

On the Testability of the Anchor-words Assumption in Topic Models

Simon Freyaldenhoven

Federal Reserve Bank of Philadelphia

Shikun Ke

Yale School of Management

Dingyi Li

Cornell University

José Luis Montiel Olea

*Cornell University**

November 2, 2023

Abstract

Topic models are a simple and popular tool for the statistical analysis of textual data. Their identification and estimation is typically enabled by assuming the existence of *anchor words*; that is, words that are exclusive to specific topics. In this paper we show that the existence of anchor words is statistically testable: There exists a hypothesis test with correct size that has nontrivial power. This means that the anchor-words assumption cannot be viewed simply as a convenient normalization. Central to our results is a simple characterization of when a column-stochastic matrix with known nonnegative rank admits a *separable* factorization. We test for the existence of anchor words in two different data sets derived from the transcripts of the meetings of the Federal Open Market Committee (FOMC)—the body of the Federal Reserve System that sets monetary policy in the United States—and reject the null hypothesis that anchor words exist in one of them.

JEL codes: C39, C55

KEYWORDS: Anchor Words, Topic Models, Nonnegative Matrix Factorization, Hypothesis Testing.

*We thank Roc Armenter, Stephane Bonhomme, Florentina Bunea, Michael Dotsey, Stephen Hansen, Tracy Ke, Aaron Schein, Marten Wegkamp, Yun Yang, and participants at numerous seminars and conferences for their comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Emails: simon.freyaldenhoven@phil.frb.org, barry.ke@yale.edu, dl922@cornell.edu, montiel.olea@gmail.com.

1 Introduction

Topic models—statistical models that aim to help uncovering the thematic structure in a collection of documents—are a simple and popular tool for the analysis of textual data; see Blei & Lafferty (2009), Blei (2012) for excellent reviews, and Boyd-Graber, Hu, Mimno et al. (2017) for a list of applications. The model assumes the existence of K latent *topics*, which are defined as probability distributions over V terms in a given vocabulary. The model also assumes that each of the D documents is characterized by a topic distribution; i.e., the share it assigns to each of the K latent topics.

An assumption that has become ubiquitous in this literature is the existence of *anchor words* (Arora, Ge & Moitra 2012), which is inspired by the notion of *separability* used in nonnegative matrix factorization problems; see Donoho & Stodden (2003) and Arora, Ge, Kannan & Moitra (2012). Broadly speaking, anchor words are defined as special terms in the vocabulary that are exclusive to each specific topic. It is well known that the existence of at least one anchor word per topic enables the identification of the parameters of the topic model. The existence of anchor words also allows the construction of estimators for the topic distributions with provable optimal statistical performance guarantees, see the recent work of Bing, Bunea & Wegkamp (2020a,b) and Ke & Wang (2022).

This paper investigates the extent to which the existence of anchor words in topic models is statistically testable. The motivation is that if the existence of anchor words is in conflict with the observed distribution of the data, then such an assumption ought to be dropped or at least relaxed.¹ The null hypothesis of interest in this paper is that the observed text data was generated by a topic model that satisfies the *anchor-words assumption*; which means that the topic distributions exhibit *at least* one anchor word per topic. The alternative hypothesis is that the anchor-words assumption does not hold. We say that the null hypothesis is testable at significance level α if there exists a test of size at most α and, in addition, the test has nontrivial power (that is, power larger than the desired significance level, for at least one parameter value in the alternative hypothesis).

Our first result (Proposition 1) identifies a necessary condition for the statistical testability of the anchor-words assumption. As is common in the literature we define the *population term-document* frequency matrix, P , as the $V \times D$ column-stochastic matrix whose (v, d) -th entry contains the probability of randomly drawing term v in document d . Our proposition shows that in order for a statistical test to have nontrivial power there must exist population term-document frequency matrices—among all of those that can be generated by a topic model with K topics—that do not admit a *separable* nonnegative matrix factorization. Our proposition simply formalizes an obvious observation: we cannot hope to test for the existence of anchor words if every population term-document frequency matrix admits a factorization for which its corresponding topic distributions have at least one anchor word per topic.

Our second result (Theorem 1) provides a characterization of when a column-stochastic matrix with known nonnegative rank admits a separable factorization. Our theorem—which builds on the seminal work of Recht, Re, Tropp & Bittorf (2012)—suggests a simple computational procedure to decide whether a separable nonnegative factorization exists for a given P . This allows us to assess, for example, how likely it is that a randomly

¹It is known that the existence of anchor words is sufficient for identification, but not necessary (Laurberg, Christensen, Plumbley, Hansen, Jensen et al. (2008), Fu, Huang, Sidiropoulos & Ma (2019)). This means that point identification of topic models can still be achieved even when this assumption is relaxed; see the recent work of Chen, He, Yang & Liang (2022) that uses the *sufficiently-scattered* condition in Huang, Sidiropoulos & Swami (2013) and Huang, Fu & Sidiropoulos (2016). Moreover, even without point identification it is still possible to use the distribution of the data to partially identify the parameters of the topic model; for example, see Ke, Montiel Olea & Nesbit (2022).

generated population term-document frequency matrix admits a separable factorization (see, for example, Figure 3a and its description). Using our theorem, we find that for $2 < K < \min\{V, D\}$ the likelihood of such an event is low.²

It is worthwhile to give a brief overview of the characterization result in Theorem 1 and explain its relation to the literature. Note that for any arbitrary matrix $P \in \mathbb{R}^{V \times D}$ that can be factorized as the product of two matrices (A, W) —with a factor $A \in \mathbb{R}^{V \times K}$ of rank K —there always exists a matrix $C \in \mathbb{R}^{V \times V}$ of rank K such that $CP = P$. Broadly speaking, the previous equation states that there are K rows of P that can be used to (linearly) generate any of its other rows. When P is a column-stochastic matrix that admits a separable factorization, it is possible to give more details on the types of linear combinations, C , that can be used to generate the rows of P . To the best of our knowledge, this observation was first made by Recht et al. (2012) and Gillis (2013). Our Theorem 1 builds on their results and shows that P has a separable nonnegative matrix factorization *if and only if* the linear program suggested by Recht et al. (2012) to find a nonnegative matrix factorization of separable matrices has a nonempty choice set. More precisely, Theorem 1 formally shows that P has a separable nonnegative matrix factorization if and only if there exists a matrix C in the set

$$\begin{aligned} \mathcal{C}_K \equiv \{ C \in \mathbb{R}^{V \times V} \mid & C \geq 0, \\ & \text{tr}(C) = K, \\ & c_{jj} \leq 1, \text{ for all } j = 1, \dots, V, \\ & c_{ij} \leq c_{jj}, \text{ for all } i, j = 1, \dots, V\}, \end{aligned} \quad (1)$$

that satisfies the equation

$$CP^{\text{row}} = P^{\text{row}}, \quad (2)$$

where P^{row} is the row-normalized version of P .

The set \mathcal{C}_K is the set of all nonnegative matrices of dimension $V \times V$ that have elements in $[0, 1]$, have trace equal to K , and have the property that the “sup-norm” of every column j is bounded by its j -th diagonal value. The set of all matrices C that satisfy (1) and (2) can be thought of as all rank- K convex combinations of the rows of P^{row} . Theorem 1 thus suggests that a reasonable test statistic for testing the anchor-words assumption given a text corpus Y is

$$T(Y) \equiv \inf_{C \in \mathcal{C}_K} \|C\hat{P}^{\text{row}} - \hat{P}^{\text{row}}\|, \quad (3)$$

where \hat{P}^{row} is a suitable estimator of the matrix P^{row} , and $\|\cdot\|$ is some matrix norm (which we will take, throughout the paper, to be the Frobenius norm).

Our third result (Theorem 2) shows that, under some high-level conditions, there exists a test of significance level α based on the test statistic (3) which has nontrivial power. Our proof is constructive, and the test we suggest rejects the anchor-words assumption whenever $T(Y)$ is large. To guarantee that the test has size at most α , we rely on a critical value that is chosen to be equal to the “worst-case” $(1 - \alpha)$ -quantile of $T(Y)$, which we

²The fact that not all nonnegative matrices of nonnegative rank K have a separable factorization with K topics should not be surprising, given well-known results in the computer science literature about the complexity of nonnegative matrix factorization. For instance, Vavasis (2010) has shown that the *exact* nonnegative matrix factorization problem is NP-hard. It is also known that finding a separable factorization (when such a factorization exists) can be done in polynomial time in (V, D, K) ; see Arora, Ge, Kannan & Moitra (2012). If every nonnegative matrix with nonnegative rank of K admitted a separable factorization, then the two previous results together would imply that the exact nonnegative matrix factorization problem is both P and NP-hard. Under the $P \neq NP$ hypothesis, an NP-hard problem cannot be in P.

denote as $q_{1-\alpha}^*$. By “worst-case” we mean the largest quantile among all those that could be obtained using a distribution for word counts generated by a model that satisfies the anchor-words assumption.

While the validity of the suggested test holds by construction, the analysis of the test’s power is more delicate. For intuition, first note that by the reverse triangle inequality,³

$$T(Y) \geq \inf_{C \in \mathcal{C}_k} \|(C - \mathbb{I}_V)(AW)^{\text{row}}\| - \sup_{C \in \mathcal{C}_k} \|(C - \mathbb{I}_V)(\hat{P}^{\text{row}} - P^{\text{row}})\|. \quad (4)$$

This means that the power of the test is lower-bounded by the probability of the event

$$\inf_{C \in \mathcal{C}_k} \|(C - \mathbb{I}_V)(AW)^{\text{row}}\| \geq \sup_{C \in \mathcal{C}_k} \|(C - \mathbb{I}_V)(\hat{P}^{\text{row}} - P^{\text{row}})\| + q_{1-\alpha}^*. \quad (5)$$

If P does not admit a separable factorization, the left-hand side of (5) is strictly positive by Theorem 1. Further, if the estimator \hat{P}^{row} is close enough to P^{row} with high probability—regardless of whether the anchor-words assumption holds—then both terms on the right-hand side of (5) will be small. Thus, one would expect (5) to hold with high probability at any point (A, W) for which the matrix $P = AW$ does not have an anchor-word factorization.

Although Theorem 2 shows that the test that rejects the null whenever “ $T(Y) > q_{1-\alpha}^*$ ” has correct size and nontrivial power, obtaining $q_{1-\alpha}^*$ is computationally infeasible. To address this issue, in Section 4.2.2 we derive a computationally tractable “bootstrap” upper bound for the critical value that allows to test for the existence of anchor words in realistic applications.

Finally, in order to illustrate the applicability of our theoretical results, we analyze the *transcripts* of the meetings of the Federal Open Market Committee (FOMC), the main body within the Federal Reserve System in charge of setting monetary policy in the United States. We focus on the FOMC transcripts during the “Greenspan period,” the meetings in which Alan Greenspan was chairman. We separate each transcript into two parts: the discussion of domestic and international economic conditions (FOMC1) and the discussion of the monetary policy strategy (FOMC2). This gives us two different corpora to analyze.⁴

The first corpus (FOMC1) allows us to illustrate the potential benefits of assuming the existence of anchor words in a concrete empirical application. Aside from the computational tractability and the theoretical identification results that become available under the anchor-words assumption, the estimated anchor words can potentially provide natural and objective labels for the estimated topics. We think this is an important point, as it has recently been argued that an inherent challenge of topic models in empirical applications is that they “*do not generate objective topic labels*” and that “*A given topic consists of many words, and words are scattered across many topics, so the outputs are often difficult to interpret.*”; see the discussion in Section 3.2.2.1 of Ash & Hansen (2023). In contrast, as we explain in detail in Section 5, the anchor words for FOMC1 are all readily interpretable (see Figure 8 and the discussion in Section 5.2.2). Moreover, the estimated topic proportions for the FOMC1 corpus seem to be consistent with historical events that shaped monetary policy decisions during the Greenspan period. In line with these results, when we apply our suggested testing procedure to this corpus, we indeed find that a nominal 5%-level test fails to reject the null hypothesis of anchor words for the FOMC1 corpus.

The results for the FOMC2 corpus are different. As we explain in Sections 5.2.2 and 5.2.3, the anchor

³Here and throughout, \mathbb{I}_H denotes the identity matrix of size H .

⁴See Chappell Jr, McGregor & Vermilyea (2004), Meade & Stasavage (2008), Meade & Thornton (2012), Hansen, McMahon & Prat (2018) for other studies using the FOMC transcript data.

words and the estimated topics for FOMC2 are difficult to interpret. Also, with the exception of two topics, it is difficult to provide a rationale for the historical evolution of the topic shares. Even without a formal statistical test, this suggests that the distribution of the data might not be compatible with the existence of anchor words, even if the topic model is assumed to be correctly specified. We then apply our suggested testing procedure to this corpus and indeed find that a nominal 5%-level rejects the anchor-words assumption for the FOMC2 corpus.

The rest of this paper is organized as follows. Section 2 presents the model. Section 3 presents the main theoretical results. This section also shows that when $K = 2$ the anchor-words assumption is not statistically testable, but gives concrete examples of statistical testability when $K = 3$. Section 4 presents numerical results. Section 5 presents the empirical application. Section 6 concludes.

2 Model

2.1 Notation

We observe documents $d = 1, \dots, D$, based on a dictionary of $v = 1, \dots, V$ terms. There is a $V \times K$ column-stochastic matrix, A , whose columns represent a probability distribution over the V terms that constitute the dictionary.⁵ We refer to each of the columns of A as a *topic*, and to A as the *term-topic* matrix. There is also a $K \times D$ column-stochastic matrix, W , collecting the probabilities that a document covers a particular topic $k = 1, \dots, K$. We refer to W as *topic-document* matrix. We assume that $K \leq \min\{V, D\}$.

It will be convenient to have specific notation to denote the v -th row, the k -th column, and the (v, k) -th entry of A . We will use $A_{v\bullet}$, $A_{\bullet k}$ and a_{vk} respectively. We use analogous notation for W and any other matrix. Further, for an arbitrary matrix B , we use \mathcal{R}_B to denote the diagonal matrix that contains the row sums of B , and use B^{row} to denote the ‘‘row-normalized’’ version of a matrix B . That is, $B^{\text{row}} = \mathcal{R}_B^{-1}B$.

We assume that the probability of a term v appearing in a given entry of document d , p_{vd} , is given by

$$p_{vd} = \sum_{k=1}^K \mathbb{P}(\text{Term } v | \text{Topic } k) \mathbb{P}_d(\text{Topic } k) = \sum_{k=1}^K a_{vk} w_{kd}. \quad (6)$$

Thus, the $V \times D$ matrix P defined by

$$P_{(V \times D)} = A_{(V \times K)} W_{(K \times D)}, \quad (7)$$

collects the terms p_{vd} . We will refer to P as the *population term-document frequency* matrix. Throughout, we maintain the assumption that both A and W are full rank and that the rows of A and P are all different from zero.⁶ We further assume that the number of topics K is known and fixed.

⁵A matrix $A \in \mathbb{R}^{V \times K}$ is column stochastic if its columns are probability distributions over \mathbb{R}^V . See p. 253 of Doeblin & Cohn (1993) for a definition.

⁶Note that, if there exists a term v with $\|A_{v\bullet}\|_0 = 0$, this term is not used in any document. Removing any unused terms from the dictionary and rewriting (7) using the smaller vocabulary V' immediately implies that $\|A_{v\bullet}\|_0 \neq 0 \forall v \in V'$.

2.2 Statistical model

The observed data consist of the number of times each term v appears in a specific document d . Denote these counts by the $V \times D$ matrix Y . Let N_d be the total number of words in document d , and $N_{\min} \equiv \min\{N_1, \dots, N_D\}$. Following the literature (e.g. Hofmann (1999)), we assume that for each document d

$$Y_{\bullet d} | (A, W) \sim \text{Multinomial}(N_d, AW_{\bullet d}). \quad (8)$$

We maintain throughout that the vectors of counts $Y_{\bullet d}$ are independent across documents, conditional on (A, W) .

It is well known that the parameters (A, W) in the statistical model (8) are not identified. This follows from the fact that any pair of parameters $(A, W) \neq (\tilde{A}, \tilde{W})$ such that $AW = \tilde{A}\tilde{W}$ will induce the same probability distribution over the data. In general, the culprit for the lack of identification is the multiplicity of solutions for the nonnegative matrix factorization problem defined by Equation (7); see Donoho & Stodden (2003), Fu et al. (2019).

The lack of identification poses statistical and computational challenges to the estimation of the parameters of the multinomial model in Equation (8). A common approach in the literature to circumvent these issues is to posit the existence of *anchor words* (Arora, Ge & Moitra (2012), Ke & Wang (2022), Bing et al. (2020a)). A term $v(k)$ in the vocabulary is an anchor word for topic k if such a term only has positive probability under topic k ; that is $A_{v(k)k} > 0$ and $A_{v(k)\tilde{k}} = 0$ for $\tilde{k} \neq k$. More formally:

Definition 1. A column stochastic, rank K matrix $A \in \mathbb{R}^{V \times K}$ is said to have anchor words if there exists a row permutation matrix Π such that

$$\Pi A = \begin{bmatrix} D \\ M \end{bmatrix}, \quad (9)$$

where $D \in \mathbb{R}^{K \times K}$ is a diagonal nonnegative matrix.

Since only the parameter $P = AW$ is identified in the multinomial model (8), it will be convenient to have an explicit definition of what it means to say that P admits a nonnegative matrix factorization with anchor words:

Definition 2. A column stochastic matrix $P \in \mathbb{R}^{V \times D}$ with nonnegative rank K is said to have a rank K , anchor-word (or separable) factorization if P can be written as

$$P = AW,$$

where $A \in \mathbb{R}^{V \times K}$ is some matrix that satisfies Definition 1, and W is a $K \times D$ column stochastic matrix.

2.3 The existence of anchor words as a statistical hypothesis

The goal of this paper is to analyze the extent to which the existence of anchor words is statistically testable. As we mentioned in the introduction, testing the conditions that enable the identification of statistical models has a long history in econometrics. Below we give a formal statement of our goal.

Let Θ denote the parameter space of the multinomial model in Equation (8). The parameter space refers to the collection of matrices (A, W) defined in (6)-(7) that could have generated the data. Define the ‘‘null set’’

Θ_0 as:

$$\Theta_0 \equiv \{(A, W) \in \Theta \mid A \text{ has anchor words as defined by Definition 1}\}. \quad (10)$$

The statistical hypothesis testing problem of interest is

$$\mathbf{H}_0 : (A, W) \in \Theta_0 \quad \text{vs.} \quad \mathbf{H}_1 : (A, W) \in \Theta_1 \equiv \Theta \setminus \Theta_0. \quad (11)$$

Let \mathcal{Y} denote the space of all possible data realizations according to the model in Equation (8). As usual, define a *statistical test* for the hypothesis testing problem in (11) as a function $\phi : \mathcal{Y} \rightarrow [0, 1]$, where $\phi(Y)$ is interpreted as the probability of rejecting the null hypothesis when the observed data is the count matrix Y .

Definition 3. *The statistical hypothesis \mathbf{H}_0 is testable at significance level α if there exists a test ϕ such that*

$$\sup_{(A, W) \in \Theta_0} \mathbb{E}_{(A, W)} [\phi(Y)] \leq \alpha, \quad (12)$$

and if there exists a parameter $(A, W) \in \Theta_1 \equiv \Theta \setminus \Theta_0$ such that

$$\mathbb{E}_{(A, W)} [\phi(Y)] > \alpha. \quad (13)$$

As usual, we refer to any test satisfying (12) as a *valid* test of significance level α . Also, for any $(A, W) \in \Theta_1$ we refer to $\mathbb{E}_{(A, W)} [\phi(Y)]$ as the *power* of the test ϕ at the parameter value (A, W) . Thus, Definition 3 says that the statistical hypothesis \mathbf{H}_0 is testable if there exists a statistical test with correct size and with nontrivial power; that is, power larger than the desired significance level at least at one parameter value in the alternative hypothesis Θ_1 .

The following simple proposition connects the statistical testability of \mathbf{H}_0 to the existence of anchor-word factorizations of the population term-document frequency matrix, P .

Proposition 1. *Let (A, W) be a parameter vector such that A does not have anchor words according to Definition 1; i.e., $(A, W) \in \Theta_1$. If the matrix $P \equiv AW$ has an anchor-word factorization—in the sense of Definition 2—then any valid test of significance level α for the hypothesis \mathbf{H}_0 has power of at most α at (A, W) .*

Proof. According to the statistical model in (8), the distribution of Y depends on the parameter (A, W) only through $P \equiv AW$. If P has an anchor-word factorization, then—by Definition 2—there exists $(\tilde{A}, \tilde{W}) \in \Theta_0$ for which $AW = P = \tilde{A}\tilde{W}$. Therefore, the power of any valid test ϕ of significance level α at (A, W) satisfies:

$$\mathbb{E}_{(A, W)} [\phi(Y)] = \mathbb{E}_{AW} [\phi(Y)] = \mathbb{E}_{\tilde{A}\tilde{W}} [\phi(Y)] \leq \alpha, \quad (14)$$

where the last inequality follows because $(\tilde{A}, \tilde{W}) \in \Theta_0$. □

The elementary result stated in Proposition 1 formalizes the observation that if any given matrix P with nonnegative rank K were to admit an anchor-word factorization, then any statistical test ϕ of significance level α for the hypothesis \mathbf{H}_0 would be trivial, in the sense that its power against any alternative $(A, W) \in \Theta_1$ is at most α . According to Definition 3 above, this makes the hypothesis \mathbf{H}_0 untestable. Consequently, Proposition 1 implies that a necessary condition for the testability of the anchor-words assumption is that not all matrices P with nonnegative rank K admit an anchor-word factorization.

A more abstract way to think about Proposition 1 is by imagining the topological structure of the null hypothesis relative to whole parameter space. For instance, it is known that if a matrix $P = AW$ for $(A, W) \in \Theta_1$ can be approximated arbitrarily well (in *total variation distance*) by elements in the set of distributions satisfying the null hypothesis (i.e., P is on the “topological boundary” of the null set), then, by continuity, the rejection probability of the test at such P must be no larger than the size of the test; see Lemma 2.1 in Canay, Santos & Shaikh (2013). When the matrix $P = AW$ for $(A, W) \in \Theta_1$ has an anchor-word factorization, then that means there is a $(A_0, W_0) \in \Theta_0$ for which $P = A_0W_0$. This means that the total variation distance between the induced data distributions for parameters (A, W) and (A_0, W_0) has to be zero. We return to this topological interpretation in the next section to argue that there are matrices P that do not admit an anchor-word factorization, and that those matrices are not on the boundary of the null set (see Remark 5, after Theorem 1).

3 Main Theoretical Results

3.1 When does P admit an anchor-word factorization?

According to Proposition 1, a necessary step to assess the testability of the anchor-words assumption is to understand whether all column-stochastic matrices P with nonnegative rank K admit an anchor-word factorization. Theorem 1 below sheds light on this issue.

Before presenting our result, we provide a brief algebraic illustration of the thought process that led to it. Note first that for any arbitrary matrix $P \in \mathbb{R}^{V \times D}$ that can be factorized as the product of two matrices (A, W) —with a factor $A \in \mathbb{R}^{V \times K}$ of rank K —there exists a matrix $C \in \mathbb{R}^{V \times V}$ such that

$$CP = P, \tag{15}$$

where C is also of rank K . Broadly speaking, the equation above says that there are K rows of P that can be used to generate any of its other rows by means of linear combinations. For example, assume w.l.o.g. that the first K rows of A , denoted A_0 , are full rank. Then, we may write

$$P = \begin{bmatrix} A_0W \\ A_1W \end{bmatrix}, \text{ and thus } C = \begin{bmatrix} \mathbb{I}_K & \mathbf{0}_{K \times (V-K)} \\ A_1A_0^{-1} & \mathbf{0}_{(V-K) \times K} \end{bmatrix} \text{ satisfies Equation (15).}$$

When P is a column-stochastic matrix that admits an anchor-word factorization, it is possible to give more details on the types of linear combinations, C , that can be used to generate the rows of P . To the best of our knowledge, this interesting observation was first made by Recht et al. (2012) and Gillis (2013).

To illustrate this point, suppose that A_0 is not just full rank, but diagonal (such that A has anchor words by Definition 1):

$$P = A^*W^* = \begin{bmatrix} A_0W^* \\ A_1W^* \end{bmatrix} = \begin{bmatrix} DW^* \\ MW^* \end{bmatrix},$$

where D is diagonal. With D diagonal, we can rewrite $A_1A_0^{-1} = MD^{-1} = \mathcal{R}_{MW^*}(\mathcal{R}_{MW^*})^{-1}M\mathcal{R}_{W^*}(\mathcal{R}_{DW^*})^{-1}$,

and all entries in $A_1 A_0^{-1}$ are nonnegative. Thus, the matrix \tilde{C} defined as

$$\tilde{C} \equiv \mathcal{R}_P C \mathcal{R}_P^{-1}, \quad \text{where} \quad C \equiv \begin{bmatrix} \mathbb{I}_K & \mathbf{0}_{K \times (V-K)} \\ (\mathcal{R}_{M W^*})^{-1} M \mathcal{R}_{W^*} & \mathbf{0}_{(V-K) \times K} \end{bmatrix}, \quad (16)$$

satisfies Equation (15). In particular, algebra shows that the matrix C in Equation (16) belongs to the set

$$\begin{aligned} \mathcal{C}_K \equiv \{ C \in \mathbb{R}^{V \times V} \mid & C \geq 0, \\ & \text{tr}(C) = K, \\ & c_{jj} \leq 1, \text{ for all } j = 1, \dots, V, \\ & c_{ij} \leq c_{jj}, \text{ for all } i, j = 1, \dots, V \}. \end{aligned} \quad (17)$$

The set \mathcal{C}_K is the set of all nonnegative matrices of dimension $V \times V$ that have diagonal elements in $[0, 1]$, have trace equal to K , and have the property that the “sup-norm” of every column j is bounded by its j -th diagonal value (which is reminiscent, but weaker, than the presence of a dominant diagonal).

Since $\tilde{C}P = P$, it follows that the matrix C in Equation (16) satisfies

$$C P^{\text{row}} = P^{\text{row}}. \quad (18)$$

The following theorem shows that the existence of an anchor-word factorization is characterized by the existence of a matrix $C \in \mathcal{C}_K$ that satisfies Equation (18).

Theorem 1. *A column-stochastic matrix $P \in \mathbb{R}^{V \times D}$ with nonnegative rank $K \leq \min\{V, D\}$ admits a rank K anchor-word factorization—in the sense of Definition 2—if and only if*

$$\mathcal{C}_K(P) \equiv \mathcal{C}_K \cap \left\{ C \in \mathbb{R}^{V \times V} \mid C P^{\text{row}} = P^{\text{row}} \right\} \neq \emptyset. \quad (19)$$

Proof. See Appendix A.1 of the online supplementary material. □

Remark 1. The set of matrices $\mathcal{C}_K(P)$ in Equation (19) can be viewed as the choice set of a linear program, where the objective function could be any arbitrary linear functional of C . To the best of our knowledge, this set was first studied by Recht et al. (2012), who use the linear program:

$$\min_{C \in \mathcal{C}_K(P)} \mathbf{b}' \text{diag}(C) \quad (20)$$

(where \mathbf{b} is any vector with distinct, non-zero entries) to factor a separable nonnegative matrix with known, nonnegative rank K . Theorem 1 shows that checking whether a column-stochastic matrix P with nonnegative rank K admits a rank K anchor-word factorization is *equivalent* to checking whether the linear program (20) has a nonempty choice set.

Remark 2. It turns out that an anchor-word factorization always exists when $K = 2$. We first use a simple geometric argument to explain the intuition behind this result. Consider a simple low-dimensional example where $V = 4$ and $K = 2$ (i.e., there are four words and only two topics). This example is depicted in Figure 1 below. Each column of the matrix P , which contains the probabilities assigned to each word in each document, can then be depicted in a tetrahedron representing the simplex in \mathbb{R}^4 . The topics themselves (the columns of A) also correspond to a set of probabilities over the four words; thus they can also be represented by points inside

the simplex. Further, because the documents are a mixture of two topics ($P = AW$), all documents will lie on the ray (depicted as a black solid line) that is spanned by the two topics, and in fact fall inside the convex hull of the two topics. Intuitively, when $K = 2$, we can always find an anchor-word factorization by intersecting the ray with the faces of the tetrahedron. This intersection is depicted by the red filled circles in the figure. It is easy to see that any matrix A with columns belonging to different faces of the tetrahedron will have the anchor-word structure.




Figure 1: Graphical representation of a topic model with $V = 4$ and $K = 2$ using the simplex in \mathbb{R}^4 . The vertices of the simplex represent the four words. The solid black line represents the ray spanned by the columns of the matrix P , which is assumed to have rank $K = 2$. The red filled circles in the intersection of the ray with the faces of the tetrahedron are the columns of a matrix A with two anchor words.

In Appendix B.3 of the online supplementary material, we complement our geometric arguments with an analytical derivation that uses Theorem 1 to formally show that when $K = 2 \leq \min\{V, D\}$, any nonnegative matrix P of rank two (and whose rows are different from zero) admits an anchor-word factorization. Our verification of Theorem 1 explicitly constructs a matrix $C \in \mathcal{C}_2(P)$ that satisfies Equation (19). Our construction also shows how to obtain the anchor words corresponding to P , starting from an arbitrary column-stochastic nonnegative matrix factorization of it.

Remark 3. We next argue that, even in simple low-dimensional problems, an anchor-word factorization need not exist. We do so through a geometric argument similar to the one discussed above (with $V = 4$) that illustrates why an anchor-word factorization frequently does not exist when $K = 3$, and to explain the differences vis-à-vis the case in which $K = 2$.

With four words ($V = 4$) and three topics ($K = 3$), we can still depict the columns of P in the tetrahedron we used in Figure 1. Further, because the documents are now a mixture of three topics, all documents will lie on the plane that is spanned by the three topics. This is illustrated in Figure 2.

We first note that if an anchor-word factorization exists, the topics must lie on the *edges* (the one-dimensional faces) of the tetrahedron. The reason is that a necessary condition for A to have anchor words is that all three topics are associated with at most two words (the word-topic matrix must have at least two zeros in each




Figure 2: Graphical representation of a topic model with $V = 4$ and $K = 3$ using the simplex in \mathbb{R}^4 . The plane represents the space spanned by the columns of the matrix P , which is assumed to have rank $K = 3$. The red filled circles are the intersection of the plane with the edges of the tetrahedron.

column).

We next note that a plane intersecting a tetrahedron will, in general, either intersect three or four of its edges. In case I (Figure 2a), the space spanned by the topics intersects three edges of the word simplex. In this case, those three edges necessarily share a common vertex. That means that the word associated with that vertex has non-zero probability under all three topics. But since the word-topic matrix has two zeros in each column, it then immediately follows that the three solid red circles provide an anchor-word factorization of P .

In case II (Figure 2b), the space spanned by the topics intersects four edges of the word simplex. No matter which three out of these four circles one selects as the columns of A , each row has at least one entry equal to zero. Thus, up to a row permutation,

$$A = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \gamma & 0 \\ 0 & 1 - \gamma & 1 - \beta \\ 1 - \alpha & 0 & \beta \end{pmatrix},$$

for $\alpha, \beta, \gamma \in (0, 1)$, and A does not have anchor words. Further, we can show using Theorem 1 that any P of the form above does not have an anchor-word factorization; see Appendix B.4 of the online supplementary material. In Section 4 we also provide numerical evidence suggesting that the probability that randomly sampled matrices P with a nonnegative rank K with $2 < K < \min\{V, D\}$ admit an anchor-word factorization could be very low.

Figure 2 is also helpful to illustrate what happens when the anchor-words assumption is erroneously imposed (and the model misspecified). Suppose P does not have an anchor-word factorization and the documents lie on the plane depicted in Case II (Figure 2b), but we estimate A under the anchor-words assumption. This restricts the set of word-topic matrices A to those that span planes which only intersect the tetrahedron at three vertices (cf. Figure 2a). Figure 2 suggests that this can lead to both misleading interpretation of the topics and a substantially poorer model fit.

Remark 4. We show in Appendix A.4 of the online supplementary material that for any matrix norm Theorem 1 is equivalent to saying that a column-stochastic matrix P with nonnegative rank K admits a rank K anchor-word factorization if and only if

$$\min_{C \in \mathcal{C}_K} \|CP^{\text{row}} - P^{\text{row}}\| = 0. \quad (21)$$

We use this simple observation to construct a statistical test for the null hypothesis of anchor words. For the remainder of the paper we let $\|\cdot\|$ denote the Frobenius norm.

Remark 5. While Theorem 1 shows that some column-stochastic matrices with nonnegative rank K do not have an anchor-word factorization, this is not yet sufficient to establish the statistical testability of the anchor-words assumption. For instance, if every matrix P that does not have an anchor-word factorization could be approximated by a sequence of matrices with an anchor-word factorization, then Lemma 2.1 in Canay et al. (2013) would imply that the power of any test of size α must also be at most α at any such P . However, intuitively, continuity of the norm in Equation (21) can be used to show that whenever P does not have an anchor-word factorization, there is no sequence of matrices with an anchor-word factorization that converges (in total variation norm) to P (see Appendix B.1 of the online supplementary material for a formal derivation). This shows that the matrices P that do not have an anchor-word factorization belong, in a sense, to the *topological interior* (with respect to the total variation norm) of Θ_1 .

Final Comment on Theorem 1. We first encountered the connection between the set $\mathcal{C}_K(P)$ and the anchor-word factorization of P in the work of Recht et al. (2012). In particular, their Theorem 3.1 on p. 4 can be viewed, *mutatis mutandi*, as showing that if an anchor-word factorization of P exists, then $\mathcal{C}_K(P)$ is nonempty.

We extend the results in Recht et al. (2012) in two ways. First, we show that it is possible for the set $\mathcal{C}_K(P)$ to be empty for some matrices P that have nonnegative rank K , provided $2 < K < \min\{V, D\}$. Second, we establish the reverse direction: if $\mathcal{C}_K(P)$ is nonempty, then an anchor-word factorization of P exists. In other words, we show that not every matrix P has an anchor-word factorization, and that the matrices P for which $\mathcal{C}_K(P)$ is empty are precisely those for which there is no anchor-word factorization.

To prove Theorem 1 we establish that—up to a permutation matrix—the construction given in our illustrative example of Equation (16) is possible if and only if P has an anchor-word factorization (see Lemma 1 in Appendix A.1 of the online supplementary material). One direction of this Lemma is implicitly used by Recht et al. (2012) in the introduction of their *hottopixxx* algorithm (see their definition of a factorization localizing matrix) and is also stated in Equation 1.1 of Gillis (2013). We formally derive this result and its reverse direction in Lemma 1.

3.2 Testing the existence of anchor words

Let \widehat{P}^{row} denote some estimator of the matrix P^{row} based on the available data Y . Consider the test statistic $T(Y)$ defined as

$$T(Y) \equiv \inf_{C \in \mathcal{C}_K} \|C\widehat{P}^{\text{row}} - \widehat{P}^{\text{row}}\|. \quad (22)$$

In Appendix B.2 of the online supplementary material we show that when $\|\cdot\|$ is the Frobenius norm, this “inf” is attained for any \widehat{P}^{row} , and thus can be replaced by a “min”. Define $\overline{N}_D = (N_1, \dots, N_D)$ to be the vector collecting the total number of words per document. Let $q_{1-\alpha}(AW, V, D, K, \overline{N}_D)$ denote the $1 - \alpha$ quantile of the test statistic $T(\cdot)$ assuming that the data was generated by the multinomial model in Equation

(8) with parameters (A, W) . Since the distribution of the data used to estimate P^{row} only depends on the parameters (A, W) through AW , then the quantiles of T only depend on the parameters through the same product. Consider then the critical value

$$q_{1-\alpha}^*(V, D, K, \bar{N}_D) \equiv \sup_{(A, W) \in \Theta_0} q_{1-\alpha}(AW, V, D, K, \bar{N}_D), \quad (23)$$

and define the test:

$$\phi^*(Y) \equiv \begin{cases} 1 & \text{if } T(Y) > q_{1-\alpha}^*(V, D, K, \bar{N}_D), \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

The next theorem shows the test in (24) has significance level α for any possible configuration (V, D, K, \bar{N}_D) of the multinomial model in Equation (8). It also gives a high-level sufficient condition under which the test has nontrivial power.

Theorem 2. *The test ϕ^* has significance level α ; i.e.,*

$$\sup_{(A, W) \in \Theta_0} \mathbb{E}_{(A, W)} [\phi^*(Y)] \leq \alpha. \quad (25)$$

Moreover, suppose there is a parameter value $(A, W) \in \Theta_1$ for which

$$\mathbb{P}_{(A, W)} \left(\inf_{C \in \mathcal{C}_K} \|(C - \mathbb{I}_V)(AW)^{\text{row}}\| - \sup_{C \in \mathcal{C}_K} \|(C - \mathbb{I}_V)(\hat{P}^{\text{row}} - (AW)^{\text{row}})\| > q_{1-\alpha}^*(V, D, K, \bar{N}_D) \right) \quad (26)$$

exceeds α . Then for such $(A, W) \in \Theta_1$ we have

$$\mathbb{E}_{(A, W)} [\phi^*(Y)] > \alpha.$$

Proof. We first establish (25). For any $(A, W) \in \Theta_0$

$$\begin{aligned} \mathbb{E}_{(A, W)} [\phi^*(Y)] &= \mathbb{P}_{(A, W)} (\phi^*(Y) = 1) \\ &= \mathbb{P}_{(A, W)} \left(\min_{C \in \mathcal{C}_K} \|C\hat{P}^{\text{row}} - \hat{P}^{\text{row}}\| > q_{1-\alpha}^*(V, D, K, \bar{N}_D) \right) \\ &\leq \mathbb{P}_{(A, W)} \left(\min_{C \in \mathcal{C}_K} \|C\hat{P}^{\text{row}} - \hat{P}^{\text{row}}\| > q_{1-\alpha}(AW, V, D, K, \bar{N}_D) \right) \\ &= \alpha, \end{aligned}$$

where the last two lines follow from the definition of $q_{1-\alpha}^*$. Thus, ϕ^* has size of at most α , regardless of the model's configuration (V, D, K, \bar{N}_D) .

Now we analyze power. The power of the test ϕ^* at $(A, W) \in \Theta_1$ is given by

$$\mathbb{P}_{(A, W)} \left(\min_{C \in \mathcal{C}_K} \|C\hat{P}^{\text{row}} - \hat{P}^{\text{row}}\| > q_{1-\alpha}^*(V, D, K, \bar{N}_D) \right).$$

Since $\|\cdot\|$ satisfies the reverse triangle inequality, then

$$\min_{C \in \mathcal{C}_K} \|C\hat{P}^{\text{row}} - \hat{P}^{\text{row}}\| \geq \inf_{C \in \mathcal{C}_K} \|(C - \mathbb{I}_V)(AW)^{\text{row}}\| - \sup_{C \in \mathcal{C}_K} \|(C - \mathbb{I}_V)(\hat{P}^{\text{row}} - P^{\text{row}})\|.$$

This means that the power of the test $\phi^*(Y)$ at any parameter values $(A, W) \in \Theta_1$ that satisfies Equation (26) is at least α . \square

The nontrivial power of the test ϕ^* in Theorem 2 is obtained under the high-level assumption in (26), which involves the following three terms:

- i) $\inf_{C \in \mathcal{C}_K} \|(C - \mathbb{I}_V)(AW)^{\text{row}}\|$,
- ii) $\sup_{C \in \mathcal{C}_K} \|(C - \mathbb{I}_V)(\widehat{P}^{\text{row}} - (AW)^{\text{row}})\|$,
- iii) $q_{1-\alpha}^*(V, D, K, \overline{N}_D)$.

Intuitively, the high-level assumption in (26) requires the term in i) to be larger than the terms ii)-iii), with probability at least α .

In Appendix A.2 of the online supplementary material we verify the high-level assumption in Theorem 2 for the estimator $\widehat{P}_{\text{freq}}^{\text{row}}$: the row-normalized version of the relative frequency counts, $\widehat{P}_{\text{freq}} \equiv n_{v,d}/N_d$. In particular, we show that, under a weak regularity condition, if $N_{\min} \equiv \min\{N_1, \dots, N_D\}$ is large enough, the high-level assumption used in Theorem 2 holds at any point $(A, W) \in \Theta_1$ such that $P = AW$ does not have an anchor-word factorization. In fact, we show in Corollary 1 in Appendix A.2 of the online supplementary material that the probability of the event in (26) (and thus the power of the test) will be arbitrarily close to one, ensuring consistency of the test at any point in the alternative for which the anchor-word factorization does not exist.

4 Numerical Results

We next present numerical results to accompany our theoretical analysis in the previous section. First, we use Theorem 1 to study how likely it is to draw a matrix P that has an anchor-word factorization for different values of (V, K, D) . Then, we illustrate Theorem 2 by showing finite sample results for a version of our test that uses a “bootstrap bound” for the critical value.

4.1 Known P

The goal of this section is to understand how likely it is for a randomly generated matrix of the form $P = AW$ to admit an anchor-word factorization for a variety of combinations of (V, K, D) . To do this, we randomly generate column-stochastic matrices $(A, W) \in \mathbb{R}^{V \times K} \times \mathbb{R}^{K \times D}$. For each realization, we then use a linear program—as the one that appears in Equation (20) in Remark 1—to check whether the set $\mathcal{C}_K(P)$ in Equation (19) is empty or not. We then report the fraction of randomly generated matrices for which the set $\mathcal{C}_K(P)$ turned out to be nonempty. By Theorem 1 this is equivalent to the fraction the sampled P that has an anchor-word factorization.

The results of this exercise are depicted in Figure 3, where we fix $D = 1000$ and vary $K \in \{2, 3, 4\}$ and $V \in \{4, 10, 100\}$. Figure 3a corresponds to the case in which the columns of A and W are sampled from independent Dirichlet distributions with concentration parameters α equal to 1 and 0.01 respectively. Note that, by construction, the probability of creating a matrix A that has anchor words is zero under this data generating process (“DGP”). We therefore refer to this data generating process for P as “No anchor words”. Figure 3b reports results for (A, W) generated as in our “No anchor words” simulation, but with all off-diagonal entries

in the first K rows of A replaced with zeros before re-normalizing the columns of A to sum to one. This ensures that under this DGP the resulting word-topic matrix A has anchor words. We refer to this data generating process as “With anchor words”.⁷

In both figures, we are reporting the fraction of simulations in which P has an anchor-word factorization, with yellow indicating an anchor factorization exists in all realizations. A blue square for a given combination of K and V indicates that P does not have an anchor factorization in any of its realizations.




Figure 3: Fraction of randomly generated matrices $P = AW$ with an anchor-word factorization for different configurations of (V, K) and $D = 1000$. Figure based on 500 simulations.

The numerical results are in line with the theoretical results discussed in Section 3. According to Remark 2 any P with rank $K = 2$ admits an anchor-word factorization. Similarly, when $K = V$ any matrix $P = AW$ admits an anchor-word factorization. This is reflected by the yellow square in the bottom left of both panels. Next, we see that for $K = 3$ and $V = 4$ some realizations of (A, W) allow an anchor-word factorization, while others do not (cf. Figure 2). In the more general case ($K > 2$, and $V \in \{10, 100\}$), we find that there does not exist an anchor-word factorization in most realizations (Figure 3a), unless we explicitly impose this structure on A (Figure 3b).

In Appendix C.1 of the online supplementary material, we study the effects of introducing varying degrees of sparsity in the word-topic matrix A on the likelihood that a randomly generated population term-document frequency matrix admits an anchor-word factorization. We find that, as the amount of sparsity in A increases, this becomes more likely.

4.2 Unknown P

In this section we conduct small scale simulations to analyze the case in which P is unknown and we observe count data generated by the multinomial model in (8). In this case, we use the count data to test for the existence of anchor words. Before presenting the results, we provide details on the construction of the test statistic and the critical value that are used in this section.

⁷We disregard P in the rare case that we obtain a word in the vocabulary that is used extremely infrequently and satisfies $\sum_d p_{vd} \leq 0.03$ to avoid numerical issues in the row-normalization step (cf. Corollary 1 of the online supplementary material).

4.2.1 Test statistic

We compute the test statistic $T(Y)$ in Equation (22) as

$$T(Y) \equiv \min_{C \in \mathcal{C}_K} \|C\hat{P}_{\text{freq}}^{\text{row}} - \hat{P}_{\text{freq}}^{\text{row}}\|_F,$$

where $\|\cdot\|_F$ denotes the Frobenius norm and $\hat{P}_{\text{freq}}^{\text{row}}$ denotes the row-normalized term-document frequency matrix. The (v, d) -entry of $\hat{P}_{\text{freq}}^{\text{row}}$ is

$$(n_{v,d}/N_d) / \sum_{d=1}^D (n_{v,d}/N_d).$$

Two remarks are in order. First, as discussed after the statement of Theorem 2, the test statistic $T(Y)$ could have been computed using a different estimator for the row-normalized population term-document frequency matrix. We use the simple row-normalized term-document frequency matrix because i) it is straightforward to implement, and ii) the uniform rates of estimation error reported in Proposition 2 (in particular, Equation 50) suggest good performance relative to the other estimators we analyzed. See Appendix (B.6) of the online supplementary material for the statistical properties of alternative estimators of P^{row} .

Second, the computation of the test statistic $T(Y)$ involves the minimization of a quadratic objective function over the set \mathcal{C}_K , which is a set of bounded, real-valued $V \times V$ matrices defined by 1 linear equality and $2V^2$ linear inequalities. We solve this optimization problem in MATLAB® (version 2022b) using the function `lsqlin`.⁸

4.2.2 Critical values

The test we presented in Theorem 2 uses the largest $1 - \alpha$ quantile of the distribution of the test statistic $T(Y)$ that can be generated by matrices (A, W) that satisfy the null hypothesis. This critical value is defined formally in Equation (23) and, in a slight abuse of notation, throughout this section we simply denote it as $q_{1-\alpha}^*$.

Theorem 2 shows that the test that rejects whenever the test statistic, $T(Y)$, exceeds $q_{1-\alpha}^*$ has correct size and nontrivial power. Although this test is useful to establish the testability of the anchor-words Assumption, obtaining $q_{1-\alpha}^*$ in our application is extremely computationally demanding. For instance, one could try to create either a deterministic or random grid of parameters (A, W) in Θ_0 , and approximate $q_{1-\alpha}^*$ from below by the largest quantile for the random variable $T(Y)$ over the grid. This will require constructing a deterministic (or random) grid over matrices of dimension $V \times D$ and $K \times V$ that satisfy the anchor-words assumption. Due to the dimension of the parameter space, it seems unlikely that one could generate a good approximation of $q_{1-\alpha}^*$ using this approach. Below, we describe two computationally feasible approaches to obtain a bound on $q_{1-\alpha}^*$.

- *Algebraic upper bound for $q_{1-\alpha}^*$.* Lemma 4 in Appendix B.5 of the online supplementary material implies

⁸The `lsqlin` function minimizes an objective function of the form $f(x) \equiv \|Cx - d\|_2$ (where x is a vector in \mathbb{R}^n and C is a matrix of dimension $m \times n$ and d is a vector of dimension $m \times 1$) subject to a set of linear equalities and inequalities. To use this function for our problem we vectorize the equation $C\hat{P}_{\text{freq}}^{\text{row}} - \hat{P}_{\text{freq}}^{\text{row}}$ as

$$(\mathbb{I}_D \otimes \hat{P}_{\text{freq}}^{\text{row}\top})\text{vec}(C) - \text{vec}(\hat{P}_{\text{freq}}^{\text{row}}),$$

and treat the choice variable x as $\text{vec}(C)$. For reference, the computation of the test statistic takes only 137 and 58 seconds respectively for the two corpora we consider in the application in Section 5.

that, under the same assumptions as in Proposition 2:

$$q_{1-\alpha}^* \leq \sup_{C \in \mathcal{C}_K} \|C - \mathbb{I}_V\|_F \cdot R_\gamma(\alpha), \quad \text{where } R_\gamma(\alpha) \equiv \sqrt{\frac{8 \left(1 - \frac{1}{V}\right)}{\gamma^2 \cdot \alpha} \cdot \frac{V^2}{N_{\min} \cdot D}},$$

and $\gamma \in (0, 1)$ is a constant such that for any $(A, W) \in \Theta$, $\sum_{d=1}^D (AW)_{vd}/D \geq \gamma/V$ for all v .

The first term in the bound has a closed-form solution and $R_\gamma(\alpha)$ can easily be computed for a chosen value of γ . However, in our simulations we find that such an algebraic bound is extremely conservative with poor power properties. We thus do not pursue this further.

- A “bootstrap bound” for $q_{1-\alpha}^*$. For any matrix $C \in \mathcal{C}_K$ we have that

$$T(Y) \leq \|C \widehat{P}_{\text{freq}}^{\text{row}} - \widehat{P}_{\text{freq}}^{\text{row}}\|_F$$

for any $C \in \mathcal{C}_K$ (by the definition of $T(Y)$). Moreover, for any $C \in \mathcal{C}_K$ we have

$$\|C \widehat{P}_{\text{freq}}^{\text{row}} - \widehat{P}_{\text{freq}}^{\text{row}}\|_F = \|(C - \mathbb{I}_V)(\widehat{P}_{\text{freq}}^{\text{row}} - P^{\text{row}}) + C P^{\text{row}} - P^{\text{row}}\|_F.$$

Theorem 1 shows that for each P such that $P = AW$ with $(A, W) \in \Theta_0$ there exists $C_P \in \mathcal{C}_K$ such that $C_P P^{\text{row}} - P^{\text{row}} = 0_{V \times D}$. Consequently,

$$T(Y) \leq \|(C_P - \mathbb{I}_V)(\widehat{P}_{\text{freq}}^{\text{row}} - P^{\text{row}})\|_F. \quad (27)$$

This means that for any $(A, W) \in \Theta_0$, the $1 - \alpha$ quantile of $T(Y)$ under P is upper bounded by the $1 - \alpha$ quantile of the random variable

$$\|(C_P - \mathbb{I}_V)(\widehat{P}_{\text{freq}}^{\text{row}} - P^{\text{row}})\|_F. \quad (28)$$

In Appendix A.3 of the online supplementary material we show that one can approximate the distribution of (28) using a parametric bootstrap that replaces C_P by $C_{\widehat{P}}$ where \widehat{P} is an estimator of P that imposes the anchor-words assumption.

In particular, let \widehat{A} and \widehat{W} denote estimators of the parameters (A, W) under the anchor-words assumption. Let $\widehat{P} \equiv \widehat{A}\widehat{W}$ denote the plug-in estimator for the population term-document frequency matrix based on \widehat{A} and \widehat{W} . Define Y_d^* as the random vector with distribution

$$Y_d^* \sim \text{Multinomial}\left(N_d, (\widehat{P})_{\bullet d}\right), \quad (29)$$

and assume that the columns of the matrix $Y^* \equiv (Y_1^*, \dots, Y_D^*)$ are generated independently according to (29).

Let $\widehat{P}_{\text{freq}}^*$ denote the matrix of frequency counts associated with Y^* . That is, $\widehat{P}_{\text{freq}}^*$ is the $V \times D$ matrix with d -th column given by Y_d^*/N_d . Consider approximating the unknown distribution in (28) by the distribution of the random vector

$$\|(C_{\widehat{P}} - \mathbb{I}_V)((\widehat{P}_{\text{freq}}^*)^{\text{row}} - \widehat{P}^{\text{row}})\|_F, \quad (30)$$

conditional on \widehat{P} . Theorem 3 in Appendix A.3 of the online supplementary material shows that the distribution of (30), conditional on the data, is *close* in P-probability to the distribution of the bounding random variable in (28). To formalize this result we use the *bounded Lipschitz metric* (see p. 394 of Dudley (2002), and also Chapter 2.2.3 and Chapter 10 in Kosorok (2007)) to measure closeness between the distributions in (28) and

(30). The bootstrap “consistency” is established under two high-level assumptions that can be readily verified when V and D are fixed and N_{\min} grows to infinity, but we think could potentially hold also in situations where V and D also grow with N_{\min} .

The bootstrap consistency result in Appendix A.3 of the online supplementary material thus suggests that the $1 - \alpha$ quantile of (30) can be used to implement a *conservative, point-wise* valid version of our test at significance level α . Note that this procedure is computationally straightforward as $C_{\hat{p}}$ is only computed once and thus there is no need to recompute the anchor-word estimates across bootstrap simulations. Note also that the bootstrap consistency in Theorem 3 essentially relies on a continuous mapping theorem; c.f., Proposition 10.7 Kosorok (2007) and, thus, there is no need for re-centering before getting the critical value.

4.2.3 Results

In the previous subsection, we showed that it is possible to use a “bootstrap bound” for the critical value of the test described in Theorem 2. We established the “consistency” of our bootstrap strategy; but, unfortunately, the consistency holds only “pointwise” at a fixed (A_0, W_0) in the null hypothesis. This means that the test based on the bootstrap upper bound need not have the correct size in finite samples. With this in mind, we next present simulation results based on the same set of DGPs (“With anchor words” and “No anchor words”) as in Section 4.1 to asses the size and power of our proposed bootstrap strategy. Throughout, we assume K is known a priori and correctly specified. To recap, the “bootstrapped” version of our test can be described as follows.

Step 1. Given the data Y , compute the test statistic $T(Y) = \min_{C \in \mathcal{C}_K} \|C \hat{P}_{\text{freq}}^{\text{row}} - \hat{P}_{\text{freq}}^{\text{row}}\|_F$.

Step 2. Obtain an estimate for P that has the anchor-word factorization.

- (a) Since K is known, we follow the recommendation in Bing et al. (2020b) and run the algorithm of Arora, Ge, Halpern, Mimno, Moitra, Sontag, Wu & Zhu (2013) on Y to obtain \hat{A}_0
- (b) Let \hat{W}_0 be the Maximum Likelihood estimator of W in the multinomial model (8) but treating \hat{A}_0 as the true unknown A (Bing, Bunea, Strimas-Mackey & Wegkamp (2022)).
- (c) Let $\hat{P}_0 = \hat{A}_0 \hat{W}_0$

Step 3. Find $C_{\hat{p}_0}$ by solving the minimization in (21).

Step 4. Estimate the quantile of the upper bound $T^*(Y) \equiv \|(C_P - \mathbb{I}_V)(\hat{P}_{\text{freq}}^{\text{row}} - P^{\text{row}})\|_F$, using the bootstrap.

- (a) Simulate n_{sim} new realizations of Y using \hat{P}_0 .
- (b) For each new realization Y_i , obtain $T_{\text{boots}}^*(Y_i) \equiv \|(C_{\hat{p}_0} - \mathbb{I}_V)((\hat{P}_{\text{freq}}^i)^{\text{row}} - \hat{P}_0^{\text{row}})\|_F$, for $i = 1 \dots, n_{\text{sim}}$, where \hat{P}_{freq}^i is the row-normalized term-document frequency matrix based on data Y_i .
- (c) Set critical value cv_α to the $(1 - \alpha)$ th percentile of $T_{\text{boots}}^*(Y_i)$.

Step 5. Reject the null hypothesis if $T(Y)$ is larger than cv_α .

Figure 4 below presents the average power under “No anchor words” and average rates of Type I error under “With anchor words” of the bootstrapped version of the test. In order to compute the average performance of the test, we generate random draws from (A, W) using the same procedures used to generate Figure 3. Then, for each of these draws, we sample the matrix of word counts, Y , from the multinomial model in Equation (8), where each document contains 10,000 words.

Figure 4a uses “No anchor words” (as described in the previous subsection) to generate draws from (A, W) . Since the probability of creating a matrix A that has anchor words is zero, the share of realizations (Y, A, W)

for which the bootstrapped test rejected the null can be interpreted as average power. For $V = 10$ and $K = 4$, the bootstrapped test rejects in all realizations. On the other hand, we note that the average power seems to deteriorate when the vocabulary size increases (e.g., for $V = 100$ and $K = 4$, we obtain a power of 42%).

Figure 4b uses “With anchor words” (as described in the previous subsection) to generate draws from (A, W) , such that the word-topic matrix A always has an anchor-word factorization. Reporting the share of realizations of (Y, A, W) for which the bootstrapped test rejects the null gives a Monte-Carlo approximation to the average rate of Type I error at a particular configuration (V, K, D) . The figure thus suggests that the bootstrapped version of the test performs well in terms of its size. Using a nominal 5%-test, the average rate of Type I error of the test ranges from 0-5% in Figure 4b.




Figure 4: Proportion of realizations (Y, A, W) in which our test rejects as we vary the number of words and the number of topics. $D = 1000$ and each document contains 10,000 words. Figure based on 500 simulations of (A, W) .

To further illustrate the power of the test under our “No Anchor words” DGP, we next increase the number of topics to $K = 6$, and vary the the vocabulary size V for two values of the document size n_d . We then compute the rejection frequency of our bootstrapped test. This is depicted in Figure 5. We again conclude that our test exhibits nontrivial power and that the power of our test deteriorates as V increases, especially for moderately sized documents. We note that the fact that for a fixed K , the power of our test deteriorates as we increase V is consistent with the results in Ding, Ishwar & Saligrama (2015). Their results essentially show that, as V increases relative to K , any matrix A generated at random by a Dirichlet distribution will be “closer” to a matrix with the anchor-word structure.

5 Empirical Application

In this section we analyze a subset of the “transcripts” of the meetings of the Federal Open Market Committee (FOMC), the main body within the Federal Reserve System in charge of setting monetary policy in the United States. We focus on the FOMC transcripts during the “Greenspan period,” the 150 meetings from August 1987 to January 2006 in which Alan Greenspan was chairman. We separate each transcript into two parts: the discussion of domestic and international economic conditions (FOMC1) and the discussion of the monetary




Figure 5: Average power of our test as we vary the size of the vocabulary. We fix $K = 6$ and simulate 1000 documents. Figure based on 100 simulations.

policy strategy (FOMC2). This gives us two different corpora to analyze.

The first corpus (FOMC1) allows us to illustrate the potential benefits of assuming the existence of anchor words in a concrete empirical application. Aside from the computational tractability and the theoretical identification results that become available under the anchor-words assumption, the estimated anchor words can potentially provide natural and objective labels for the estimated topics. We think this is an important point, as it has recently been argued that an inherent challenge of topic models in empirical applications is that they “*do not generate objective topic labels*” and that “*A given topic consists of many words, and words are scattered across many topics, so the outputs are often difficult to interpret.*”; see the discussion in Section 3.2.2.1 of Ash & Hansen (2023). In contrast, the anchor words for FOMC1 are all relatively easy to interpret. Moreover, the estimated topic proportions for the FOMC1 corpus seem to be consistent with historical events that shaped monetary policy decisions during the Greenspan period.

On the other hand, we find the estimates we obtain under the anchor-words assumption for FOMC2 harder to interpret: Anchor words for different topics have very similar meanings, and thus it becomes difficult to understand the difference between topics. Further, with the exception of two topics, we found it difficult to provide a rationale for the historical evolution of the topic shares. We would like to argue that this is not a flaw of the method; instead we think it may be a warning about the compatibility of the anchor-words Assumption and the true data generating process.

We then apply our suggested testing procedure to these two corpora and indeed find that a nominal 5%-level test fails to reject the null hypothesis of anchor words for the FOMC1 corpus, but rejects for the FOMC2 corpus.

The rest of this section is organized as follows. Section 5.1 presents a broad description of the FOMC transcripts, along with some descriptive statistics for the FOMC1 and FOMC2 corpora. Section 5.2 presents the estimation results for the parameters of the topic model, assuming the existence of anchor words. This section also provides a detailed interpretation of the results. In Section 5.3 we then test the anchor-words assumption in both corpora. Finally, Section 5.4 discusses the finite-sample properties of the test.

5.1 FOMC transcripts

The nineteen participants of the FOMC meetings—seven members of the Board of Governors of the Federal Reserve System and the presidents of the twelve regional Reserve Banks—convene regularly to discuss domestic and international economic conditions, conditions in financial markets, and other factors considered relevant for monetary policy. The purpose of this discussion is to make key decisions on the stance of monetary policy. The FOMC Secretariat typically prepares a verbatim *transcript* of the FOMC meeting proceedings and conference calls after their occurrence.⁹ This is the most detailed record of the FOMC meeting and it is currently released with a lag of five years.

We focus on the FOMC transcripts during the “Greenspan period,” the 150 meetings from August 1987 to January 2006 in which Alan Greenspan was chairman. The transcripts can be obtained directly from the website of the Federal Reserve. This dataset has been used recently in the work of Hansen et al. (2018) (henceforth HMP) to study the effects of increased ‘transparency’ on the discussion inside the FOMC when deciding monetary policy. We followed HMP in merging the transcripts for the two back-to-back meetings in September 2003 and dropping the meeting on May 17, 1998.¹⁰ As a result, we ended up with 148 transcripts.

We removed non-alphabetical words, words with a length of one, and common stop words. We also constructed the 150 most frequent bigrams (combinations of two words) and 50 most frequent trigrams (three words). We then stemmed all the words using a standard approach¹¹.

We separate each transcript into two parts: the discussion of domestic and international economic conditions (FOMC1) and the discussion of the monetary policy strategy (FOMC2). These sections are not sign-posted, so we manually separated each transcript (we tried to match closely the separation rules used by HMP and discussed in their work). At the end, we construct two separate term-document matrices, one for each section. To reduce the size of the vocabulary, we follow Ke et al. (2022) and further rank the remaining terms by their term frequency-inverse document frequency (tf-idf) score and keep those with the highest tf-idf score (we also manually looked at these terms to ensure that they were meaningful for our analysis). At the end we are left with 200 terms for FOMC1 and 150 for FOMC2. The final two term-document matrices that we use for estimation have dimension 200×148 and 150×148 each.


We start by providing a high level overview of our data. First, Figure 6 plots the document size for each of the meetings included in our sample. The figure shows that documents in the FOMC1 corpus are typically larger: The average document size in the FOMC1 corpus is 2309, but only 853 for FOMC2. We also note that the number of words per meeting for FOMC1 exhibits a positive time trend, while the size of the FOMC2 documents remained relatively stable over time.

Second, Figure 7 presents the “word cloud” corresponding to the vocabulary used in each corpus. A word cloud is a convenient graphical representation of the frequency of each term in a corpus. Terms that appear more frequently are depicted with a larger font size. The five highest terms in each corpus are depicted in orange. Although the two corpora have a number of overlapping terms (e.g., data, concern, expect, inflat, growth to name but a few), the word clouds suggest that the term distributions in the two corpora are markedly different. This is consistent with the fact that the FOMC1 corpus focuses mainly on the description of the domestic and


⁹The speakers’ original words are lightly edited to facilitate the reader’s understanding. In addition, a very small amount of information received on a confidential basis is subject to deletion.

¹⁰The beginning of the transcript for the May 17, 1998 meeting states: “No transcript exists for the first part of this meeting, which included staff reports and a discussion of the economic outlook.”

¹¹We used the Natural Language Toolkit (nltk) library in Python, its PorterStemmer package for word stemming, and its Collocation package for the bigrams and trigrams.



(a) FOMC1



(b) FOMC2

Figure 6: Number of words per document in the FOMC1/FOMC2 corpora. The solid horizontal line represents the average number of words per meeting. For reference, the grey bars represent recession dates, as reported by the National Bureau of Economic Research.

foreign economic conditions that are relevant for monetary policy decisions, while FOMC2 focuses on the discussion of monetary policy alternatives.



(a) FOMC1

(b) FOMC2

Figure 7: Word Cloud for the FOMC1/FOMC2 corpora. The five highest terms in each corpus are colored in orange.

5.2 Anchor words in FOMC1 and FOMC2 corpora

5.2.1 Choosing K

Although the theory presented in Section 2 assumed the number of topics in the model to be known, in practice K needs to be selected (a priori or a posteriori) by the researcher. As noted by Blei & Lafferty (2009) “choosing the number of topics is a persistent problem in topic modeling and other latent variable analysis. In some cases,

the number of topics is part of the problem formulation and specified by an outside source. In other cases, a natural approach is to use cross validation on the error of the task at hand (e.g., information retrieval, text classification).”

Bing et al. (2020a) have recently shown that the anchor-words assumption allows the researcher to estimate K and, under some regularity assumptions, guarantee that the estimator is consistent (it coincides with the true number of topics with high probability).¹² We thus estimate the number of topics for the FOMC1 and FOMC2 corpus separately using the algorithm suggested by Bing et al. (2020a), and obtain $\hat{K}_{\text{FOMC1}} = 4$ and $\hat{K}_{\text{FOMC2}} = 5$. In the remaining part of the application we estimate the remaining parameters of the topic model using these numbers of topics as given.

5.2.2 Estimation of A

We start by reporting the estimates of A and W based on state-of-the-art algorithms that assume the existence of anchor words.

To the best of our knowledge, the FOMC corpus has only been analyzed using the Latent Dirichlet Allocation model of Blei, Ng & Jordan (2003) and the robust Bayes version of the algorithm recently suggested by Ke et al. (2022).¹³ By reporting the model’s estimated parameters under the anchor-words assumption, we provide a novel estimate of the topics discussed in FOMC meetings and their distributions.

Our results, however, suggest that—even without a formal statistical test—the estimates obtained from imposing the anchor-words assumption may appear more reasonable in some contexts than in others. To us, this means that the anchor-words assumption may not always be appropriate, and that a statistical test for the existence of anchor words is a valuable tool for practitioners.

Estimated matrix A for FOMC1: Figure 8 presents word clouds summarizing the estimator of A obtained from the FOMC1 corpus under the anchor-words assumption. Terms that have a higher estimated probability under a given topic are depicted in larger font sizes, and the five terms with the highest probability appear in orange. Our baseline results are for the estimator suggested in Bing et al. (2020b), which adapts to unknown sparsity of A , and is minimax optimal under some assumptions.¹⁴ The caption that appears below each sub-figure presents the anchor words corresponding to each topic; that is, the words that are exclusive to the topic represented by the word cloud.

A practical advantage of using the anchor-words assumption in the estimation of A is that the anchor words, along with the most frequent words in each topic, usually provide a simple interpretation for the latent topic (and thus, a simple interpretation of the thematic structure in the corpus). For example, we think that, without much controversy, we could label Topic 1 as “foreign conditions.” The anchor word for this topic is “foreign” and the most frequent words on this topic —“export,” “dollar,” “import”— can be associated to developments in foreign markets (such as changes in the exchange rate, foreign demand, etc).

Topics 2 and 3 (which, using their anchor words, we can label “recoveri” and “uncertainty” respectively)

¹²We would like to thank the authors for kindly sharing their code to implement Algorithm 2 in Bing et al. (2020a).

¹³e.g., see Hansen et al. (2018) and Fligstein, Brundage & Schultz (2017)

¹⁴Appendix C.2 of the online supplementary material presents results for the estimators suggested in Arora, Ge & Moitra (2012), Ke et al. (2022), as well as the Latent Dirichlet Allocation. Note that Arora, Ge & Moitra (2012)’s algorithm outputs a unique anchor word for each topic, whereas Bing et al. (2020b)’s algorithm can output multiple anchor words for a topic. The topic estimates from Arora, Ge & Moitra (2012) are similar to our baseline result, giving anchor words “wage,” “uncertainty” and “recoveri,” which are also anchor words in our baseline results. Ke & Wang (2022) and the LDA implementation don’t explicitly impose anchor-words Assumption, and give estimates different from Bing et al. (2020b).

also have a straightforward interpretation. Topic 3 is an interesting finding given anecdotal evidence on the importance that the themes of “risk and uncertainty” played on Alan Greenspan’s framework for monetary policy.¹⁵

It is worth mentioning that the anchor words for each topic need not coincide with its most frequent terms. For example, the anchor words in Topic 4 could, in principle, all be linked to goal of *maximum employment* in the Federal Reserve’s policy mandate. However, none of the anchor words appears in the five most frequent terms in the topic. In fact, the most frequent terms —“inflat,” “price,” “increase”— are evocative of the goal of price stability, which is the other part of the Federal Reserve’s dual mandate. Thus, one could label Topic 4 as the “dual mandate” topic.

In summary, we think that the four topics found in FOMC1 —“foreign conditions,” “recoveri,” “uncertainty,” and “dual mandate”— indeed uncover a reasonable thematic structure in the FOMC1 corpus.

Estimated matrix A for FOMC2: Figure 9 presents word clouds summarizing the estimator of A obtained from the FOMC2 corpus under the anchor-words assumption. Recall that FOMC2 corpus covers the discussion of the monetary policy strategy. While it is again possible to interpret and label the topics using a combination of its anchor words and its most likely terms, we think that the results are not as clear-cut as in FOMC1.

Before giving an interpretation of the word clouds, it is worthwhile to make a few comments about i) the policy instruments that the FOMC has available to conduct monetary policy, and ii) the way in which policy choices are usually communicated to both the public and the Open Market Trading Desk at the Federal Reserve Bank of New York. Understanding both of these components is important for the interpretation of the estimated FOMC2 topics.

- *FOMC’s Policy instruments.* Traditionally, the Federal Reserve’s policy actions referred mainly to *open market operations* (buying or selling securities issued or backed by the U.S. government in the open market) in order to keep a key short-term money market interest rate, called the *federal funds rate*, at or near a desired target. It is common to think about this desired target for the federal funds rate as the policy variable selected by the Federal Reserve. Currently the Federal Reserve sets and announces a range for the target rate (for example, 5.00% to 5.25%), provides “forward guidance” to markets, and makes choices regarding balance sheet policies.¹⁶

- *FOMC’s Communication of Monetary Policy.* At the conclusion of each FOMC meeting, the Committee issues operating instructions to the Open Market Trading Desk at the Federal Reserve Bank of New York (Thornton, Wheelock et al. (2000)). Also, after each meeting, the FOMC currently communicates its decision about the stance of monetary policy to the public. The format in which the FOMC communicates the outcome of the meeting has changed over time. For example, before 1994, the monetary policy decision of the FOMC was not immediately communicated to the public. Instead, market participants had to infer the Federal Reserve’s actions from conditions in the money market. Beginning in 1994 the Federal Reserve started issuing a *statement* immediately after its meetings, but only if policy had changed. Starting in June 1999 such a statement was released for every scheduled meeting, regardless of whether or not there was a policy change. Also, from 1983 through 1999, the instructions to the Open Market Trading Desk included a statement about the Committee’s

¹⁵See, for example, Alan Greenspan’s famous 2003 speech in Jackson Hole, WY entitled “Monetary Policy Under Uncertainty,” available at the Federal Reserve’s website: <https://www.federalreserve.gov/boarddocs/speeches/2003/20030829/default.htm>.

¹⁶It is worth mentioning that the Federal Reserve has not always had an explicit operating target for the federal funds rate, and has not always provided explicit forward guidance to markets participants. While the exact point in time at which the Federal Reserve started using an explicit federal funds target rate is subject to some debate (Thornton 2006), it is common to assume that the target for the federal funds rate summarized the FOMC’s deliberations about the monetary policy stance throughout the Greenspan period.



(a) Topic 1: foreign



(b) Topic 2: recoveri



(c) Topic 3: invest, neg, uncertainty



(d) Topic 4: acceler, exampl, hear, impact, job, labor, moder, pressur, slow, trend, unemploy_rate, wage, worker

Figure 8: Bing et al. (2020b)'s estimator of A in the FOMC1 corpus. Each panel shows the word cloud of words of a topic (column in A matrix), where the font size is proportional to term's weight in the topic, and the top 5 terms with largest weights are colored in orange. The estimated anchor words for each topic are in the caption.



(a) Topic 1: asymmetr



(b) Topic 2: rang, target



(c) Topic 3: sentenc



(d) Topic 4: announc



(e) Topic 5: basi_point

Figure 9: Bing et al. (2020b)'s estimator of A in the FOMC2 corpus. Each panel shows the word cloud of words of a topic (column in A matrix), where the font size is proportional to term's weight in the topic, and the top 5 terms with largest weights are colored in orange. The estimated anchor words for each topic are in the caption.

expectations for future changes in the stance of monetary policy, in addition to instructions for current policy. From Thornton (2006), “the statement pertaining to possible future policy was known as the “symmetry,” “tilt,” or “bias,” of the policy directive. The directive was said to be symmetric if it indicated that a tightening or an easing of policy were equally likely in the future. Otherwise, the directive was said to be asymmetric toward either tightening or easing.”

Based on the discussion above, we can assign the label “asymmetric policy directive” to Topic 1, given that the anchor word for Topic 1 is “asymmetr” and the top five words associated with this topic are “asymmetr,” “move,” “policy,” “inflation,” and “data”. The estimated W for FOMC2 confirms this topic is important in the meetings between 1987 and 1999 (cf. Figure 10), which seems quite reasonable as the policy directive was explicitly communicated to the Open Market Trading Desk (and was plausibly an important part of the FOMC deliberations).

Topics 3 and 4 also seem to be related to the FOMC communication (and their corresponding anchor words are “sentenc” and “announc”), but their interpretation is less clear (beyond the fact that they clearly relate to the communication of the policy choice to the public). We would expect these topics to increase after the year 2000, when the statements became more detailed. We come back to this point in the subsequent subsection when we discuss the estimated W for FOMC2. It is not quite clear to us why Topics 3 and 4 are considered different by the model.

A similar point can be made about Topic 2 and Topic 5. Topic 2 includes both “target” and “rang” as anchor words (thus suggesting explicit targeting of the federal funds rate), while Topic 5 has the anchor word “basi point” (which again is suggestive of explicit discussions about the target federal funds rate).

In summary, we think that the interpretation of the FOMC2 topics is not very transparent, which informally suggests that the anchor-words assumption may not be appropriate for this corpus.

5.2.3 Estimation of W

We next report estimates of the matrix W , which contains the topic proportions in each document, again estimating W separately in the FOMC1 and FOMC2 corpus. Our estimates of W are based on the recent work of Bing et al. (2022), and correspond to the Maximum Likelihood estimator of W in the multinomial model (8) but treating \hat{A} as the true unknown A .

Figure 10 presents the estimated topic proportions using a stacked bar graph. Since each FOMC transcript is indexed by the day of its associated FOMC meeting, the x -axis in each graph is simply a date stamp. At each of these dates, the stacked bars give the proportion that each of the meetings assigned to each of the K latent topics (with the proportions adding to one by construction).

Estimated matrix W for FOMC1: Panel a) in Figure 10 presents the topic proportions corresponding to the FOMC1 documents. The evolution of the topic proportions over time, and the label of the topics, are consistent with historical events that shaped monetary policy decisions during the Greenspan period. For example, it is well-known that Greenspan faced at least five periods of economic turbulence during his tenure as chairman of the Federal Reserve: the October 1987 stock market crash, the Asian financial crisis of 1997, the 9/11 terrorist attacks, and two US recessions (one in the early 90’s and one in the early 2000’s, cf. Figure 6). The estimated matrix W shows that the “uncertainty” topic increases around these dates. The “recoveri” topic also seems to become larger after these events. Further, the share of the “foreign conditions” topic gets close to zero from 1992 to 1996, corresponding to the period between the Gulf War and the 1997 Asian Financial Crisis.

Estimated matrix W for FOMC2: Panel a) in Figure 10 presents the topic proportions corresponding to the FOMC2 documents. The evolution of the topic proportions over time seems to be more erratic than what we reported for FOMC1.

As we expected, Topic 1 (“asymmetric policy directive”) is very important before January 2000, but practically disappears after this date. This is consistent with the fact that the FOMC decided to stop communicating explicitly the likely direction or the timing of future policy moves to the public (and instead decided to include the “Committee’s assessment of the balance of risks between heightened inflation pressure and economic weakness over the foreseeable future;” see Thornton et al. (2000)). Relatedly, Topic 3 (which has “sentence” as its anchor word, and “statement” as its most likely term) has a very small share before January 2000, but it is the most important topic in the transcripts at the end of the sample. We found it difficult to provide a rationale for the shares of the other topics in FOMC2.




Figure 10: Bing et al. (2022)’s estimator of W for FOMC1 and FOMC2. The topic labels are based on the anchor words as explained in Section 5.2.2.

5.3 Testing the anchor-words Assumption

In the previous subsection we argued that the estimated parameters for FOMC1 admit a straightforward interpretation. The estimated anchor words provide a clear distinction between the topics, and the estimated topic proportions are consistent with historical events that shaped monetary policy decisions during the Greenspan period. We also noted that results for the FOMC2 corpus are markedly different: both the anchor words and the topics are difficult to interpret. With the exception of two topics, we found it difficult to provide a rationale for the historical evolution of the topic shares. Motivated by these results, in this section, we test the assumption of the existence of anchor words in both the FOMC1 and FOMC2 corpus. Our main finding is that the assumption of anchor words is rejected by a nominal 5%-level test in the FOMC2 corpus, but not in the FOMC1 corpus.

5.3.1 Test statistic

As we mentioned before, the computation of the test statistic $T(Y)$ involves the minimization of a quadratic objective function over the set C_K , which is a set of bounded, real-valued $V \times V$ matrices defined by 1 linear equality and $2V^2$ linear inequalities. We solve this optimization problem in MATLAB® (version 2022b) using the function `lsqlin`. The computation of the test statistic in our application takes only 137 seconds for FOMC1 and 58 seconds for FOMC2. The test statistics we obtain for the FOMC1 and FOMC2 corpus are

$$T(Y_{\text{FOMC1}}) = .4938, \quad T(Y_{\text{FOMC2}}) = .6401. \quad (31)$$

5.3.2 Critical values

As discussed in Section 4.2.2, obtaining the critical value used in the test of Theorem 2 in our application is extremely computationally demanding. For instance, one could try to create either a deterministic or random grid of parameters (A, W) in Θ_0 , and approximate $q_{1-\alpha}^*$ from below by the largest quantile for the random variable $T(Y)$ over the grid. In the FOMC1 corpus, this will require constructing a deterministic (or random grid) over matrices of dimension 200×4 and 4×148 that satisfy the anchor-words assumption. Due to the dimension of the parameter space, it seems unlikely that one could generate a good approximation of $q_{1-\alpha}^*$ using this approach. Below, we report the critical values based on the two computationally feasible approaches discussed in Section 4.2.2.

- *Algebraic Upper Bound for $q_{1-\alpha}^*$.* Lemma 4 in Appendix B.5 of the online supplementary material implies that, under the same assumptions as in Proposition 2:

$$q_{1-\alpha}^* \leq \sup_{C \in \mathcal{C}_K} \|C - \mathbb{I}_V\|_F \cdot R_\gamma(\alpha),$$

where

$$R_\gamma(\alpha) \equiv \sqrt{\frac{8 \left(1 - \frac{1}{V}\right)}{\gamma^2 \cdot \alpha} \cdot \frac{V^2}{N_{\min} \cdot D}},$$

and $\gamma \in (0, 1)$ is a constant such that for any $(A, W) \in \Theta$, $\sum_{d=1}^D (AW)_{vd} / D \geq \gamma / V$ for all v . The first term in the bound has a closed-form solution. Its value for FOMC1 is 31.50, and for FOMC2 is 29.83. To compute the second term that appears in the upper bound, we just need to choose a value of γ . The value of γ controls the magnitude of the row sums of the matrix AW uniformly in our parameter space. We pick the value of γ using the estimated values of A and W under the anchor-words assumption. More precisely, we set

$$\hat{\gamma} \equiv \frac{V}{2D} \cdot \min_{v \in V} \left\{ \sum_{d=1}^D (AW)_{vd} \right\},$$

which is guaranteed to be smaller than or equal to $1/2$.¹⁷

¹⁷Note that for any $v \in V$

$$\hat{\gamma} \leq \frac{V}{2D} \cdot \sum_{d=1}^D (AW)_{vd}.$$

Thus, adding both sides over $v \in V$ implies

$$V\hat{\gamma} \leq \frac{V}{2},$$

Using this formula, the bound for $q_{1-\alpha}^*$ in FOMC1 becomes $87.84/\sqrt{\alpha}$ and the bound in FOMC2 becomes $197.65/\sqrt{\alpha}$. This means that, using this conservative critical value, we fail to reject the null hypothesis of anchor words in both FOMC1 and FOMC2 for any significance level. This suggests that the algebraic upper bound is overly conservative.

- A “bootstrap bound” for $q_{1-\alpha}^*$. Finally, we compute the “bootstrap bound” for $q_{1-\alpha}^*$ discussed in Section 4.2.2. In our application, computing the critical value using 1,000 simulations takes 182 seconds for FOMC1 and 113 seconds for FOMC2. The 5%-critical values for FOMC1 and FOMC2 are 0.6310 and 0.6038 respectively. Comparing these critical values to our test statistics in (31), our test rejects the null hypothesis of the existence of anchor words for FOMC1, but fails to reject it for FOMC2.

5.4 Finite-sample properties of the test

We have shown that it is possible to use a “bootstrap bound” for the critical value of the test described in Theorem 2. We established the “consistency” of our bootstrap strategy; but, unfortunately, the consistency holds only “pointwise” at a fixed (A_0, W_0) in the null hypothesis. This means that the test based on the bootstrap upper bound need not have the correct size in finite samples. With this in mind, this section presents a small simulation study to analyze both the rate of Type I error and Type II error of our test. The simulation is based on the setup of FOMC2 data. This means that we set $V = 150$, $D = 148$, and we consider document sizes equal to each of the FOMC2.

- *Type I error.* We first analyze the rate of Type I error of the test that uses the test statistic described in Section 4.2.1 and the critical value based on the “bootstrap” upper bound described in Section 4.2.2. To guarantee that the true data-generating process has anchor words and is comparable to the Type II error discussed later, we do the following. We generate 1,000 arbitrary matrices, $\{P_i\}_{i=1}^{1,000}$, by sampling $D = 148$ independent columns from the Dirichlet distribution in \mathbb{R}^V and with concentration parameter $\alpha = 1/200$. We then generate multinomial counts according to P_i with a large number of trials, and use the data to construct estimates A_{0i} and W_{0i} (according to our discussion in Sections 5.2.2 and 5.2.3 based on Arora et al. (2013), Bing et al. (2020b) and Bing et al. (2022)). Specifically, we use the STM-TOP algorithm described in Bing et al. (2020b) with $K = 5$. In the remaining part of this section, we use A_{0i} , W_{0i} , and K_0 to denote the true model parameters used in the simulation.

Using $P_{0i} = A_{0i}W_{0i}$, we generate $i = 1, \dots, 1000$ new matrices of counts Y_i (of dimension $V \times D$) based on the multinomial model in (8), where each of these multinomial trials uses the true size of the documents in the application. For each of these new matrices Y_i , we compute our test statistic in Equation (22) (as we have explained before, computing this statistic takes around 58 seconds for each new dataset).

We then get, for each Y_i , the “bootstrap bound” suggested in Section 4.2.2. Denote this critical value by c_i . The average rate of Type I error using this critical value (the share of simulations for which $T(Y_i) > c_i$) is 3.7% for the nominal 5% test. Thus, the simulations suggest the critical value based on the “bootstrap bound” is conservative at certain parameter space under the setup of the FOMC2.

- *Type II error/Power.* Our claim in Theorem 2 is that the test suggested therein will have non trivial power against at least one alternative. As we have discussed before, the critical value for this test is not computationally feasible, so we analyze the power of the test that uses a critical value based on the bootstrap upper bound discussed in Section 4.2.2.

which implies $\hat{\gamma} \leq 1/2$.

We extract nonnegative matrix factorizations of $\{P_i\}_{i=1}^{1,000}$ using the standard nonnegative matrix factorization routine in Matlab (which uses the KL-divergence as objective function, see the documentation of MATLAB®’s function `nnmf`). We use the nonnegative factors as the true data generating process (after normalizing the matrices to be column stochastic) and we denote them as A_{1i} and W_{1i} . Letting $P_{1i} \equiv A_{1i}W_{1i}$, we compute the value of $\inf_{C \in \mathcal{C}_K} \|CP_{1i}^{\text{row}} - P_{1i}^{\text{row}}\|_F$ (to confirm that P_{1i} does not have an anchor-word factorization). The average value of this statistic is 0.0885, and the 5% lower quantile is 0.0064. The average value of $\inf_{C \in \mathcal{C}_K} \|CP_{1i}^{\text{row}} - P_{1i}^{\text{row}}\|_F$ for concentration parameters $\alpha = 1$ and $\alpha = 0.1$ are 0.0410 and 0.0585, respectively. These values also suggest that using a concentration parameter equal to $\alpha = 1/200$ will lead to a larger average power than $\alpha = 1$ and $\alpha = .1$. We now take A_{1i} and W_{1i} as the true data-generating process. The average power of the test (the share of simulation draws for which $T(Y_i) > c_i$) that uses the critical value based on the “bootstrap bound” is close to 71.2% for the 5% nominal test.

6 Conclusion

In this paper we show that the existence of anchor words in topic models where $2 < K < \min\{V, K\}$ is statistically testable: There exists a test for the null hypothesis that anchor words exist, that has correct size and nontrivial power. This means that imposing the anchor-words assumption to identify the parameters of a topic model cannot be viewed simply as a convenient normalization. A key result to establish the statistical testability of the anchor-words assumption is Theorem 1. This theorem shows that a column-stochastic matrix (with known nonnegative rank K) admits a *separable* factorization if and only if the linear program suggested by Recht et al. (2012) to find a nonnegative matrix factorization of separable matrices has a nonempty choice set.

We establish the statistical testability of the anchor-words assumption by constructing an explicit test that has correct size in finite samples. Our Theorem 2 shows that our suggested test has nontrivial power, provided a certain high-level condition is verified. We also show that our high-level condition can be verified in settings where the size of the available documents is large enough. In fact, Corollary 1 in Appendix A.2 of the online supplementary material provides primitive conditions under which our test is consistent (its power approaches one) at any (A, W) for which the corresponding matrix $P = AW$ does not have an anchor-word factorization.

An unsatisfactory aspect about our constructive results is that the critical value we suggest for the test in Theorem 2 is computationally infeasible in any realistic application. The computational difficulties we face are in part due to the fact that testing whether there exists a nonnegative solution to a large-scale system of linear equations—whose coefficients and ordinates may depend on the unknown data distribution—is a difficult statistical problem. It is known that guaranteeing size control while remaining computationally feasible is challenging; see Kitamura & Stoye (2018), Fang, Santos, Shaikh & Torgovitsky (2023) and Bai, Santos & Shaikh (2022). In fact, Fang et al. (2023) have recently devised a procedure for testing the abstract hypothesis that the unknown distribution of an i.i.d sample satisfies a linear system of equations of the form $Ax = \beta$, where x is a nonnegative (high-dimensional) vector and β depends on the distribution of the data. Unfortunately, their results do not seem to be directly applicable to our problem as the characterization provided in Theorem 1 involves the linear equation $CP^{\text{row}} = P^{\text{row}}$ (which implies that both sides of the linear system of interest depend on the true distribution of the data). An interesting question for future research is whether some extension of their recommended testing procedure can be used to construct a test for the existence of anchor words. Another question of interest is whether the “bootstrap bound” for the critical value we suggest in Section 4.2.2 of this

paper could be used for the problems considered in Fang et al. (2023).

In order to show the applicability of our results, we test for the existence of anchor words in two different datasets derived from the transcripts of the meetings of the Federal Open Market Committee (FOMC). One corpus discusses domestic and international economic conditions, and one corpus discusses possible monetary policy strategies. In the latter, we reject the null hypothesis that anchor words exist. For this case, it would be an interesting exercise for future work to estimate a topic model replacing the anchor-words assumption by some weaker condition that yields point identification, and leads to a computationally tractable statistical procedure; for example, some version of the *sufficiently-scattered* assumption discussed in Huang et al. (2013), Huang et al. (2016), and more recently in Chen et al. (2022)

Finally, it is worth mentioning that the scope of the theoretical results established in this paper may extend beyond applications of topic models to textual data. We discuss two concrete examples below.

First, in a very interesting recent paper, Moran, Sridhar, Wang & Blei (2021) have shown that the popular “deep generative models” (which are used for conducting unsupervised representation learning in high-dimensional data) can be identified by assuming the existence of “anchor features.” We think it would be interesting to study whether such an assumption (which is analogous to the anchor-words assumption in topic models) has testable implications (and is, therefore, incompatible with certain distributions of the data).

Second, Wang, Zhu & Wang (2023) have recently argued that the “in-context learning” capabilities of large language models (LLMs) can be explained by viewing LLMs as topic models that implicitly infer “task-related” information from a small number of examples. It would be interesting to think about whether the reported sensitivity of the in-context learning capabilities of LLMs to “choice, format, and even the order of the demonstrations used” reported in Wang et al. (2023) can be linked to the identification (or the lack thereof) of topic models.

References

- Arora, S., Ge, R., Halpern, Y., Mimno, D., Moitra, A., Sontag, D., Wu, Y. & Zhu, M. (2013), A practical algorithm for topic modeling with provable guarantees, *in* ‘International conference on machine learning’, PMLR, pp. 280–288.
- Arora, S., Ge, R., Kannan, R. & Moitra, A. (2012), Computing a nonnegative matrix factorization—provably, *in* ‘Proceedings of the forty-fourth annual ACM Symposium on Theory of Computing’, pp. 145–162.
- Arora, S., Ge, R. & Moitra, A. (2012), Learning topic models—going beyond SVD, *in* ‘2012 IEEE 53rd Annual Symposium on Foundations of Computer Science’, IEEE, pp. 1–10.
- Ash, E. & Hansen, S. (2023), ‘Text algorithms in economics’, *Annual Review of Economics* **15**, 659–688.
- Bai, Y., Santos, A. & Shaikh, A. M. (2022), ‘On testing systems of linear inequalities with known coefficients’, *Working Paper* .
- Bing, X., Bunea, F., Strimas-Mackey, S. & Wegkamp, M. (2022), ‘Likelihood estimation of sparse topic distributions in topic models and its applications to wasserstein document distance calculations’, *The Annals of Statistics* **50**(6), 3307–3333.

- Bing, X., Bunea, F. & Wegkamp, M. (2020a), ‘A fast algorithm with minimax optimal guarantees for topic models with an unknown number of topics’, *Bernoulli* **26**(3), 1765–1796.
- Bing, X., Bunea, F. & Wegkamp, M. (2020b), ‘Optimal estimation of sparse topic models’, *The Journal of Machine Learning Research* **21**(1), 7189–7233.
- Blei, D. M. (2012), ‘Probabilistic topic models’, *Communications of the ACM* **55**(4), 77–84.
- Blei, D. M. & Lafferty, J. D. (2009), ‘Topic models’, *Text mining: classification, clustering, and applications* **10**(71), 71–89.
- Blei, D. M., Ng, A. Y. & Jordan, M. I. (2003), ‘Latent dirichlet allocation’, *Journal of Machine Learning Research* **3**, 993–1022.
- Boyd-Graber, J., Hu, Y., Mimno, D. et al. (2017), ‘Applications of topic models’, *Foundations and Trends® in Information Retrieval* **11**(2-3), 143–296.
- Canay, I. A., Santos, A. & Shaikh, A. M. (2013), ‘On the testability of identification in some nonparametric models with endogeneity’, *Econometrica* **81**(6), 2535–2559.
- Chappell Jr, H. W., McGregor, R. R. & Vermilyea, T. (2004), *Committee decisions on monetary policy: Evidence from historical records of the Federal Open Market Committee*, MIT Press.
- Chen, Y., He, S., Yang, Y. & Liang, F. (2022), ‘Learning topic models: Identifiability and finite-sample analysis’, *Journal of the American Statistical Association* pp. 1–16.
- Ding, W., Ishwar, P. & Saligrama, V. (2015), Most large topic models are approximately separable, in ‘2015 Information Theory and Applications Workshop (ITA)’, IEEE, pp. 199–203.
- Doebelin, W. & Cohn, H. (1993), *Doebelin and Modern Probability*, Vol. 149, American Mathematical Soc.
- Donoho, D. & Stodden, V. (2003), ‘When does non-negative matrix factorization give a correct decomposition into parts?’, *Advances in neural information processing systems* **16**.
- Dudley, R. (2002), *Real Analysis and Probability*, Vol. 74, Cambridge University Press.
- Fang, Z., Santos, A., Shaikh, A. M. & Torgovitsky, A. (2023), ‘Inference for large-scale linear systems with known coefficients’, *Econometrica* **91**(1), 299–327.
- Fligstein, N., Brundage, J. S. & Schultz, M. (2017), ‘Seeing like the Fed: Culture, cognition, and framing in the failure to anticipate the financial crisis of 2008’, *American Sociological Review* **82**(5), 879–909.
- Fu, X., Huang, K., Sidiropoulos, N. D. & Ma, W.-K. (2019), ‘Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications’, *IEEE Signal Process. Mag.* **36**(2), 59–80.
- Gillis, N. (2013), ‘Robustness analysis of hottopixx, a linear programming model for factoring nonnegative matrices’, *SIAM Journal on Matrix Analysis and Applications* **34**(3), 1189–1212.

- Hansen, S., McMahon, M. & Prat, A. (2018), ‘Transparency and deliberation within the FOMC: A computational linguistics approach’, *The Quarterly Journal of Economics* **133**(2), 801–870.
URL: <https://doi.org/10.1093/qje/qjx045>
- Hofmann, T. (1999), Probabilistic latent semantic indexing, in ‘Proceedings of the 22nd annual international ACM SIGIR Conference on Research and Development in Information Retrieval’, pp. 50–57.
- Horn, R. A. & Johnson, C. R. (2012), *Matrix analysis*, Cambridge university press.
- Huang, K., Fu, X. & Sidiropoulos, N. D. (2016), ‘Anchor-free correlated topic modeling: Identifiability and algorithm’, *Advances in Neural Information Processing Systems* **29**.
- Huang, K., Sidiropoulos, N. D. & Swami, A. (2013), ‘Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition’, *IEEE Transactions on Signal Processing* **62**(1), 211–224.
- Ke, S., Montiel Olea, J. L. M. & Nesbit, J. (2022), ‘A robust machine learning algorithm for text analysis’, *Working paper* .
- Ke, Z. T. & Wang, M. (2022), ‘Using SVD for topic modeling’, *Journal of the American Statistical Association* pp. 1–16.
- Kitamura, Y. & Stoye, J. (2018), ‘Nonparametric analysis of random utility models’, *Econometrica* **86**(6), 1883–1909.
- Kosorok, M. R. (2007), *Introduction to empirical processes and semiparametric inference*, Springer Science & Business Media.
- Laurberg, H., Christensen, M. G., Plumbley, M. D., Hansen, L. K., Jensen, S. H. et al. (2008), ‘Theorems on positive data: On the uniqueness of NMF’, *Computational intelligence and neuroscience* **2008**.
- Levin, D. A. & Peres, Y. (2017), *Markov chains and mixing times*, Vol. 107, American Mathematical Soc.
- McRae, A. D. & Davenport, M. A. (2021), ‘Low-rank matrix completion and denoising under poisson noise’, *Information and Inference: A Journal of the IMA* **10**(2), 697–720.
- Meade, E. E. & Stasavage, D. (2008), ‘Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve’, *The Economic Journal* **118**(528), 695–717.
- Meade, E. E. & Thornton, D. L. (2012), ‘The Phillips curve and US monetary policy: What the FOMC transcripts tell us’, *Oxford Economic Papers* **64**(2), 197–216.
- Moran, G. E., Sridhar, D., Wang, Y. & Blei, D. M. (2021), ‘Identifiable deep generative models via sparse decoding’, *arXiv preprint arXiv:2110.10804* .
- Munkres, J. R. (2000), *Topology: International edition*, Pearson Prentice Hall.
- Recht, B., Re, C., Tropp, J. & Bittorf, V. (2012), ‘Factoring nonnegative matrices with linear programs’, *Advances in neural information processing systems* **25**.

Thomas, L. B. (1974), 'Problem 73-14', *SIAM Review* **16**(3), 393–394.

URL: <http://www.jstor.org/stable/2029161>

Thornton, D. L. (2006), 'When did the FOMC begin targeting the federal funds rate? What the verbatim transcripts tell us', *Journal of Money, Credit and Banking* **38**(8), 2039–2071.

Thornton, D. L., Wheelock, D. C. et al. (2000), *History of the Asymmetric Policy Directive*, Inter-university Consortium for Political and Social Research.

Vavasis, S. A. (2010), 'On the complexity of nonnegative matrix factorization', *SIAM Journal on Optimization* **20**(3), 1364–1377.

Wang, X., Zhu, W. & Wang, W. Y. (2023), 'Large language models are implicitly topic models: Explaining and finding good demonstrations for in-context learning', *arXiv preprint arXiv:2301.11916* .