



Individuals, institutions, and innovation in the debates of the French Revolution

Alexander T. J. Barron^a, Jenny Huang^{a,b}, Rebecca L. Spang^c, and Simon DeDeo^{b,d,1}

^aSchool of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN 47408; ^bSanta Fe Institute, Santa Fe, NM 87501; ^cDepartment of History, College of Arts and Sciences, Indiana University, Bloomington, IN 47405; and ^dDepartment of Social and Decision Sciences, Dietrich College, Carnegie Mellon University, Pittsburgh, PA 15213

Edited by Danielle S. Bassett, University of Pennsylvania, Philadelphia, PA, and accepted by Editorial Board Member Michael S. Gazzaniga March 19, 2018 (received for review October 9, 2017)

The French Revolution brought principles of “liberty, equality, fraternity” to bear on the day-to-day challenges of governing what was then the largest country in Europe. Its experiments provided a model for future revolutions and democracies across the globe, but this first modern revolution had no model to follow. Using reconstructed transcripts of debates held in the Revolution’s first parliament, we present a quantitative analysis of how this body managed innovation. We use information theory to track the creation, transmission, and destruction of word-use patterns across over 40,000 speeches and a thousand speakers. The parliament as a whole was biased toward the adoption of new patterns, but speakers’ individual qualities could break these overall trends. Speakers on the left innovated at higher rates, while speakers on the right acted to preserve prior patterns. Key players such as Robespierre (on the left) and Abbé Maury (on the right) played information-processing roles emblematic of their politics. Newly created organizational functions—such as the Assembly president and committee chairs—had significant effects on debate outcomes, and a distinct transition appears midway through the parliament when committees, external to the debate process, gained new powers to “propose and dispose.” Taken together, these quantitative results align with existing qualitative interpretations, but also reveal crucial information-processing dynamics that have hitherto been overlooked. Great orators had the public’s attention, but deputies (mostly on the political left) who mastered the committee system gained new powers to shape revolutionary legislation.

cultural evolution | political science | cognitive science | computational social science | digital history

The French Revolution was a turning point in European history. Revolutionary commitments to individual liberty collided with ideals of social equality, while the rejection of Divine-Right monarchy and the embrace of laws based on reason opened a host of practical questions about how to govern the most populous state in Europe. The first parliament of the Revolution, the National Constituent Assembly (NCA), was a picture of upheaval from its outset.

Over the course of 2 years, the thousand or more individuals in that Assembly took it upon themselves to propose and argue the previously unimaginable: the revocation of Old-Regime privilege and the reinvention of the relationship between individual and state. But this parliament was more than a debate society for ambitious young men. It was also the origin of a system of rule. In the years that followed, successive legislative bodies declared war on most of Europe, dissolved the French monarchy, declared a Republic, and sentenced the former king to death—all while simultaneously writing constitutions and passing ordinary legislation. Many of their procedures and some of their personnel were drawn from the experience of the NCA.

As a parliament, the NCA confronted the problems that come with managing massive flows of information—problems still faced by the modern deliberative political bodies that, in many cases, are its direct descendants (1). But as the first

parliament, the body itself had little precedent to follow. Its members therefore faced a double challenge: how to convey points in a way familiar enough to be intelligible by others, while nonetheless making claims that were in many cases substantially novel (“revolutionary,” even). The NCA was a site, therefore, of both epistemic and political innovation. Conceiving it as such suggests two sets of questions. First, how did new ideas enter that parliament room; how were they adopted, adapted, or discarded by the men who heard them? Second: What institutions did the parliament evolve to manage the onslaught of novelty and reaction, optimism and grievance, philosophical argument and organizational minutiae that characterized the day-to-day tasks of governance and nation-building?

The digitization of historical archives allows us to answer these questions in a fundamentally new way. Using latent Dirichlet allocation (2) and new techniques in information theory drawn from the cognitive sciences (3), we track the emergence and persistence of word-use patterns in over 40,000 speeches made in the NCA and later reconstructed in the *Archives Parlementaires* (AP) from detailed records kept at the time. Two critical measures—novelty (how unexpected a speech’s patterns are, given past speeches) and transience (the extent to which those patterns fade or persist in future speeches)—allow us to trace both new manners of speech and the emergence of new institutions. Our mapping of the French Revolution’s turbulent early days in terms of the creation, sharing, and destruction of word-use patterns complements existing studies of specific ideas

Significance

How do democracies make decisions? We can read transcripts from parliament houses and legislative halls to see how particular ideas are introduced and debated, but we understand very little about the general principles of how these systems deal with information, or the origins of those principles. Here we study the parliamentary assembly of the first 2 years of the French Revolution, a model for democracies and revolutions across the globe, and show how patterns of speaking are created, picked up, and ignored or propagated. Political ideology, top-down rules, and individual charisma all affect how word patterns survive and thrive or, conversely, disappear and drop away.

Author contributions: A.T.J.B., R.L.S., and S.D. designed research; A.T.J.B., J.H., R.L.S., and S.D. performed research; A.T.J.B., J.H., R.L.S., and S.D. contributed new reagents/analytic tools; A.T.J.B., J.H., R.L.S., and S.D. analyzed data; and A.T.J.B., J.H., R.L.S., and S.D. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. D.S.B. is a guest editor invited by the Editorial Board.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence should be addressed. Email: sdedeo@andrew.cmu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1711729115/-DCSupplemental.

(much as evolutionary biology analyzes both mechanisms of transmission/selection and the particular phenotypes for which an environment selects).

We find, at high significance, a bias in favor of the propagation of novel patterns. In the framework of cultural evolution, the flow of ideas through NCA is out of equilibrium: the system reveals itself as having preferentially selected for what violated prior expectations. This effect was driven in part by charismatic political radicals such as Robespierre and Pétion de Villeneuve, who not only introduced new patterns more often than their peers but did so in a way such that others followed. By contrast, influential conservative figures such as Abbé Maury and Cazalès acted as inertial dampeners: their speeches maintained past patterns and carried them forward, despite the Assembly's overall bias toward innovation. Conservatives of the French Revolution “conserved”: not only did they refer to past traditions, but they did so with familiar discursive strategies and inherited word-use patterns.

In parallel with these individual-level differences, our methods reveal a major transition in how the parliament as a whole processed novelty. Roughly halfway through the NCA's existence, committees—which met outside the parliament but reported to it—gained new power to raise and resolve questions. Orators on the left and right continued to confront each other in public speeches from the floor, but those on the left also captured this new institutional mechanism and used it to their own advantage. The consolidation of this structural shift was accompanied by radicalization of the left and an accelerating flight of conservatives from both the parliament and the country itself.

Results

Social systems are characterized by heteroglossia: the coexistence, sharing, and competition of different patterns of speech. Heteroglossia makes linguistics and rhetoric (the reception, influence, and propagation of language within a community) (4, 5) core components in the quantitative study of culture. Tracking changes in speech patterns within a social body allows us to examine cultural evolution: the circulation, selection, and differential propagation of speech patterns in the group as a whole (rather than, say, tracking the ideas of a single individual). Patterns of heteroglossia demonstrate existing power relations, create new ones, and are a key method for the definition of both institutions and genres (6–8). Our methods here quantify a key aspect of cultural evolution: the extent to which

one agent's language patterns are used and copied by another (9, 10). To study the flow of rhetorical influence and attention in the NCA over time, we characterize how patterns of language use, uncovered by topic modeling, are propagated from speech to speech. We do so using Kullback–Leibler Divergence (KLD) (11): KLD, or “surprise,” measures the extent to which the expectations of an optimal learner, trained on one pattern, are violated by later patterns. Other work has demonstrated that surprise (in the Kullback–Leibler sense) is a cognitive as well as an information-theoretic quantity. It predicts what a subject will look at in a dynamically evolving visual scene (12) and can be used to map an individual's higher level activities (detecting, for example, biographically significant transitions in a subject's intellectual life) (3). Methodologically, this paper extends that work by considering surprise in relation to both past and future.

We use surprise to analyze a corpus of speeches by many different individuals. Surprise here measures both the deviation of one speech from the patterns of prior ones (novelty) and from patterns that appear in the future (transience). High surprise compared with the past indicates the topic mixture is new compared with previous speeches, hence the term “novelty”; high surprise compared with the future indicates that later speeches do not retain that pattern very strongly, hence the term “transience”. We provide a detailed introduction to these methods in *SI Appendix*.

Novelty and transience track a number of different effects. In addition to capturing intuitive notions of influence—a speaker with high novelty and low transience may have successfully shifted the terms of a discussion—they also track strategic effects (speakers who angle to speak early in a debate may have higher novelty) and external common causes (the first speech after a major event outside the chamber will have higher novelty). We test for a number of these latter effects (see *SI Appendix*), through day-level fixed effects and debate-level position analysis, to which our results below are robust. Our methods track the replications of patterns of speech and subject matter, rather than (for example) agreement: A speaker may introduce a new subject (high novelty) that is discussed by others (low transience), only to have his position contradicted or rejected.

Innovation Bias. Speeches in the NCA span a wide range in both novelty and transience; Fig. 1 summarizes the system at the level of individual speeches, in this case at the relatively rapid time scale (window width, w) of seven speeches. While the majority

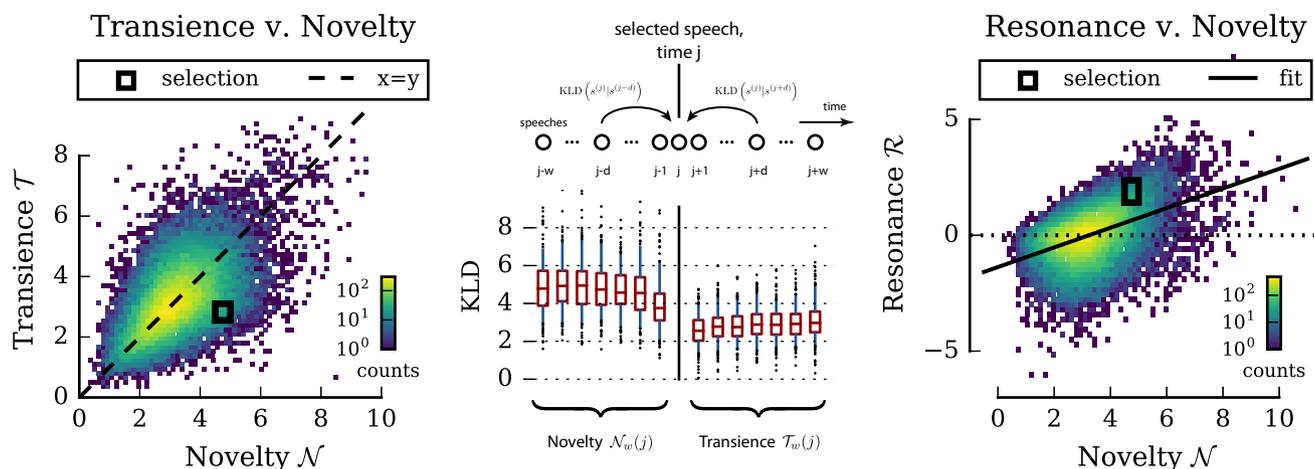


Fig. 1. Novelty, transience, and resonance in the French Revolution. (Left) A density plot of transience vs. novelty per speech at scale $w = 7$. Resonant speeches, with low transience compared with their novelty, fall below the identity ($x = y$) line. Resonant speeches at any time j are more surprising compared with preceding speeches (time $j - d$, $1 \leq d \leq w$) than successors (time $j + d$). This temporal asymmetry can be seen in the center plot of surprise for speech delay d surrounding highly resonant speeches from the selection at Left. (Right) Resonance vs. novelty, with regression line. Although novelty is tied to transience, it is also necessary to achieve resonance.

of speeches concentrate near the symmetry line—speeches with high novelty are likely to have similarly high transience—two results stand out. First, the scatter is large: many speeches lie far off the novelty–transience line of equality, and it is easy (for example) to find speeches with top-quartile novelty that have bottom-quartile transience. “What is new is quickly forgotten” is a useful heuristic but holds only in the average. Below, we consider two potential drivers of these diverse receptions: speaker and context.

Second, novel speeches were unexpectedly influential. We quantify this with “resonance,” the imbalance between novelty and transience (see *Materials and Methods*). Resonance, the quality of at once differing from the past and leaving traces on the future, increases with novelty, as shown in the rightmost plot of Fig. 1. We refer to this positive relationship as innovation bias: penalties to high novelty speeches are lower than expected in a system at equilibrium. This bias is measured by Γ , the slope of the novelty–resonance line; positive Γ indicates innovation bias. We find this bias in place from the most rapid time scales (one speech to the next, $w = 1$), to $w \approx 1000$, those on the order of days (see *SI Appendix*). This innovation bias lasts for at least the course of a day, as speakers deliberately turn to new topics, but fades away on longer time scales.

Organizational Roles and Individual Strategies. In choosing when to speak and what to say, speakers had some control over the relative novelty of their speeches. A speaker’s attitude toward the Assembly as an institution, like other political or philosophical commitments, would contribute to his willingness to create new word patterns or copy earlier speakers. Conversely, speakers had much less control over the reception of what they said. Idiosyncratic properties of the speaker, ranging from the political (e.g., faction membership) and the social (e.g., demeanor or prestige) to the rhetorical (everything from word choice to pitch and volume of speech) could have altered the reception of their words.

Patterns at the level of the individual could, of course, temporarily break the system-level trend that favored innovation. An unskilled but ambitious deputy might have tried to introduce new ideas that his colleagues tended to dismiss, registering in our data as high novelty with low, or even negative, resonance. Conversely, the prestige or social power of another deputy might have meant he combined low novelty with high resonance, thereby keeping the conversation on track.

Decisions at the individual level cannot account for all deviations from system-level patterns, however. From its first days, the Assembly organized itself in such a way as to assign explicit roles to particular speakers. These information-processing functions overruled an individual’s personal characteristics. An example from the NCA, still common to many parliaments today, is the role of president, who served as point of contact for the King and enforcer of the daily agenda (13). The president’s role—one in which only 49 of the NCA’s more than a thousand members ever served—was largely functional and organizational.

The NCA also created another specialized entity: the committee. Committees, whose members were notionally selected on the basis of expertise, deliberated in private. They developed content outside the debate process and then presented it to the full body for public review. While a speaker on the Assembly floor might play to the audience in the visitors’ galleries, committee members addressed only each other.

Lexical markers in our data identify when a committee proxy was speaking and allow us to classify his speech into two categories: “new-item” speech and “in-debate” committee speech (see *SI Appendix*). New-item speeches introduced official content to the Assembly floor (typically draft legislation to be debated by the body) and mark transitions in attention from one topic to another. In-debate speeches occurred when a committee mem-

ber or spokesman engaged with other delegates following the item’s introduction.

To understand how individual-level differences and system-imposed roles affected the production and reception of new patterns, we consider the novelty and resonance of speeches given by those in three distinct Assembly roles: the 40 most common orators, the President (regardless of who held the position), and committee proxies. We calculate the average novelty and resonance for each category (scaled by z score), $z(\mathcal{N})$ and $z(\mathcal{R})$, identifying both potentially idiosyncratic speech-pattern innovations and their reception by the system as a whole, at scales of w from 1 (one speech compared with the next, and prior, speech) to 5,000 (one speech compared with 3 mo of speeches before or after).

The system’s overall bias in favor of innovation would predict that the category of speakers with the highest novelty would also have the highest resonance. To determine if this expectation is validated in our sources, we compare the overall novelty–resonance relationship of the system (depicted by the fit line in Fig. 1) to the measured resonance for speeches from each of the three categories (orator, President, committee proxy). Specifically, for each category we report $\Delta z(\mathcal{R})$, defined as $z(\mathcal{R}) - \mathbb{E}[z(\mathcal{R})|z(\mathcal{N})]$, the difference between the measured mean resonance of speeches and the expected mean resonance under the OLS model $z(\mathcal{R}) \sim z(\mathcal{N})$. While $z(\mathcal{R})$ measures the effect speakers had on later discourse, $\Delta z(\mathcal{R})$ measures the extent to which they broke system-level trends in achieving those effects. For example, a speaker with high novelty may have high resonance but negative $\Delta z(\mathcal{R})$, indicating that his adventurousness was rewarded less than expected. Full results for $w = 36$ (roughly half a day) are shown in Table 1; results for novelty are stable on all time scales, while resonance at $w = 36$ is strongly correlated from $w = 3$ to $w = 100$ (see *SI Appendix*).

Institutional roles—speaking on behalf of a committee, presiding over the Assembly—did not follow these system-level trends. High novelty and high resonance together characterize committees as gatherers of new information that they injected into debate in ways that defined downstream discussion. In contrast, the president’s role as agenda enforcer led to lower than average resonance: he acted, at best, to summarize what had come before, while having less influence on patterns of speech that followed. Though he might break from conversation to further the agenda, the content he introduced tended not to persist. The overall novelty bias cannot be explained by the taking-in of new information from committees: while committees show above average resonance ($z(\mathcal{R})$ greater than zero), they have lower resonance than expected given their novelty ($\Delta z(\mathcal{R})$ less than zero).

The data similarly show individual orators departing from system-wide trends. Of the top 40 speakers in the assembly, 27 show significant deviations from aggregate patterns in either novelty or resonance at at least the $p < 0.05$ level, with 22 speakers showing deviations at $p < 10^{-3}$. Speakers deviate in both directions, with some showing anomalously high tendencies to break with past patterns and others showing similarly strong tendencies to preserve them. High-novelty speakers are overwhelmingly associated with the left wing and the bourgeoisie, while all of our right-wing speakers, and the vast majority of nobility, are low-novelty. The order in which speakers joined a debate appears to be driving some, but not all, of these effects (*SI Appendix*).

While more than half of the top 40 orators show surprisingly high or low novelty, fewer are distinguished by their resonance. The latter were, however, among the key players of the Revolution. The celebrated radicals Robespierre and Pétion achieved not only the highest resonance but also significantly higher resonance than even that due to the system-wide novelty bias [positive $\Delta z(\mathcal{R})$]. In contrast, speakers such as Armand-Gaston Camus and Théodore Vernier, called on primarily for their specialized knowledge in canon law and finances, show high

Table 1. Mean novelty and resonance by speaker at scale 36, for role and type (in bold) and the top 40 orators

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).

novelty but low resonance: they presented information that either failed to make an impact or (more likely) settled questions so conclusively that the room moved on to completely different discussions. Finally, prominent political conservatives such as Jean-Siffrein Maury and Jacques de Cazalès appear in the low-novelty, high-resonance quadrant. They break the system-level novelty bias and are notable not only for keeping the conversation on track (low novelty) but for speaking in ways that persist forward (high resonance). In this, Maury and Cazalès are characteristic of the right-wing overall: while the novelty-biased left was composed of both high- and low-resonance speakers, right-wing speech patterns persist, with positive $z(\mathcal{R})$ and $\Delta z(\mathcal{R})$ despite their anomalously low novelty.

The Emergence and Evolution of the Committee. Committees were a key NCA innovation, allowing the system to manage vast amounts of information without overwhelming legislative debate. Committees in the NCA did not appear overnight. Our previous section establishes their unusual functional role, but the AP's comprehensive coverage allows study of their role's emergence as well. In this section, we show how returns to novelty, Γ , were modulated by committee roles over time. We fit, separately, two terms that quantify the additional boost (or decrement) to the novelty-resonance relationship when speeches either introduce new committee items (Γ_n) or advocate on behalf of a committee during debate (Γ_d). A speech of novelty \mathcal{N} , for example, achieves on average a resonance \mathcal{R} equal to $(\Gamma + \Gamma_n)(\mathcal{N} - \mathcal{N}_0)$ when made by a committee member introducing a new item, compared with $\Gamma(\mathcal{N} - \mathcal{N}_0)$ when the speaker acts on his own behalf.

We look for discrete shifts in committee function, doing change-point detection with a maximum-likelihood model of the novelty-resonance relationship where Γ , Γ_n , and Γ_d are allowed to vary in time. Following ref. 3, we consider a two-epoch model, where all three quantities are fixed to constant values in each epoch, with a single discrete change at a particular time point whose position is a free parameter. The two-epoch model is preferred to a single-epoch model, as well as to a linear (secular shift) model under AIC. Our maximum likelihood change-point in the nature of committee functions occurs in late 1790; the modal best fit date across all scales is October 31, 1790. Allowing the intercepts of the new-item and in-debate speeches, as well as their slopes, to vary produces nearly identical results. A separate frequentist analysis rejects a randomly ordered null model at $p < 10^{-2}$ (see *SI Appendix*).

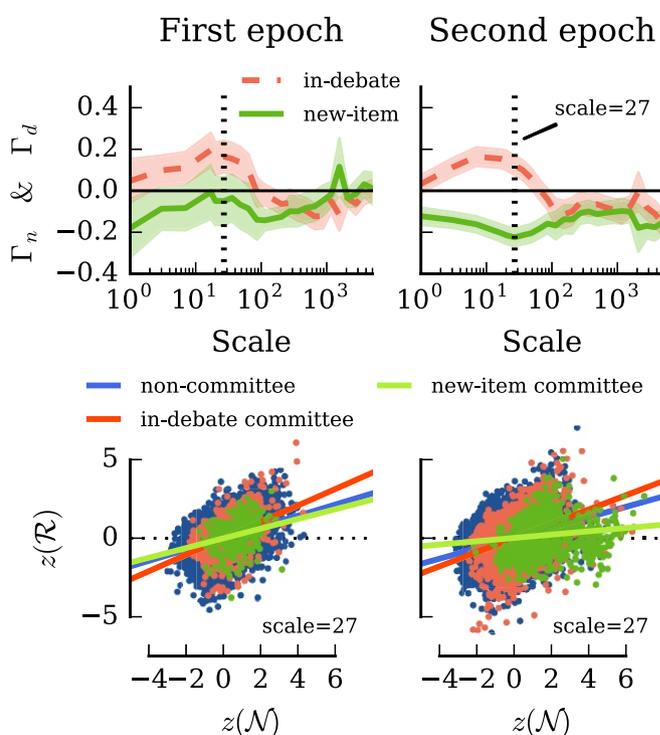


Fig. 2. Information-processing functions of NCA committees before (first column) and after (second column) the late-1790 change-point. (Top) The shift in the novelty-resonance relationship for new-item and in-debate committee speech, with 99% confidence intervals. (Bottom) Scatter plots and fit lines at scale 27 for these speech types, compared with all other speeches. The "undebated tail" appears in the second epoch as a new cloud of green points along the dotted line, generated by committees with new powers to propose and dispose.

