How Surprising is the French Revolution? Insights and Information Theory

REBECCA L. SPANG AND SIMON DEDEO

Eighteenth-century studies has long understood itself as interdisciplinary, but the disciplines involved have almost exclusively been Humanities. This paper, co-authored by a historian and a physicist-turned-cognitive scientist, both presents some initial results of our work on the French Revolution's equivalent of the *Congressional Record* and reflects on the challenges and satisfactions the collaboration presented.

Along with Alexander Barron and Jenny Huang in the Lab for Social Minds, we worked with the digitized corpus of the first 33 volumes of the *Archives parlementaires* (*AP*): an anthology of the major speeches and debates delivered in the French National Assembly from summer 1789 to September 1791 (when that body dissolved and was replaced by the Legislative Assembly). Each volume of the *AP* runs to 800 pages of very small type and for more than a hundred years the only navigational tools were nineteenth-century indexes. Now, however, the *AP* volumes published before 1911 are digitally searchable and can be manipulated using computational methods (not those edited since 1961, which remain under copyright).[1]

Much of this collaboration consisted of making our disciplinary assumptions—our prior beliefs about the proper goals and methods of academic research—intelligible to each other (and, thereby, conscious to ourselves). To begin with: What is "information theory" and why use it to study the French Revolution? We have, after all, already plenty of what is normally called "information" on the French Revolution and no shortage of theories about that information and what it might mean.[2] Indeed, even the most apparently fact-driven account of revolutionary events relies on some implicit theory about the past (a set of assumptions, a way of seeing), even if that theory is so widely shared as to no longer be recognized as such. A historian works with a theory of information just as much as with the information itself.

While a statistician often speaks in terms of probability, historians deal in plausibilities. Consider R.R. Palmer's now classic *Twelve Who Ruled* (first published in 1941, reissued in 1973, 1989, and 2005). When Palmer wrote "As if the Hébertist uprising of September 5 were not enough to occupy the Committee [of Public Safety], it was on the same day that a depressing message came from Houchard, general in command of the Army of the North"[3] what he was in some sense really saying was: "I think it much more likely that the twelve men on this Committee got these two pieces of news on the same day, than that they did not." When a historian praises or criticizes Palmer, his or her reac-

[1] We worked with the version available from the French Revolution Digital Archive, a collaborative project between Stanford University and the Bibliothèque Nationale de France.

[2] This is not the place for a full historiographical essay on the French Revolution; see Rebecca L. Spang, "Paradigms and Paranoia: How Modern is the French Revolution?" *American Historical Review* 108 (February 2003), 119-147 and "Self, Field, Myth: What We Will Have Been," *H-France Salon* 1:1 (November 2009), 24-32.

[3] R. R. Palmer, *Twelve Who Ruled: The Year of Terror in the French Revolution* (Princeton: Princeton University Press, 2005), 78.

tion is not usually couched in terms of probabilities—we do not tend to say "there is only a 60% likelihood that at least eleven of twelve committee members learned about both these developments on 5 September 1793 and a 75% chance they were saddened by them."

Nonetheless, any historian's work depends on implicit attention to "what was probable" because, as a discipline, History is strongly committed to *un*certainty. History is the domain of both the contingent (the past did not have to happen as it did) and the fragmentary (our sense of what indeed did happen is reconstructed from available evidence). When historians produce an account of why the past happened as it did, they concern themselves with what might more or less plausibly—and thus, more or less probably—fill those gaps and resolve those possibilities.

That history as a discipline is always already implicitly probabilistic is important because "information theory" is an explicitly probabilistic way of knowing. The probabilities with which it deals are not like those we encountered in the Center's 2012 Workshop on Play, however.[4] That Workshop dealt with probability as we often consider it today, as something best understood in terms of gambling and games. Approached through a question such as "How many times will a coin come up heads?" probability is a measure of frequency and tells us about the behavior of objects.

But an epistemic, even moral, understanding of probability is both older and more relevant here. In the eighteenth century, most uncertainty resulted from not knowing the logic or properties of objects: hand-crafted dice, unlike the machine-made ones of a later era, did not all roll the same way and it was possible (even probable) that one die would regularly come up "two" more often than another.[5] And if we today assume (rightly or wrongly) that we live in a world of uniform objects, we nonetheless do not predicate that same uniformity of subjects.

We do not expect people to repeat their behaviors in a statistically reliable, predictable way and indeed, we consider a certain kind of unreliability—the possibility of gaps and changes of mind—as the very sign of subjecthood. (When people become so regular that they remind us of clocks, they become the stuff of legend, as with Immanuel Kant's late afternoon walk through the streets of Königsberg).[6] Still, we do expect human behavior to be plausible and those expectations frame and limit our speculations on others' actions (in past, present, and fiction alike).

Thinking in terms of plausibility gives us an intuitive, qualitative framework for measuring human probabilities; information theory aspires to make this thinking into a science. It brings to the fore a philosophical difference at the heart of probability itself. *Frequentist* probability theory focuses on the behavior of objects: it counts past outcomes and derives a set of regularities in the world that hold for the future as well. Information theory, in contrast, draws on an understanding of probability first formulated by the English non-conformist clergyman Thomas Bayes in the mid-eighteenth century and presented to the Royal Society by his Unitarian friend Richard Price (known to all *dix-*

[4] For paper summaries and transcripts of discussions—including the wonderful debate on the rationality of gambling—see *The Workshop* 1 (June 2013)

[5] Lorraine Daston, *Classical Probability in the Enlightenment* (Princeton: Princeton University Press, 1988).

[6] See Frédéric Gros, *A Philosophy of Walking*, trans. John Howe (Verso, 2015) and Manfred Kuehn, *Kant: A Biography* (Cambridge: Cambridge University Press, 2001).

huitièmistes for his 1789 sermon, "A Discourse on the Love of our Country," the provocation for Edmund Burke's *Reflections on the Revolution in France*).[7] The Bayesian approach used in information theory introduces a subject into the equation and makes probability into a dyadic relation. It asks: given what I knew (or thought I knew), how surprised am I by this or that new piece of information? Bayesian statistics (like much Enlightenment thought) understands this "I" not as a particular embodied subject but as an ideal, rational learner. We can hence program a computer to be this subject and to recognize surprises.

"Surprise" as we are using it here is not simply a byproduct of (or synonym for) variability, the way we might say we are surprised, for example, by the outcome of a dice roll. It is fundamentally relational, and does not have a constant: two poker players may, for example, have different levels of surprise about the same event. Consider the punter cheated by a card shark: the pair of aces is an (unfortunate) surprise for the former, and a foregone conclusion for the latter. We weigh probabilities against each other—and not against an independent, standard measure (such as a meter, pound, or gallon)—because surprise always depends on expectations and expectations ("prior beliefs") are neither constant between individuals nor for any given individual at different points in time. As we learn about the world—or about a datatset or corpus—our expectations of it shift.

Information theory studies "stochastic" (that is to say, aleatory or apparently random) processes as a model for how communication works. Information theorists—or, rather, mathematicians, cosmologists, computer scientists, and others working in the area known as "information theory"—look at some dataset (about whose contents they know very little) and try to find a "signal" (i.e., evidence that the dataset does communicate *something* even though they may not understand what is being said).[8] Their goal is to identify how the data communicate, not explain what is being said or craft a response.

Games are a common metaphor for the interplay of expectations and data, but the origins of information theory were more martial than recreational. Two of the founding names in the field (Solomon Kullback and Richard Leibler) worked as United States Army cryptologists during World War Two (using methods in part devised by the Coast Guard to decode messages being sent by "rum runners" under Prohibition).[9] Kullback did not know a word of Japanese, but he and a colleague (Frank Rowlett) could still "break the code" used in Japanese diplomatic communications. For the translator who then read the decrypted message, those communications were not stochastic (except insofar as any sign system is just that) but coded Japanese had no meaning for Kullback and his fellow codebreakers.[10] Claude Shannon, often called the "father of the information age," had a similar background in signal transmission and decoding, writing one paper on how to separate signals from noise and another (in telephonics) on how to distinguish a message from static. As a Digital Humanities method, Information Theory is a kind of distant

[7] In addition to Daston, *Classical Probability* see Stephen M. Stigler, "Richard Price, the First Bayesian," *Statistical Science* 33:1 (2018), 117-125 and Sharon Bertsch McGrayne, *The Theory that Would not Die* (New Haven: Yale University Press, 2011).

[8] For a basic textbook introductions see Thomas M. Cover and Joy A. Thomas, *Elements of Information Theory* (Wiley, 1991; 2006). See also, Brian Skyrms, *Signals: Evolution, Learning, and Information* (Oxford: Oxford University Press, 2010).

[9] https://www.nsa.gov/public_info/_files/oral_history_interviews/nsa_oh_17_82_kullback.pdf, 7

[10] https://www.nsa.gov/about/cryptologic_heritage/hall_of_honor/1999/kullback.shtml

reading, but unlike many forms of text mining, it measures not the content but the structure of communication.

Information theory lets us calculate how much information there is in any utterance. The more unexpected an utterance, the more information the perceiving subject gets from it. (Imagine waiting for a coin to be tossed versus waiting for two dice to be rolled—you will always get more information from the latter outcome, because there are more ways it could have turned out.) There is more "information" (also wonderfully called "surpisal") in a single unlikely occurrence than in multiple likely ones. In text mining, an "information rich" word is one that helps distinguish one author, text, or chapter from another, but its informative-ness should not be confused with its *meaning*.

From the point of view of information theory, it doesn't matter what the word or communication is: if your expectation was that a text could say only *a*, *b*, or *c*, and then it said *f* or it said *q*, the amount of information gained would be identical (regardless of whether *f* and *q* are synonyms, antonyms, or otherwise unrelated). In terms of information or surprisal, common or "function" words such as *the*, *and*, *he*/*she*/*it*, *under* etc. (pronouns, conjunctions, prepositions) are the equivalent of an "s" or "e" when playing Scrabble. Very common and useful to have, but low score. (Stylometry is the limit case here, using subtle differences in the weak surprisals of these function words to separate the patterns of one writer, or genre, from another.[11])

Our first surprise in this research arose when we trained the computer to read the volumes of the *AP* that cover the Constituent Assembly with the expectation that it would be able to distinguish conservatives from radicals or representatives of the First Estate (Catholic clergy) from those of the Second (nobility) and Third (everybody else) on the basis of their vocabularies and word-choice patterns. We formed this hypothesis after reading a paper by colleagues in Informatics at IU that demonstrates clear vocabulary polarization within the United States Congress over the last thirty years (1994-2012). Correia, Chan, and Rocha show that Democrats in this period were significantly more likely than Republicans to say "Africa," "tax cut," or "human right": Republicans were similarly much more likely to say "federal government," "bureaucracy," or "death tax."[12]

Using analogous methods, we wanted to know what words distinguished reactionaries from Jacobins, royalists from future republicans in the French Revolution's first tumultuous years. To our great surprise, no such clear polarization emerged. Instead, clustered together in the center of "discourse space" were the leading figures on *both sides*: Mirabeau and Cazalès, Robespierre and the abbé Maury. That is, given a word such as "king," "law," or "nation" in the *AP*, the probability of it having been uttered by any of those four speakers (as well as others such as Barnave, Garat, or the Lameth brothers) was much more equal than we expected.[13] (see Figure 1) At the edges of "discourse space"—that is,

[11] For an introduction to basic methods, see François Dominic Laramée, "Introduction to stylometry with Python," *The Programming Historian* (April 2018); for one interesting analysis of a contemporary author in the context of debates about the eighteenth-century "rise of the novel," see Simon Fuller and James O'Sullivan, "Structure over Style: Collaborative Authorship and the Revival of Literary Capitalism," *Digital Humanities Quarterly* 11: 1 (2017).

[12] R. B. Correia, K.N. Chan, and L.M. Rocha, "Polarization in the U.S. Congress," *The 8th Annual Conference of the Comparative Agendas Project* (Lisbon, June 2015).

[13] For the purposes of this first hypothesis, we limited our analysis to the 96 "major orators" in the Assembly (as identified first by Alphonse Aulard and then updated in Edna Hindie Lemay, ed.

where the vocabularies are least likely to be shared—we find those representatives who were the Assembly's policy nerds. Whether it was Vernier reporting on financial matters, Merlin or Target on detailed questions of jurisprudence, they used specialist languages that had comparatively little overlap with the rhetoric central to Revolutionary politics.

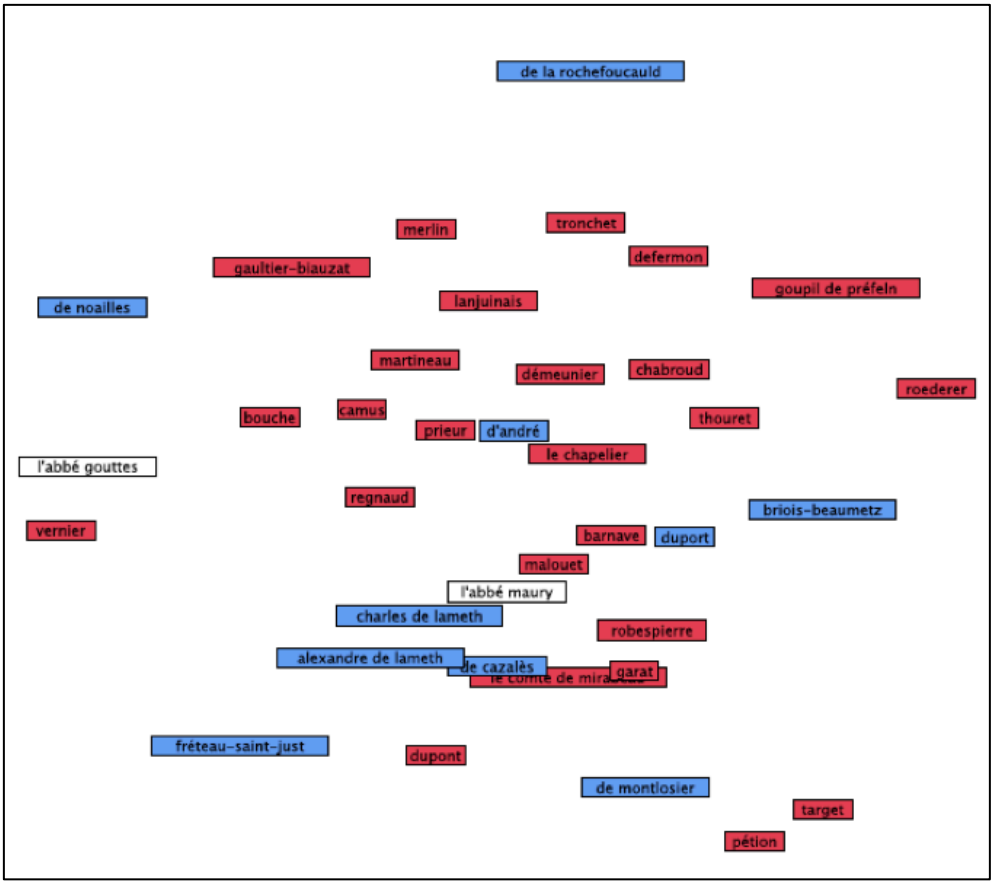

Figure 1: Discourse space in the *AP* (July 1789-September 1791). White=First Estate deputies; Blue=Second Estate; Red=Third.

The clustering of speakers from Right and Left alike in the center of discourse space did not conform to our prior expectation, but we learned something very important from it nonetheless. Discourse space in the first years of the French Revolution does not correspond to (it does not adequately represent or explain) political space. In this respect, the Assembly's "discourse space" as we define and analyze it today is remarkably different from its physical space as experienced by the deputies in 1789-1791—where those deputies pushing for change clustered on the presiding member's left and those resisting it on

---

Dictionnaire des constituants 1789-1791 (Paris: Voltaire Foundation and Oxford University Press, 1991) and Timothy Tackett, *Becoming a Revolutionary* (Pennsylvania State University Press, 1996). Our analysis was of course also limited to those utterances actually included in the *AP*—whether Robespierre and Maury were equally likely to say these words when speaking "off the record" remains necessarily unknown.

his right (hence our metaphorical use of Left and Right to describe political alignments today).

A crucial takeaway is that the computational analysis of the Assembly's discourse seems to disprove the so-called "discursive explanation" of the French Revolution. There were not distinct discourses of justice, will, and reason (as Keith Baker postulated thirty years ago).[14] Instead, there were shared words and on-going, fierce contestation over what those words meant. Information theory cannot tell us about meaning. If we want to understand those debates over meaning, we will have to read them the old-fashioned way.

Having discovered that the French Revolution's National Assembly of 1789 was far less polarized (at least in vocabulary) than the United States Congress today (a finding that should give us all pause), we turned then to a diachronic analysis. If computational methods and distant reading would not allow us to distinguish one speaker from another in the *AP*, what could they tell us about different periods in the Assembly's history? How did speakers' vocabularies and word patterns change (if at all) from the beginning of the Assembly to its end?

Comparing each speech in the *AP* to those that preceded and followed it (both immediately, and over the entire run of the 1789-1791 volumes), we employed two measures: *novelty*, or how much one speech's vocabulary and word patterns differed from those before it, and *resonance*, or how much those coming after resembled it. Surprisal as an information-theory category has the non-historical property of being something one can weigh with reference to both past and future. Our ahistorical gambler can look backwards, and judge the relative surprise of a new opponent's strategies to those seen before or, look forward, and ask how long these new strategies are likely to persist.

In the *AP*, high surprise in relation to the past indicates some sort of discursive innovation; to be surprising in relation to the future is to be transient and forgotten. In general, we find that most surprising speeches in the Assembly were surprising in relation to *both* past and future—in other words, most novelty did not resonate (the Revolution generated a lot of new-ness that was lost almost as quickly as it was created). A few speakers did, however, innovate in a way that others then imitated (this is what is sometimes called "influence," but the key behavior was of course not that of the so-called "influencer" but of those who followed).[15] Near the top of these influencers was Robespierre. (see Figure 2, next page).

[14] Keith Michael Baker, *Inventing the French Revolution* (Cambridge: Cambridge University Press, 1990).

[15] See Bruno Latour, "War and Peace of Microbes," in his *The Pasteurization of France*, trans. Alan Sheridan and John Law (1984; Cambridge: Harvard University Press, 1988).

Name	$z(\mathcal{N})$	$z(\mathcal{R})$	$\Delta z(\mathcal{R})$	Type
High novelty, high resonance				
Jérôme Pétion de Villeneuve	0.10	0.28***	+0.25***	3g
Maximilien Robespierre	0.11	0.18**	+0.14*	3g
Jean-Denis Lanjuinais	0.06	0.16***	+0.15**	3g
Alexandre Lameth	0.17*	0.14	+0.09	2g
Charles Antoine Chasset	0.31***	0.13	+0.04	3g
Committee (new item)	1.31***	0.12***	−0.27***	—
Philippe-Antoine Merlin	0.27***	0.05	−0.03	3g
Pierre-François Gossin	0.65***	0.03	−0.17*	3g
Jacques François Menou	0.40***	0.02	−0.10	2g
Committee (in debate)	0.29***	0.02	−0.07***	—
Left wing	0.07***	0.02*	0.00	(g)
3rd estate	0.06***	0.01	−0.02*	—
High novelty, low resonance				
Jacques Guillaume Thouret	0.16**	0.00	−0.05	3g
Jacques-Joseph Defermon	0.35***	−0.03	−0.13*	3-
François Denis Tronchet	0.24***	−0.04	−0.11*	3g
Armand-Gaston Camus	0.29***	−0.04	−0.13***	3g
President	0.02	−0.07***	−0.08***	—
Théodore Vernier	0.55***	−0.14	−0.31***	3g
Low novelty, high resonance				
Guillaume Goupil-Préfelne	−0.21***	0.13	+0.20***	3g
Jean-François Reubell	−0.18***	0.11	+0.16**	3g
Jacques Antoine de Cazalès	−0.44***	0.08	+0.21***	2d
Pierre Victor Malouet	−0.27***	0.08	+0.16***	3d
Jean-Siffrein Maury	−0.46***	0.07	+0.20***	1d
Pierre-Louis Prieur	−0.27***	0.05	+0.13**	3g
1st and 2nd estates	−0.10***	0.03***	+0.05***	—
Jean-François Gaultier de Biauzat	−0.13*	0.03	+0.06	3g
Right wing	−0.32***	0.03*	+0.10***	(d)
Low novelty, low resonance				
Antoine de Folleville	−0.44***	−0.01	+0.12	2d
Michel Le Peletier de Saint-Fargeau	−0.20***	−0.01	+0.05	2g
François-Dominique de Montlosier	−0.61***	−0.02	+0.17*	2d
Louis Foucauld de Lardimalie	−0.53***	−0.05	+0.11	2d
Charles Lameth	−0.15*	−0.06	−0.02	2g
Pierre François Bouche	−0.09*	−0.10	−0.07	3g
Antoine Barnave	−0.04	−0.12**	−0.11	3g

Bolded categories include all speeches by speakers who match either the type (estate or political affiliation; based on ref. 14), or role (committee or president; defined in text). $z(\mathcal{N})$: novelty compared with system average; $z(\mathcal{R})$: resonance compared with system average; $\Delta z(\mathcal{R})$: resonance relative to predicted resonance given novelty. "Type" codes for estate (3: bourgeoisie; 2: nobility; 1: clergy) and political affiliation (g: *gauche*, left-; d: *droit*, right-wing). *p* values corrected for multiple comparisons using Holm–Bonferroni (15).

Figure 2 (source: Barron, Huang, Spang, DeDeo, "Individuals, Innovations and Institutions"…)

A common and valid criticism of Digital Humanities projects is that they use sophisticated technological means to tell us something we already know (e.g., "Macbeth is a tragedy").[16] That Robespierre was important for the French Revolution is hardly a new finding, but that our methods allowed us to learn this "at a distance"—with a team of researchers led by a physicist and of whom several do not read much (or any) French—confirms that our methods "work" and opens the way to using them for further analysis.[17]

It also brings us to new questions and confessions (disciplinary as well as empirical, practical as well as conceptual). Humanities researchers share work by e-mailing attachments and using the "track changes" function in Word; scientists prefer LaTeX (an online system a bit like an early version of GoogleDoc, but where all editing is done in plain text). Historians argue among themselves about the value of endnotes over footnotes; scientists use parenthetical, in-text citations. The common model in the sciences is to publish a series of short papers (very short, by humanists' standards), whereas history remains based on the monograph or definitive, much longer, journal article. (In population biology terms, it's the difference between *r*-selection and *K*-selection strategies, where the first has many offspring that grow quickly and receive comparatively little care and the latter has only one or two on which it dotes.)

A historian's questions about this data—does the vocabulary of the 58 Third Estate deputies with some personal or hereditary claim to nobility align more closely with that of the other Third-Estate representatives, or of the Second Estate? do we see vocabulary differences between deputies from the *pays d'etats* and those from the *pays d'élection*?—may interest fellow scholars of 1780s-1790s France, but even very clear answers (if we could get them) are not likely to make much of an impact in the field of Social and Decision Sciences. (In our own vocabulary, those findings would be high novelty and low resonance.) Historians also know that most of the Revolution's debates and conversations did not happen on the floor of the Assembly, but until the era's pamphlets and newspapers have been digitized in an equally clear and standard format, our methods will remain applicable to the *AP* alone.

[16] Michael Witmore, Jonathan Hope, and Michael Gleicher, "Digital Approaches to the Language of Shakespearean Tragedy," in Michael Neill and David Schalkwyk, eds., *The Oxford Handbook of Shakespearean Tragedy* (Oxford: Oxford University Press, 2016)

[17] See our later paper: Alexander Barron, Jenny Huang, Rebecca Spang, and Simon DeDeo, "Individuals, Innovation, and Institutions in the Debates of the French Revolution," *Proceedings of the National Academy of Sciences* (*PNAS*) 115:18 (May 1, 2018), 4607-4612.