

Drivers of English Syntactic Change in the Canadian Parliament

Liwen Hou

Khoury College of Computer Sciences
Northeastern University
Boston, MA
hou.l@northeastern.edu

David A. Smith

Khoury College of Computer Sciences
Northeastern University
Boston, MA
dasmith@ccs.neu.edu

Abstract

Corpus linguists have long noted the “colloquialization” of many genres of English. While the average decline in many features of formal speech is obvious in aggregate, we are better able to disentangle drivers of change by examining Canadian parliamentary speeches coded for characteristics of individual speakers across more than 100 years—much longer than previous studies of individuals’ language change in a common environment. While many language changes proceed by cohort replacement and often originate with female speakers, the Canadian Hansard shows that most speakers employed increasingly colloquial language over their careers and that gender effects are mostly explained by the stronger effect of political power. Using multilevel regression, we investigate growing colloquialization by tracing declining passives, modals, and pied-piping relative clauses, and increasing numbers of semi-modals and progressives in the speech patterns of individuals.

1 Introduction

Although it is well known that languages change over time, the exact mechanisms of such change are less understood. In this work, we investigate how syntactic change propagated over the past century in the Parliament of Canada. We focus on instances of language change that correspond to the colloquialization of parliamentary speech.

Leech (2004) defines *colloquialization* as “a tendency for the written language gradually to acquire norms and characteristics associated with the spoken conversational language”, but the meaning that we attribute to the term is the one found in Collins and Yao (2013), where it is a “process involving a spreading of colloquial features from casual face-to-face conversation to other genres, not only written”. Although this can manifest as more than one

type of language change, we focus on syntactic aspects of colloquialization in this work.

Specifically, we find that the *be*-passive became rarer, that the progressive became more common, that pied-piping became less frequent, that modals such as “shall” also decreased in frequency, and that semi-modals such as “have to” rose in frequency. See Section 3 (and in particular Table 1) for examples of the five aforementioned features.

We are particularly interested in the forces driving language change. One natural hypothesis is that newer parliamentary members repeatedly introduce more colloquial language in their speeches, so that language shifts whenever an older member of Parliament retires and is replaced by a young member (cohort replacement). Another hypothesis is that the entry of women into Parliament was a driver of the observed changes, as when Labov (2001) concludes that the “leaders of linguistic change have been located as women”.

By analyzing variation at the individual level, we show that neither of these hypotheses can explain the increasing average colloquialization we observe in Canadian parliamentary speeches. In a given year, variation in speakers’ ages affects colloquialization measures of their speeches little. Seniority of service in Parliament has a small effect in the wrong direction: new members of Parliament are slightly less colloquial. As for women, they speak more formally than men rather than less. We show that it is individual members of Parliament who change their speech style to become more colloquial over time. We also find that whether a speaker belongs to the party in power has a large effect on colloquialization, which reverses when a speaker’s party loses power, indicating that analyses of language change need to account for synchronic variation.

Previous studies of colloquialization with large datasets (e.g. Kruger et al., 2019; Hiltunen et al.,

2020) did not analyze language at the level of individuals, for which we show that the various colloquial features are correlated. Crucially, having data at the level of individuals allows us to analyze the social variables that cause change.

While other individually-labeled datasets (e.g. Reddit comments) have previously been used to study language change (e.g. Del Tredici and Fernández, 2018; Stewart and Eisenstein, 2018; Barron et al., 2018; Danescu-Niculescu-Mizil et al., 2013), there are not many such sources of data that also span a significant number of years. The study that is the most similar to ours is that of Arnaud (1998), which claims to have found a correlation between the progressive and two sociolinguistic factors (gender and intimacy) from a survey of private letters of 22 people in the 1787-1880 period.

According to Collins and Yao (2014), “studies of grammatical change in the Late Modern English period have concentrated almost exclusively on British and American English”. Moreover, previous studies of Canadian colloquialization (e.g. Collins and Yao, 2013; Tagliamonte and D’Arcy, 2007) were based on datasets that do not span many years. In addition to the novel study of individual syntactic changes and their social correlates using computational methods, another contribution of this work is thus that we present the first study of colloquialization for Canadian English using data covering many decades.

2 Data

The Canadian Hansard dataset ¹ contains over 650 million words and consists of the digitized proceedings of the Canadian House of Commons from 1901 to 2017 (with missing data in 1906 and 1907); the digitization process is described in Beelen et al. (2017). This dataset is notable for the fact that it contains information not only about what was spoken but also about who spoke it; this individual-level labeling of the data is key in allowing us to correlate colloquialization features across individuals. Furthermore, biographical information (e.g. date of birth) is linked to the identifier assigned to each member of Parliament.

3 Colloquialization Features

We examine five features hypothesized to be related to colloquialization according to the literature

¹Version 1.0.2 of the dataset, which we used for this project, was downloaded from www.lipad.ca/data/

(e.g. Leech, 2004; Mair, 2006; Biber and Finegan, 1989): the passive, the progressive, modals, semi-modals, and pied-piping in relative clauses. We use the modals (*will, would, can, could, may, might, should, shall, must, ought, need*) found in Leech et al. (2009), p. 72. While semi-modals have no consistent definition, we use the “emergent modals” listed in Leech et al. (2009), p. 98: *have got to, be going to, want to, have to, need to* (which are also the top five most frequent semi-modals listed on p. 100 of the same source). Examples of these features are shown in Table 1.

3.1 Processing and Identification of Features

We process the data using Stanford CoreNLP (Manning et al., 2014) to obtain labeled dependency parses, part-of-speech tags, and lemmas for all tokens. We then identify the relevant linguistic features as follows.

Passives: We classify passive verb phrases as those with dependency labels “auxpass” and “nsubjpass”.

Progressives: We find progressive constructions by looking for a POS tag beginning with “VB” that corresponds to the lemma “be” and that is followed by the POS tag “VBG”.

Pied-piping: We detect pied-piping relative clauses by looking for a preposition (tagged “IN”) immediately followed by “who”, “whom”, “whose”, or “which”.

Modals: We identify modals by checking to see if the part-of-speech tag for a word is “MD” and if its lemma matches an element from our list of modals.

Semi-modals: Semi-modals are identified in the following way: For “have got to”, we look for instances of “got” that is followed by “to” plus a verb and preceded (either immediately or with an intervening word to cover cases such as “smuggling has apparently got to be a science”) by a word whose lemma is “have”. For “be going to”, we identify instances of “going” where the lemma of the preceding word is “be” or “not” and where the following word is “to”. For our other three semi-modals, we look for words whose lemma is “want”, “have” or “need”, and we then check whether it is a verb and whether the word after it is “to”.

Feature	Example
passive	This provision is modified by a time limit.
progressive	We are seeing these valuable resources go to other countries.
pied-piping	It was a question in which everyone was interested.
modal	The RCMP could show the file to an insurance company that will increase my insurance.
semi-modal	That has got to be good, when we consider the years when it did not grow.

Table 1: Examples found in parliamentary speech of the five features that we study

4 Summary Statistics

We have confirmed for Canadian English the trends observed in British and American English. First, the rate of passivization, pied-piping, and modals has decreased over time, and the progressive and semi-modal constructions have become more common. The average rate of these constructions per year is depicted in Figures 1a to 1e. The number of verb phrases is the denominator of these rates.

We also find that these colloquialization features correlate across individuals. In other words, some individuals speak more colloquially than others. For the top 100 individuals (based on the number of verb phrases spoken), the correlation matrix for these features is given in Table 2; observe that the signs are consistent with what one would expect them to be in all cases.

5 Colloquialization Metric

We define a colloquialization score based on the aforementioned linguistic features. For each feature, we calculate its mean and standard deviation over the top 100 speakers. For example, for passives, we look at the number of passive constructions divided by the number of verbs spoken by each of the top 100 speakers, and we calculate the mean and standard deviation of this list of 100 passivization rates. This gives us *pass-mean* and *pass-stdev*, which we use as baselines to normalize all other blocks of text. We repeat this procedure to get baselines for our other linguistic features.

The top 100 speakers were chosen because they have enough speech data (at least 350,000 words each) to get reasonably clean estimates of the rates at which they use various linguistic constructions; using more speakers would risk overestimating the standard deviation because the estimate of the rates of usage of their linguistic features will be noisy.

Next, for any block of text t , we define its colloquialization score as follows. First, we count the number of passive constructions $pass_t$ in the text and the number of verbs V_t in the text, and divide them to get the passivization rate $pass-rate_t :=$

$pass_t/V_t$. Next, we normalize this passivization rate with respect to the baselines *pass-mean* and *pass-stdev* that we calibrated for the top 100 speakers; this gives us a Z score for passivization:

$$pass-Z_t := \frac{pass-rate_t - pass-mean}{pass-stdev}.$$

We repeat this procedure for our other linguistic features to get a Z score for each.

Finally, we combine these Z scores linearly, with a fixed weight for each feature. In other words, we define

$$colloq-score_t := w_{pass}pass-Z_t + w_{prog}prog-Z_t + w_{mod}mod-Z_t + w_{semi}semi-Z_t + w_{pied}pied-Z_t.$$

We pick the weights $w_{pass} = -1/5$, $w_{prog} = 1/5$, $w_{mod} = -1/5$, $w_{semi} = 1/5$, and $w_{pied} = -1/5$, where the signs correspond to whether the feature goes up or down over time. These even weights are chosen for simplicity; the purpose of this colloquialization score is primarily to aggregate together our various linguistic features. In practice, we observe that these features usually move in tandem. In Section 8, we introduce a more principled model that represents colloquialization as a latent variable.

6 Explanatory Features

We consider a variety of independent variables that may help explain the increase in colloquialization in our dataset. We focus on 1961-2010, a 50-year time period during which there was a marked increase in colloquialization. In this section, we describe these independent variables and provide some summary statistics about them.

6.1 Party and Power Effects

One notable pattern is that speakers who belong to the governing political party tend to speak less colloquially compared to speakers whose party is not in power as shown in Figure 2. (For the purpose of this analysis, we merge the Conservative party with its predecessors and compare their members to the Liberals and New Democrats. The latter party,

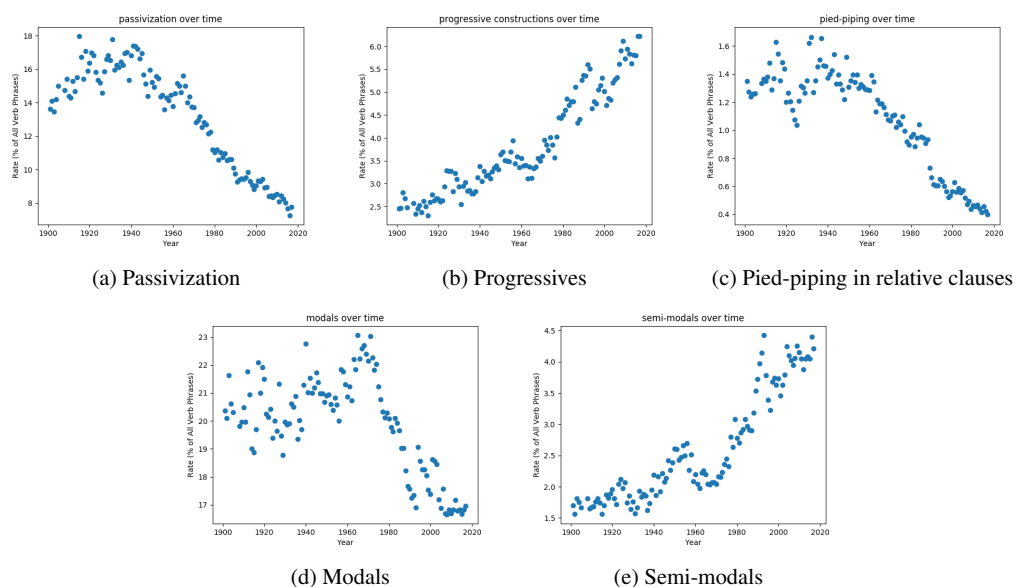


Figure 1: Rate of constructions in verb phrases by year

	passives	progressives	modals	semi-modals	pied-piping
passives	1	-0.787	0.690	-0.793	0.795
progressives	-0.787	1	-0.687	0.752	-0.739
modals	0.690	-0.687	1	-0.620	0.548
semi-modals	-0.793	0.752	-0.620	1	-0.683
pied-piping	0.795	-0.739	0.548	-0.683	1

Table 2: Correlation between colloquialization features of top 100 individuals

founded in 1961, has been seen as having “no hope of ever forming government itself” according to [Hausman and Turnbull, 2014](#).)




Figure 2: Speeches by members of the governing party are on average less colloquial

For the years 1961 to 2010, the average colloquialization score of the members belonging to the party in power was 0.07. In contrast, the average colloquialization score of the members not in power was 0.30 (restricting only to Liberal and Conservative members of Parliament for a fair com-

parison). The difference between these scores is 0.23. For comparison, in the same 50-year period, the average colloquialization score of the New Democrats was 0.48, significantly higher than both Liberals and Conservatives regardless of whether they were in or out of power. Between Conservatives and Liberals, the party holding power spoke more formally than its out-of-power counterpart by an amount averaging approximately 6 years of change. In other words, the party in power between 1961 and 2010 appears to revert to the speech style of around 6 years ago in terms of formality. In contrast to studies that show lower-status speakers accommodating to higher-status ones ([Danescu-Niculescu-Mizil et al., 2012](#)), out-of-power speakers continue to use more colloquial features while in-power speakers eschew them. In contrast to the study by [Barron et al. \(2018\)](#) on the first two years of the French Revolutionary assembly, political power in the Canadian Parliament over the 20th century is associated with linguistic conservatism. In Section 8, we quantify these differences by fitting a model which treats colloquialization as a latent variable.

6.2 Cohort and Age Effects

The colloquialization scores of new members (at most 4 years in Parliament, corresponding to the bottom third) as opposed to more senior members (at least 10 years in Parliament, corresponding to the top third) do not show visible differences.

Similarly, if we restrict the analysis to only the speakers who are below age 47 at the time of the speech (corresponding to the bottom third) and those above age 56 (corresponding to the top third), we do not find differences between the two groups.

However, in Section 8, we use our model to give some evidence that there are modest age and seniority effects in opposite directions: older members of Parliament speak more formally, while more veteran members of Parliament speak more colloquially. Since age and seniority are correlated, these effects partially cancel each other out.

6.3 Gender Effects

When women first joined Parliament, their speeches were less colloquial compared to the speeches of their male counterparts, but the gap between the genders has narrowed over time as we can see in Figure 3. In Section 8, we show that this narrowing can be explained by power effects.




Figure 3: Colloquialization scores by gender

Labov (2001) states that “in the great majority of cases, it is women who are ahead”. Our observation that women are not the leaders of change for colloquialization is thus a surprising result.

7 Individual Changes

As in the previous section, we restrict this analysis to the period from 1961 to 2010, a 50-year period in which colloquialization consistently increased in our dataset. Our goal in this section is to discern to what extent this language change represents changing speech styles of individual members of

Parliament, rather than changes in the pool of parliamentary speakers (the latter category includes cohort effects, the effect of women joining Parliament in increasing numbers, and the effects of party representation in Parliament, among others).

First, we fit a line of best fit to the yearly colloquialization scores of each of the top 100 speakers in this time period (as defined in Section 5). To do this, we use weighted least squares, where the weight of each year is the number of words the speaker spoke in that year. Of the top 100 speakers, 83 of them had positive slopes, meaning they became more colloquial over this time period. Moreover, the slopes of the individual lines of best fit approximate the overall colloquialization change observed in the dataset well: the aggregate colloquialization score of the dataset increased at a rate of 0.0378 per year, while the average speaker’s colloquialization score increased 0.0325 per year (with a standard deviation of 0.0442).

One potential problem with this analysis is that it may be confounded by the power effect described above: individual speakers (belonging to the Liberal party or one of the Conservative parties) change their colloquialization rates when their party attains or loses power, becoming less colloquial when in power. This effect may cause some speakers to appear to have declining colloquialization scores simply because they first joined Parliament when out-of-power and later came into power. To account for this, we adjust the colloquialization score of each Liberal or Conservative speaker by adding or subtracting a parameter δ_y , which we set to equal half the “colloquialization power gap” for the year y ; the latter gap is the overall colloquialization score of the Liberals or Conservatives who are not in power minus the overall colloquialization score of the Liberals or Conservatives who are in power (restricted to that specific year). In other words, in a year when the governing party has an aggregate colloquialization score of X and the party that is not in power has a colloquialization score of Y , we set $\delta_y = (Y - X)/2$. We then add δ_y to the colloquialization scores of the in-power members as well as subtract δ_y from the scores of the out-of-power members. (Only Liberals and Conservatives are adjusted this way.) After this adjustment, 86 of the top 100 speakers have positive slopes, which means that their colloquialization scores increased over time.

7.1 A Differential Analysis

To examine whether there are truly individual changes in speech rather than a changing pool of speakers, we do an individual differential analysis: we average all the individual year-to-year changes.

To do so, we first compute the passivization rates $pass-rate_{p,y}$ for each speaker p and year y (and similarly for the other linguistic features). We turn these rates into colloquialization scores $colloq_{p,y}$ as described in Section 5. We then subtract the rates for consecutive years for a given speaker:

$$\Delta_{p,y} := colloq_{p,y+1} - colloq_{p,y}.$$

We also keep track of the amount of speech spoken by person p in year y , which we measure by the number of verbs spoken. We denote this by $V_{p,y}$. We would like to take a weighted average of the differences $\Delta_{p,y}$; however, using $V_{p,y}$ as the weight of $\Delta_{p,y}$ neglects the possibility that $V_{p,y+1}$ is very small and hence that $colloq_{y+1,p}$ is poorly estimated. Instead, we define the weight $W_{p,y}$ as the harmonic mean of the number of verbs spoken in years y and $y + 1$:

$$W_{p,y} := \frac{2V_{p,y}V_{p,y+1}}{V_{p,y} + V_{p,y+1}}.$$

This weight scheme gives the right behavior of being small if one of the two counts $V_{p,y}$ and $V_{p,y+1}$ is small. We then take the weighted average of all the differences $\Delta_{p,y}$ to get the average individual colloquialization change:

$$average-change := \frac{\sum_{p,y} W_{p,y} \Delta_{p,y}}{\sum_{p,y} W_{p,y}}.$$

We exclude from the sum the years y in which the party in power changes from year y to $y + 1$; this way, we average the individual changes while keeping the “in power” feature fixed to remove any colloquialization effect of being in or out of power.

Note that this definition of *average-change* makes it the average yearly *individual* colloquialization change. This is deliberately defined to exclude any effects of changing population; for example, if the proportion of female members of Parliament increases over time, and if women speak differently than men, this will not affect *average-change* at all, since it only measures the average change for a single individual at a time. Similarly, this variable is not affected by changing

parliamentary composition (e.g. if the Conservatives or New Democrats win or lose seats) or cohort effects.

We compare *average-change* to the slope of the weighted line of best fit of colloquialization over time; the latter is the naïve measurement of the colloquialization increase in our data, corresponding to the amount of effect to be explained. We measure both of these over the 1961–2010 period in which we see a sharp colloquialization increase.

As mentioned earlier, the slope of the weighted line of best fit was 0.0378; this is the amount of colloquialization increase per year in our data to be explained. This can be interpreted to mean that on average across our linguistic features, the feature increases or decreases 0.0378 standard deviations per year, where the standard deviation is measured with respect to the top 100 speakers. In contrast, *average-change* was 0.0348 for the relevant time period. This means that of the yearly change of 0.0378, as much as 0.0348 can be explained by individual changes, rather than changes in the composition of the pool of speakers. This is 92% of the effect, which means that individual change constitutes the vast majority of the dataset’s change.

8 Regression Model

We model the language change over the 50-year time period (1961–2010) as being driven by a single underlying measure of colloquialization. More explicitly, we write the following model:

$$\begin{aligned} colloq &= \alpha_0 + \alpha_{year}year + \alpha_{power}power + \alpha_{age}age \\ &\quad + \alpha_{senior}senior + \alpha_{gender}gender + \dots \\ pass &\sim \text{Binomial}(V, \ell(\beta_{pass}colloq + \gamma_{pass})) \\ prog &\sim \text{Binomial}(V, \ell(\beta_{prog}colloq + \gamma_{prog})) \\ mod &\sim \text{Binomial}(V, \ell(\beta_{mod}colloq + \gamma_{mod})) \\ semi &\sim \text{Binomial}(V, \ell(\beta_{semi}colloq + \gamma_{semi})) \\ pied &\sim \text{Binomial}(V, \ell(\beta_{pied}colloq + \gamma_{pied})) \end{aligned}$$

In the above model, colloquialization is a linear function of explanatory features such as year, gender, party, seniority (number of years in Parliament), and age. The function $\ell(\cdot)$ is the logistic function, which maps $(-\infty, \infty) \rightarrow (0, 1)$ to convert unconstrained real numbers to probabilities. We model the number of passive verbs in a passage of text, denoted *pass*, as coming from a binomial distribution with parameters V and $\ell(\beta_{pass}colloq + \gamma_{pass})$; in other words, in our generative model, each verb phrase in the text becomes passivized with independent probability $\ell(\beta_{pass}colloq + \gamma_{pass})$. The passivization rate is

therefore allowed to depend linearly on the underlying colloquialization, with coefficient β_{pass} and intercept γ_{pass} . The number of progressives, modals, and semimodals are modeled similarly (with their own parameters β_* and γ_* dictating their dependence on the underlying colloquialization parameter). In other words, we model the linguistic features *pass*, *prog*, *mod*, *semi*, *pied* as dependent variables and *year*, *gender*, *party*, *power*, *seniority*, *age* as independent variables, and we set the dependence to be funnelled through a single latent parameter *colloq*. In the above model, the intercept α_0 of the colloquialization parameter is redundant, as we can compensate for any change in α_0 by changing the γ_* parameters. For this reason, we set $\alpha_0 = 0$ without loss of generality. Similarly, the scale of *colloq* is arbitrary; if we scale all the α_* parameters by a constant factor, we can compensate for this by scaling the β_* parameters in the opposite direction. For this reason, we set $\alpha_{year} = 1$ without loss of generality. This will allow the other α_* parameters to have a natural interpretation: for example, if $\alpha_{gender} = 5$, it means that the colloquialization difference between male and female speakers is equal to around 5 years worth of change.

We use Stan (Carpenter et al., 2017; Stan Development Team, 2018) to fit the model to our data. For priors we use 0-centered Cauchy distributions:

$$\begin{aligned} \alpha_{gender}, \alpha_{power}, \alpha_{age}, \dots &\sim \text{Cauchy}(0, \sigma_1) \\ \beta_{pass}, \beta_{prog}, \beta_{mod}, \beta_{semi}, \beta_{pied} &\sim \text{Cauchy}(0, \sigma_2) \\ \gamma_{pass}, \gamma_{prog}, \gamma_{mod}, \gamma_{semi}, \gamma_{pied} &\sim \text{Cauchy}(0, \sigma_3) \end{aligned}$$

The values used for the standard deviations are $\sigma_1 = 3$ (corresponding to the expectation that a typical feature, such as gender, would affect colloquialization by an amount equivalent to around 3 years of change), $\sigma_2 = 0.01$ (corresponding to the expectation that the rate of a feature such as passivization would change 0.01 per year), and $\sigma_3 = 3$. To justify this choice of σ_3 , we need to mention another implementation detail: we shift the *year* feature so that *year* = 0 corresponds to 1985, the middle year of our range. Setting $\sigma_3 = 3$ then corresponds to the expectation that in 1985, a typical linguistic feature would be around $\ell(\pm 3)$, where ℓ is the logistic function. We chose the Cauchy distribution for its heavy tail.

8.1 Party Effects

To check for the colloquialization of various parties, we first group all conservative parties together,

as they share many parliamentary members who switch between them over time. As we are restricting to the years 1961–2010 for this analysis, the relevant parties are the modern Conservative party and its predecessors Progressive Conservative, Canadian Alliance, and Reform (Bauman and Kahana, 2006). Next, we discard the parties that have few parliamentary members, as they do not have enough speakers from which to draw meaningful conclusions. Over the relevant time span, the major parties were the Liberal Party, the Conservative Party, and the New Democratic Party (NDP); we therefore restrict our attention to these three.

Starting with the NDP, we fit the above model with the features *year* and *gender*, as well as an indicator feature for NDP membership. This shows a very large colloquialization boost for NDP membership: NDP members are more colloquial than non-NDP members by an equivalent of 7.5 years of change. We also extract posterior uncertainty intervals from the model, also known as credible intervals (CI); in this case, the 95% interval was 7.3–7.8. This effect is much larger than any difference between Liberals and Conservatives. Note that the NDP has never held power. In order to better focus on the difference between Liberals and Conservatives and on the effects of being in power, we remove the NDP and apply the remaining regressions to only Liberals versus Conservatives (and with *year* and *gender* as additional features).

The results of this analysis indicate that Liberals are more formal than Conservatives by an average of 3.7 years worth of colloquialization change, but there is a potential confounder here: the party in power tends to speak more formally than the party out of power, and the Liberals held power for the majority of the relevant time period (Blais, 2005). Adding a *power* feature, the Conservative-Liberal colloquialization gap shrinks from 3.7 years to 0.7 (95% CI: 0.5–1.0). The effect of holding power is much larger: the model indicates that the party in power speaks more formally than the party out of power by an equivalent of 5.9 years of colloquialization change (95% CI: 5.7–6.1). This result stays fairly robust with respect to adding additional features; for example, after adding *gender*, the Conservative-Liberal colloquialization gap changes to 0.8, and the *power* gap changes to 6.0. *Seniority* and *age* features also don't seem to have a large effect on these parameters.

After NDP membership, belonging to the govern-

ing party is the most significant predictor of colloquialization (when controlling for year). The party in power (the Liberals or Conservatives) speaks the most formally, followed by the other party (between Liberals and Conservatives), and then the NDP, whose members speak the most colloquially.

8.2 Gender Effects

To examine the effect of gender, we first run the model using only the features *year*, *gender*, and the product $gender \times year$ (plus the linguistic features as dependent variables: *pass*, *prog*, *mod*, *semi*, and *pied*). The goal of this is to investigate how the colloquialization of women in Parliament compares to men's, and how this changes over time.

The results of this run indicate that women speak more formally than men but that this effect diminishes over time: the coefficient of $gender \times year$, which can be interpreted as the rate at which the male and female colloquialization gap changes, was 0.07 (95% CI: 0.04 to 0.10), and the coefficient of *gender* was -2.8 (95% CI: -2.3 to -3.4). The interpretation of this is that at the beginning of our time period in 1961, women spoke more formally than men by around 4.5 years, meaning that a woman's speech would on average have colloquialization similar to a man's from 4.5 years prior; by the end of our time period in 2010, this gap shrank to 1 year. This might initially suggest that women started speaking less formally as they became more accepted in Parliament. However, this result disappears once we add in other parameters: with *power* added as a feature, the coefficient of $gender \times year$ decreases to 0.02, (95% CI: -0.02 to 0.05), meaning that most or all of the catch-up effect of women's colloquialization rates can be attributed to women belonging primarily to the party in power in the early time periods (1960s and 1970s). Adjusted for the effect of party in power, women's speeches appear to be only 1.6 years more formal than men's (95% CI: 1.0 to 2.1), and this gap remains relatively stable over time.

Adding other features, such as age, seniority, and party membership, changes the numbers slightly but does not affect the main story: women consistently speak more formally than men (equivalent to around 1.5 to 2 years of colloquialization change), and this gender colloquialization gap does not significantly change over time.

8.3 Age and Seniority Effects

We add features for age and seniority (number of years in Parliament) in order to see whether they give evidence of cohort effects. To simplify the analysis, we group members of Parliament into two age groups: the oldest third (56 years old or older) and the youngest third (47 years old or younger), omitting the people in the middle age range. We similarly categorize people into two groups by their seniority: veteran members (with 10 or more years since their first year in Parliament, corresponding to the top third after sorting the speeches by the number of years their speakers spent in Parliament before making that speech) and new members (who spent at most 4 years since their first year in Parliament, corresponding to the lowest third).

Year, power, party, and gender are included as features. Because power is a feature, we exclude NDP from this analysis. We then fit the model with the age and/or seniority indicators added. Since age and seniority are correlated, we should be careful when interpreting the results. Adding only age, the resulting coefficient indicates that the older third of Parliament speaks more formally by an equivalent of 0.4 years (95% CI: 0.2-0.7). This is trending in the expected direction, but it is a small effect.

Adding only seniority, the resulting coefficient indicates that the more senior third of Parliament speaks *less* formally by an equivalent of 2.3 years of colloquialization change (95% CI: 2.0-2.6). This is the opposite direction of what we expected: speakers newly entering Parliament are moving the average in the direction of being *less* colloquial while the overall colloquialization is increasing.

We also fit the analysis with both age and seniority features in place at the same time. The resulting coefficients indicate that the oldest third are more formal than the youngest third by an equivalent of 1.9 years of colloquialization change (95% CI: 1.6-2.2), and that the more senior third of Parliament speakers are less formal than the newest members by 3.2 years (95% CI: 2.9-3.5). Both of these are stronger effects than before. This indicates that new members' formality was hiding an age effect in the opposite direction, where younger members are slightly more colloquial but temporarily speak more formally when new to Parliament.

Note that the age effects we observe leave room for a small but non-zero cohort effect. However, the average age of the two age groups differs by around 19 years, which suggests that the age effect

(in which younger people speak more colloquially) only explains an increase of colloquialization of around 0.1 per year, out of the observed increase of 1 per year. 90% of the effect appears to come from non-cohort factors, primarily individual change.

9 Conclusion

We show that the English spoken in the Parliament of Canada has become more colloquial over many decades and that various features hypothesized to be related to colloquialization are correlated with each other at the level of individuals, indicating the existence of a latent colloquialization parameter.

We investigate several social variables as explanatory factors for the observed change in colloquialization. We find that cohort replacement cannot explain the change, that women did not lead this change, but that whether a member of Parliament belongs to the governing party has a large effect. The primary driver of the aggregate change appears to be an increase in the use of colloquial features in the speech styles of individuals.

Future directions The combination of a large dataset covering a long time span and information about individuals was important to our analysis but occurs infrequently in prior work. It would be interesting to identify other such sources of data, especially ones where social class could be extracted for investigation, as Labov (1990) has noted that social class is an important correlate of change.

Another direction for future work is to examine the same social variables for other kinds of syntactic change, as social correlates have most commonly been analyzed for sound changes. One could also compare syntactic change to other types of change, such as semantic shifts, and study whether the latter display a pattern that is different from the one observed here for changes in syntax.

References

René Arnaud. 1998. [The development of the progressive in 19th century English: A quantitative survey](#). *Language Variation and Change*, 10(2):123–152.

Alexander T. J. Barron, Jenny Huang, Rebecca L. Spang, and Simon DeDeo. 2018. [Individuals, institutions, and innovation in the debates of the French Revolution](#). *PNAS*, 115(18):4607–4612.

Richard W Bauman and Tsvi Kahana. 2006. *The least examined branch: The role of legislatures in the constitutional state*. Cambridge University Press.

Kaspar Beelen, Timothy Alberdingk Thijm, Christopher Cochrane, Kees Halvemaan, Graeme Hirst, Michael Kimmins, Sander Lijbrink, Maarten Marx, Nona Naderi, Ludovic Rheault, Roman Polyanovsky, and Tanya Whyte. 2017. [Digitization of the Canadian parliamentary debates](#). *Canadian Journal of Political Science/Revue canadienne de science politique*, 50(3):849–864.

Douglas Biber and Edward Finegan. 1989. [Drift and the evolution of English style: A history of three genres](#). *Language*, pages 487–517.

André Blais. 2005. [Accounting for the electoral success of the Liberal party in Canada presidential address to the Canadian political science association London, Ontario June 3, 2005](#). *Canadian Journal of Political Science/Revue canadienne de science politique*, 38(4):821–840.

Bob Carpenter, Daniel Lee, Marcus A Brubaker, Allen Riddell, Andrew Gelman, Ben Goodrich, Jiqiang Guo, Matt Hoffman, Michael Betancourt, and Peter Li. 2017. [Stan: A probabilistic programming language](#). *Journal of Statistical Software*, 76(1).

Peter Collins and Xinyue Yao. 2013. [Colloquial features in world Englishes](#). *International Journal of Corpus Linguistics*, 18(4):479–505.

Peter Collins and Xinyue Yao. 2014. [Grammatical change in the verb phrase in Australian English: A corpus-based study](#). *Australian Journal of Linguistics*, 34(4):506–523.

Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. [Echoes of power: Language effects and power differences in social interaction](#). In *Proceedings of the 21st international conference on World Wide Web*, pages 699–708.

Cristian Danescu-Niculescu-Mizil, Robert West, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. [No country for old members: User lifecycle and linguistic change in online communities](#). In *Proceedings of the 22nd international conference on World Wide Web*, pages 307–318.

Marco Del Tredici and Raquel Fernández. 2018. [The road to success: Assessing the fate of linguistic innovations in online communities](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1591–1603, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Melissa Haussman and Lori Turnbull. 2014. [Legislatures and parties: Heightened divisions since the 1990s](#). In David Martin Thomas and David Biette, editors, *Canada and the United States: Differences that count*, chapter 8, pages 163–185. University of Toronto Press.

- Turo Hiltunen, Jenni Räikkönen, and Jukka Tyrkkö. 2020. Investigating colloquialization in the British parliamentary record in the late 19th and early 20th century. *Language Sciences*.
- Haidee Kruger, Bertus van Rooy, and Adam Smith. 2019. Register change in the British and Australian Hansard (1901-2015). *Journal of English Linguistics*, 47(3):183–220.
- William Labov. 1990. The intersection of sex and social class in the course of linguistic change. *Language variation and change*, 2(2):205–254.
- William Labov. 2001. *Principles of linguistic change, Volume 2: Social factors*, volume 29. Blackwell.
- Geoffrey Leech. 2004. Recent grammatical change in English: data, description, theory. In *Advances in corpus linguistics*, pages 61–81. Brill Rodopi.
- Geoffrey Leech, Marianne Hundt, Christian Mair, and Nicholas Smith. 2009. *Change in contemporary English: A grammatical study*. Cambridge University Press.
- Christian Mair. 2006. *Twentieth-century English: History, variation and standardization*. Cambridge University Press.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- Stan Development Team. 2018. Pystan: the Python interface to Stan, version 2.17.1.0. <http://mc-stan.org>.
- Ian Stewart and Jacob Eisenstein. 2018. Making “fetch” happen: The influence of social and linguistic context on the success of lexical innovations. In *Proceedings of the 2018 conference on empirical methods in natural language processing (EMNLP)*, pages 4360–4370.
- Sali A Tagliamonte and Alexandra D’Arcy. 2007. The modals of obligation/necessity in Canadian perspective. *English World-Wide*, 28(1):47–87.