Transparency as an Index for Monetary Policy Shocks

Research Proposal

Katie Rha

November 14, 2017

Economics Department: ECON 691H

University of North Carolina at Chapel Hill

Thesis Advisor: Professor Neville Francis

Faculty Advisor: Professor Lutz Hendricks

1. INTRODUCTION AND MOTIVATION

The primary goal of this paper is to create a quantitative and procedural method for measuring transparency. I aim to quantify the effectiveness of this current method of policy communication by analyzing the reactions, and thus relationship, between publically available FOMC meeting materials and federal funds futures rates and euro-dollar fluctuations around FOMC meeting dates. I will be analyzing the language used in these documents and news articles through natural language processing techniques, primarily by way of text scraping/mining for linguistic variations of key variable terms, such as "inflation" and "unemployment". Not only is this in an effort to understand the history and chronology of circulating, policy/economic-related conversations between the Fed and the public, but is also to understand to what magnitude the release of information can contribute to reactionary effects in the aggregate economy. Studying fed funds futures rates and euro-dollar fluctuations will help in isolating and identifying such monetary policy shocks.

Given the importance of the topics contained in these documents, it is not surprising that, over the years, both the public and the FOMC itself have favored greater transparency. The transparency considered here is concerned with document publication, which, in practice, can take two forms: the first is a matter of timing, the second is one of translation. uestions of translation consider the process of distilling detailed information into a concise summary; for example, which words and topics from FOMC meetings should be included in the meeting minutes? Despite the importance of these documents, topics of timing and, to an even lesser extent, translation have received little empirical attention. This is the first quantified measure of transparency of which is constructed from documents to which both the public, and the FOMC itself, pay a great deal of attention: the meeting minutes and transcripts—the main vehicles through which the public can ascertain the content of FOMC meetings and the nature of the deliberations that underlie monetary policy decisions as well as news sources.

2. RELATED LITERATURE

The view that more transparency may lead to more conformity and hence less information revelation is formalized in the career concerns literature. Greater disclosure can induce experts who are concerned with their professional reputation to pool on actions that are optimal given available public signals even when their private signals would suggest that other actions are optimal (Prat 2005). In such circumstances, the principal benefits from committing to a policy of limited transparency. Of course, it is possible that both effects—discipline and conformity—operate simultaneously, in which case one should ask whether on balance more disclosure improves or worsens information aggregation. We are able to explore these issues by exploiting the natural experiment that led to the release of the FOMC transcripts. Since the 1970s, FOMC meetings were tape recorded to help prepare minutes. Unknown to committee members, though, these tapes were transcribed and stored in archives before being recorded over. They only learned this when Greenspan, under pressure from the US Senate Committee on Banking, Housing, and Urban Affairs (Senate Banking Committee hereafter), discovered and revealed their existence to the politicians and the rest of the FOMC. To avoid accusations of hiding information, and to relieve potential pressure to release information in a more timely fashion, the Fed quickly agreed to publish the past transcripts and all future transcripts with a five-year lag. We thus have a complete record of deliberation both when policymakers did not know that their verbatim discussions were being kept on file let alone that such information would be made public (prior to November 1993), and when they knew with certainty that their discussions would eventually be made public. Meade and Stasavage (2008) have previously used this natural experiment to analyze the effect of transparency on members' incentives to dissent in voice. This dissent data, recorded in Meade (2005), is a binary measure based on whether a policymaker voiced disagreement with Chairman Greenspan's policy proposal during the policy debate. Their main finding, which they interpret as conformity, is that the probability that members dissent declines significantly after transparency. We instead generate communication measures based on basic text counts and on topic models, a class of machine learning algorithms for natural language processing that estimates what fraction of time each speaker in each section of each meeting spends on a variety of topics. This approach allows one to construct several measures of communication relating to both discipline and conformity, and also to compare which effect is stronger. The wealth of data also allows us to extend Meade and Stasavage (2008) in another direction.

3. CONSTRUCTION OF TRANSPARENCY INDEX

The construction of a transparency index involves use of natural language processing techniques, particularly text scraping and text mining, on two types of materials. The first includes publicly available FOMC materials, primarily transcripts and minutes. The second involves news sources, which include articles in Bloomberg, Wall Street Journal, the NY Times, and the Chicago Tribune containing information pertinent to Federal Reserve actions. By text scraping these two sources, we should be able to understand the translation of Fed policy actions as described in the FOMC statements into readership by the public at-large.

3.1. Application of LDA to Central Bank Policy

The primary algorithm we use to conduct topic-modelling of the FOMC transcripts and minutes is Latent Dirichlet Allocation (LDA) introduced by Blei, Ng, and Jordan (2003). LDA is widely used in linguistics, computer science, and other fields and has been cited over 8,000 times in ten years. While topic modelling approaches are beginning to appear in the social science literature, there use so far is mainly descriptive. For example, Quinn, Monroe, Colaresi, Crespin, and Radev (2010) apply a topic model similar to LDA to congressional speeches to identify which members of Congress speak about which topics. An innovation of our paper is to use communication measures constructed from LDA output as dependent variables in an econometric model explicitly motivated by economic theory (more specifically, career concerns). We believe this illustrates the potential fruitfulness of combining traditional economic tools with those from the increasingly important world of "Big Data" for empirical research in economics more broadly. Fligstein, Brundage, and Schultz (2014)—developed independently6 from this paper—also apply LDA to FOMC transcripts focusing on the period 2000-2007. They describe the topics that the meeting as a whole covers rather than the topics of individuals, and verbally argue they are consistent with the sociological theory of "sense making". They claim that the standard models that macroeconomists use led them to fail to connect topics related to housing, financial markets and the macroeconomy. In contrast, this paper uses LDA applied to all data from the Greenspan era (1987-2006) to construct numerous measures of communication patterns at the meeting-section-speaker level and embeds them within a difference-in-differences regression framework to identify how transparency changes individual incentives. Bailey and Schonhardt-Bailey (2008) and Schonhardt-Bailey (2013) also use text analysis to examine the FOMC transcripts. They emphasize the arguments and persuasive strategies adopted by policymakers (measured using a computer package called "Alceste") during three periods of interest (1979-1981, 1991-1993, and 1997-1999). Of course, many others have analyzed the transcripts without using computer algorithms; for example, Romer and Romer (2004) use the transcripts to derive a narrative-based measure of monetary policy shocks.

Our current method of LDA as applied to transcripts include a topic clustering of the top ten words most frequently associated together in four-year intervals. For example, the most frequently cited topic from 1976-1979 includes "rate", "burns", "time", "range", "miller", "market", "growth", "m1", "funds", and "point". The most frequently cited topic from 2006-2009 includes "inflation", "rate", "growth", "policy", "year", "prices", "markets", "forecast", and "time". An in-depth analysis and summary statistic table will be provided in the next draft of this proposal.

3.2. News-Source Construction and Sensitivity

Similar to the process outlined by Husted et al. (2016), within news sources we search for articles containing variations on key variable terms, such as (i) "inflate" or "unemploy," (ii) "Federal fund(s) rate" or "Fed fund(s) rate," and (iii) "Federal Reserve" or "the Fed" or "Federal Open Market Committee" or "FOMC". We do this for all articles surrounding the FOMC meeting dates. It is worth noting that this daily index is aggregated into inter-meeting intervals (as well as monthly levels), while Baker et al. (2015) constructs a monthly measure. Since FOMC decisions are made on pre-specified meeting dates, our index allows us to incorporate information arrival following each FOMC meeting and capture any effects that FOMC (in)actions have on the transparency between the Fed and the public.

I consider the accuracy and sensitivity of the baseline transparency index by adding several adjustments into its construction. With the news-based approach specifically, our search is narrowed to only search for articles which have both key variable terms, such as "inflate" and "unemploy" and the phrases: "Federal Reserve," "The Fed," or "monetary policy" existing in the same space. More deeply, I restrict this universe, by limiting our search to word zones of a 20-word maximum. Husted et al. (2016) takes a similar approach with proximity refinement, and concludes that this restriction has a smaller type II error, as it filters our more of both "false" articles and "correct" articles. In order to better understand the trade-offs associated with using the proximity refinement, we plan to extract and read several articles. However, while this proximity search does filter out articles that mention all the keywords, it is difficult to analyze these articles, as the explicit term "transparency" is not indicative of describing Fed policy action by these news sources. There are certain lexicons that may describe Fed policy actions, with words that contain clear-cut terms, such as action verbs and explicit adjectives. A direct approach such as this is foreseeable to be problematic and extremely subjective. Thus, it is more important to look upon the news-based approach as a robustness check after the transparency index has been applied to monetary policy data. By doing so, we will be able to identify only articles that pertain to reactionary effects of the FOMC's released policy decisions and assess the level of readership and understanding afterwards.

3 EMPIRICAL PLAN: RESPONSE OF INDEX TO MONETARY POLICY SHOCKS

As a test of our index, we examine whether the FOMC was more or less likely to surprise the markets with their interest rate decisions. To do this, we plan to use the surprise data as developed in both Gurkaynak, Sack, and Swanson (2005), which measures the effect of both short and longer term movements in the yield curve, and Kuttner (2001) (which was subsequently used in Bernanke and Kuttner (2005)) which captures the shorter end movements 41 only. These data are derived from traded futures securities which allow one to decompose movements in the Fed Funds target rate into expected and unexpected moves.

A visual of our current summary statistics of Fed Funds Futures Rate and Euro-Dollar fluctuations around FOMC meeting dates from 1990-2012 and 1985-2012, respectively, we plan to use are presented in Appendix A. It is worth noting that this high frequency data of futures and euro-dollar rates, is scaled to account for the day of the month in which the meeting took place.

WORKS CITED

Kuttner, Kenneth N., “Monetary policy surprises and interest rates: Evidence from the Fed funds futures market,” Journal of Monetary Economics Volume 47, Issue 3, June 2001, Pages 523–544.

Gürkaynak, Refet S., Brian Sack and Eric Swanson, “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” American Economic Review, Vol. 95, No. 1 (Mar., 2005), pp. 425-436

Campbell, Jeffrey R., Charles L. Evans, Jonas DM Fisher, and Alejandro Justiniano. "Macroeconomic effects of federal reserve forward guidance." Brookings Papers on Economic Activity (2012): 1-80.

Cochrane, John and Monika Piazzesi, 2002. "The Fed and Interest Rates - A High-Frequency Identification." American Economic Review, 92(2): 90-95.

Faust, Jon, Eric T. Swanson, and Jonathan H. Wright (2004), “Identifying VARS based on high frequency futures data,” Journal of Monetary Economics, Volume 51, Issue 6, September 2004, Pages 1107–113.

Barakchian, S. Mahdi and Christopher Crowe, “Monetary Policy Matters: Evidence from New Shocks,” Journal of Monetary Economics, Volume 60, Issue 8, November 2013, Pages 950–966.

Gertler, Mark and Peter Karadi, “Monetary Policy Surprises, Credit Costs, and Economic Activity,” American Economic Journal: Macroeconomics, 7(1) (January 2015) 44–76.

Nakamura, Emi, and Steinsson, Jón, “High Frequency Identification of Monetary Non-Neutrality,” October 2015 working paper.

Acosta, Miguel (2015). “FOMC Responses to Calls for Transparency,” Finance and Economics Discussion Series 2015-060. Washington: Board of Governors of the Federal Reserve System, http://dx.doi.org/10.17016/FEDS.2015.060.

Geraats, Petra (2014). “Monetary Policy Transparency,” CESifo Working Paper No. 4611. Category 7: Monetary Policy and International Finance (January 2014).

Husted, Roger, Sun (2017). “Monetary Policy Uncertainty,” Federal Reserve Working Paper.

Geraats, Petra, “Central Bank Transparency,” The Economic Journal, 112 (November 2002), F532–F565.

Faust, Swanson, Wright, “Identifying VARS Based on High Frequency Futures Data” Journal of Monetary Economics 51 (2004) 1107–1131.

Barakchian, Crowe, “Monetary policy matters: Evidence from new shocks data” Journal of Monetary Economics (2013) 950-966.

Bernanke, Ben S. and Frederic S. Mishkin. "Inflation Targeting: A New Framework For Monetary Policy?," Journal of Economic Perspectives, 1997, v11(2,Spring), 97-116.

Hamilton, James D., “Daily Changes in Fed Funds Rates” Journal of Money, Credit and Banking, Vol. 41, No. 4 (June 2009)

Hamilton, James D., “Daily Monetary Policy Shocks and New Home Sales,” Journal of Monetary Economics 55 (2008) 1171–1190

Thornton, D. L. (2003), Monetary policy transparency: transparent about what?. The Manchester School, 71: 478–497. doi:10.1111/1467-9957.00363

Faust, Svensson, “Transparency and Credibilit: Monetary Policy with Unobservable Goals,” International Economic Review, Vol. 42, no. 2 (May 2001): 369-397

Appendix A: Cited Facts and Figures

Figure 1. Monetary Policy Shocks

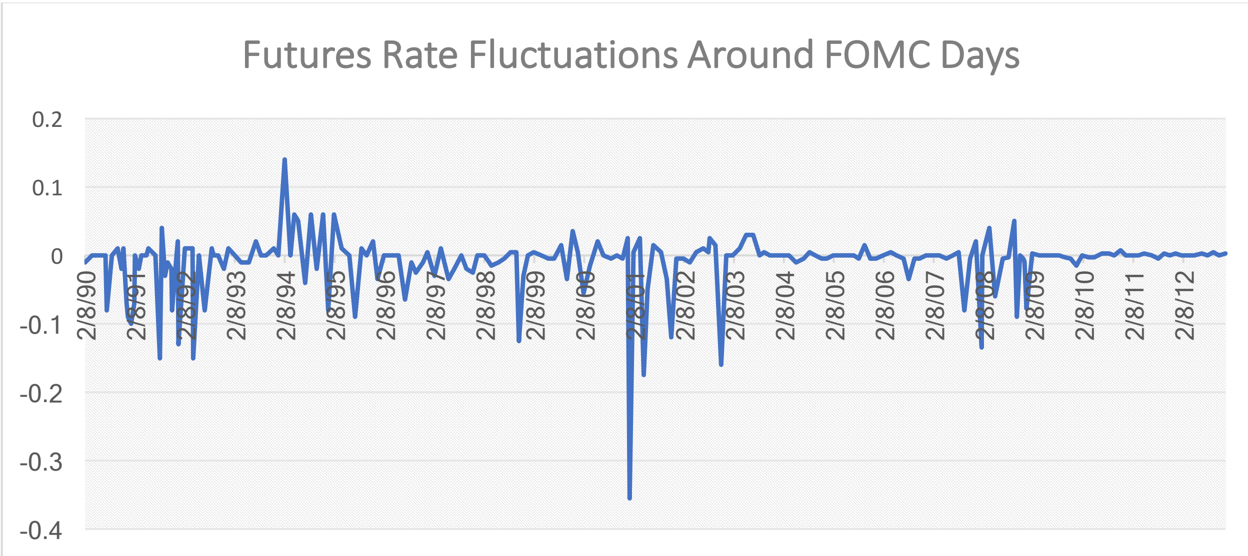

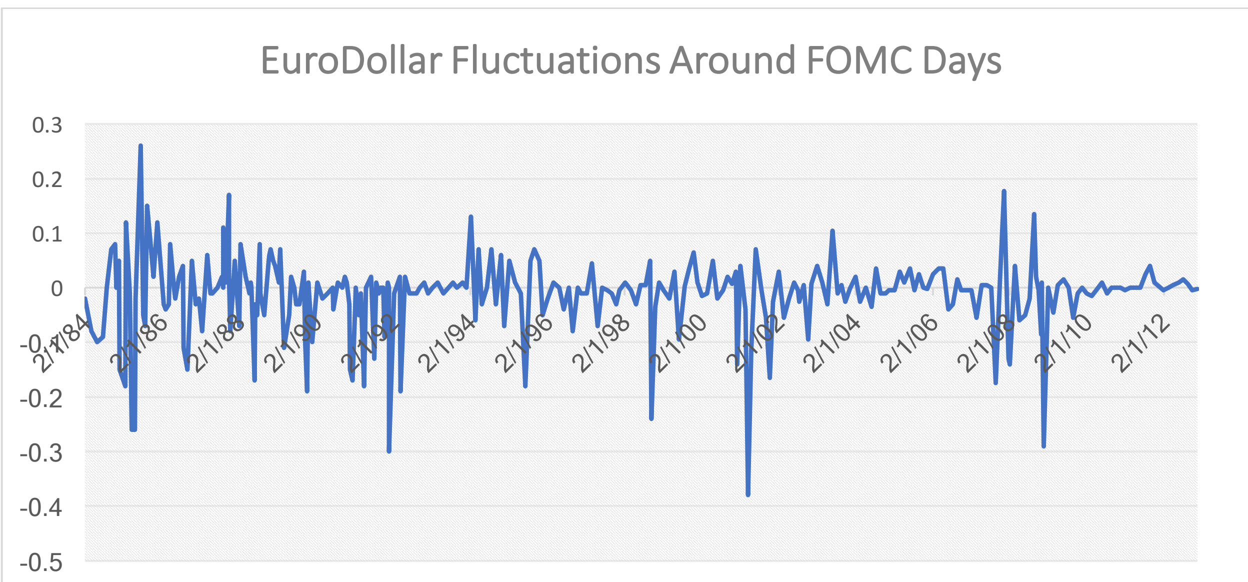

Data provided by Michael T. Kiley, Board of Governors Division of Financial Stability.

Table 1. Stop Word List

a	aside	clearly	etc	happens	indicates	may	off	really	sometimes	they'll	v	whom
able	ask	co	even	hardly	inner	maybe	often	reasonably	somewhat	they're	value	whose
about	asking	com	ever	has	insofar	me	oh	regarding	somewhere	they've	various	why
above	associated	come	every	hasn't	instead	mean	ok	regardless	soon	think	very	will
according	at	comes	everybody	have	into	meanwhile	okay	regards	sorry	third	via	willing
accordingly	available	concerning	everyone	haven't	inward	merely	old	relatively	specified	this	viz	wish
across	away	consequentl	everything	having	is	might	on	respectively	specify	thorough	vs	with
actually	awfully	consider	everywhere	he	isn't	more	once	right	specifying	thoroughly	w	within
after	b	considering	ex	he's	it	moreover	one	s	still	those	want	without
afterwards	be	contain	exactly	hello	it'd	most	ones	said	sub	though	wants	won't
again	became	containing	example	help	it'll	mostly	only	same	such	three	was	wonder
against	because	contains	except	hence	it's	much	onto	saw	sup	through	wasn't	would
ain't	become	correspondi	f	her	its	must	or	say	sure	throughout	way	would
all	becomes	could	far	here	itself	my	other	saying	t	thru	we	wouldn't
allow	becoming	couldn't	few	here's	j	myself	others	says	t's	thus	we'd	x
allows	been	course	fifth	hereafter	just	n	otherwise	second	take	to	we'll	y
almost	before	currently	first	hereby	k	name	ought	secondly	taken	together	we're	yes
alone	beforehand	d	five	herein	keep	namely	our	see	tell	too	we've	yet
along	behind	definitely	followed	hereupon	keeps	nd	ours	seeing	tends	took	welcome	you
already	being	described	following	hers	kept	near	ourselves	seem	th	toward	well	you'd
also	believe	despite	follows	herself	know	nearly	out	seemed	than	towards	went	you'll
although	below	did	for	hi	knows	necessary	outside	seeming	thank	tried	were	you're
always	beside	didn't	former	him	known	need	over	seems	thanks	tries	weren't	you've
am	besides	different	formerly	himself	l	needs	overall	seen	thanx	truly	what	your
among	best	do	forth	his	last	neither	own	self	that	try	what's	yours
amongst	better	does	four	hither	lately	never	p	selves	that's	trying	whatever	yourself
an	between	doesn't	from	hopefully	later	nevertheles:	particular	sensible	thats	twice	when	yourselves
and	beyond	doing	further	how	latter	new	particularly	sent	the	two	whence	z
another	both	don't	furthermore	howbeit	latterly	next	per	serious	their	u	whenever	zero
any	brief	done	g	however	least	nine	perhaps	seriously	theirs	un	where	
anybody	but	down	get	i	less	no	placed	seven	them	under	where's	
anyhow	by	downwards	gets	i'd	lest	nobody	please	several	themselves	unfortunate	whereafter	
anyone	c	during	getting	i'll	let	non	plus	shall	then	unless	whereas	
anything	c'mon	e	given	i'm	let's	none	possible	she	thence	unlikely	whereby	
anyway	c's	each	gives	i've	like	noone	presumably	should	there	until	wherein	
anyways	came	edu	go	ie	liked	nor	probably	shouldn't	there's	unto	whereupon	
anywhere	can	eg	goes	if	likely	normally	provides	since	thereafter	up	wherever	
apart	can't	eight	going	ignored	little	not	q	six	thereby	upon	whether	
appear	cannot	either	gone	immediate	look	nothing	que	so	therefore	us	which	
appreciate	cant	else	got	in	looking	novel	quite	some	therein	use	while	
appropriate	cause	elsewhere	gotten	inasmuch	looks	now	qv	somebody	theres	used	whither	
are	causes	enough	greetings	inc	ltd	nowhere	r	somehow	thereupon	useful	who	
aren't	certain	entirely	h	indeed	m	o	rather	someone	these	uses	who's	
around	certainly	especially	had	indicate	mainly	obviously	rd	something	they	using	whoever	
as	changes	et	hadn't	indicated	many	of	re	sometime	they'd	usually	whole	

Total: 569

Originated from Cornell's SMART stopword list

Figure 2. Latent Dirichlet Allocation Explained

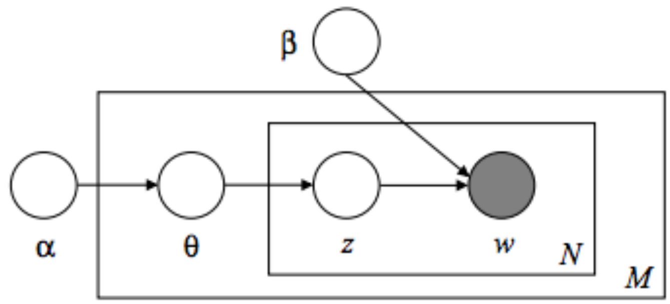

Graphical model representation of LDA. The boxes are "plates" representing replicates. The outer plate represents documents, while the inner plate represents the repeated choice of topics and words within a document.

Latent Dirichlet allocation (LDA) is a generative probabilistic model of a corpus. The basic idea is that documents are represented as random mixtures over latent topics, where each topic is characterized by a distribution over words. Several simplifying assumptions are made in this basic model, some of which we remove in subsequent sections. First, the dimensionality k of the Dirichlet distribution (and thus the dimensionality of the topic variable z) is assumed known and fixed. Second, the word probabilities are parameterized by a $k \times V$ matrix β where $\beta_{i j} = p(w_j = 1|z_i = 1)$, which for now we treat as a fixed quantity that is to be estimated. Finally, the Poisson assumption is not critical to anything that follows and more realistic document length distributions can be used as needed. Furthermore, note that N is independent of all the other data generating variables (θ and z). It is thus an ancillary variable and we will generally ignore its randomness in the subsequent development. A k-dimensional Dirichlet random variable θ can take values in the $(k-1)$-simplex (a k-vector θ lies in the $(k-1)$-simplex if $\theta_i \geq 0$, $\sum_{i=1}^{k} \theta_i = 1$), and has the following probability density on this simplex:

$$p(\theta|\alpha) = \frac{\Gamma\left(\sum_{i=1}^{k} \alpha_i\right)}{\prod_{i=1}^{k} \Gamma(\alpha_i)} \theta_1^{\alpha_1 - 1} \cdots \theta_k^{\alpha_k - 1},$$

where the parameter α is a k-vector with components $\alpha_i > 0$, and where $\Gamma(x)$ is the Gamma function. The Dirichlet is a convenient distribution on the simplex — it is in the exponential family, has finite dimensional sufficient statistics, and is conjugate to the multinomial distribution. These properties will facilitate the development of inference and parameter estimation algorithms for LDA. Given the parameters α and β, the joint distribution of a topic mixture θ, a set of N topics z, and a set of N words w is given by:

$$p(\theta, \mathbf{z}, \mathbf{w} | \alpha, \beta) = p(\theta | \alpha) \prod_{n=1}^{N} p(z_n | \theta) p(w_n | z_n, \beta),$$

where p(zn |θ) is simply θi for the unique i such that zi n = 1. Integrating over θ and summing over z, we obtain the marginal distribution of a document:

$$p(\mathbf{w} \mid \alpha, \beta) = \int p(\theta \mid \alpha) \left(\prod_{n=1}^{N} \sum_{z_n} p(z_n \mid \theta) p(w_n \mid z_n, \beta) \right) d\theta.$$

Finally, taking the product of the marginal probabilities of single documents, we obtain the probability of a corpus:

$$p(D \mid \alpha, \beta) = \prod_{d=1}^{M} \int p(\theta_d \mid \alpha) \left(\prod_{n=1}^{N_d} \sum_{z_{dn}} p(z_{dn} \mid \theta_d) p(w_{dn} \mid z_{dn}, \beta) \right) d\theta_d.$$

The LDA model is represented as a probabilistic graphical model in Figure 1. As the figure makes clear, there are three levels to the LDA representation. The parameters α and β are corpuslevel parameters, assumed to be sampled once in the process of generating a corpus. The variables θd are document-level variables, sampled once per document. Finally, the variables zdn and wdn are word-level variables and are sampled once for each word in each document. It is important to distinguish LDA from a simple Dirichlet-multinomial clustering model. A classical clustering model would involve a two-level model in which a Dirichlet is sampled once for a corpus, a multinomial clustering variable is selected once for each document in the corpus, and a set of words are selected for the document conditional on the cluster variable. As with many clustering models, such a model restricts a document to being associated with a single topic. LDA, on the other hand, involves three levels, and notably the topic node is sampled repeatedly within the document. Under this model, documents can be associated with multiple topics. Structures similar to that shown in Figure 1 are often studied in Bayesian statistical modeling, where they are referred to as hierarchical models (Gelman et al., 1995), or more precisely as conditionally independent hierarchical models (Kass and Steffey, 1989). Such models are also often referred to as parametric empirical Bayes models, a term that refers not only to a particular model structure, but also to the methods used for estimating parameters in the model (Morris, 1983).