

Adjusting for Confounding with Text Matching

Margaret E. Roberts University of California, San Diego
Brandon M. Stewart Princeton University
Richard A. Nielsen Massachusetts Institute for Technology

Abstract: We identify situations in which conditioning on text can address confounding in observational studies. We argue that a matching approach is particularly well-suited to this task, but existing matching methods are ill-equipped to handle high-dimensional text data. Our proposed solution is to estimate a low-dimensional summary of the text and condition on this summary via matching. We propose a method of text matching, topical inverse regression matching, that allows the analyst to match both on the topical content of confounding documents and the probability that each of these documents is treated. We validate our approach and illustrate the importance of conditioning on text to address confounding with two applications: the effect of perceptions of author gender on citation counts in the international relations literature and the effects of censorship on Chinese social media users.

Verification Materials: The materials required to verify the computational reproducibility of the results, procedures, and analyses in this article are available on the *American Journal of Political Science* Dataverse within the Harvard Dataverse Network, at: <https://doi.org/10.7910/DVN/HTMX3K>.

Social media users in China are censored every day, but it is largely unknown how the experience of being censored affects their future online experience. Are social media users who are censored for the first time flagged by censors for increased scrutiny in the future? Is censorship “targeted” and “customized” toward specific users? Do social media users avoid writing after being censored? Do they continue to write on sensitive topics or do they avoid them?

Experimentally manipulating censorship would allow us to make credible causal inferences about the effects of experiencing censorship, but this is impractical

and unethical outside of a lab setting. Inferring causal effects in *observational* settings is challenging due to confounding. The types of users who are censored might have different opinions that drive them to write differently than the types of users who are not censored. This in turn might affect both the users’ rate of censorship as well as future behavior and outcomes. We argue that conditioning on the text of censored social media posts and other user-level characteristics can substantially decrease or eliminate confounding and allow credible causal inferences with observational data. Intuitively, if we can find nearly identical posts—one of which is censored while the

Margaret E. Roberts is Associate Professor, Department of Political Science, University of California, San Diego, Social Sciences Building 301, 9500 Gilman Drive, #0521, La Jolla, CA 92093-0521 (meroberts@ucsd.edu). Brandon M. Stewart is Assistant Professor and Arthur H. Scribner Bicentennial Preceptor, Department of Sociology, Princeton University, 149 Wallace Hall, Princeton, NJ 08544 (bms4@princeton.edu). Richard A. Nielsen is Associate Professor, Department of Political Science, Massachusetts Institute for Technology, 77 Massachusetts Avenue, E53 Room 455, Cambridge, MA 02139 (rnielsen@mit.edu).

We thank the following for helpful comments and suggestions on this work: David Blei, Naoki Egami, Chris Felton, James Fowler, Justin Grimmer, Erin Hartman, Chad Hazlett, Seth Hill, Kosuke Imai, Rebecca Johnson, Gary King, Adeline Lo, Will Lowe, Chris Lucas, Walter Mebane, David Mimno, Jennifer Pan, Marc Ratkovic, Matt Salganik, Caroline Tolbert, and Simone Zhang; audiences at the Princeton Text Analysis Workshop, Princeton Politics Methods Workshop, the University of Rochester, Microsoft Research, the Text as Data Conference, and the Political Methodology Society and the Visions in Methodology conference; and some tremendously helpful anonymous reviewers. We especially thank Dustin Tingley for numerous insightful conversations on the connections between STM and causal inference and Ian Lundberg for extended discussions on some technical details. Dan Maliniak, Ryan Powers, and Barbara Walter graciously supplied data and replication code for the gender and citations study. The JSTOR Data for Research program provided academic journal data for the international relations application. This research was supported, in part, by the Eunice Kennedy Shriver National Institute of Child Health and Human Development under grant P2-CHD047879 to the Office of Population Research at Princeton University. The research was also supported by grants from the National Science Foundation RIDIR program, award numbers 1738411 and 1738288. This publication was made possible, in part, by a grant from the Carnegie Corporation of New York, supporting Richard Nielsen as an Andrew Carnegie Fellow. The statements made and views expressed are solely the responsibility of the authors.

American Journal of Political Science, Vol. 00, No. 0, xxxx 2020, Pp. 1–17

other is not—from similar users, we can compare downstream online behavior to obtain credible estimates of the effects of censorship.

Traditional matching methods are not suited to the task of conditioning on the text of documents. Quantitative analysts typically represent text using thousands or even millions of dimensions (e.g., the columns of a document term matrix, collections of word embeddings). Common matching techniques such as propensity score matching (Rosenbaum and Rubin 1983) and coarsened exact matching (CEM; Iacus, King, and Porro 2011) were developed for applications with fewer variables (dimensions) than observations in the data set. For example, Rubin and Thomas (1996, 249) note that in “typical examples” of matching, the number of variables is “between 5 and 50,” whereas the number of observations is much larger. As the number of variables increases, the “curse of dimensionality” makes it difficult to find units similar on all dimensions. This poses a well-known problem for exact matching, which requires observations to match on all covariates and often fails to find any matches in high-dimensional settings (Rubin and Thomas 1996, 250). Existing methods that relax the requirement of exact matching (Iacus, King, and Porro 2011) or perform dimension reduction (e.g., by matching on Mahalanobis distances or propensity scores) can suffer from poor efficiency or fail entirely with high-dimensional data.

We propose a text-matching strategy that allows analysts to effectively and transparently address confounding captured in text. We make four contributions. First, we introduce a framework for using the content of text to address confounding in observational data. Second, we propose a general text-matching adjustment strategy that involves balancing both a low-dimensional density estimate of the data and a metric that captures the probability of treatment. This approach produces matches that capture aspects of text related to treatment and facilitates qualitative comparison and evaluation. Third, we design a specific procedure, topical inverse regression matching (TIRM), to match on a jointly estimated propensity for treatment and density estimate. We show that this procedure has strong performance in a simulation study.¹ Finally, we demonstrate how to apply text matching through two applications.

A strength of matching relative to other conditioning strategies is that analysts can evaluate the quality of

the adjustment by reading treated documents alongside their matches. Comparing documents allows analysts to use substantive knowledge to recognize improved textual similarity in the matched sample even if they cannot formalize that knowledge *a priori* in a balance metric to be minimized in the matching procedure.² This type of human validation is an essential part of making comparisons in a high-dimensional and complex setting such as text (Grimmer and Stewart 2013).

Our first application extends the work of Maliniak, Powers, and Walter (2013), who find that perceived gender of international relations scholars affects citations to their articles. To address confounding, the authors controlled for article content with hand-coded variables. We show that applying TIRM recovers similar effect estimates without using the hand-coded data, suggesting that TIRM can be a viable alternative to measuring text confounders by hand. In our second application, we estimate how censorship affects the online experience of social media users in China.³ We match censored social media users to uncensored social media users who write similar social media posts and have very similar censorship histories. We find that censored social media users are more likely to be censored again in the future, suggesting that either censors are flagging users when they are censored, that social media users write about more sensitive topics after censorship, or both.

In each application, we address confounding by conditioning on text, a solution we believe is broadly applicable across social science. Scholars of American politics could match legislative bills with similar content to estimate the effect of veto threats on repositioning in Congress. Scholars of race and ethnicity might match students with similar college admissions profiles and essays to estimate the effect of perceived race on college admissions. And scholars of international relations might condition on the content of international agreements when estimating the determinants of international cooperation. Our approach could also apply to nontext data in computer vision, population genetics, biological microarrays, and other areas where a generative model of pretreatment covariates can be reliably estimated, or when the observed

¹Our primary contribution is to pose the problem of text-based confounding and offer TIRM as one possible solution. Since our paper started circulating in July 2015, Mozer et al. (2020) and Veitch, Sridhar, and Blei (2019) have introduced alternative approaches to text confounding. We hope that there will be further developments in this area.

²If the analyst can confidently define an *a priori* balance metric that captures confounding in text, they can directly optimize it using standard methods (Diamond and Sekhon 2013; Hainmueller 2011; Imai and Ratkovic 2014) while trading off between balance and sample size (King, Lucas, and Nielsen 2017). Our approach cannot obviate the need to (implicitly) choose a balance metric, but it does weaken reliance on the balance metric by facilitating human validation.

³A similar analysis is performed in Roberts (2018), who uses exact matches.

data are noisy measures of a latent confounder (Kuroki and Pearl 2014).

The article proceeds as follows. In the second section, we describe text-based confounding adjustment, explain why we opt for a matching approach, and define basic notation. We highlight the importance of conditioning on both a density estimate and a propensity score. In the third section, we present topical inverse regression matching as a way to jointly estimate the density and a propensity score. We also offer an approach to balance checking and discuss the method’s strengths, limitations, and relation to prior work. In the fourth section, we detail our two applications: a validation study demonstrating the effect of perceived author gender on academic article citations and a study estimating the effect of being censored on the reactions of Chinese social media users. The last section concludes with a discussion of future directions.

Using Text to Address Confounding

We begin by describing the setting for which we develop our approach. To fix ideas, we use one of our applications—the effects of experiencing government censorship on Chinese social media users—as a running example. In this example, our goal is to answer two questions. First, are Chinese social media users who have a post censored more likely to be censored in subsequent posts? Second, does censorship decrease the number of future posts by a user? To answer both questions, we match censored bloggers to uncensored bloggers with similar posts and similar histories of censorship and posting. We use matching to identify censorship mistakes: similar posts by different authors where one post was censored and the other was not. We find that censorship increases the probability of future censorship, but it does not have an effect on the number of posts the user writes. This provides evidence that censorship is targeted toward users who are recently censored, but that it does not induce a chilling effect on the number of posts written.


We adopt the following notation for the confounding problem. We start with a data set of n units. Each unit i is assigned treatment T_i , which takes a value of 1 for treated units and 0 for control. Under the potential outcomes framework, the outcome variable Y_i takes on the value $Y_i(1)$ when unit i is treated and $Y_i(0)$ when unit i is a control. In the censorship case, the units are individual Chinese social media users, the treatment T_i is censorship, and the outcome Y_i is the subsequent censorship rate of the social media user.

Because we have observational data, T_i is not randomly assigned and treated and control groups may not be comparable. A common practice is to match on p pretreatment covariates $\mathbf{X} = (X_1, X_2, \dots, X_p)$ to improve similarity in the distribution of covariates within treatment and control groups, a condition called balance. If we assume conditional ignorability, $T_i \perp\!\!\!\perp Y_i(0), Y_i(1) | \mathbf{X}$, then balancing the distribution of observed covariates $\mathbf{X}$ across treatment groups provides a way to estimate the average causal effect of T on Y . If we were able to exactly match each treated unit to control units, we would estimate the average treatment effect on the treated by averaging the difference between the value of Y_i for the treated unit and the value of Y_i for the matched control units.⁴

In most matching applications, $\mathbf{X}$ is low-dimensional, with $p \ll n$. We consider cases where the potential dimensionality of p is very large. For example, censorship in social media posts may be a function of thousands of particular words, particular combinations of words, hierarchies of words, and so on. As is common in the text analysis literature (Grimmer and Stewart 2013), we represent each document in a sparse count matrix $\mathbf{W}$ whose typical element, W_{ij} , contains the number of times the j th word appears within the text associated with unit i . The $\mathbf{W}$ matrix has dimension n (number of documents) by v (number of unique words in the corpus). The data are high-dimensional in the sense that v is always large relative to n . This approach largely ignores word order and combinations, though $\mathbf{W}$ can be modified to include these features.

As with conditioning in low-dimensional settings, conditioning on the text represented by $\mathbf{W}$ is appropriate when $\mathbf{W}$ is pretreatment and conditioning on it will address confounding. However, if the text is a result of treatment, rather than its cause, conditioning on the text will induce posttreatment bias. Even when the text is pretreatment, conditioning on it can cause problems. For example, if the text is an instrument, meaning it explains the outcome only through the treatment, conditioning on the text can amplify the bias of unconditioned confounders (Pearl 2011). Consider the following directed acyclic graph, which contains some small, but unavoidable, unobserved confounding U .

⁴Typically, exact matches on all observed covariates $\mathbf{X}$ are not possible, so we match treated units to the closest control units and then estimate the average treatment effect on the treated (ATT) within matches. In situations where some treated units are not similar enough to any control units, researchers typically remove the unmatchable treated units. If so, the quantity estimated is the feasible sample average treatment effect on the treated (FSATT)—the ATT for the set of sampled treated units for which matches can be found.



Conditioning on V reduces bias, but conditioning on Z would amplify bias from U . If the text contains some combination of V and Z that cannot be separated, we want to condition on the text only if the confounding (V) was large relative to the instrument (Z). We do not recommend our approach when text is posttreatment or text is unrelated to the outcome except through treatment.

In nontext settings with a transparent treatment assignment mechanism, the researcher may be confident about which p variables of $\mathbf{X}$ are necessary for the assumption of selection on observables to hold. This is not generally the case for text matching; some words in the text affect treatment, but we are unsure of which ones. If we attempt to match all words, only texts with *identical* word frequencies will match, a possibility that becomes vanishingly small as v grows large.

Our approach is to create a low-dimensional density estimate of the variables in $\mathbf{W}$ that we can use to address confounding (in concert with any nontext confounders in $\mathbf{X}$). Our function for creating this low-dimensional summary of $\mathbf{W}$ is called g , and we refer to the resulting summary as $g(\mathbf{W})$. We make the standard conditional ignorability assumption with respect to $g(\mathbf{W})$ and $\mathbf{X}$, along with standard assumptions of positivity⁵ and stable unit treatment values (SUTVA):

Assumption 1 (Conditional Ignorability). $T_i \perp\!\!\!\perp Y_i(0), Y_i(1) | g(\mathbf{W}), \mathbf{X}$.

Assumption 2 (SUTVA). For all individuals i , $Y_i(T) = Y_i(T_i)$.

Assumption 3 (Positivity). For all individuals i , $\Pr(T_i = t) > 0$ for all $t \in T$.

For Assumption 1 to be plausible, $g(\mathbf{W})$ must retain the aspects of the text that confound causal inference (much in the way we might assume that a “years of education” variable blocks confounding through education). Yet, to be useful, $g(\mathbf{W})$ must be a dramatic simplification of $\mathbf{W}$. What should this g function look like? The answer depends on the application and the way in which text confounds the treatment and outcome. To develop a useful general method, we note that in many applications, text confounds treatment assignment in two ways: through topics that affect treatment assignment and through indi-

vidual words that affect treatment assignment, regardless of topic. A useful g function for many applications is one that retains information about both of these channels and eliminates other information about the text.

The first likely channel of text confounding is through topics; treatment assignment might be related to the amount that a document discusses one or more topics. In the censorship application, for example, censorship might be related to the topic(s) that the document discusses, which could also be correlated with the subsequent behavior of the social media user. If topics of documents confound treatment assignment, then methods that summarize the topical content using density estimates become strong candidates for our g function.

The second likely channel of text confounding is through individual words that affect the probability of treatment, regardless of topic. This channel of confounding is especially likely if treatment assignment is made by humans using forms of word searching. For example, if Chinese censors sift through large numbers of posts for the keyword “protest,” they might flag posts that use this word incidentally, but are not about political protest. A more subtle version arises if treatment assignment is a function of the specific terms in which a particular topic is discussed, rather than merely the amount of that topic. For example, Chinese censors might encounter two posts with equal attention to the topic of large public marches. However, if one post talks about a public march in terms of a political protest while the other talks in terms of a holiday parade, the first is more likely to be censored. To account for this channel of text confounding, we might turn to methods such as propensity scores that can summarize how individual features affect the probability of treatment.

We prefer g functions that address text confounding from both of these channels. In the next section, we introduce a method called topical inverse regression matching (TIRM) that addresses confounding from both topics and individual words by combining elements of a topic model with elements of a propensity score model.

A distinguishing advantage of using text to adjust for confounding is that reading documents provides the analyst with a richer depiction of the units under study than what is encoded in the data matrix. This is in contrast to, for example, survey data where all known information about each respondent is encoded in the data matrix. Our matching approach leverages this advantage: Reading matched documents is a powerful tool for assessing the suitability of our g function.⁶ Analysts should read

⁵D’Amour et al. (forthcoming) show that this assumption may fail in high-dimensional data, presenting another challenge for current matching approaches that we do not take on here.

⁶Although matching has many strengths and is a popular approach (Ho et al. 2007; Sekhon 2009), we acknowledge that it lacks some

matched documents to evaluate two potential points of failure: (1) the g function has obscured an important distinction (e.g., the difference between a parade and a protest), and (2) the matching is too permissive.

Matching on both topics and propensity scores aids in this manual evaluation. Theoretically, the propensity score alone is a sufficient balancing score (Rosenbaum and Rubin 1983). However, propensity score matches are difficult to manually assess because propensity score matching approximates a fully randomized experiment rather than a block-randomized experiment (King and Nielsen 2019); documents with similar treatment probabilities may look very different. For example, propensity scores might match posts about protests to posts about pornography because they have equal probability of censorship. However, it may be difficult to distinguish this from a failure of the g function or the matching algorithm, even for an expert analyst, so we seek topically similar matches.

Text matching is only useful when conditioning on $g(\mathbf{W})$ (and other available variables $\mathbf{X}$) is sufficient to block confounding. Not only does the text need to contain relevant information about confounding, but also the summary g needs to capture that information. We expect that TIRM will often be a sufficiently rich representation of the text, but if confounding is based on textual aspects besides topics and words (e.g., sentiment, word order, punctuation, spacing, rhyme, font), then analysts should modify $\mathbf{W}$ or g accordingly. We prefer topic models for interpretability, but analysts with large data sets might find it easier to fit a simpler model, such as principal components analysis, instead. This would then require fitting a separate model for the propensity score (e.g., regularized logistic regression).

Topical Inverse Regression Matching

There are many matching approaches and many approaches to modeling text, each with strengths and weaknesses. Here, we propose one method for matching on text, topical inverse regression matching (TIRM), that includes the two attributes we believe should be present in most high-dimensional matching. First, it matches on a *coarsened* representation of the text to ensure that the resulting matches are substantively similar. Second, it uses information about how the text relates to *treatment assignment*. Our general methodology relies only on the ability to extract these two quantities: a low-dimensional

attractive theoretical properties. For example, Abadie and Imbens (2006) show that matching estimators are not $N^{1/2}$ consistent and do not attain the Hahn (1998) semiparametric efficiency bound.

TABLE 1 Overview of the TIRM Method

Step	Rationale
1. Estimate a structural topic model including the treatment vector as a content covariate.	Reduces dimension of the text
2. Extract each document's topics calculated as though treated (part of $g(\mathbf{W})$).	Ensures semantic similarity of matched texts
3. Extract each document's projection onto the treatment variable (part of $g(\mathbf{W})$).	Ensures similar treatment probability of matched texts
4. Use a low-dimensional matching method to match on $g(\mathbf{W})$ and estimate treatment effects using the matched sample.	Standardizes matching

representation of the text and the treatment propensity. TIRM uses a structural topic model (STM; Roberts, Stewart, and Airolidi 2016) as the measurement model to jointly estimate the topics (the low-dimensional representation) and a projection of information about treatment (the treatment assignment model). Although we explain our methodology in terms of STM below, other methodology could be substituted.

Estimation

STM will always estimate the topic distribution of each document. By including the treatment indicator as a content covariate in STM, we show below how to calculate a projection that captures information about treatment propensity not captured in the topics. Matching on this projection and the topic profile of the documents ensures that we will find documents that are topically similar to each other and have a similar probability of receiving treatment. Table 1 provides an overview of the complete procedure.

TIRM Step 1: Estimate STM. STM is a logistic-normal topic model that can incorporate document-specific covariates affecting both *topic prevalence* and *topical content*. Whereas prior work has focused on topic prevalence, we leverage the topical content covariate to capture the relationship between individual words and propensity to

treatment. We provide a basic overview of the model but refer readers to details in Roberts, Stewart, and Airolidi (2016).

As with the latent Dirichlet allocation (LDA) model (Blei, Ng, and Jordan 2003), we adopt a “bag of words” language model: For each token l in a document i , first sample a topic $z_{i,l}$ from the document-specific distribution θ_i and then sample the observed word from a topic-specific distribution over the vocabulary. Unlike LDA, STM allows a distribution over the vocabulary that is document-specific; B_i is a $k \times v$ matrix. We represent $z_{i,l}$ as a one-hot-encoding column vector so that $z_{i,l} B_i$ returns a v -length vector giving the distribution over the vocabulary for the particular token’s topic. Thus, each token is generated by

$$z_{i,l} \sim \text{Multinomial}_k(\theta_i), \quad \text{for } l = 1 \dots L_i; \quad (1)$$

$$w_{i,l} \sim \text{Multinomial}_v(z_{i,l} B_i), \quad \text{for } l = 1 \dots L_i. \quad (2)$$

STM allows each document to have an individual prior for θ_i based on topic prevalence covariates, but for notational simplicity we consider a shared global prior:

$$\theta_i \sim \text{LogisticNormal}_{k-1}(\mu, \Sigma), \quad (3)$$

where θ_i is on the $k - 1$ dimensional simplex and Σ is a global $k - 1 \times k - 1$ covariance matrix. Including topic prevalence covariates allows the model to share information about the value θ_i across different documents with similar covariate values. This enters the model by parameterizing μ (see Roberts, Stewart, and Airolidi 2016).

We form the document-specific topic distributions over words by a combination of a baseline word prevalence and sparse deviations from the baseline due to the topic, the content covariate, and the topic–content covariate interaction. For TIRM, we use the treatment status T_i as the content covariate and model row r (the topic) and column c (the vocabulary word) of the matrix B_i as

$$B_{i,r,c} = \frac{\exp\left(m_c + \kappa_{r,c}^{(\text{topic})} + \kappa_{T_i,c}^{(\text{cov})} + \kappa_{T_i,r,c}^{(\text{int})}\right)}{\sum_c \exp\left(m_c + \kappa_{r,c}^{(\text{topic})} + \kappa_{T_i,c}^{(\text{cov})} + \kappa_{T_i,r,c}^{(\text{int})}\right)}, \quad (4)$$

where m_c is the baseline propensity of the word, $\kappa_{r,c}^{(\text{topic})}$ is the sparse topic-specific deviation from the baseline for vocabulary entry c , $\kappa_{T_i,c}^{(\text{cov})}$ is the deviation from the baseline due to the treatment status, and $\kappa_{T_i,r,c}^{(\text{int})}$ is the deviation from the baseline due to the treatment and topic interaction. The parameters in STM are estimated using variational expectation-maximization.

The analyst selects the number of topics k that sets the granularity of the summary. As we increase the number of topics, matches will be harder to find, but more

substantively similar. It is not strictly necessary that the topics be interpretable as long as they provide an accurate density estimate. However, interpretable topics can be useful because they allow the analyst to understand the dimensions along which matched documents are similar. Interpretable topics can also facilitate selective matching on substantively important topics in confounding known to be captured in a limited subset of the document.

In the next two steps, we extract quantities of interest from this model that capture topic similarity and propensity to receive treatment. We can then apply standard matching techniques on those summary measures.

TIRM Step 2: Extract Topic Proportions. The parameter θ_i provides a measure of the document’s topical content. To ensure topics are comparable irrespective of treatment and control differences, we reestimate the topics for all control documents as though they were treated. This choice is consistent with an estimand of the (feasible-sample) average treatment effect on the *treated*. Two different words with the same topic are stochastically equivalent under the model and thus can be matched together. In this way, matching on topics can be seen as a high-dimensional analog to the coarsening step in coarsened exact matching (Iacus, King, and Porro 2011). Where coarsened exact matching coarsens within a variable (e.g., treating years 9–12 of schooling as *high school*), topics coarsen across variables (e.g., treating *tax*, *tariff*, and *economic* as part of an economy topic). When we have estimated more topics, θ_i will be longer, facilitating a more fine-grained match of substantive content. On the other hand, fewer topics will result in a shorter θ_i , which will treat more words as equivalent and create a coarser match.

TIRM Step 3: Extract Treatment Projection. While the refit θ_i captures information about the topic of the document, we want to extract the model-based information about whether or not the unit is treated. The core insight is to treat the topic model with content covariates as a multinomial inverse regression conditional on the latent variables and derive a projection in the style of Taddy (2013; see also Rabinovich and Blei 2014). Appendix B in the supporting information (SI) provides additional details, properties, and alternative strategies. Here, we overview the form of the projection for document i , which we denote p_i .

For each treatment level, the STM content covariate model learns a weight for each word ($\kappa^{(\text{cov})}$ in Equation 4) and for each topic–word combination ($\kappa^{(\text{int})}$). The projection for a given document is the sum of the document’s weighted word counts normalized by document length.

For a given level of the content covariate denoted by t ,

$$\rho_{i,t} = \frac{1}{L_i} \left(\sum_{l=1}^L \left(\underbrace{w'_{i,l} \kappa_{t,c}^{(cov)}}_{\text{weight}} + \sum_r \underbrace{w'_{i,l} I(z_{i,l} = r) \kappa_{t,r,c}^{(int)}}_{\text{topic-specific weight}} \right) \right), \quad (5)$$

where each $w_{i,l}$ is a one-hot encoding that indicates the observed word at token l . In practice, we do not observe the true topic indicators $z_{i,l}$, so we use their posterior means. Each document has a projection value for each level of the treatment that can then be included along with the topics in the matching. Given the model and the latent variables, this term captures the information about treatment assignment not contained in the topics.

TIRM Step 4: Matching and Treatment Effect Estimation. In the final step, we match on both the STM projection and the estimated topic proportions, which ensures that matches are both topically similar and have similar within-topic probabilities of treatment. Though our framework is compatible with most matching algorithms at this final stage, we generally prefer using CEM if pruning treated units is acceptable. Researchers can also match on other pretreatment covariates that are thought to be confounders in this step. In our Chinese censorship example, we also match on the day the post was written, previous posting rate, and previous censorship rate.

Most matching algorithms require analyst input beyond the choice of matching variables. For CEM, analysts must choose the values at which to “coarsen” or “bin” the matching variables. Popular CEM software uses histogram binning algorithms to make automated choices, but we prefer coarsening that is meaningful for text. Our default is to coarsen in a way that roughly corresponds to the presence or absence of each topic; for example, by using two bins—one ranging from 0 to 0.1 and another ranging from 0.1 to 1. This may seem overly permissive, but we have found that this coarsening creates a relatively demanding criterion with many topics because CEM requires matches to share the same bin across all variables. We generally use automated binning with five levels or more for the projections. A larger number of bins with more fine-grained distinctions allows for closer matches but prunes more units. These choices are necessarily application-specific. They depend on the similarity and number of documents available for matching.

Using the matched sample, we fit a model predicting the outcome as a function of treatment status and possibly other controls to estimate the effect of treatment. Inference following matching procedures has been subject to substantial debate, and we do not break new ground on this issue. Standard practice in political science is to use the standard errors from the analyst-preferred post-

matching analysis model without correction (Ho et al. 2007), which we do here. Iacus, King, and Porro (2019) show that this results in accurate inference, provided analysts are willing to change their axiom of inference from simple random sampling to stratified sampling. Some analysts have tried to construct alternative estimates of uncertainty using the bootstrap method; Abadie and Imbens (2008) show that the bootstrap method; results are inconsistent estimates of the treatment effect standard error, and they propose an asymptotically consistent alternative for some matching settings (Abadie and Imbens 2006). None of these approaches account for uncertainty in the estimation of the representation of the text. We encourage further work on accounting for uncertainty in text matching, but it is beyond the scope of this article.

Balance Checking

Balance checking—confirming that matched units are in fact similar on pretreatment confounders—is important for assessing whether matching is successful. However, we see no reason to believe there is a universally best balance metric for text similarity, and therefore, checking balance after text matching is not straightforward. We recommend several procedures.

First, we check whether words that predict treatment in the unmatched sample are balanced in the matched sample. Second, we verify that the distribution of topics in treated and control documents is similar in the matched sample. TIRM is designed to jointly minimize both of these, so if these checks reveal that matches are not adequately similar, then technical problems may be to blame, or else good matches may not exist in the data.

The TIRM procedure is designed to maximize balance on the term frequencies.⁷ However, our hope is that the procedure has picked up more general similarities in the texts. We assess this in two ways: first, by automated balance checking using a metric not directly optimized by the procedure, and second, by manual evaluation of document pairs.

For the automated balance checking we turn to string kernels, which measure similarities in sequences of characters (Spirling 2012). String kernels retain word order information that we typically discard in text models, so confirming that matching improves the string kernel similarity of treated and control texts builds confidence that even though our procedure omits word order, it still

⁷Two documents of very different lengths can be matched together if they allocate a similar proportion of their length to the same topics. If document length is a confounder, it can be included as a covariate in the matching algorithm.

improves balance in a metric that respects that ordering. In this project, we use simple graphical diagnostics (see SI Figure 5), but future work could develop formal hypothesis tests.

Finally, we check balance manually by reading matched documents. A crucial advantage of our matching approach is that it allows experts to directly scrutinize the claim that matched texts are sufficiently similar. Evaluation through reading is subjective, but it can help analysts judge whether they believe the texts are sufficient for identification and whether balance has been achieved in practice. SI Tables 4 and 6 provide some examples, though we recommend examining more pairs than we have space to present there.

Strengths and Limitations

TIRM is just one solution to the text-matching problem, but it satisfies our desiderata of producing human-verifiable matches. TIRM also estimates both document topic proportions and within-topic propensities for treatment, and, as a result, increased weight is given to words that predict treatment while the resulting matches are topically similar. This allows TIRM to prioritize variables that are related to treatment assignment while approximating a blocked design on the full set of confounders. The method is easy to apply and can be estimated with the existing `stm` software (Roberts, Stewart, and Tingley 2019) or through a new R package called `textmatching`.

One limitation of TIRM is that it requires an adequate density estimate for a complex data-generating process. Loosely speaking, the matches are only useful if the topic model is a sufficiently accurate summary of the confounding in the documents. We have found topic models to work well for this purpose. Analysts can always evaluate the quality of their model by substantively interpreting the topics, verifying that they are coherent, and considering whether documents with similar topic proportions are really similar upon close reading. The density estimate of STM can also be replaced by other density estimators that are more attuned to particular types of data (e.g., network or genetic data) or simply because better alternatives to topic models are developed in future years.

TIRM also inherits limitations that are common to other matching methods. In general, matching for causal inference requires the stable unit treatment value assumption (SUTVA; Rubin 1980), which requires any interference between units to be properly modeled. Interference between units is especially likely in applications of high-dimensional matching involving text because the purpose

of writing is often to influence or respond to the writing of others. Violations of SUTVA should be carefully considered based on the context of the analyst’s application. Like other conditioning approaches to causal inference, TIRM also requires that the selection on observables assumption is met. The core idea of text matching is that the documents themselves contain information about confounding and that the learned representation is sufficient to capture this confounding. If there are other pathways for confounding beyond the text that are not adjusted for, the procedure could be badly biased. Finally, the estimand can change as matching drops observations, particularly treated observations; if so, it is important to characterize the group to which the estimated effect applies (King, Lucas, and Nielsen 2017; Rubin 2006, 221–230). Dropping too many units can also result in a loss of efficiency.

Related Work

Before moving to applications of TIRM, we briefly mention how it relates to other approaches for similar problems. The matching literature has considered the problems of high-dimensional data, but mainly for estimating propensity scores (Belloni, Chernozhukov, and Hansen 2014; Hill, Weiss, and Zhai 2011). We do not find these approaches useful for text matching because they produce matches that have high probabilities of treatment for very different textual reasons. Our approach of augmenting propensity scores with information about topic balance is most similar to the covariate-balancing propensity scores of Imai and Ratkovic (2014). Mozer et al. (2020) and Veitch, Sridhar, and Blei (2019) directly build on our framework to propose alternative text adjustment approaches, and the related literature is reviewed in Keith, Jensen, and O’Connor (2020).

There is little work in the matching framework that proposes matching on a density estimate, as we do here. Price et al. (2006) reduce the dimensionality of genotype data with an eigen decomposition, but follow it with regression adjustment. A recent working paper by Kallus (2018) balances covariate representations based on a deep neural network with applications to image data (which shares some structural similarities to text data). Johansson, Shalit, and Sontag (2016) and Louizos et al. (2017) consider learned representations optimized for causal inference.

Finally, Egami et al. (2018) provide a framework for causal inference with text as treatment or outcome, complementing our discussion of text-based confounding. See the supporting information for further related work.

Applications and Simulations

To demonstrate the effectiveness of our text-matching approach, we present a two-part validation study in the next two subsections that builds on previous work by Maliniak, Powers, and Walter (2013). Their original study estimates the effect of perceived author gender on the citation counts of academic journal articles in the discipline of international relations. They condition on the text of the articles using variables hand-coded by research assistants based on close reading. First, we use this hand-coded data to produce a simulated data set that we use to study the performance of our proposed estimator. Next, we show that we recover a result using text matching that is similar to the Maliniak, Powers, and Walter (2013) analysis without their hand-coded data. Finally, we demonstrate the use of our methods in our motivating example of studying the effects of government censorship on Chinese social media users.

The Gender Citation Gap: Data and Simulation

If an international relations (IR) article published under a woman's name were instead published in the same venue under the name of a man with the same scholarly credentials, would it be cited more?⁸ Maliniak, Powers, and Walter (2013) say yes.⁹ Obtaining credible answers to this question is not straightforward with observational data. On average, authorial teams of different gender compositions tend to write on different topics within IR, use different methods, and have different epistemological commitments. Because these factors may affect citation counts, it is possible that lower citation counts for all-female authorial teams reflect bias against certain topics and approaches, rather than against perceived gender of the authors. Maliniak, Powers, and Walter (2013) address this challenge using information from the Teaching, Research, and International Policy (TRIP) Journal Article Database to control for the broad subfield of each article, the issue areas covered, the general methodology, paradigm,¹⁰ and epistemology. They find that academic

articles by female authors in IR have lower citation counts than articles by men or mixed-gender author teams, even after accounting for a range of potential text and non-text confounding.

We revisit the question of whether a gender citation gap exists in IR to illustrate the benefits to text matching with TIRM. With the help of JSTOR's Data for Research Program, we supplement the data from Maliniak, Powers, and Walter (2013) with the full text of 3,201 articles in the IR literature since 1980, 333 of which are authored solely by women.¹¹ We have two goals. The first is to show that TIRM can recover treatment effects in simulated data. Because simulating realistic text data is hard, we base our simulation off of the real text of articles in the IR literature, but simulate treatment effects and confounding. This subsection reports the results of these simulations. Our second goal is to demonstrate how text matching would have allowed comparable adjustment for text-based confounding without the time-consuming process of hand-coding the articles.

In order to design a simulation that has a credible joint distribution of confounder, treatment, and outcome, we use the observed text of the articles and simulate both treatment and outcomes. To avoid assuming that the topics themselves exactly capture the necessary confounding, we use one of the observed hand-coded categories, "quantitative methodology," as the true unobserved confounder in our simulation. Using the *real* article text and hand-coding, we *simulate* a treatment and outcome using the following simple model.

$$T_i \sim \text{Bernoulli}(\pi = .1X_i + .25(1 - X_i)), \quad (6)$$

$$Y_i \sim \text{Normal}(\mu = .5X_i - .2T_i, \sigma^2 = .09), \quad (7)$$


where Y , T , and X are, respectively, the outcome, the treatment, and a binary confounder indicating whether the article is hand-coded as using quantitative methodology. This defines a joint distribution over the outcome (Y_i), treatment (T_i), unobserved binary confounder (X_i taking on the value of whether each article is coded as using quantitative methodology), and observed text (W_i). We then use only information about the text (W_i) in TIRM to adjust for the unobserved confounder (X_i) to see whether TIRM can use the text to identify the form of confounding coming from the human-labeled category "quantitative methodology." This is a difficult test of TIRM's performance because it must recover a process of

⁸We are estimating the effect of *perceived* author gender in the minds of other authors making citation decisions.

⁹The finding is critiqued in Zigerell (2017) and defended in Maliniak, Powers, and Walter (2017).

¹⁰Scholarship in international relations is sometimes organized into "paradigms," or schools of thought about which factors are most crucial for explaining international relations. The predominant paradigms are realism, liberalism, and constructivism, though others exist.

¹¹We analyze more articles than Maliniak, Powers, and Walter (2013) because the TRIP database has coded more articles since 2013. However, we are missing data for a few articles used by Maliniak, Powers, and Walter (2013) because they are not in JSTOR's data set.

FIGURE 1 Gender Citation Gap Simulation

Note: The plot shows 100 simulations generated using real texts, the real hand-coded “Quantitative Methodology” variable, a simulated treatment, and a simulated outcome (according to Equations 6 and 7). Each panel shows the estimates and 95% confidence intervals from one of five different estimators, including the benchmark correct linear model specification using the unobserved confounder (True Model), our proposed estimator (TIRM), matching on only the topic proportions from the TIRM procedure (Topic Matching), matching on only the projection from the TIRM procedure (Matching on Projection), and the completely unadjusted estimator (Naive Difference-in-Means). Line types indicate whether the interval covers the truth (denoted by a dashed gray line at -0.2). TIRM achieves the best coverage out of the models we evaluate.

unobserved confounding from real texts. SI Appendix C offers additional details on the simulation and discusses some of the strengths and weaknesses of this design.

We produce 1,000 simulated data sets for analysis. In Figure 1, we show the results of each model for a subset of 100 simulations. For each simulated data set, we plot the treatment effect estimate and 95% confidence interval for five estimators: the true linear model using the unobserved quantitative methodology variable,

the TIRM model,¹² matching on the topics only from the TIRM model, matching on the projections from the TIRM model, and the unadjusted difference-in-means estimator. Table 2 provides summary statistics for all 1,000 simulations.

¹²For each simulated data set, we apply the TIRM procedure using 15 topics and matching with CEM on the projection score for treatment using eight automatically generated bins and the topics using two bins each ($0-0.1$ and $0.1-1$).

TABLE 2 Demonstration of TIRM Estimator Performance and Relevant Benchmark Estimators across 1,000 Simulations

	MSE	Bias	Coverage	Avg. # Treated	Avg. # Matched
True Model	0.00018	0.00032	0.955	619.6	3,201.0
TIRM	0.00099	−0.02214	0.826	392.8	1,447.2
Topic Matching	0.00144	−0.03377	0.547	548.8	2,142.8
Matching on Projection	0.00312	−0.03220	0.456	496.2	2,693.4
Naive Difference-in-Means	0.01299	−0.11278	0.000	619.6	3,201.0

Note: Results are compared on mean squared error (MSE), bias, coverage of the 95% confidence interval, average number of treated units in the matched set, and average number of treated and control units in the matched set.

The TIRM model performs substantially better than matching only on the topics or the projection, particularly in terms of bias and coverage. Although TIRM’s 82% coverage rate on this example does not match the nominal 95%, it performs well given that the true confounder has been withheld from the model.

The Gender Citation Gap: Female IR Scholars Are Cited Less Often

We now return to the substantive question motivating our reanalysis of the data from Maliniak, Powers, and Walter (2013): Are female IR scholars cited less often due to perceived gender? Maliniak, Powers, and Walter (2013) control for differences in the writing of men and women using qualitative variables painstakingly hand-coded by research assistants over more than a year. We recover similar effects with TIRM using far less manual effort.


Women’s and men’s writings are not the same, on average, in the unmatched data set. Among other differences, women are more likely to write about gender (*women*, *gender*, and *children*), whereas men or male/female coauthored articles are more likely to use words associated with statistical methods (*model*, *estimate*, and *data*). We also find substantial imbalance in the human-coded variables measuring each article’s research approach, paradigm, methodology, and issue area.

We apply TIRM to the data, specifying 15 topics in the STM portion of the algorithm, using all-female authorship as the treatment variable. This is a relatively small number of topics to span the diversity of the IR literature, so this model recovers relatively broad topics. We use CEM to match on the TIRM output, along with three nontext variables from Maliniak, Powers, and Walter (2013): (1) whether at least one author is tenured; (2) whether the article appeared in the *American Political Science Review*, *American Journal of Political Sci-*

ence, or *Journal of Politics*; and (3) article age. We drop treated units with no available matches, so it is important to note how the treated pool changes after matching. On average, women’s articles in the *matched* sample are 3 years older and 8 percentage points more likely to have a tenured author than women’s articles in the *full* sample. The matched sample is also heavier on topics related to security and political economy and lighter on international organizations and institutions.

We compare the performance of TIRM to alternative approaches: matching only on propensity scores, matching only on topics, and exact matching on the original human-coded data. We first look for improvements in balance among the words that are most imbalanced in the unmatched data set. We identify the words that have high *mutual information* with author gender in the raw data set as well as the matched data from TIRM, matching only on topics, and human matching.¹³ Figure 2 shows the relationship between the difference in word occurrence by gender and the mutual information of each word in each data set. If perfect balance on all words were possible, we would hope to see every word lined up vertically on the $x = 0$ line (and in sans-serif typeface accordingly). However, since not all words can be balanced, balance on words with high mutual information is most important. TIRM—shown in the bottom-right panel—outperforms the others in balancing the high mutual information words. Many high mutual information words such as *interview* that were previously imbalanced are now lined up down the center. TIRM makes the imbalance substantially worse on words with low mutual information, but because these words have low imbalance

¹³We calculate the mutual information for an individual word w as the difference between the entropy of category k , $H(k)$, and the entropy of category k when conditioning on a word’s appearance in the document, $H(k|w)$. $H(k) - H(k|w)$ can be calculated as follows: $H(k) - H(k|w) = \sum_{t=0}^1 \sum_{s=0}^1 P(k=t, w=s) \log_2 \frac{P(k=t, w=s)}{P(k=t)P(w=s)}$. See Grimmer (2010) and Manning, Raghavan, and Schütze (2008) for a longer treatment of mutual information.

FIGURE 2 Relationship between Mutual Information and Difference in Word Occurrence

Note: (all female - male/coed) in a) the full data set; b) topic matched; c) matched on human coding; and d) TIRM. In panels b, c, and d, words for which matching decreased the absolute value of the difference in word appearance are in sans-serif typeface and words for which matching increased the absolute value of the difference in word appearance are in italicized serif typeface. Darker words changed more in either direction. Word size is proportional to each word's mutual information

to begin with, we anticipate that they will not substantially increase bias. This analysis highlights the benefits of the treatment model because it can identify and address the most imbalanced words—exact matching on human coding and matching only on topics do not perform as well as TIRM.

We also check balance in other ways. We manually examine pairs of matched documents to confirm these pairs match our intuitions about which articles in the IR

literature are similar enough for comparison. We read more pairs than we can conveniently present, but see SI Table 4 for a few examples. We also evaluate TIRM's success at balancing the human-coded variables of article substance and the estimated topics from the topic model. We find that TIRM performs reasonably well at balancing the human-coded variables, particularly the most imbalanced ones. This is reassuring because in most applications, the purpose of TIRM is to substitute for painstaking

human coding of texts. We also find that TIRM balances the STM topics well. Matching on topics only (without the TIRM projection) performs slightly better at balancing the STM topics, but that is to be expected; TIRM is trying to simultaneously balance the estimated topics *and* the probability of treatment. Finally, we check string kernel similarity of the matched data sets and find that TIRM outperforms the alternatives and offers substantial improvement over the raw data. Details of these balance checks are in SI Appendix D.

Maliniak, Powers, and Walter (2013, 906) report that women's articles receive 4.7 fewer citations on average. Our estimates from similar specifications using the TIRM matched sample show a slightly more pronounced gender gap of 6.5 citations, most of which seems concentrated among the highest citation-earning articles. Because this article is focused on the method, we refer readers to SI Appendix D for further results and a sensitivity analysis.

Government Censorship of Chinese Social Media Users

The Chinese government oversees one of the most sophisticated censorship regimes in the world (Esarey and Qiang 2008; MacKinnon 2011; Marolt 2011), with technologies ranging from manipulating search results to blocking foreign websites. Even as we learn more about the types of content that are censored and the technical infrastructure enabling censorship (Bamman, O'Connor, and Smith 2012; King, Pan, and Roberts 2013, 2014), we still know little about the subsequent impacts of censorship on individual social media users in China. Is censorship completely determined by the text of a particular post, or does censorship become more targeted toward users based on their previous censorship history? This issue is particularly important; the targeted use of artificial intelligence for censorship has become a cause of concern because it further complicates the ability of the public to hold the government accountable for censorship decisions (Morozov 2012; Roberts 2018; Tufekci 2014).

Our goal is to estimate the causal effect of experiencing censorship on the subsequent censorship and posting rates of social media users in China. The Chinese censorship regime appears to attempt to thoroughly remove social media posts it deems sensitive (King, Pan, and Roberts 2013, 2014), but the censors miss some. We exploit these mistakes by using TIRM to identify pairs of nearly identical social media posts written by nearly identical users, where one is censored and the other is not. We can then observe the subsequent censorship rates of both


users to estimate the causal effect of censorship on the treated units who remain in our sample.

We use data on 4,685 social media users on Weibo from the Weiboscope data set (Fu, Chan, and Chau 2013).¹⁴ Fu, Chan, and Chau (2013) collected posts from Weibo in real time and then revisited these posts later to observe whether they were censored. After processing the text to account for linguistic features of Chinese, we convert the text of social media posts into a matrix of document word counts and estimate TIRM using 100 topics. There may be confounding factors not captured in the text if, for example, the government uses nontext information to make censorship decisions. To address these possibilities, we also match on nontext confounders: previous post rate, previous censorship experience, and date. We assume these covariates capture pretreatment differences in user targeting. Our final matched sample contains 879 posts. Since we drop treated units without matches, we describe the differences between our matched and unmatched samples. On average, the matched sample has a slightly lower previous censorship rate, is more likely to discuss environmental protests such as those in Shifang and Qidong in 2012, and is less likely to discuss human rights lawyers and central politics. Additional details are in the supporting information.

For illustration, we compare matching with TIRM to matching only on the topics or only on the projection. We find that the TIRM is effective at identifying social media posts about the same topic or event, but with different censorship statuses. Figure 3 shows that TIRM matching outperforms other matching strategies in reducing the difference between topics in censored and uncensored posts. Topic matching is a close second, and matching only on the projection performs poorly. TIRM also outperforms the others at improving the string kernel similarity of documents in the matched data set (details in SI Figure 7). TIRM allows us to manually evaluate the similarity of the matched documents, some of which we show in Table 3.

Our first outcome measure is the censorship rate of each blogger after the matched post. Our second outcome is the rate of posting after censorship. The results show that censored users are likely to experience more censorship in the future as a result of their censorship experience. Having a post censored increases the

¹⁴Due to data size, we selected only users censored at least once in the first half of 2012 and study their behavior in the last half of 2012, and we restrict the control donation sample to posts made on the same day and with a cosine similarity to a censored post greater than 0.5.

FIGURE 3 Topic Balance Comparison

Note: The top panel shows the balance for unmatched, projection-only matched, topic-only matched, and TIRM-matched for the most unbalanced topics in the unmatched data set. The bottom panel shows the distribution of the absolute values of topical differences for all 100 topics under each type of matching. TIRM outperforms the other matching methods.

probability of future censorship significantly,¹⁵ but it does not decrease the number of posts written by the censored user. This suggests one of two scenarios. This evidence is

¹⁵The probability of a “Permission denied” post is not different between the two groups in the pretreatment period, but it is 0.009 for the censored group and 0.004 for the uncensored group in the period after treatment. The probability of a “Weibo does not exist” post is not different between the two groups in the pretreatment period, but it is 0.25 for the censored group and 0.20 for the uncensored group in the period after treatment.

consistent with algorithmic targeting of censorship, where social media users are more likely to be censored after censorship because they are flagged by the censors. Alternatively, social media users may chafe against censorship and respond by posting increasingly sensitive content that is more likely to be censored. Either way, these results are consistent with Roberts (2018), who finds a similar pattern using exact matches, and indicate that users do not seem to experience a chilling effect of censorship. In

TABLE 3 Translations of example social media posts that were censored (left) with matched uncensored social media posts selected by TIRM (right)

Censored Post	Uncensored Post
There may be even bigger plans: When the chaos escalate to a certain degree, there will be military control, curfew, Internet cutoff, and then complete repression of counterrevolution. There are precedent examples.	The person on the right (refers to the previous comment) knew too much. Bang (Sound effect for gunshot)! You knew too much! Bang! There may be even bigger plans: When the chaos escalate to a certain degree, there will be military control, curfew, Internet cutoff, and then complete repression of counterrevolution. There are precedent examples.
#Weitianxia#Shifang netizen's disclose: I saw police officers looking for restaurants to eat on the street and the intestine vermicelli restaurant owner immediately said they don't sell it to the police. Then everyone on the street came out and yelled, which was very impressive. Now many stores have signs saying that police tactical units are not allowed to enter. Shifang people said: F*k you, you beat us up, bombed us, and still ask us to feed you, why don't you eat sh*t?	Due to the lack of prior publicity procedures, some people are unfamiliar, uncomprehending and unsupportive of this program. To respond to the general public's request, the municipal party committee and the government researched and decided to stop the project. Shifang will not ever develop the Molybdenum Copper project in the future.
[17-year-old young athlete fails 3 attempts to lift The media calls it a shame of Chinese female weightlifters] According to Sina: Chinese female weightlifters faced a shameful failure of its Olympic history last night! During the female 53kg weightlifting competition, joined as the black horse, Zhou Jun, a 17-year-old young athlete from Hubei, failed in all 3 of her attempts and ended with no result, which ends her Olympic journey. Many media reported this using "the most shameful failure of Chinese female Olympic weightlifters" as the title.	[17-year-old young athlete fails 3 attempts to lift The media calls it a shame of Chinese female weightlifters] According to Sina: Chinese female weightlifters faced a shameful failure of its Olympic history last night! During the female 53kg weightlifting competition, joined as the black horse, Zhou Jun, a 17-year-old young athlete from Hubei, failed in all 3 of her attempts and ended with no result, which ends her Olympic journey. Many media reported this using "the most shameful failure of Chinese female Olympic weightlifters" as the title. I personally think, it is not a shame of Zhou Jun, but a shame of Chinese media!

the supporting information, we explore sensitivity to unobserved confounding and the choice of CEM coarsening.

Conclusion

Scholars across the social sciences are finding an increasing number of ways to use text as data. In this article, we have proposed conditioning on text to address confounding using matching. We identify the core concerns for addressing confounding from text, provide a method for text matching, and introduce approaches to balance checking. Matching text is difficult because it is inherently high-dimensional; we address this concern with a simple approach matching on a density estimate and a projection that captures propensity to treatment. To assist applied researchers wishing to make causal inferences with high-

dimensional data, we provide the `textmatching` package in R, which uses the `stm` package (Roberts, Stewart, and Tingley 2019) to implement the matching procedures described in this article. This general strategy may have applications in other types of high-dimensional data.

Our interest in text matching is born out of a practical necessity from the applications we present. There are an enormous number of research problems in which the content of texts potentially confounds causal inference in observational studies, and the different characteristics of our case studies reflect this diversity. We have used these methods in several languages, with corpora of varying size, and typical document lengths as short as a couple of sentences to roughly 10,000 words. These applications suggest that our solution has the flexibility to address the tremendous variety characteristic of social science data.

Text matching is not a panacea for observational causal inference, and perfectly balancing on text does

not guarantee unbiased causal inference. In both our applications, we had to condition on additional nontext variables to block other paths of confounding. Credible causal inference requires careful design and domain knowledge. This domain knowledge may suggest other approaches to text adjustment, such as supervised measurement of a specific quantity. Regardless, we encourage applied users to leverage the greatest strength of text-based adjustment: the ability to carefully read documents to assess what the model can and cannot measure.

Social science has made tremendous progress in developing text-as-data methods, but we are only at the beginning of developments at the intersection of text and causal inference. We see several opportunities for pushing forward the text-based confounding literature. We hope that scholars will extend our work to a proposed alternative to TIRM, a task already started by Mozer et al. (2020) and Veitch, Sridhar, and Blei (2019). A central challenge is developing general-purpose methods for evaluating these new models. Future work could also introduce additional approaches to model criticism and balance checking to fill out the text-matching workflow. In addition to these aspects of practice, key aspects of theory remain to be addressed, including rigorous approaches to uncertainty, consistency results that allow our model to function as only an approximation to g , and implications of positivity assumptions in high dimensions.

References

- Abadie, Alberto, and Guido W. Imbens. 2006. "Large Sample Properties of Matching Estimators for Average Treatment Effects." *Econometrica* 74(1): 235–67.
- Abadie, Alberto, and Guido W. Imbens. 2008. "On the Failure of the Bootstrap for Matching Estimators." *Econometrica* 76(6): 1537–57.
- Bamman, David, Brendan O'Connor, and Noah Smith. 2012. "Censorship and Deletion Practices in Chinese Social Media." *First Monday* 17(3). <https://journals.uic.edu/ojs/index.php/fm/article/view/3943>
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014. "Inference on Treatment Effects after Selection among High-Dimensional Controls." *Review of Economic Studies* 81(2): 608–50.
- Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 3(jan): 993–1022.
- D'Amour, Alexander, Peng Ding, Avi Feller, Lihua Lei, and Jasjeet Sekhon. Forthcoming. "Overlap in Observational Studies with High-Dimensional Covariates." *Journal of Econometrics*.
- Diamond, Alexis, and Jasjeet S. Sekhon. 2013. "Genetic Matching for Estimating Causal Effects: A General Multivariate Matching Method for Achieving Balance in Observational Studies." *Review of Economics and Statistics* 95(3): 932–45.
- Egami, Naoki, Christian J. Fong, Justin Grimmer, Margaret E. Roberts, and Brandon M. Stewart. 2018. "How to Make Causal Inferences Using Texts." arXiv preprint arXiv:1802.02163.
- Esarey, Ashley, and Xiao Qiang. 2008. "Political Expression in the Chinese Blogosphere: Below the Radar." *Asian Survey* 48(5): 752–72.
- Fu, King-wa, Chung-hong Chan, and Michael Chau. 2013. "Assessing Censorship on Microblogs in China: Discriminatory Keyword Analysis and the Real-Name Registration Policy." *IEEE Internet Computing* 17(3): 42–50.
- Grimmer, Justin. 2010. "A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases." *Political Analysis* 18(1): 1–35.
- Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3): 267–97.
- Hahn, Jinyong. 1998. "On the Role of the Propensity Score in Efficient Semiparametric Estimation of Average Treatment Effects." *Econometrica* 66(2): 315–31.
- Hainmueller, Jens. 2011. "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies." *Political Analysis* 20(1): 25–46.
- Hill, Jennifer, Christopher Weiss, and Fuhua Zhai. 2011. "Challenges with Propensity Score Strategies in a High-Dimensional Setting and a Potential Alternative." *Multivariate Behavioral Research* 46(3): 477–513.
- Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15(3): 199–236.
- 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): 345–61.
- Iacus, Stefano M., Gary King and Giuseppe Porro. 2019. "A Theory of Statistical Inference for Matching Methods in Causal Research." *Political Analysis* 27(1): 46–68.
- Imai, Kosuke, and Marc Ratkovic. 2014. "Covariate Balancing Propensity Score." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76(1): 243–63.
- Johansson, Fredrik, Uri Shalit, and David Sontag. 2016. "Learning Representations for Counterfactual Inference." In *Proceedings of the 33rd International Conference on Machine Learning*, New York, NY, USA.
- Kallus, Nathan. 2018. "DeepMatch: Balancing Deep Covariate Representations for Causal Inference Using Adversarial Training." arXiv preprint arXiv:1802.05664.
- Keith, Katherine, David Jensen, and Brendan O'Connor. 2020. "Text and Causal Inference: A Review of Using Text to Remove Confounding from Causal Estimates." arXiv preprint arXiv:2005.00649, University of Massachusetts Amherst.
- King, Gary, Christopher Lucas, and Richard A. Nielsen. 2017. "The Balance-Sample Size Frontier in Matching Methods

- for Causal Inference.” *American Journal of Political Science* 61(2): 473–89.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. “How Censorship in China Allows Government Criticism but Silences Collective Expression.” *American Political Science Review* 107(1): 1–18.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2014. “Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation.” *Science* 345(6199): 1251722. <https://gking.harvard.edu/files/gking/files/chinasci2.pdf>
- King, Gary, and Richard Nielsen. 2019. “Why Propensity Scores Should Not Be Used for Matching.” *Political Analysis* 27(4): 435–54.
- Kuroki, Manabu, and Judea Pearl. 2014. “Measurement Bias and Effect Restoration in Causal Inference.” *Biometrika* 101(2): 423–37.
- Louizos, Christos, Uri Shalit, Joris M. Mooij, David Sontag, Richard Zemel, and Max Welling. 2017. “Causal Effect Inference with Deep Latent-Variable Models.” In *Advances in Neural Information Processing Systems*. 6446–56. 31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.
- MacKinnon, Rebecca. 2011. “China’s ‘Networked Authoritarianism.’” *Journal of Democracy* 22(2): 32–46.
- Maliniak, Daniel, Ryan Powers, and Barbara F. Walter. 2013. “The Gender Citation Gap in International Relations.” *International Organization* 67(4): 889–922.
- Maliniak, Daniel, Ryan Powers, and Barbara Walter. 2017. “A Reply to ‘Reducing Political Bias in Political Science Estimates.’” *PS: Political Science & Politics* 50(1): 184–85.
- Manning, Christopher D., Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. New York, NY: Cambridge University Press.
- Marolt, Peter. 2011. “Grassroots Agency in a Civil Sphere? Rethinking Internet Control in China.” In *Online Society in China: Creating, Celebrating, and Instrumentalising the Online Carnival*, ed. David Herold and Peter Marolt. New York: Routledge, 53–68.
- Morozov, Evgeny. 2012. *The Net Delusion: The Dark Side of Internet Freedom*. New York: PublicAffairs.
- Mozer, Reagan, Luke Miratrix, Aaron Russell Kaufman, and L. Jason Anastasopoulos. 2020. “Matching with Text Data: An Experimental Evaluation of Methods for Matching Documents and of Measuring Match Quality.” *Political Analysis*, 1–24. <https://10.1017/pan2020.1>
- Pearl, Judea. 2011. “Invited Commentary: Understanding Bias Amplification.” *American Journal of Epidemiology* 174(11): 1223–27.
- Price, Alkes L., Nick J. Patterson, Robert M. Plenge, Michael E. Weinblatt, Nancy A. Shadick, and David Reich. 2006. “Principal Components Analysis Corrects for Stratification in Genome-Wide Association Studies.” *Nature Genetics* 38(8): 904–09.
- Rabinovich, Maxim, and David Blei. 2014. “The Inverse Regression Topic Model.” In *Proceedings of the 31st International Conference on Machine Learning*. 199–207.
- Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China’s Great Firewall*. Princeton, NJ: Princeton University Press.
- Roberts, Margaret E., Brandon M. Stewart, and Edoardo M. Airoldi. 2016. “A Model of Text for Experimentation in the Social Sciences.” *Journal of the American Statistical Association* 111(515): 988–1003.
- Roberts, Margaret E., Brandon M. Stewart, and Dustin Tingley. 2019. “stm: An R Package for Structural Topic Models.” *Journal of Statistical Software* 91(2): 1–40.
- 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.
- Rubin, Donald B. 1980. “Discussion of ‘Randomization Analysis of Experimental Data in the Fisher Randomization Test.’” *Journal of the American Statistical Association* 75: 591–93.
- Rubin, Donald B. 2006. *Matched Sampling for Causal Effects*. New York: Cambridge University Press.
- Rubin, Donald B., and Neal Thomas. 1996. “Matching Using Estimated Propensity Scores: Relating Theory to Practice.” *Biometrics* 52: 249–64.
- Sekhon, Jasjeet S. 2009. “Opiates for the Matches: Matching Methods for Causal Inference.” *Annual Review of Political Science* 12: 487–508.
- Spirling, Arthur. 2012. “U.S. Treaty Making with American Indians: Institutional Change and Relative Power, 1784–1911.” *American Journal of Political Science* 56(1): 84–97.
- Taddy, Matt. 2013. “Multinomial Inverse Regression for Text Analysis.” *Journal of the American Statistical Association* 108(503): 755–70.
- Tufekci, Zeynep. 2014. “Engineering the Public: Big Data, Surveillance and Computational Politics.” *First Monday* 19(7).
- Veitch, Victor, Dhanya Sridhar, and David M. Blei. 2019. “Using Text Embeddings for Causal Inference.” arXiv preprint arXiv:1905.12741.
- Zigerell, L. J. 2017. “Reducing Political Bias in Political Science Estimates.” *PS: Political Science & Politics* 50(1): 179–83.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix A: Additional Related Work

Appendix B: Treatment Projection

Appendix C: Simulation Details

Appendix D: The Gender Citation Gap: Balance Checking, Results, and Sensitivity Analysis

Appendix E: Details of Chinese Social Media User Analysis

Supplemental Information (Online Only)

Adjusting for Confounding with Text Matching

January 14, 2020

Table of Contents

A	Additional Related Work	1
B	Treatment Projection	3
C	Simulation Details	8
D	The Gender Citation Gap: Balance Checking, Results, and Sensitivity Analysis	10
E	Details of Chinese Social Media User Analysis	18

A Additional Related Work

In this short section we highlight some additional related work that we were unable to cite in the main text because of space constraints.

One of the most exciting frontiers in causal inference right now is the use of machine learning methods for causal inference including many methods which could be used to adjust for high-dimensional covariates in experimental and observational settings (Van der Laan and Rose, 2011; Bloniarz et al., 2015; Hazlett, Forthcoming; Sales, Hansen and Rowan, 2015; Athey and Imbens, 2016; Athey, Imbens and Wager, 2016; Hartford et al., 2016; Chernozhukov et al., 2017; Ratkovic and Tingley, 2017). Most of these approaches leverage models of the outcome (see also, Rubin and Thomas, 2000; Hansen, 2008). By contrast our approach is focused on the analysis of observational data, falls in the matching framework, and does not use a separate regression model of the outcome data. There has been some work on the particular problem of using high-dimensional data for estimating propensity scores (Schneeweiss et al., 2009; Westreich, Lessler and Funk, 2010; Hill, Weiss and Zhai, 2011; Belloni, Chernozhukov and Hansen, 2014).

Although not about matching, Taddy (2013a) offers an approach that is conceptually related to ours: considering how to select documents for manual coding in supervised learning. Ideally, manual coding should use an optimally *space filling design*, but this is impractical in high-dimensions. Taddy proposes a topic model followed by a *D-optimal space filling design* in the lower-dimensional topic space. Both of these approaches share our intuition that if two features in a high-dimensional covariate set commonly co-occur, then they can be treated interchangeably to identify appropriate counterfactual cases.

There are several recent lines of work considering embeddings and density estimates as proxies for unobserved confounding (Veitch, Wang and Blei, 2019; Sridhar and Getoor, 2019; Wang and Blei, 2019; Yao et al., 2019; Tran and Blei, 2017). These are also connected to broader examinations of the role of text representations in causal inference (Egami et al., 2017; Wood-Doughty, Shpitser and Dredze, 2018).

In balance checking with string kernels we primarily use visual diagnostics here. However, there is a framework for formal hypothesis tests using the Minimum Mean Discrep-

ancy framework developed in Gretton et al. (2012) (see also, Gretton et al., 2007; Sejdinovic et al., 2013; Szabó et al., 2015). For practical purposes these tests would need to be made more computationally efficient (Zhang et al., N.d.) and altered to reflect the appropriate null hypothesis for a balance test (Hartman and Hidalgo, 2018). Those developments are beyond the scope of this paper.

B Treatment Projection

In this appendix section we derive the the treatment projection discussed in Section 3.1.3. As a reminder of the notation: $w_{i,l}$ is a one-hot-encoding vector indicating the observed word in document i at token l . $z_{i,l}$ is a categorical variable indicating the topic of that token and κ are word weights (with parenthetical superscripts indicating whether they correspond to parameters for the topics, content covariate, or interaction between topics and content covariates). Equation 5 from Section 3.1.3 provides the projection which is reproduced here:

$$\rho_{i,t} = \frac{1}{L_i} \left(\sum_{l=1}^L \left(w'_{i,l} \underbrace{\kappa_{t,c}^{(cov)}}_{\text{weight}} + \sum_r^k w'_{i,l} \underbrace{\overbrace{I(z_{i,l} = r)}^{\text{topic indicator}} \kappa_{t,r,c}^{(int)}}_{\text{topic-specific weight}} \right) \right) \quad (8)$$

Although notationally dense, the projection has the straightforward interpretation of summing up two weights for each word appearing in the document and normalizing by document length. The first weight is specific to the entry of the vocabulary (e.g. **parade** and **protest** have different weights) while the second weight is specific to the entry of the vocabulary and that token’s topic (e.g. **parade** has one weight under Topic 3 but a different weight under Topic 2).

Connections to Inverse Regression We arrive at this projection by noting that conditional on the token-level latent variables z (which denote the topic of a given word token), the structural topic model with content covariates has the form of the multinomial inverse regression (MNIR) in Taddy (2013b). In that work, Taddy derives a projection for the MNIR model and proves that it satisfies classical sufficiency for the outcome such that given the model and the parameters, the treatment indicator is independent of the words given the projection (to use our example). Given this low-dimensional representation, Taddy (2013b) then fits the low-dimensional forward regression which predicts the treatment indicator using the projection. We don’t actually need this final step because we are matching on the projection itself (this is roughly analogous to the practice of matching on the linear predictor in the propensity score).

Rationale for Projection The rationale for using the projection is the same as in Taddy’s work: efficiency. As explained in the rejoinder to the original paper (Taddy, 2013c), using the inverse regression provides efficiency gains relative to the forward regression which derive from assuming a generative model for the words. Given the generative model, the variance on the projection decreases in the number of words rather than the number of documents. Even when the generative model does not hold, this can provide substantial gains in practice.

Advantages of Joint Estimation In our setting, the inverse regression formulation affords us two additional advantages. First, we can allow words to have different weights depending on their context (as captured through the topics). For example, in the censorship example we can allow for the possibility that certain words may always increase your odds of censorship while others are only sensitive in particular situations. Second, we avoid redundant information between the topics and the words.

Properties In Taddy (2013b) there are no topics or interactions between topics and covariates. Here we observe that his Propositions 3.1 and 3.2 establishing sufficiency of the projection conditional on the parameters of the model (including the document level random effects), extend to our setting as well by conditioning on the token-level latent variables $z_{i,l}$. We show that our model can be written in that form, following closely on Taddy (2013b, page 758)

$$\begin{aligned}
w_{i,l} &\sim \text{Multinomial}(q_{i,l}, 1) \\
q_{i,l,c} &= \frac{\exp(\eta_{i,l,c})}{\sum_c \exp(\eta_{i,l,c})} \\
\eta_{i,l,c} &= \underbrace{m_c}_{\text{baseline}} + \underbrace{\left(\sum_r^k I(z_{i,l} = r) \kappa_{r,c}^{(\text{topic})} \right)}_{\text{topic}} + \underbrace{\sum_a^b I(T_i = a) \kappa_{a,c}^{(\text{cov})}}_{\text{treatment}} + \\
&\quad \underbrace{\sum_a I(T_i = a) \left(\sum_r^k I(z_{i,l} = r) \kappa_{a,r,c}^{(\text{int})} \right)}_{\text{treatment-topic-interaction}}
\end{aligned}$$

The data $w_{i,l}$ and $q_{i,l}$ are v -length column vectors representing the one-hot encoding vector of the observed data in token l and the probability vector that draws it respectively. All other terms are scalars. Entries in the vocab are indexed by c and runs to v , the levels of the content covariate is indexed by a and runs to b , the topics are indexed by r and runs to k .

We can rewrite this in a more compact notation by suppressing the dependence on i and writing m as a v -length column vector, T as a b -length column vector (one-hot encoding), and z_l as a k -length column vector (one-hot encoding). We use $\kappa^{(\text{topic})}$ to denote the v -by- k matrix of topic parameters and $\kappa^{(\text{cov})}$ to denote the v -by- a matrix of covariate parameters and $\kappa_r^{(\text{int})}$ to indicate the v -by- a matrix of interaction parameters for the r -th topic. We can now write the v -length vector η_l as

$$\begin{aligned}\eta_l &= m + \kappa^{(\text{topic})'} z_l + \kappa^{(\text{cov})'} T + \sum_{c=1}^k \kappa_r^{(\text{int})'} T \\ &= m + \kappa^{(\text{topic})'} z_l + \Phi' T\end{aligned}$$

where $\Phi = \kappa^{(\text{cov})} + \sum_{r=1}^k \kappa_r^{(\text{int})}$ collects the v -by- a matrix of word weights.

Rewriting the components that do not depend on T as $\alpha = m + \kappa^{(\text{topic})'} z_l$, we can write the likelihood in exponential family form

$$\begin{aligned}\exp(w_l' \eta_l - A(\eta_l)) &= \exp(w_l' \alpha) \exp((w_l' \Phi) T - A(\eta)) \\ &= h(w) g(\Phi' w, T)\end{aligned}$$

where $A(\eta) = \log(\sum_c \exp(\eta_c))$ is the log-partition function. The form of the model is now the same as in Taddy (2013b) and the remainder of his proof follows, with standard sufficiency results for the exponential family implying that $p(w_l | \Phi' w_l, T) = p(w_l | \Phi' w_l)$. Proposition 3.2 follows analogously and establishes that the reduction still holds when we normalize for document length.

Limitations As in Taddy (2013b), it is worth emphasizing that sufficiency only holds conditional on the latent variables (in our setting $z_{i,l}$). We subsequently include the

document level topics θ_i when we are doing our matching. This is to say that we don't invest too heavily in the sufficiency result, rather the results simply provide an intuition for why these word weights would be helpful in understanding the propensity to be treated.

Alternatives Taddy (2013b) leaves open the question of the best way to use the latent structure in the projection and there is likely still further work to be done on this point. Rabinovich and Blei (2014) introduce a simpler projection in the context of their inverse regression topic model (IRTM). The model structure is similar to the Structural Topic Model with a content covariate and the predecessor of both models, the Sparse Additive Generative model of text (Eisenstein, Ahmed and Xing, 2011). In IRTM, the probability of observing word c from topic r in document i is given by

$$\beta_{i,r,c} \frac{\beta_{r,c} \exp(\Phi_c y_i)}{\sum_c \beta_{r,c} \exp(\Phi_c y_i)}$$

where β_r is a draw from a Dirichlet distribution, Φ_c is a draw from a Laplace distribution and y_i is a *continuous* covariate (NB: we have adopted their notation from equation 1 of their paper except for the indices we have changed to match ours). Thus, their representation of a topic is a background topic distribution (β_r) multiplied by a distortion factor ($\exp(\Phi_c y_i)$) where the argument of $\exp()$ is sparse, leading the distortion factor to often be 1. We can rewrite the STM model to look more like this in order to clarify the connections,

$$B_{i,r,c} \propto \overbrace{\exp(m_c + \kappa_{r,c}^{(\text{topic})})}^{\beta_{r,c}} \underbrace{\exp(\kappa_{T_{i,c}}^{(\text{cov})})}_{\exp(\Phi_c y_i)} \overbrace{\exp(\kappa_{y_d,k,v}^{(\text{int})})}^{\text{no equivalent}}.$$

The first section is an alternate way of representing the background topic distribution. In IRTM, each topic r is a draw from a Dirichlet distribution, in STM it is a log-linear model with a dense vector shared by all topics and a sparse, topic-specific deviation. The second section is mostly the same with the distinction that in IRTM the covariate is continuous in 1-dimension and in STM it is categorical. The third chunk in STM which captures the topic-covariate interaction has no equivalence in the IRTM. The IRTM performs MAP

estimation using a stochastic variational EM algorithm.

Rabinovich and Blei (2014) compute the analogous projection from their model ($\frac{w'\Phi}{L}$) and find that using the projection *alone* is not very effective for prediction. They instead opt to compute the MAP estimate of the content covariate. Performing the joint optimization of the content covariate and the topics is complicated and Rabinovich and Blei (2014) employ a coordinate ascent estimation strategy that would be even more complicated (and slow) in our setting, but it is a direction to possibly explore in future work. Because we are already conditioning on the topics, we would expect better performance than the initial Rabinovich and Blei (2014) tests on the simple projection.

C Simulation Details

In this appendix we provide the technical details of the simulation which we summarized in Section 4.1. To review the motivating logic, we wanted to avoid simulating new documents because simulating text from the TIRM model itself would make strong and unrealistic assumptions about the world. We know of no way to provide a realistic model for treatment assignment, the outcome, and the confounder at the same time while still also knowing the true causal effect. We opted instead to use the structure of the female IR scholarship application in Section 4.2 to simulate a confounded treatment assignment and outcome using the real texts and hand-coding of a plausible confounder (a binary variable indicating quantitative research methodology). In this appendix we provide additional details on the simulation including the rationale behind our choices and interpretation of our results. We also describe the factors that make the simulation challenging enough to be interesting and those that are still unrealistically simplistic.

We conducted a thousand simulations with each simulation taking approximately 40 minutes to run. Each of the 1000 simulations is linked to a unique seed and can be run independently using the code in our replication archive.

Simulating the Data We preprocessed the 3201 articles in the JSTOR data using the default settings of the `stm` package’s command `textProcessor`. We then limited to words which appear in at least 25 documents in order to shrink the vocabulary to a manageable size. We then construct a model where articles on quantitative methodology are treated 10% of the time and non-quantitative articles are treated 25% of the time. We then simulate an outcome using the actual hand-coded variable on quantitative methodology (X_i) as an unobserved confounder which along with the treatment generates the outcome using a linear model.

$$T_i \sim \text{Bernoulli}(\pi = .1X_i + .25(1 - X_i))$$

$$Y_i \sim \text{Normal}(\mu = .5X_i - .2T_i, \sigma^2 = .09).$$

Estimation We largely mirror the estimation choices in our application on this data, fitting a structural topic model with 15 topics using the treatment as a content covariate and matching topics in two bins (less than .1 topic proportion in the document or more than .1 topic proportion in the document) and eight automatically generated bins for the treatment projection.

Strengths and Weaknesses of the Simulation Strategy The structure of the simulation preserves both the real documents and hand-coding of a category from those documents which might plausibly represent a confounder. The simulation is hard (and thus meaningful) because the confounder is not available to our matching model and there is no guarantee that it is recoverable from the word-count data. We also induced a substantial amount of confounding and noise (as evidence by the strikingly poor performance of the unadjusted estimator). We have relatively few treated units (about 20% of the sample on average).

The simulation is also easy in some important ways that might give us pause in generalizing. The treatment and outcome model are quite simple. We use a binary unobserved confounder because it makes the simulation straightforward and easier to describe but one could imagine constructing a more complicated basis for unobserved confounding. The linear model for the outcome means that the treatment effect is constant. This helps remove any complications from the estimand changing as units are pruned, but also makes things substantially easier for models like TIRM which drop a large number of treated units.

For computational reasons, we compare performance of topic matching and matching on the projection using the fitted TIRM model. This effectively demonstrates what each of these components is contributing to our overall estimate but does not necessarily reflect what would happen if a topic model or balancing score were fit separately to the data. We conducted smaller tests (of 100 simulations) using separately estimated topic models and balancing scores alone and did not observe substantially different results from what we have reported here. Because we use real texts, we are also limited to only the 3201 articles in our database and thus were not able to study performance as sample size increases.

D The Gender Citation Gap: Balance Checking, Results, and Sensitivity Analysis

This section reports details about the balance checking and analysis for the application estimating whether there is a gender citation gap in the IR literature. For balance checking in this application we are able to include comparisons based on human-coding of the articles from the Teaching Research and International Policy (TRIP) article database. The collection of this data involved a large-scale human coding effort that will not be generally available in other applications, but provides us a chance to probe the performance of our method. We do however note that this human coding is not necessarily the gold-standard for this application because the coding system was not designed to facilitate estimation of our causal estimand. However if TIRM is performing well, then we believe it should improve balance on the human-coded variables measuring article content. When presenting comparisons to matches based on human coding we are matching on 35 variables which includes the methods variables (see Fig 6), issue areas (see Fig 6), paradigms (realist, liberal, constructivist, atheoretical, nonparadigmatic, marxist) and epistemic orientation variables (positivist orientation, material variables, ideational variables) from the TRIP dataset.

Balance Checking: Balancing Estimated Topics

We check whether TIRM balanced the estimated topics adequately in Figure 4. Our primary comparison is between TIRM, projection matching (propensity score only), and topic matching. We find that topic matching performs best at balancing the estimated topics, as we would expect. Matching only on the projection generally does not improve balance on the topics – in fact balance on many topics gets notably worse. TIRM performs almost as well as topic matching, despite also balancing on the propensity score, which we see as evidence of TIRM’s effectiveness in combination with the balance checks below. We also include a secondary comparison to the balance we obtain when we use matching on the human-coded variables. Matching on human-coded variables generally improves




Figure 4: Balancing Estimated Topics

balance on the topic model topics, but not in all cases.

Balance Checking: Kernel Similarity

We also compare the matching based on human coding to TIRM using a string kernel similarity metric. Figure 5 shows the similarity between matched documents in the corpus matched using TIRM and corpus matched exactly on human codes. Overall, TIRM performs as well to the human-coding matching in producing semantically similar documents when measured with string kernel similarity.

Balance Checking: Comparing TIRM and Human Coding

Figure 6 evaluates balance on the human coded categories. The rows along the y-axis correspond to non-mutually exclusive, human-coded categories from the article text: methodological categories on top and issue-area categories below. To the right of each category label, we plot a bar showing the imbalance of this category by gender of article author




Figure 5: String Kernel Similarity Comparison

in the raw data set. TIRM performs reasonably well at balancing the human-coded categories, especially on the variables that were initially most imbalanced in the full sample: qualitative methodology, formal methodology, quantitative methodology, and the issue area of international security, suggesting that TIRM comes closest to mimicking the human coding process. Topic matching also substantially reduces imbalance in many of the human-coded categories. In contrast, projection-only matching makes balance worse on several of the most imbalanced human-coded categories (formal methods and quantitative methods) and does not improve balance much on several others (qualitative methods and international security).

It's true that the other methods also occasionally make imbalance worse on some of the human-coded variables. TIRM and topic matching both increase imbalance in articles using descriptive methodology, for example. However, we are more concerned about correcting extreme imbalances present in the full data set because we *ex ante* believe they will produce the greatest bias in our estimates. On these, TIRM and topic matching outperform projection matching.




Figure 6: Automated Matching Comparison and Human Categories

Balance Checking: Comparing matched pairs

Treated Document	Matched Control Document
“Democratic Synergy and Victory in War, 1816-1992.” Ajin Choi. 2004. <i>International Studies Quarterly</i>. Abstract: “This study investigates the question of why democracies are more likely to win wars than non-democracies. I argue that due to the transparency of the politics, and the stability of their preferences, once determined, democracies are better able to cooperate with their partners in the conduct of wars, and thereby are more likely to win wars. In support of my argument, the main findings in this study show that, other things being equal, the larger the number of democratic partners a state has, the more likely it is to win; moreover, democratic states are more likely to have democratic partners during wars. These results are in contrast with those in current literature about the high likelihood of prevailing by democracies in wars, which emphasize, on the one hand, the superior capacity of democratic states to strengthen military capabilities and, on the other hand, to select wars in which they have a high chance of winning.”	“Third-Party Interventions and the Duration of Intrastate Conflicts.” Patrick M. Regan. 2002. <i>Journal of Conflict Resolution</i>. Abstract: “Recent research has begun to focus on the role of outside interventions in the duration of civil conflicts. Assuming that interventions are a form of conflict management, ex ante expectations would be that they would reduce a conflict’s expected duration. Hypotheses relating the type and timing of outside interventions to the duration of civil conflicts are tested. The data incorporate 150 conflicts during the period from 1945 to 1999, 101 of which had outside interventions. Using a hazard analysis, the results suggest that third-party interventions tend to extend expected durations rather than shorten them. The only aspect of the strategy for intervening that reduces the likelihood that a conflict will end in the next month is that it be biased in favor of either the opposition or the government. In effect, neutral interventions are less effective than biased ones.”
“The Present As Prologue: Europe and Theater Nuclear Modernization.” Catherine McArdle Kelleher. 1981. <i>International Security</i>. Introduction: “More than a year after formal Alliance decision, controversy still surrounds NATO’s plan for new long-range theater nuclear force (LRTNF) deployments. The controversy focuses both on the substance and the process involved. Proponents, in Washington as elsewhere, see the decision as an Alliance success, and the process as a model for future decision-making. The Alliance has now demonstrated it can meet new Soviet challenges; the exhaustive consultation procedures did lead to a genuinely informed NATO consensus despite the inherent political risks. Although not designed to match Soviet LRTNF capabilities, ground-launched cruise missiles (GLCMs) and Pershing IIs do provide a new element in the overall East-West military balance. And there has been due, measured common attention to the problems both of reinforcing American strategic linkage, and of pursuing opportunities for East-West limitations on LRTNF deployments. ”	“Counterforce: Illusion of a Panacea.” Henry A. Trofimenko. 1981. <i>International Security</i>. Introduction: “In recent years, American military strength has been moving in a vicious cycle. It has been unable to get out of the impasse created by Washington’s desire to outstrip the Soviet Union in strategic arms and by the practical impossibility of achieving this aim. The Soviet Union is fully resolved not to fall behind the United States, nor to permit such U.S. preponderance. Recent evidence of this strategic merry-go-round was provided by the Carter Administration’s Directive No. 59, which returned U.S. strategy to the concepts of counterforce nuclear targeting. This policy concludes the long campaign in the U.S. media and academic press that was intended to frighten the Americans with a purported Soviet counterforce and “war-winning” threat. ”
“Locating “Authority” in the Global Political Economy.” A. Claire Cutler. 1999. <i>International Studies Quarterly</i>. Abstract: “This article addresses the problematic nature of “authority” in the global political economy. Focusing on the rules governing international commercial relations, which today form part of the juridical conditions of global capitalism, the location and structure of political authority are argued to be historically specific. They have changed with the emergence of different historic blocs and as a result of consequent alterations in state-society relations. The article emphasizes the significance of private corporate power in the construction of the global political economy and hegemonic authority relations. However, the significance of private authority is obscure and little understood by students of international relations. This gives rise to analytical and normative grounds for adopting a historical materialist approach to the analysis of global authority that incorporates national, subnational, and transnational influences.”	“South Korean and Taiwanese Development and the New Institutional Economics.” David C. Kang. 1995. <i>International Organization</i>. Introduction: “The publication of books by both Alice Amsden and Robert Wade provide an opportune moment to reflect on the study of East Asian development. After an initial surge of interest beginning in the 1970s, the field has reached a plateau, and scholars recently have cast a wide net searching for ways to extend the field. In assessing the “state of the art” regarding economic development of the East Asian newly industrialized countries (NICs), this review will treat three themes. First, I will argue that the focus on states versus markets is becoming stale and much scholarly interest lies in the politics behind the economics. Second, I argue that political scientists have underexplored the historical origins of Korean and Taiwanese capitalism and that such attention promises to strengthen both theories and explanations of development. Third, I argue that the international system has been more important in promoting development in East Asia than accounts in the “first wave” have recognized.”

Table 4: Matched pairs of documents (female authorial teams on left).

Results

We estimate negative binomial regression models using similar specifications to those in Table 2 of Maliniak, Powers and Walter (2013). Our results are in Table 5. Column 1 of Table 5 shows the results of a specification with no conditioning variables other than the TIRM matching. Columns 2 and 3 add progressively more control variables, but only includes those that could be coded easily, without human reading. Column 4 adds the human-coded variables for research orientations, paradigms, and methodology. Column 4 corresponds closely to the “kitchen sink” model in Maliniak, Powers and Walter (2013), but omits a handful of variables that do not vary in our matched sample.

These results support the conclusion of Maliniak, Powers and Walter (2013) that there is a detectable gender citation gap in the IR literature. Maliniak, Powers and Walter (2013, 906) estimate the gap to be about 4.7 citations, while our equivalent model in column 4 estimates this gap to be 6.5 citations. Our estimate may not be directly comparable to theirs because we use coarsened exact matching which discards treated units and changes the quantity of interest from the SATT to FSATT (King, Lucas and Nielsen, 2017). Our result show that the gender citation gap persists in the subset of women’s and men’s articles that are textually similar according to TIRM. We cannot extrapolate our estimate to the broader population of IR articles without additional assumptions, and we don’t undertake that extrapolation exercise here. This means our findings are not definitive evidence of the magnitude of the gender citation gap in the overall population of IR articles. However, they do show that evidence of a gap persists even when we condition on the text of the articles in ways Maliniak, Powers and Walter (2013) could not.

Sensitivity Analysis

Regression and matching approaches on observational data rely on the assumption that all confounding is due to the observable factors included in the matching and regression procedures. Sensitivity analysis offers a way to test how robust the findings are to violations of this assumption. We use Rosenbaum’s sensitivity analysis framework based on randomization inference (Rosenbaum, 2002), implemented by Keele (2010). This pro-

Table 5: Maliniak, Powers, and Walter 2013 with Text Matching

	<i>Dependent variable:</i>			
	Citation Count (source: SSCI)			
	(1)	(2)	(3)	(4)
Female author(s)	-1.02*** (0.25)	-0.91*** (0.30)	-0.92*** (0.28)	-0.59*** (0.22)
Mixed gender authors		-0.49 (0.69)	0.36 (0.69)	0.46 (0.53)
Article age		0.08 (0.08)	0.05 (0.08)	0.07 (0.06)
Article age ²		-0.003 (0.002)	-0.003 (0.002)	-0.003** (0.002)
Tenured		1.28*** (0.27)	0.96*** (0.27)	0.33 (0.23)
Tenured female		-0.55 (0.50)	-0.29 (0.47)	0.34 (0.38)
Coauthored		-0.12 (0.37)	-0.84** (0.40)	-0.56* (0.31)
R1		0.38* (0.22)	0.45** (0.22)	0.09 (0.18)
AJPS			0.70 (0.81)	1.14* (0.65)
APSR			1.56** (0.68)	1.39*** (0.54)
IO			1.08*** (0.28)	0.40* (0.23)
IS			-0.48 (0.33)	-0.26 (0.28)
ISQ			-0.49 (0.40)	-0.11 (0.34)
JCR			0.62* (0.37)	0.76** (0.36)
Positivist				0.60* (0.31)
Materialist				0.86 (0.65)
Ideational				-0.22 (0.19)
Paradigm: Atheoretical				-0.69* (0.40)
Paradigm: Constructivist				1.53*** (0.56)
Paradigm: Liberal				-0.52** (0.21)
Paradigm: Realist				0.07 (0.41)
Method: Qualitative				1.27*** (0.37)
Method: Quantitative				0.82** (0.37)
Method: Formal theory				0.67* (0.36)
Method: Analytical				1.42*** (0.49)
Method: Description				0.82* (0.45)
Method: Counterfactual				1.47*** (0.57)
Constant	3.60*** (0.13)	2.79*** (0.72)	2.88*** (0.72)	0.37 (0.98)
Observations	181	181	181	181
Log Likelihood	-749.84	-738.62	-723.23	-665.45

Note:

Negative binomial generalized linear models. *p<0.1; **p<0.05; ***p<0.01

cedure compares the differences in outcomes between matched pairs in a data set as if treatment were randomly assigned and then calculates how large a confounder would be necessary to eliminate the observed difference. This is done by positing an odds ratio Γ that corresponds to the magnitude of the potential unobserved confounder. When $\Gamma = 1$, there is no confounding. As Γ increases from 1, the odds of units being treated increase. We do not know the true Γ , so we instead posit increasing values and see whether our results could be overturned by a relatively small unobserved confounder, or only by a very large one. Matched samples for which Γ is higher are considered less sensitive to potential confounding.

In the matched sample for the Maliniak data, we find that the result we report would be overturned with $\Gamma > 1.9$. This means that an unobserved confounder associated with female authorship by a factor of 1.9 would overturn our result. This is a middling value: it is possible to imagine an unobserved factor that might be this strongly correlated with the gender of authors and citation counts but the effect of this unobserved factor would have to be somewhat large. By comparison, the sensitivity analysis of the unmatched data indicates that the difference in citations between treated and control documents would be overturned with $\Gamma > 1.24$, indicating that the unmatched result is sensitive to very modest levels of unobserved confounding.

E Details of Chinese Social Media User Analysis

This section contains details of the analysis of Chinese social media users. To create our sample, we identify all users within the Weiboscope dataset (Fu, Chan and Chau, 2013) who were censored at least once in the first half of 2012. We subset the data to contain only these 4,685 users. We identify all posts that are censored for these users. Fu, Chan and Chau (2013) identify two types of error that could be censorship – one is posts that have a “permission denied” error after being removed and others that have a “Weibo does not exist” error after being removed. Fu, Chan and Chau (2013) determine from their own experiments that the “permission denied” error indicates censor removal all the time, but while the “Weibo does not exist” error is usually censorship, it could also be the user removing the post at their own discretion. To ensure that match posts indicate censorship, we only use “permission denied” posts as treated units and match to control posts that are neither “permission denied” nor “Weibo does not exist”.

After identifying posts that were censored with a “permission denied” message by these users, we subset to posts that are longer than 15 characters. We use 15 characters in order to ensure that there is enough information in the censored post to match the content of the post – many of the posts only include the text “reposting” which does not include enough information to ensure that matches are substantively similar. Because the post volume of the 4,685 users is sufficiently large, we restrict our pool of control posts to those that have a cosine similarity with the censored post of greater than 0.5 and were posted on the same day as the censored post. Thus, the potential control donation sample is all posts that are greater than 15 characters in length that have a cosine similarity to a censored post posted on the same day of greater than 0.5. This leaves us with 75,641 posts from 4,160 users across the last 6 months of 2012, with 21,503 censored posts and 54,138 uncensored posts from which to create matches.

We run a topic model on these 75,641 posts with 100 topics in order to ensure a close match and an indicator of censorship as the content covariate. We extract from this: 1) the estimated topic proportion for each post within our dataset and 2) the estimated projection for each post within our dataset.

To estimate the effect of censorship on future post rate and future censorship, we extract all posts for the users within the dataset for the four weeks before and the four weeks after censorship. For the four weeks before censorship, we calculate the number of posts each user wrote and the number of these that were censored or went missing. For the four weeks after censorship, we calculate the number of posts each user wrote and the number of these that were censored and went missing.

We then proceed with matching. Using coarsened exact matching, we match censored posts to uncensored posts written on the same day that have similar topics and projection scores. In addition, we make sure that matched users have a similar previous posting rate and a similar previous censorship rate by matching on these variables. We also ensure that users are not matched to themselves within a strata.

Balance in terms of topics is described in the main text of the paper. Here we show that string kernel similarity results. Matched posts were randomly sampled within each matched dataset and their string kernel similarity was calculated. Posts matched via TIRM had the highest level of similarity, followed by topic matching, then propensity score matching, and last the unmatched dataset (Figure 7).

Further, we conducted a qualitative comparison of matched posts to ensure that TIRM was retrieving posts that were qualitatively similar. We provide some examples of matched posts in Table 3.

After matching with TIRM, the matched data had 305 censored posts matched to 574 uncensored posts, with a total of 879 matched posts. There was no difference in the censorship rate before the match for matched users – both users who were censored in the matched post and not censored had a previous censorship rate of 0.003. Further, we also matched on the previous history of missing posts, and there was not difference in the history of “Weibo does not exist” messages between users who were censored in the matched post and not censored – both had a previous missingness rate of 0.22. Last, there was no difference between the number of posts treated and control users posted before the match – on average, treated users post 661 posts in the four weeks before the match, while untreated users post 649. Balance tests for these pre-censorship non-text covariates




Figure 7: Mean String Kernel Similarity for Matched Posts Randomly Sampled Within Each matched Dataset

Censored Post	Uncensored Post
There may be even bigger plans: When the chaos escalate to a certain degree, there will be military control, curfew, Internet cutoff, and then complete repression of counterrevolution. There are precedent examples.	The person on the right (refers to the previous comment) knew too much. Bang (Sound effect for gunshot)! You knew too much! Bang! There may be even bigger plans: When the chaos escalate to a certain degree, there will be military control, curfew, Internet cutoff, and then complete repression of counterrevolution. There are precedent examples.
#Weitianxia#Shifang netizen’s disclose: I saw police officers looking for restaurants to eat on the street and the intestine vermicelli restaurant owner immediately said they don’t sell it to the police. Then everyone on the street came out and yelled, which was very impressive. Now many stores have signs saying that police tactical units are not allowed to enter. Shifang people said: F*k you, you beat us up, bombed us, and still ask us to feed you, why don’t you eat sh*t?	Due to the lack of prior publicity procedures, some people are unfamiliar, uncomprehending and unsupportive of this program. To respond to the general public’s request, the municipal party committee and the government researched and decided to stop the project. Shifang will not ever develop the Molybdenum Copper project in the future.
[17-year-old young athlete fails 3 attempts to lift The media calls it a shame of Chinese female weightlifters] According to Sina: Chinese female weightlifters faced a shameful failure of its Olympic history last night! During the female 53kg weightlifting competition, joined as the black horse, Zhou Jun, a 17-year-old young athlete from Hubei, failed in all 3 of her attempts and ended with no result, which ends her Olympic journey. Many media reported this using “the most shameful failure of Chinese female Olympic weightlifters” as the title.	[17-year-old young athlete fails 3 attempts to lift The media calls it a shame of Chinese female weightlifters] According to Sina: Chinese female weightlifters faced a shameful failure of its Olympic history last night! During the female 53kg weightlifting competition, joined as the black horse, Zhou Jun, a 17-year-old young athlete from Hubei, failed in all 3 of her attempts and ended with no result, which ends her Olympic journey. Many media reported this using “the most shameful failure of Chinese female Olympic weightlifters” as the title. I personally think, it is not a shame of Zhou Jun, but a shame of Chinese media!

Table 6: Translations of example social media posts that were censored (left) with matched uncensored social media posts selected by TIRM (right).

are provided in Table 7.

Even though the matched users posted very similar posts, and had very similar previous censorship rates, we find that their experience with censorship diverges after the match. The “Permission denied” rate of the treated users is approximately twice as large

	Mean Treated	of	Mean of Con- trol	Difference	P-Value
Previous Censorship Proportion	0.003		0.003	<0.0001	0.223
Previous Not Exist Proportion	0.220		0.220	<0.0001	0.984
Previous Number of Posts	661.987		649.560	12.427	0.773

Table 7: Balance Tests (T-tests) For Pre-Treatment Covariates

as the censorship rate of the untreated users after the match. Further, the rate of “Weibo does not exist” messages also increases for treated users after the match – treated users have on average 25% of their posts missing in the four weeks after the match, in comparison to control users who only have 20% of their posts missing in the four weeks after the match. This suggests that either treated users are affected by their experience with censorship in a way that inspires them to post more sensitive material after the match or that they are put on a list after being censored that increases their likelihood of future censorship.

Users are more likely censored after experiencing censorship; however, we do not see them posting less after experiencing censorship. On average, treated users post 600 posts in the four weeks after censorship, while untreated users post 588, an insignificant difference. This indicates that despite experiencing more censorship, treated users are not deterred in posting online after being censored.

We use a sensitivity analysis (Rosenbaum, 2002) to estimate the magnitude of unobserved confounding that would overturn the findings. We estimate Γ , the factor by which a hypothetical unobserved confounder would have to be associated with treatment to erase the effect. For both outcomes with statistically detectable results, we find that an unobserved factor that increased the odds of treatment by roughly 1.6 would overturn the result ($\Gamma = 1.62$ and $\Gamma = 1.64$ respectively). It is possible to imagine such a factor, so our results are somewhat sensitive.

We conduct another sensitivity analysis to ensure that our results are not peculiar to the tuning parameters of the matching algorithm. As shown in Figure 8, we compare

the results across many matches of topical bin size – from 2 bins for each topic to 15 bins for each topic.¹⁶ The plots show a high level of consistency across bins – while the rate of censorship consistently increases for both “permission denied” and “Weibo does not exist” posts, the total number of posts written by censored and uncensored users is not different after the match.

¹⁶We allow CEM to automatically construct the bins.




Figure 8: Sensitivity Analysis for Censorship Results, Varying Topical Bin Size. Numbers next to each confidence interval indicate sample size. Top Panel: Effect of Censorship on “Permission Denied” rate. Middle Panel: Effect of Censorship on “Weibo Does Not Exist” rate. Bottom Panel: Effect of Censorship on Posting Rate.

References

- Athey, Susan and Guido Imbens. 2016. “Recursive partitioning for heterogeneous causal effects.” *Proceedings of the National Academy of Sciences* 113(27):7353–7360.
- Athey, Susan, Guido W Imbens and Stefan Wager. 2016. “Approximate residual balancing: De-biased inference of average treatment effects in high dimensions.” *arXiv preprint arXiv:1604.07125* .
- Belloni, Alexandre, Victor Chernozhukov and Christian Hansen. 2014. “Inference on treatment effects after selection among high-dimensional controls.” *The Review of Economic Studies* 81(2):608–650.
- Bloniarz, Adam, Hanzhong Liu, Cun-Hui Zhang, Jasjeet Sekhon and Bin Yu. 2015. “Lasso adjustments of treatment effect estimates in randomized experiments.” *arXiv preprint arXiv:1507.03652* .
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen and Whitney Newey. 2017. “Double/Debiased/Neyman Machine Learning of Treatment Effects.” *arXiv preprint arXiv:1701.08687* .
- Egami, Naoki, Christian J Fong, Justin Grimmer, Margaret E Roberts and Brandon M Stewart. 2017. “How to Make Causal Inferences Using Texts.”
URL: <https://scholar.princeton.edu/sites/default/files/bstewart/files/ais.pdf>
- Eisenstein, Jacob, Amr Ahmed and Eric P. Xing. 2011. Sparse Additive Generative Models of Text. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*. ICML’11 USA: Omnipress pp. 1041–1048.
URL: <http://dl.acm.org/citation.cfm?id=3104482.3104613>
- Fu, King-wa, Chung-hong Chan and Michael Chau. 2013. “Assessing censorship on microblogs in China: Discriminatory keyword analysis and the real-name registration policy.” *IEEE Internet Computing* 17(3):42–50.
- Gretton, Arthur, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf and Alexander Smola. 2012. “A kernel two-sample test.” *Journal of Machine Learning Research* 13(Mar):723–773.
- Gretton, Arthur, Karsten M Borgwardt, Malte Rasch, Bernhard Schölkopf and Alex J Smola. 2007. A kernel method for the two-sample-problem. In *Advances in neural information processing systems*. pp. 513–520.
- Hansen, Ben B. 2008. “The prognostic analogue of the propensity score.” *Biometrika* 95(2):481–488.
- Hartford, Jason, Greg Lewis, Kevin Leyton-Brown and Matt Taddy. 2016. “Counterfactual Prediction with Deep Instrumental Variables Networks.” *arXiv preprint arXiv:1612.09596* .

- Hartman, Erin and F Daniel Hidalgo. 2018. “An Equivalence Approach to Balance and Placebo Tests.” *American Journal of Political Science* 62(4):1000–1013.
- Hazlett, Chad. Forthcoming. “Kernel Balancing: A flexible non-parametric weighting procedure for estimating causal effects.” *Statistica Sinica* .
- Hill, Jennifer, Christopher Weiss and Fuhua Zhai. 2011. “Challenges with propensity score strategies in a high-dimensional setting and a potential alternative.” *Multivariate Behavioral Research* 46(3):477–513.
- Keele, Luke. 2010. “An overview of rbounds: An R package for Rosenbaum bounds sensitivity analysis with matched data.” *White Paper. Columbus, OH* pp. 1–15.
URL: <https://pdfs.semanticscholar.org/c0b1/823186cf5e0869ee35657d6809f803c8e35c.pdf>
- King, Gary, Christopher Lucas and Richard A Nielsen. 2017. “The Balance-Sample Size Frontier in Matching Methods for Causal Inference.” *American Journal of Political Science* 61(2):473–489.
- Maliniak, Daniel, Ryan Powers and Barbara F Walter. 2013. “The gender citation gap in international relations.” *International Organization* 67(04):889–922.
- Rabinovich, Maxim and David Blei. 2014. The inverse regression topic model. In *Proceedings of The 31st International Conference on Machine Learning*. pp. 199–207.
- Ratkovic, Marc and Dustin Tingley. 2017. “Causal Inference through the Method of Direct Estimation.” *arXiv preprint arXiv:1703.05849* .
- Rosenbaum, Paul R. 2002. *Observational Studies*. Springer.
- Rubin, Donald B. and Neal Thomas. 2000. “Combining propensity score matching with additional adjustments for prognostic covariates.” *Journal of the American Statistical Association* 95(450):573–585.
- Sales, Adam C, Ben B Hansen and Brian Rowan. 2015. “Rebar: Reinforcing a Matching Estimator with Predictions from High-Dimensional Covariates.” *arXiv preprint arXiv:1505.04697* .
- Schneeweiss, Sebastian, Jeremy A Rassen, Robert J Glynn, Jerry Avorn, Helen Mogun and M Alan Brookhart. 2009. “High-dimensional propensity score adjustment in studies of treatment effects using health care claims data.” *Epidemiology (Cambridge, Mass.)* 20(4):512.
- Sejdinovic, Dino, Bharath Sriperumbudur, Arthur Gretton and Kenji Fukumizu. 2013. “Equivalence of distance-based and RKHS-based statistics in hypothesis testing.” *The Annals of Statistics* pp. 2263–2291.
- Sridhar, Dhanya and Lise Getoor. 2019. “Estimating Causal Effects of Tone in Online Debates.” *arXiv preprint arXiv:1906.04177* .

- Szabó, Zoltán, Arthur Gretton, Barnabás Póczos and Bharath Sriperumbudur. 2015. Two-stage sampled learning theory on distributions. In *Artificial Intelligence and Statistics*. pp. 948–957.
- Taddy, Matt. 2013a. “Measuring Political Sentiment on Twitter: Factor Optimal Design for Multinomial Inverse Regression.” *Technometrics* 55(4):415–425.
- Taddy, Matt. 2013b. “Multinomial inverse regression for text analysis.” *Journal of the American Statistical Association* 108(503):755–770.
- Taddy, Matt. 2013c. “Rejoinder: Efficiency and Structure in MNIR.” *Journal of the American Statistical Association* 108(503):772–774.
- Tran, Dustin and David M Blei. 2017. “Implicit causal models for genome-wide association studies.” *arXiv preprint arXiv:1710.10742* .
- Van der Laan, Mark J and Sherri Rose. 2011. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media.
- Veitch, Victor, Yixin Wang and David M Blei. 2019. “Using embeddings to correct for unobserved confounding.” *arXiv preprint arXiv:1902.04114* .
- Wang, Yixin and David M Blei. 2019. “The blessings of multiple causes.” *Journal of the American Statistical Association* (just-accepted):1–71.
- Westreich, Daniel, Justin Lessler and Michele Jonsson Funk. 2010. “Propensity score estimation: neural networks, support vector machines, decision trees (CART), and meta-classifiers as alternatives to logistic regression.” *Journal of clinical epidemiology* 63(8):826–833.
- Wood-Doughty, Zach, Ilya Shpitser and Mark Dredze. 2018. “Challenges of Using Text Classifiers for Causal Inference.” *arXiv preprint arXiv:1810.00956* .
- Yao, Liuyi, Sheng Li, Yaliang Li, Hongfei Xue, Jing Gao and Aidong Zhang. 2019. On the estimation of treatment effect with text covariates. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. AAAI Press pp. 4106–4113.
- Zhang, Qinyi, Sarah Filippi, Arthur Gretton and Dino Sejdinovic. N.d. “Large-scale kernel methods for independence testing.” *Statistics and Computing*. Forthcoming.