself. We find that the policies that diffuse with the least adaptation in the diffusion process are those that use simple language, that do not require binding actions (e.g., material disclosures), and that target firms operating in specific industries. Our findings therefore speak to those in the broader literature about how policymakers (and other actors) may be able to communicate most *effectively*. Rauh (2022) uses similar NLP methods to show how overly complex language used in European Commission press releases reinforces public attitudes about the institution's aloof and technocratic nature (and may even negatively impact processes of European integration). As noted by Tweedie (2022, 1853), complex language insulates and distances texts from their audiences. Both studies underscore the importance of textual attributes – in particular, readability – in diffusion processes.

While the text of more targeted, less complex, and less restrictive policy documents is adapted less as it travels, that does not necessarily mean that the impact of those policies in shaping firm-level behavior is greater. As Linos and Pegram (2016) showed empirically, the language of policies has important consequences for state behavior, even when the policy is non-binding. Less binding policies may result in less impact on firms' behavior, which is the aim of policies, especially in the CSP empirical area. Hence, there is a tension between the diffuse-ability of policies and the desirability of those policies; the text of CSPs that are diffusing most may not be the ones that most affect firms' behavior.

Going further, our study does not analyze the impact of governmental CSPs on firms' activities and commitments. An important avenue for future research would be to build on our dataset by adding sustainability reports issued by a sample of firms, as well as industry-led CSPs operating in different sectors. The same NLP techniques employed in this analysis will be useful for teasing out the attributes of firms' sustainability policies and reports. Using a measure of text similarity, we could assess the extent to which businesses are copying international, regional, or state-level CSPs and, vice versa, how firms' language informs governmental CSP texts. This could reveal important insights on the implications of the relative diffusion of policy documents based upon the text's complexity, flexibility, and restrictiveness in firm-level (reported) activities.[23] Shaping firm actions is, after all, the aim of international, regional, and national government CSPs.

[23] We are aware that reporting does not always equal behavior. Firm CSP reporting might reflect actions or aspirations. These types of reporting should, nevertheless, not be underestimated as they can have a performative effect on long-term action (Christensen, Morsing, and Thyssen 2013; van den Broek 2022). As such, establishing if governmental CSPs diffuse to firm reporting can be valuable on its own.

References

Abbott, Kenneth W., Robert Keohane, Anne-Marie Slaughter, and Duncan Snidal. 2000. “The Concept of Legalization.” *International Organization* 54 (3): 401–20.

Abbott, Kenneth W., and Duncan Snidal. 2000. “Hard and Soft Law in International Governance.” *International Organization* 54 (3): 421–56.

Acharya, Amitav. 2004. “How Ideas Spread: Whose Norms Matter? Norm Localization and Institutional Change in Asian Regionalism.” *International Organization* 58 (2): 239–75.

Acquier, Aurélien, Jean-Pascal Gond, and Jean Pasquero. 2011. “Rediscovering Howard R. Bowen’s Legacy: The Unachieved Agenda and Continuing Relevance of Social Responsibilities of the Businessman.” *Business & Society* 50 (4): 607–46.

Aizenberg, Ellis, and Moritz Müller. 2020. “Signaling Expertise Through the Media? Measuring the Appearance of Corporations in Political News Through a Complexity Lens.” *Journal of European Public Policy*. DOI:10.1080/13501763.2020.1797144

Allee, Todd M., and Manfred Elsig. 2019. “Are the Contents of International Treaties Copied and Pasted? Evidence from Preferential Trade Agreements.” *International Studies Quarterly* 63 (3): 603–13.

Alschner, Wolfgang, and Dmitriy Skougarevskiy. 2016. “Mapping the Universe of International Investment Agreements.” *Journal of International Economic Law* 19 (3): 561–88.

Al-Ubaydli, Omar, and Patrick McLaughlin. 2017. “RegData: A Numerical Database on Industry-Specific Regulations for all United States Industries and Federal Regulations, 1997–2012.” *Regulation & Governance* 11: 109–23.

Ansari, Shahzad M., Peer C. Fiss, and Edward J. Zajac. 2010. “Made to Fit: How Practices Vary as They Diffuse.” *Academy of Management Review* 35 (1): 67–92.

Beaulieu-Guay, Louis-Robert, Marc Tremblay-Faulkner, and Éric Montpetit. 2020. “Does Business Influence Government Regulations? New Evidence from Canadian Impact Assessments.” *Regulation & Governance* 15 (4): 1419–35.

Benoit, Kenneth, Kohei Watanabe, Haiyan Wang, Paul Nulty, Adam Obeng, Stefan Müller and Akitaka Matsuo. 2018. “Quanteda: An R Package for the Quantitative Analysis of Textual Data.” *Journal of Open Source Software* 3 (30): 1–4.

Bos, Thomas, and Flavius Frasincar. 2021. “Automatically Building Financial Sentiment Lexicons While Accounting for Negation.” *Cognitive Computation* 14: 442–60.

Boudt, Kris, Peter de Goeij, James Thewissen, and Geert Van Campenhout. 2015. "Analysts' forecast error: a robust prediction model and its short-term trading profitability." *Accounting & Finance* 55(3): 683-715.

Bowen, Howard R. 1953. *Social Responsibilities of the Businessman*. Iowa City: University of Iowa Press.

Breen, Michael, Dermot Hodson, and Manuela Moschella. 2019. “Incoherence in Regime Complexes: A Sentiment Analysis of EU–IMF Surveillance.” *Journal of Common Market Studies (JCMS)* 58 (2): 419–37.

Breuer, Wolfgang, and Bushra Ghufran. 2020. "The Predictive Power of Managerial Tone: A Text-Based Analysis of Takeover Performance." Available at SSRN 3763744.

Carroll, Archie B. 2008. “A History of Corporate Social Responsibility: Concepts and Practices.” In *The Oxford Handbook of Corporate Social Responsibility*, edited by A. Crane, A McWilliams, D Matten, J Moon, and DS Siegel, 19–46. Oxford: Oxford University Press.

Christensen, Lars T., Mette Morsing, and Ole Thyssen. 2013. “CSR as Aspirational Talk.” *Organization* 20 (3): 372–93.

Colyvas, Jeannette A., and Stefan Jonsson. 2011. “Ubiquity and Legitimacy: Disentangling Diffusion and Institutionalization.” *Sociological Theory* 29 (1): 27–53.

Du, Shuili, and Kun Yu. 2021. “Do Corporate Social Responsibility Reports Convey Relevant Information? Evidence from Report Readability and Tone.” *Journal of Business Ethics* 172: 253–74.

Elkins, Zachary, and Beth Simmons. 2005. “On Waves, Clusters, and Diffusion: A Conceptual Framework.” *Annals of the American Academy of Political and Social Science* 598: 33–51.

Ferrari, Silvia, and Francisco Cribari-Neto. 2004. “Beta Regression for Modeling Rates and Proportions.” *Journal of Applied Statistics* 31 (7): 799–815.

Flesch, Rudolf. 1948. “A New Readability Yardstick.” *Journal of Applied Psychology* 32 (3): 221–33.

Gibson, James. 1979. *The Ecological Approach to Perception*. London: Houghton Mifflin.

Gilardi, Fabrizio. 2010. “Who Learns from What in Policy Diffusion Processes?” *American Journal of Political Science* 54 (3): 650–66.

Giroux, Hélène. 2006. “‘It was such a handy term’: Management fashions and pragmatic ambiguity.” *Journal of management studies 43* (6): 1227-1260.

Grimmer, Justin, and Brandon M. Stewart. 2013. “Text as Data: The Promise and Pitfalls of Automatic Content Analysis for Political Texts.” *Political Analysis* 21 (3): 267–97.

Grimmer, Justin, Margaret E. Roberts, and Brandon M. Stewart. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences.* Princeton: Princeton University Press.

Gunning, Robert. 1969. "The Fog Index After Twenty Years." *International Journal of Business Communication* 6 (2): 3–13.

Hager, Anselm, and Hanno Hilbig. 2020. "Does Public Opinion Affect Political Speech?" *American Journal of Political Science* 64 (4): 921–37.

Hinkle, Rachael K. 2015. "Into the Words: Using Statutory Text to Explore the Impact of Federal Courts on State Policy Diffusion." *American Journal of Political Science* 59 (4): 1002–21.

Huang, Ann. 2008. "Similarity Measures for Text Document Clustering." *Proceedings of the New Zealand Computer Science Research Student Conference* X: 49–56

Jansa, Johsua M., Eric R. Hansen, and Virginia H. Gray. 2019. "Copy and Paste Law-Making: Legislative Professionalism and Policy Reinvention in the States." *American Politics Research* 47 (4): 739–67.

Karsh, Andrew, Sean C. Nicholson-Crotty, Neal D. Woods, and Ann O'M, Bowman. 2016. "Policy Diffusion and the Pro-Innovation Bias." *Political Research Quarterly* 69 (1): 83–95.

Kern, Kristine, Ingrid Kissling-Näf, Ute Landmann, Corine Mauch, and Tina Löffelsend. 2001. "Policy convergence and policy diffusion by governmental and non-governmental institutions: An international comparison of eco-labeling systems." No. FS II 01-305. WZB Discussion Paper.

Klingler-Vidra, Robyn. 2018. *The Venture Capital State: The Silicon Valley Model in East Asia*. Ithaca, NY: Cornell University Press.

Klingler-Vidra, Robyn, and Philip Schleifer. 2014. "Convergence More or Less: Why Do Practices Vary as They Diffuse?" *International Studies Review* 16 (2): 264–74.

Knudsen, Jette Steen, and Jeremy Moon. 2017. *Visible Hands: Government Regulation and International Business Responsibility*. Cambridge, UK: Cambridge University Press.

Knudsen, Jette Steen, Jeremy Moon, and Rieneke Slager. 2015. "Government Policies for Corporate Social Responsibility in Europe: A Comparative Analysis of Institutionalisation." *Policy & Politics* 43 (1): 81–99.

Kourula, Arno, Jeremy Moon, Marie-Laure Salles-Djelic, and Christopher Wickert. 2019. "New Roles of Government In the Governance of Business Conduct: Implications for Management and Organizational Research." *Organization Studies* 40 (8): 1101–23.

Linder, Fridolin, Bruce Desmarais, Matthew Burgess, and Eurgenia Giraudy. 2020. "Text as Policy: Measuring Policy Similarity through Bill Text Reuse." *Policy Studies Journal* 48 (2): 546–73.

Linos, Katerina, and Tom Pegram. 2016. "The Language of Compromise in International Agreements." *International Organization* 7 (3): 587–621.

Loughran, Tim, and Bill McDonald. 2011. “When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks.” *Journal of Finance* LXVI (1): 35–65.

Makse, Todd, and Craig Volden. 2011. “The Role of Policy Attributes in the Diffusion of Innovations.” *Journal of Politics* 73 (1): 108–24.
Möller, Rouven, and Doron Rechmann. 2021. “ECB Language and Stock Returns – A Textual Analysis of ECB Press Conferences.” *Quarterly Review of Economics and Finance* 80: 590–604.

Mooney, Christopher Z., and Mei-Hsien Lee. 1995. “Legislative Morality in the American States: The Case of Pre-Roe Abortion Regulation Reform.” *American Journal of Political Science* 39 (3): 599–627.

Mooney, Christopher Z., and Mei-Hsien Lee. 2000. “The Influence of Values on Consensus and Contentious Morality Policy: U.S. Death Penalty Reform, 1956–82.” *Journal of Politics* 62 (1): 223–39.

Moschella, Manuela, and Nicola N. Diodati. 2020. “Does Politics Drive Conflict in Central Banks’ Committees? Lifting the Veil on the European Central Bank Consensus.” *European Union Politics* 21 (2): 183–203.

Mucko, Przeyslaw. 2021. “Sentiment Analysis of CSR Disclosures in Annual Reports of EU Companies.” *Procedia Computer Science* 192: 3351–9.

Nicholson-Crotty, Sean. 2009. “The Politics of Diffusion: Public Policy in the American States.” *Journal of Politics* 71 (1): 192–205.

Pagliari, Stefano, and Meredith Wilf. 2020. “Regulatory Novelty after Financial Crises: Evidence from International Banking and Securities Standards, 1975–2016.” *Regulation & Governance* 15 (3): 933-951.

Pagliari, Stefano, and Kevin Young. 2020. “Exploring Information Exchange among Interest Groups: A Text-Reuse Approach.” *Journal of European Public Policy* 27 (11): 1698–1717.

Papke, Leslie E., and Jeffrey M. Wooldridge. 1996. “Econometric Methods for Fractional Response Variables with an Application to 401 (K) Plan Participation Rates.” *Journal of Applied Econometrics* 11: 619–32.

Peacock, Claire, Karolina Milewicz, and Duncan Snidal. 2019. “Boilerplate in International Trade Agreements.” *International Studies Quarterly* 63: 923–37.

Pope, Shawn, and Alwyn Lim. 2020. “The Governance Divide in Global Corporate Responsibility: The Global Structuration of Reporting and Certification Frameworks, 1998–2017.” *Organization Studies* 41 (6): 821–54.

Radaelli, Claudio. 2005. “Diffusion without Convergence: How Political Context Shapes the Adoption of Regulatory Impact Assessment.” *Journal of European Public Policy* 12 (5): 924–43.

Rauh, Christian. 2022. “Clear Messages to the European Public? The Language of European Commission Press Releases 1985–2022.” *Journal of European Integration*. DOI:10.1080/07036337.2022.2134860.

Rogers, Everett. 2003 [1962]. *The Diffusion of Innovations*. Glencoe, New York: Free Press.

Scarbrough, Harry, Maxine Robertson, and Jacky Swan. 2015. “Diffusion in the Face of Failure: The Evolution of a Management Innovation.” *British Journal of Management* 26: 365–87.

Shipan, Charles R., and Craig Volden. 2006. “Bottom-Up Federalism: The Diffusion of Antismoking Policies from U.S. Cities to States.” *American Journal of Political Science* 50 (4): 825–43.

Shipan, Charles R., and Craig Volden. 2008. “The Mechanisms of Policy Diffusion.” *American Journal of Political Science* 52 (4): 840–57.

Simmons, Beth A., Frank Dobbin, and Geoffrey Garrett. 2006. “Introduction: The International Diffusion of Liberalism.” *International Organization* 60 (4): 781–810.

Simmons, Beth A., Paulette Lloyd, and Brandon M. Stewart. 2018. “The Global Diffusion of Law: Transnational Crime and the Case of Human Trafficking.” *International Organization* 72 (2): 249–81.

Solingen, Etel. 2012. “Of Dominoes and Firewalls: The Domestic, Regional, and Global Politics of International Diffusion.” *International Studies Quarterly* 56: 631–44.

Sommerer, Tomas, and Jonas Tallberg. 2019. “Diffusion across International Organizations: Connectivity and Convergence.” *International Organization* 73 (2): 399–433.

Strang, David. 1991. “Adding Social Structure to Diffusion Models – An Event History Framework.” *Sociological Methods and Research* 19 (3): 324–53.

Tews, Kerstin. 2005. “The Diffusion of Environmental Policy Innovations: Cornerstones of an Analytical Framework.” *European Environment* 15: 63–79.

Tweedie, Jonathan. 2022. “Against Mystifying Complexity: On Asking Simple, Burning Questions.” *Organization Studies* 43 (11): 1853–6.

van den Broek, Onna M. 2022. “How Political Actors Co-Construct CSR and Its Effect on Firms’ Political Access: A Discursive Institutionalist View.” *Journal of Management Studies* https://doi.org/10.1111/joms.12871.

van den Broek, Onna M., and Robyn Klingler-Vidra. 2022. “The SDGs as a North Star: How an Intermediary Network Makes, Takes and Retrofits the Meaning of the SDGs.” *Regulation & Governance* 16 (4): 1306–24.

Westphal, J. D., & Zajac, E. J. (1997). Defections from the inner circle: Social exchange, reciprocity, and the diffusion of board independence in US corporations. *Administrative science quarterly*, 161-183

Weyland, Kurt. 2007. *Bounded Rationality and Policy Diffusion: Social Sector Reform in Latin Americ*a. Princeton, NJ: Princeton University Press.