

Amano, Robert A.; Gosselin, Marc-André; Mc Donald-Guimond, Julien

Working Paper

Evolving temperature dynamics in Canada: Preliminary evidence based on 60 years of data

Bank of Canada Staff Working Paper, No. 2021-22

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Amano, Robert A.; Gosselin, Marc-André; Mc Donald-Guimond, Julien (2021) : Evolving temperature dynamics in Canada: Preliminary evidence based on 60 years of data, Bank of Canada Staff Working Paper, No. 2021-22, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2021-22>

This Version is available at:

<https://hdl.handle.net/10419/241245>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Evolving Temperature Dynamics in Canada: Preliminary Evidence Based on 60 Years of Data

by Robert Amano, Marc-Andre Gosselin and Julien Mc Donald-
Guimond

Canadian Economic Analysis Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

RAmano@BankofCanada.ca, MGosselin@BankofCanada.ca, JMcDonald-Guimond@BankofCanada.ca

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.



Acknowledgements

We thank Tatjana Dahlhaus, Glenn Rudebusch and our colleagues at the Bank of Canada for helpful comments and suggestions. We are also grateful to Lucie Vincent for her guidance with the daily temperature data.

Abstract

Recent discussions on climate change have led to an interest in its potential impact on economic phenomena and public policy. In this paper, we focus on one aspect of the climate change question by documenting the time-series properties of temperatures across Canada. In particular, we examine the evolving dynamics of daily average temperature and diurnal temperature range (the difference between the daily maximum and minimum temperatures at a given location) for select Canadian cities using data from the past 60 years. While rising mean temperature levels in Canada and elsewhere has been well documented, research exploring the other elements of temperature dynamics using modern econometric methods and rich model specifications are sparse. To fill in this gap, we extend the work of Diebold and Rudebusch (2019) and examine the evolution of daily temperature averages, volatility, seasonality and duration. This new evidence provides economists exploring issues related to climate change with a better understanding of the nature of Canadian temperature dynamics and their magnitudes.

Topics: topics: Climate change; Econometric and statistical methods

JEL codes: C22, Q54

1. Introduction

Research aimed at better understanding climate change dynamics has been gaining momentum (see, for example, Hsiang and Kopp 2018, Carney 2019, Bolton et al. 2020 and Diebold et al. 2020). Despite this increasing interest, surprisingly little work has analyzed the dynamics of climate metrics using modern time-series methods. We take up this task by exploring, in a systematic fashion, one margin of the climate change issue: The evolution of temperature dynamics across a range of Canadian locations. We do not attempt to isolate drivers behind changing temperature dynamics nor do we advance a specific model of temperature dynamics or have one in the background. At this relatively early stage of climate change economic research, we believe it is sufficient to provide a factual background of changing temperature dynamics from which theoretical climate models and climate simulation models can be developed and tested.

As a result, we follow the work reported in Diebold and Rudebusch (2019). Diebold and Rudebusch conduct a formal econometric analysis of the evolving dynamics of daily average temperatures and diurnal temperature ranges (the difference between the daily maximum and minimum temperatures at a given location) for 15 large cities across the United States using data from the 1960 to 2017 sample period. We conduct a similar examination for select large Canadian cities but also include two northern and relatively less populated areas: Eureka, Nunavut and Yellowknife, Northwest Territories. The latter areas allow us to explore whether the results presented in the current paper for more southern and populated Canadian cities and those reported in Diebold and Rudebusch are robust to more northern and less populated areas. We also extend their work by examining the evolution of the frequency and duration of so-called extreme temperature events. This allows us to shed light on a common narrative in the Canadian media regarding increasingly longer “heat waves” as well as provide researchers with empirical findings on an additional dimension of temperature dynamics in Canada.

While the rise of average temperatures has been well documented, our analysis differs from previous research along three margins. First, we use modern econometric methods, which allows us to perform valid hypothesis testing in the presence of autocorrelated or heteroscedastic regression residual terms. We find that “typical” temperature regressions admit such non-spherical residuals and earlier work did not properly account for them in their inferential procedure. As demonstrated in den Haan and Levin (1997), this type of omission can lead to gross differences between nominal and empirical test sizes. Second, we use richer temperature dynamics specifications, which may help to sharpen our parameter estimates. Third, the effects of fixed and time-varying seasonality, a potentially important element of

temperature dynamics, are examined. This analysis should improve our understanding of the dynamics of an important aspect of climate change, which, in turn, should help inform weather modelling and, eventually, public policy. Indeed, Dietz et al. (2020) emphasize the importance of properly modelling climate change dynamics for providing informed recommendations for climate change public policy.

The remainder of the paper is as follows. The next section describes the temperature data used in this study. Section 3 demonstrates our empirical strategy using one city, Toronto, as an example. Section 4 expands the analysis described in the previous section to our full sample of locations. The penultimate section examines the evolution of “extreme” temperature events since 1960. The final section offers concluding remarks and directions for future research.

2. Data

This paper leverages daily surface air temperature data from the third generation of the Adjusted and Homogenized Canadian Climate (AHCC) database. The information collected includes minimal, maximal and average daily temperatures measured in Celsius (C). This dataset has many desirable properties, such as its spatial and temporal coverage; but more importantly, its data have been adjusted for several types of non-climatic breaks. These include corrections for station relocation, changes in observing practices and automation, among others. According to the dataset’s authors, Vincent et al. (2018), the AHCC database “provides the best data sets for temperature trends in Canada.”

We focus the analysis on large metropolitan areas that provide us with broad regional coverage across Canada over 1960–2019.¹ Concentrating on these areas allows us to study temperature dynamics with near complete data for regions with an important share of both population and economic activity. Importantly, the majority of the data was collected at airport weather stations, where temperature could differ from that of the cities themselves, an issue to which we will come back later. More specifically, the analysis focuses on temperature data from 11 large metropolitan areas, *viz.*, Victoria, Vancouver, Edmonton, Calgary, Winnipeg, Toronto, Ottawa, Montreal, Quebec, Moncton and St. John’s. As well, we include Eureka, Nunavut and Yellowknife, Northwest Territories, in our sample to examine whether our results for more populated Canadian cities and those reported in Diebold and Rudebusch (2019) apply to

¹ Missing values are imputed by taking an exponentially weighted average of the same month-day pairs in the year preceding and following the missing observation.

more northern and less populated areas.² Indeed, Eureka lies north of the Arctic Circle (79.9 N latitude versus 66.5 N) and had a population of eight according to the most recent data. To provide a sense of distance, Eureka is about 3,844 kilometres further north (holding the longitude constant) than the most northern U.S. city in Diebold and Rudebusch (*viz.*, Portland, Oregon) and 2,956 kilometres from the most northern large Canadian city in our sample, Edmonton, Alberta.

For the purposes of the current analysis, we create two statistics for each city: (i) daily average temperature or $AVG = (MAX + MIN) / 2$, where MAX is the maximum temperature during a particular day and MIN is the minimum temperature during the same day; and (ii) diurnal temperature range $DTR = MAX - MIN$, which is a natural measure of intraday volatility or variability. According to Diebold and Rudebusch (2019), DTR is an intuitive estimator of daily volatility and, as well, a highly efficient statistic for estimating underlying quadratic variation from discretely sampled data.

3. Preliminary Data Analysis

3.1 Graphs and Trends

In order to provide an initial sense of the data, we follow Diebold and Rudebusch (2019) and present some simple statistics and graphs based on temperature data from one city. In our case, we use Toronto, Canada's most populated city, as our base case. We expand the analysis to the full sample of locations in the next section, but a focus on one city helps to lay out our analytical approach.

In Figure 1, we plot the unconditional distributions of AVG and DTR . As in Diebold and Rudebusch, AVG admits a bimodal distribution reflecting the strong seasonality patterns found in Canada. In particular, the winter mode is around 2°C, while the summer mode is about 21°C. The DTR distribution, in contrast, is unimodal and approximately symmetric, and centred around 7°C with much less dispersion than the AVG density.

Next, we explore the temporal evolution of AVG and DTR by regressing them on a series of deterministic variables via ordinary least squares. Modelling processes as complex as average temperatures and diurnal temperature range with a simple constant and trend will almost ensure that the residuals will be serially

² An obvious omission from the list of northern and less populated areas is Whitehorse, Yukon. Unfortunately, consistent temperature data for Whitehorse is difficult to obtain, with, in one case, 651 consecutive missing observations.

correlated, autoregressively heteroscedastic or both. As a result, in this section we merely present the modelling approach and defer the important question of statistical significance to section 4.

Figure 2 presents plots of *AVG* and *DTR* along with their fitted trend line from a simple regression of the form:

$$y_t = \alpha + \beta \cdot T_t + v_t, \quad (1)$$

where y_t is either *AVG* or *DTR*, α and β are parameters to be estimated, T_t is a linear time trend, and v_t is an error term.³

When y_t is set equal to *AVG*, β is positive, supporting the result that average temperatures have been increasing over time (or an increasing trend line in the left panel of Figure 2). The parameter estimate implies that the average temperature in Toronto has increased by about 1.5°C over the past 60 years, which is greater than the average increase observed globally over the same period (see Rudebusch 2019). As noted in Diebold and Rudebusch, there are two reasons for greater warming of cities: (i) the increasing urban heat island effect; and (ii) the fact that average land temperatures generally grow faster than the global average, which includes ocean areas that are slower to warm.⁴

Setting y_t equal to *DTR*, the estimated slope term in the above regression is negative. This result indicates that the daily range of temperatures in Toronto has been narrowing (represented by a decreasing trend line in the right panel of Figure 2). Specifically, *DTR* has fallen by about 0.3°C over the sample period with the *MIN* rising faster than the *MAX* and, thereby, narrowing the spread between the two measures.

Thus far, the picture for temperatures in Toronto is a diurnal asymmetry characterized by increasing average temperatures and declining intraday temperature ranges. The latter is owing to nighttime temperatures increasing faster than daytime ones.

³ We acknowledge that a linear trend is a simple approach to modelling changes over time, but we believe it is a reasonable approximation to capture general changes in a variable's dynamics over a 60-year horizon. Accordingly, we do not attempt to forecast the future path of any of the variables under consideration.

⁴ To obtain a sense of impact of the urban heat island effect, we tested the equality of trends between the cities of Toronto and Burlington. Burlington is located approximately 50 kilometres southeast of Toronto, also next to Lake Ontario. Despite Toronto being more than 14 times more populous than Burlington, the null hypothesis of equal trends could not be rejected.

3.2 Seasonality

In this subsection, we continue working with temperature data from Toronto to consider another important feature of temperature dynamics—seasonality. We examine both fixed and time-varying seasonality and start the analysis with the former. More specifically, the detrended residuals from the previous *AVG* and *DTR* specifications, \hat{v}_t , are regressed, respectively, on 12 monthly dummy variables:

$$\hat{v}_t = \sum_{i=1}^{12} \gamma_i D_{it} + u_t, \quad (2)$$

where D_{it} is equal to one if day t resides in month i , otherwise it is equal to zero; the γ 's are parameters to be estimated; and u_t is an error term. The regression results point to the existence of seasonal effects for both *AVG* and *DTR* and explain about 81 percent of the variation in detrended *AVG* and 14 percent in detrended *DTR*. Figure 3 plots the fitted and actual values from equation (2) and offers a visual means to glean these results. That is, the figure shows a pronounced seasonality pattern in *AVG* (left panel) and the inability of monthly dummy variables to capture a lot of variation in *DTR* (right panel).

Figure 4 provides another view of the fixed seasonal patterns in the data. In Figure 4, the 12 estimated monthly seasonal parameters from detrended *AVG* and *DTR* regressions are plotted in the left and right panels, respectively. The seasonal pattern for *AVG* is as expected: high in the summer months and low in the winter period. In particular, estimated seasonality reaches its maximum in July for *AVG* and its trough in January. The variation in seasonality for *DTR* is much less than that of *AVG* but also culminates over the summer months.⁵

Since it is possible that seasonal patterns have also been evolving over the sample period, we explore this possibility in the remainder of this subsection. To this end, we augment equation (2) with 12 monthly dummy variables interacted with time. That is:

$$\hat{v}_t = \sum_{i=1}^{12} \delta_i D_{it} + \sum_{i=1}^{12} \lambda_i D_{it} T_t + e_t, \quad (3)$$

where δ and λ are coefficients to be estimated and e_t is an error term. The interaction term, $D_{it} T_t$, should allow us to capture any linearly evolving changes in the monthly seasonal pattern found above. We present the results graphically in Figure 5. The figure displays the estimated seasonal parameters for the first (1960) and last (2019) years of our sample. Overall, we see that, whereas the broad seasonality pattern remains the same, there are changes in the magnitude of some monthly seasonality effects. In

⁵ Interestingly, the seasonality of *DTR* in Toronto does not appear to be clearly bimodal, as is the case for most other cities in our sample as well as for most locations studied by Diebold and Rudebusch (2019).

particular, the estimated seasonal effects on average daily temperatures seem to have declined in October and November while, for *DTR*, seasonality appears to have evolved with decreases in June, October and December and increases in March and November.

3.3 Extreme Temperature Events

An idea that is often mentioned in the popular media and elsewhere is that “heat waves” have become more common events. In this subsection, we explore a more general question of whether the frequency and duration of “extreme temperature” events in Toronto has evolved over time and, if so, we estimate the extent of the change. To operationalize the idea of “extreme temperature events,” we follow Vincent et al. (2018) and define the following five categories: (i) summer days = days with $MAX > 25^{\circ}C$; (ii) hot days = days with $MAX > 30^{\circ}C$; (iii) hot nights = days with $MIN > 22^{\circ}C$; (iv) frost days = days with $MIN \leq 0^{\circ}C$; and (v) ice days = days with $MAX \leq 0^{\circ}C$. Applicable daily *MIN* and *MAX* temperature observations for Toronto are then placed into these categories.

Looking across these categories over the sample period, it is readily apparent that there have been large shifts in the frequency of these extreme temperature events. For instance, the decadal average for summer and hot days was about 57 and 10 in the 1960s and jumped to average about 72 and 14.5 days for the 2010s. Similar shifts are seen for colder temperatures: The 1960s decadal average for frost and ice days was about 121 and 53 days and fell to 104 and 44 days, on average, for the 2010s.

To obtain more formal evidence of evolution of extreme temperature events, the annual frequency in each category is regressed on a constant and linear trend. The estimated coefficient associated with the trend terms then provides us with a sense of the degree of change in frequency of extreme temperature events. The results are presented in the upper panel of Table 1 and support the casual observations from the previous paragraph. In particular, summer and hot days have increased by about 18.5 and 7 days, respectively, since 1960, and frost and ice days have dropped by almost 23 and 16 days, respectively.

Another, perhaps, more interesting dimension of extreme temperatures is the duration of these events since it helps to formalize the notion of longer “heat waves.” To begin, a simple examination of decadal averages suggests the consecutive hot and summer temperature days are becoming longer. In the 1960s, for example, the average duration of summer and hot days was 3 and 1.5 days which, in the 2010s, increased to 4.2 and 1.9 days. Again, to provide more formal evidence of changes in the duration of extreme temperature events, the annual average durations of each type of temperature event is regressed on a constant and a trend. These results can be found in the lower panel of Table 1. The results

suggest that since 1960, the durations of summer and hot days have increased by about 1.2 and 0.3 days. Moreover, durations of frost and ice days in Toronto have declined by almost 1.7 and 0.8 days.

To summarize the results from this section, temperature dynamics for Toronto can be characterized by rising average temperatures, narrowing intraday temperature ranges, and evolving seasonal effects. As well, extreme temperature events (as defined above) have become more frequent and their durations longer lasting. In the next section, the analysis is expanded to the full sample, which includes 11 large Canadian cities and two northern and less populated areas.

4. Full-Sample Data Analysis

In this section, we continue to follow DR and estimate specifications for both mean and variance dynamics that control for trend, time-varying seasonality and serial correlation. Thus, in contrast to the previous section, where trends and seasonality were estimated sequentially, the empirical approach in this section starts immediately with the specification that estimates all these effects jointly. The analysis is conducted with data from 11 large metropolitan areas, *viz.*, Victoria, Vancouver, Edmonton, Calgary, Winnipeg, Toronto, Ottawa, Montreal, Quebec, Moncton and St. John's, along with two more northern and less populated areas, Eureka and Yellowknife.

As mentioned in section 3, modelling processes as complex as average temperatures and diurnal temperature range will almost ensure that the residuals will be serially correlated, autoregressively heteroscedastic or both. Therefore, the standard errors of all following regressions are estimated using the Newey and West (1987) heteroscedasticity and autocorrelation consistent (HAC) estimator, which allows us to conduct valid statistical inference even in the presence of non-spherical residuals.⁶ As well, given the Monte Carlo results reported in den Haan and Levin (1997), we re-estimate each regression's statistics using the Andrews (1991) quadratic spectral HAC estimator augmented with the Andrews and Monahan (1992) pre-whitening filter and automatic bandwidth selection procedure to examine the robustness of our results.⁷

⁶ For all regressions in this analysis, the Newey and West estimator uses a Bartlett kernel and a truncation parameter selection approach as recommended in Newey and West (1994).

⁷ In addition, every linear trend was also estimated using the non-parametric Sen's slope estimator (Sen 1968). These results together with those from the QS HAC estimator are not reported in this paper since the conclusions are robust to their application.

4.1 Conditional Mean Dynamics

To examine the conditional mean dynamics for the 11 cities and two northern areas, we estimate the following specification for each city or area:

$$y_t = \alpha + \beta T_t + \sum_{i=1}^{11} \delta_i D_{it} + \sum_{i=1}^{11} \lambda_i D_{it} T_t + \rho y_{t-1} + \varepsilon_t, \quad (4)$$

where y_t is either *AVG* or *DTR*; α, β, δ 's, λ 's and ρ are parameters to be estimated; and T_t is a linear time trend. The dummy variable, D_{it} , is defined slightly differently than above owing to the presence of the constant term. The variable is a monthly seasonal dummy series where July has been omitted, and which equals to one if day t is in month i and zero otherwise. The interaction term, $D_{it} T_t$, is intended to capture any potential time variations in seasonal factors. In effect, the constant and non-interacted trend terms reflect the month of July, and the seasonal dummies and interacted trends reflect marginal effects for the 11 other months relative to July. Finally, ε_t is a disturbance term, and y_t is a lagged dependent variable included to provide us with an indication of the degree of autocorrelation in *AVG* and *DTR* after controlling for the effects of trend and seasonal components.

The estimation results for *AVG* for each city are presented in the upper panel of Table 2, while the comparable results for Eureka and Yellowknife are displayed in the lower panel. Table 3 reports analogous statistics for *DTR*. For both tables, the first (left-most) column presents the direction and magnitude of the trend movement over the whole sample based on a simple regression on a constant and a trend. The statistical significance of this trend is assessed via a Wald test that all trend terms in equation (4) are jointly equal to zero, for which the p-values are reported in column $p(nt)$.⁸ The reason for this two-step procedure is that while a simple model can provide a reasonable approximation of a variable's trend over such a long period of time, its specification is arguably not rich enough to provide precise estimates because of the large amount of noise left. A joint significance test of all trend terms in a richer model is thus better suited to determine the significance of an overall trend in the data. The third and fourth columns report p-values for Wald tests corresponding to the null hypotheses of no seasonality $p(ns)$ and no time-varying seasonality $p(nts)$, respectively. The former jointly tests the significance of the seasonal dummies, the idea being that if the marginal effects for those months are not different from zero, then there is no seasonality in the sense that all months share the same average, represented by the constant. Similarly, testing for no time-varying seasonality involves the joint significance test of the interacted trend

⁸ A single asterisk indicates statistical significance at, at least, the 5 percent level.

terms, where a non-rejection of the null hypothesis implies that all months share the same linear trend, and thus that their relation to each other has remained constant through time. The penultimate column shows the coefficient on the lagged dependent variable, and the last column reports the coefficient of determination as an indicator of the goodness-of-fit of each regression specification. Finally, an estimate of the average trend in southern locations is computed using the mean group estimator (MGE) of Pesaran and Smith (1995).

Turning to Table 2, it is readily apparent that the parameter estimates associated with the trend term are all positive and statistically significant. For the whole sample, the MGE reports an average warming of 1.47°C, which is also statistically significant.⁹ An especially notable result is that the increase in average temperature in the more northern climates represented by Eureka and Yellowknife (lower panel) is almost twice as large as the average increase of the 11 cities, consistent with previous studies (e.g., IPCC 2018 and Vincent et al. 2018). Indeed, the point estimate for Eureka indicates that its average temperature has been growing at more than three times the global warming rate of about 1°C since 1960 (see Rudebusch 2019). In other words, our results imply that average temperatures in northern Canadian climates are increasing faster than southern regional temperatures and this conclusion is consistent with the evidence suggesting the ice sheets are melting much faster than previously anticipated based on data from the 1990s (see Shepherd et al. 2019).¹⁰ Fixed seasonality in *AVG* is a common feature across every location with all the *p(ns)*'s being less than 0.01. In contrast, there is less evidence of evolving seasonality with a median *p*-value for southern cities over 10 percent.

Table 3 presents evidence supporting the idea, found for Toronto, that the intraday temperature range has been narrowing in some cities. Yet, while a majority of cities have seen significant declines in *DTR*, others have experienced important increases, such as Edmonton with a cumulative trend of 1.18 degrees. Indeed, Davy et al. (2017) find that changes in *DTR* trends are not homogeneous across locations and are specific to geographical areas. The MGE trend for *DTR* is about -0.1 but is not significant, presumably

⁹ We checked again for the impact of urban heat island effects by testing the equality of trends between Montreal (for which data are recorded at the Trudeau airport) and the Saint-Hubert airport, which is located on the south shore of Montreal, but we could not reject the null hypothesis. This result, together with the one contrasting Toronto and Burlington, suggests that urban heat island effects might not be playing a large role when assessing long-run temperature trends.

¹⁰ A simple exercise where temperatures are averaged across Eureka and Yellowknife and the more southern cities separately shows that the trends for each of those areas are statistically different at the 5 percent level.

because of the relatively large span of results (-0.69 for Montreal and 1.18 for Edmonton). Consistent with the results for *AVG*, we find statistically significant evidence of fixed seasonality in *DTR* for all locations. In contrast to *AVG*, however, indications of time-varying seasonality are discernable at the 5 percent level of significance or lower. Notably, Victoria, Calgary, Winnipeg, Toronto, Ottawa, Montreal, St. John's, Eureka, and Yellowknife show statistically significant signs of evolving seasonality.

4.2 Conditional Variance Dynamics

As argued earlier, weather data typically exhibit serial correlation and other time-related forms of heteroskedasticity. We thus follow Diebold and Rudebusch and test for the presence of heteroskedasticity in our data by modelling the conditional variance of *AVG* and *DTR*. Specifically, we estimate the following model for each city or area:

$$\varepsilon_t^2 = \alpha + \beta T_t + \sum_{i=1}^{11} \delta_i D_{it} + \sum_{i=1}^{11} \lambda_i D_{it} T_t + \rho \varepsilon_{t-1}^2 + v_t, \quad (5)$$

where ε_t^2 is the squared residuals from the *AVG* or *DTR* conditional mean equations, and the other terms are analogous to those of the conditional mean equation.

The ε^2 's represent the variance of the data conditional on the trend and seasonality. Because the latter are deterministic components, the ε^2 's can be interpreted as capturing the volatility *beyond* the systematic fluctuations induced by the seasons and the linear trend. In contrast to *DTR*, which captures *intraday* volatility, the conditional variance equations model *interday* volatility. If the data are homoscedastic, the coefficients from the conditional variance equation should be jointly statistically indiscernible from zero. Otherwise, it would indicate that we have heteroskedastic shocks.

The estimation results for *AVG* for each city are presented in the upper panel of Table 4, and the comparable results for Eureka and Yellowknife are displayed in the lower panel. Table 5 reports analogous statistics for *DTR*. Each column is defined as in Tables 2 and 3.

Trend estimates for the conditional variances are statistically significant and are generally negative across locations for both *AVG* and *DTR*. Evidence of declining volatility is particularly strong in the case of average temperatures, with a significant MGE of -0.09 degrees. On the other hand, the picture for *DTR* is less clear, with some locations having important positive trends, such as Victoria and Edmonton. In our northern locations, however, volatility seems to have declined for both *AVG* and *DTR*. Seasonal patterns are also found to be highly significant, further supporting the presence of heteroskedasticity. Similar to the conditional mean equation, the evidence of trending seasonality is relatively weak. Worth noting is

the general significance of the lagged dependent variable, which suggests the existence of ARCH effects. We defer the examination of the latter to future research.

So far, the conditional variance dynamics indicate that average temperature volatility has actually been declining over the past 60 years. This is consistent with the estimated decline in intraday volatility across most locations captured by the conditional mean results for *DTR*, and it contrasts with the general impression that temperature swings are more frequent. The next two sections provide insights into how the conditional variance results can be reconciled with this narrative.

4.3 Frequency and Duration of Extreme Temperatures

The decline in temperature volatility estimated in the previous sections may appear counterintuitive given the numerous warnings of increasing extreme weather events over time. However, it is important to determine whether temperature change is assessed relative to a reference climate or an evolving climate. Indeed, results from the conditional variance analysis merely indicate that around the trending, constantly evolving, temperature distribution volatility has decreased. Yet, climate change is typically defined from the point of view of a reference climate, usually the one prevailing over the 1850–1900 period (Vincent et al. 2018). In that sense, extreme events may have increased despite the fall in volatility.

Figure 6 illustrates this using a stylized climate distribution changing over time. Clearly, an increase in average temperatures, as the one documented above, implies a decrease in the likelihood of cold events and an increase in that of hot events. On the other hand, a fall in volatility implies declines in the likelihood of both cold *and* hot temperatures. For cold events, the result is unequivocal: a decline in frequency. Hot events, however, may become more or less likely depending on which one of the mean effect or variance effect dominates. To shed light on this issue, we expand the analysis in subsection 3.3 and examine whether the frequency and duration of extreme temperature events has changed over the sample period. We use the same extreme temperature definitions as the subsection. To reiterate, we follow Vincent et al. (2018) and define the following five categories: (i) summer days = days with $MAX > 25^{\circ}C$; (ii) hot days = days with $MAX > 30^{\circ}C$; (iii) hot nights = days with $MIN > 22^{\circ}C$; (iv) frost days = days with $MIN \leq 0^{\circ}C$; and (v) ice days = days with $MAX \leq 0^{\circ}C$. Applicable daily *MIN* and *MAX* temperature observations for each location are then placed into these categories.

Table 6 displays the trend parameter estimates of specifications that regress the annual frequency of different types of extreme temperature events on a constant and a linear trend for each location under consideration. Looking across the table, we see some general patterns. Temperatures consistent with

summer days are becoming more frequent events with an MGE increase of 11 days since 1960. The effect is especially notable in Toronto, Ottawa and Montreal. The frequency of hot days is also rising with statistically significant increases ranging from more than 1 day in Victoria to about 7 days in Toronto. Consistent with the idea of global warming, the frequency of frost and ice days is estimated to be declining. The location-specific effects are, on the whole, statistically significant with an MGE decrease of 15 days for frost days, and 11 days for ice days. The 18-day decline in frost days in Eureka stands out as its location is more than 1,000 kilometres north of the Arctic Circle.

Estimates of changing extreme temperature duration are presented in Table 7. The estimates are based on equations regressing the annual average durations of each type of temperature event on a constant and a trend. The statistically significant summer trend estimates suggest that the average duration has increased by about 0.4 days. There is no evidence, however, of changes in the duration of hot days. On the other hand, the average duration of frost and ice spells appears to have fallen on average across locations. Interestingly, durations in the northern areas do not seem to have changed over the estimation sample, despite some changes in frequency.

One pattern that emerges from the extreme events analysis is that the typical change in cold events seems to be greater, in absolute value, than that of warm events. The reduction in the frequency of cold events is more widespread across the country than the increase in the frequency of warm events, and the absolute magnitude is larger. The change in average duration is also more important for cold events, especially for ice days.

This feature provides a first possible explanation as to why the conditional variances might be declining despite the frequent references to higher weather volatility. Two effects seem to be at play in Canada: extreme cold temperatures have weakened, in terms of both frequency and duration, while extreme warm events have become more common and persistent. However, the evidence suggests that the latter effect might be of lesser importance than the former.¹¹ If this is the case, then one could observe a rise in extreme warm events together with an *overall* fall in (average) volatility across all seasons. This would be consistent with the fact that warming has particularly affected winters in Canada (Bush et al. 2019).

¹¹ This is also supported by an examination of the distribution of *AVG* over time. The probability of observing a summer or a hot day has increased by less, in absolute value, than the decline in the probability of observing a frost or ice day. The fall in the 5th percentile of the distribution is also more important than the increase in the 95th percentile.

5. Concluding Remarks

Climate change is a multidimensional problem. In this paper, we examine a small subset of the problem by studying the evolution of temperature dynamics in Canada along a number of different margins. Our contribution to the climate literature is to bring modern econometric methods and richer temperature dynamics specifications to this issue. We find evidence to confirm the well-known result that average temperatures have been increasing over time, and that the daily average temperature in more northern areas such as Eureka has been increasing at much faster rates than the Canadian city or global average. As well, the estimation results indicate changes in the frequency and duration of extreme temperature events. Finally, we document the presence of important changes in fixed and time-varying seasonality not examined in earlier research.

In addition to documenting interesting temperature dynamics for Canadian locations, we believe we have estimated statistics to help develop theoretical models of climate change as well as condition climate simulation models that are used for public policy advice. For instance, when providing policy advice, these models will need to account for or incorporate the result of changing seasonal temperature patterns. As well, the modelling of temperature distributions in climate models could be enriched by the addition of more complex patterns in the evolution of their higher moments.

Looking ahead, we plan on conducting a statistical analysis exploring whether precipitation patterns have also been evolving over time. In addition, we intend on measuring the impact of weather-related events, such as extreme temperatures and heavy precipitation, on economic activity.

6. References

- Andrews, D.W.K. (1991) "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica* 59:817–58.
- Andrews, D.W.K., and J.C. Monahan (1992) "An Improved Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimator." *Econometrica* 60:953–66.
- Bolton, P., M. Despres, L.A. Pereira da Silva, F. Samama, and R. Svartzman (2020) "The Green Swan: Central Banking and Financial Stability in the Age of Climate Change." Bank for International Settlements (<https://www.bis.org/publ/othp31.pdf>).
- Carney, M. (2019) "Remarks Given during the UN Secretary General's Climate Action Summit 2019." (<https://www.bankofengland.co.uk/-/media/boe/files/speech/2019/remarks-given-during-the-un-secretary-generals-climate-actions-summit-2019-mark-carney.pdf>)
- Davy, R., I. Esau, A. Chernokulsky, S. Outten, and S. Zilitinkevich (2017) "Diurnal Asymmetry to the Observed Global Warming." *International Journal of Climatology* 37:79–93.
- den Haan, W.J., and A.T. Levin (1997) "A Practitioner's Guide to Robust Covariance Matrix Estimation." In G.S. Maddala and C.R. Rao (Eds.), *Handbook of Statistics: Robust Inference* 299-342. Elsevier.
- Diebold, F.X., P. Goulet Coulombe, M. Goebel, G.D. Rudebusch and B. Zhang (2020) "Optimal Combination of Arctic Sea Ice Extent Measures: A Dynamic Factor Modeling Approach." University of Pennsylvania (<https://arxiv.org/pdf/2003.14276.pdf>).
- Diebold, F.X. and G.D. Rudebusch (2019) "On the Evolution of U.S. Temperature Dynamics." University of Pennsylvania (<https://arxiv.org/pdf/1907.06303.pdf>).
- Dietz, S., F. van der Ploeg, A. Rezai, and F. Vermans (2020) "Are Economists Getting Climate Dynamics Right and Does it Matter?" London School of Economics and Political Science.
- Hsiang, S., and R.E. Kopp (2018) "An Economist's Guide to Climate Change Science." *Journal of Economic Perspectives* 32:3–32.
- Jahn, M. (2015) "Economics of Extreme Weather Events: Terminology and Regional Impacts Models." *Weather and Climate Extremes* 10:29-39
- Newey, W.K., and K.D. West (1987) "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55:703–708.
- Newey, W.K., and K.D. West (1994) "Automatic Lag Selection in Covariance Matrix Estimation." *Review of Economic Studies* 61:631–53.
- Pesaran, M.H., and R. Smith. (1995) "Estimating Long-Run Relationships from Dynamic Heterogeneous Panels." *Journal of econometrics* 68(1):79–113.
- Rudebusch, G.D. (2019) "Climate Change and the Federal Reserve." *FRBSF Economic Letter* 2019-09.
- Sen, P.K. (1968) "Estimates of the Regression Coefficient Based on Kendall's Tau." *Journal of the American Statistical Association* 63(234):1379–89.

Shepherd, A., E. Ivins, E. Rignot, B. Smith, M. van den Broeke, I. Velicogna, P. Whitehouse, K. Briggs, I. Joughin, G. Krinner, S. Nowicki, T. Payne, T. Scambos, N. Schlegel, A. Geruo, C. Agosta, A. Ahlstrøm, G. Babonis, V.R. Barletta, A.A. Bjørk, A. Blazquez, J. Bonin, W. Colgan, B. Csatho, R. Cullather, M.E. Engdahl, D. Felikson, X. Fettweis, R. Forsberg, A.E. Hogg, H. Gallee, A. Gardner, L. Gilbert, N. Gourmelen, A. Groh, B. Gunter, E. Hanna, C. Harig, V. Helm, A. Horvath, M. Horwath, S. Khan, K.K. Kjeldsen, H. Konrad, P.L. Langen, B. Lecavalier, B. Loomis, S. Luthcke, M. McMillan, D. Melini, S. Mernild, Y. Mohajerani, P. Moore, R. Mottram, J. Mouginot, G. Moyano, A. Muir, T. Nagler, G. Nield, J. Nilsson, B. Noël, I. Ootosaka, M.E. Pattle, W.R. Peltier, N. Pie, R. Rietbroek, H. Rott, L. Sandberg Sørensen, I. Sasgen, H. Save, B. Scheuchl, E. Schrama, L. Schröder, K-W. Seo, S.B. Simonsen, T. Slater, G. Spada, T. Sutterley, M. Talpe, L. Tarasov, W.J. van de Berg, W. van der Wal, M. van Wessem, B. Dutt Vishwakarma, D. Wiese, D. Wilton, T. Wagner, B. Wouters, and J. Wuite (2019) "Mass Balance of the Greenland Ice Sheet from 1992 to 2018." *Nature* 579:233–39. <https://www.nature.com/articles/s41586-019-1855-2>

Vincent, L.A., X. Zhang, E. Mekis, H. Wan, and E.J. Bush (2018) "Changes in Canada's Climate: Trends in Indices Based on Daily Temperature and Precipitation Data." *Atmosphere-Ocean* 56:332–49.

Appendix A – Figures

Figure 1: Unconditional distributions of AVG and DTR for Toronto




Figure 2: Temperature trends for Toronto




Figure 3: Temperature seasonality for Toronto




Figure 4: Seasonal factors for Toronto

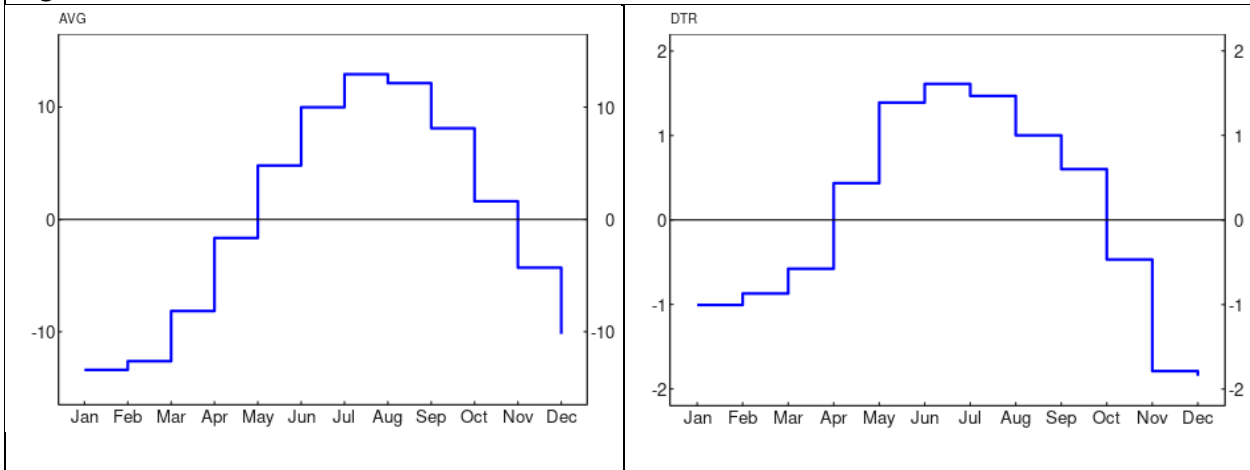


Figure 5: Evolving seasonal factors for Toronto






Figure 6: Stylized climate distribution



Source: Jahn (2015)

Figure 7: Selected locations



Appendix B – Tables

Table 1: Trends in extreme temperatures for Toronto

Summer days	Hot days	Hot nights	Frost days	Ice days
<i>Trends in the number of days</i>				
18.48	7.06	2.20	-22.32	-15.98
<i>Trends in the average duration</i>				
1.21	0.32	-0.13	-1.65	-0.82

Table 2: AVG, conditional mean dynamics

Location	Province	$\Delta trend$	$p(nt)$	$p(ns)$	$p(nts)$	p	R^2
South							
Victoria	BC	1.40*	0.00	0.00	0.00	0.70*	0.88
Vancouver	BC	1.36*	0.00	0.00	0.09	0.76*	0.91
Calgary	AB	1.16*	0.05	0.00	0.08	0.79*	0.87
Edmonton	AB	1.16*	0.04	0.00	0.04	0.79*	0.91
Winnipeg	MB	1.52*	0.02	0.00	0.11	0.76*	0.92
Toronto	ON	1.51*	0.00	0.00	0.84	0.72*	0.91
Ottawa	ON	1.89*	0.00	0.00	0.67	0.71*	0.91
Montreal	QC	2.11*	0.00	0.00	0.41	0.70*	0.91
Quebec	QC	1.27*	0.00	0.00	0.39	0.68*	0.91
Moncton	NB	1.56*	0.00	0.00	0.36	0.63*	0.88
St John's	NL	1.20*	0.00	0.00	0.19	0.61*	0.85
MGE		1.47	0.00				
North							
Yellowknife	NT	2.78*	0.00	0.00	0.05	0.81*	0.95
Eureka	NU	2.98*	0.00	0.00	0.13	0.85*	0.97

Note:

All results are based on daily data, 1960–2019.

Column 3 reports the estimated trend movement over the 60-year sample in degrees Celsius.

$p(nt)$ is the robust p -value for a Wald test of no trend. $p(ns)$ is the robust p -value for a Wald test of no seasonality. $p(nts)$ is the robust p -value for Wald a test of no trend in seasonality. p is the estimated autoregressive coefficient.

R^2 is the adjusted coefficient of determination.

Asterisks denote significance at the 5 percent level.

Table 3: DTR, conditional mean dynamics

Location	Prov	$\Delta trend$	$p(nt)$	$p(ns)$	$p(nts)$	p	R^2
South							
Victoria	BC	-0.46*	0.00	0.00	0.00	0.42*	0.42
Vancouver	BC	-0.67*	0.00	0.00	0.15	0.35*	0.29
Calgary	AB	-0.30*	0.01	0.00	0.01	0.28*	0.13
Edmonton	AB	1.18*	0.00	0.00	0.13	0.35*	0.21
Winnipeg	MB	-0.20*	0.01	0.00	0.02	0.27*	0.20
Toronto	ON	-0.30*	0.00	0.00	0.05	0.22*	0.18
Ottawa	ON	-0.15*	0.00	0.00	0.00	0.23*	0.15
Montreal	QC	-0.69*	0.00	0.00	0.00	0.24*	0.15
Quebec	QC	0.13	0.22	0.00	0.22	0.28*	0.18
Moncton	NB	-0.19	0.17	0.00	0.19	0.25*	0.14
St John's	NL	0.24*	0.00	0.00	0.01	0.27*	0.17
MGE		-0.13	0.41				
North							
Yellowknife	NT	-0.35*	0.01	0.00	0.05	0.23*	0.28
Eureka	NU	-0.09*	0.00	0.00	0.00	0.27*	0.16

Note:

All results are based on daily data, 1960–2019.

Column 3 reports the estimated trend movement over the 60-year sample in degrees Celsius. $p(nt)$ is the robust p -value for a Wald test of no trend. $p(ns)$ is the robust p -value for a Wald test of no seasonality. $p(nts)$ is the robust p -value for Wald a test of no trend in seasonality. p is the estimated autoregressive coefficient.

R^2 is the adjusted coefficient of determination.

Asterisks denote significance at the 5 percent level.

Table 4: AVG, conditional variance dynamics

Location	Province	$\Delta trend$	$p(nt)$	$p(ns)$	$p(nts)$	p	R^2
South							
Victoria	BC	0.08*	0.00	0.00	0.00	0.14*	0.04
Vancouver	BC	-0.12*	0.00	0.00	0.06	0.09*	0.05
Calgary	AB	-0.24*	0.03	0.00	0.11	0.10*	0.10
Edmonton	AB	0.04	0.19	0.00	0.14	0.08*	0.10
Winnipeg	MB	-0.17	0.33	0.00	0.45	0.01	0.08
Toronto	ON	-0.15*	0.00	0.00	0.15	0.03*	0.05
Ottawa	ON	-0.21*	0.00	0.00	0.27	0.01	0.08
Montreal	QC	-0.15*	0.00	0.00	0.04	0.00	0.09
Quebec	QC	-0.06*	0.00	0.00	0.00	0.01	0.11

Moncton	NB	-0.04	0.06	0.00	0.04	0.03*	0.10
St John's	NL	0.02*	0.05	0.00	0.03	0.05*	0.03
MGE		-0.09	0.00				
North							
Yellowknife	NT	-0.15*	0.02	0.00	0.01	0.01	0.10
Eureka	NU	-0.16*	0.00	0.00	0.00	0.13*	0.11

Note:

All results are based on daily data, 1960–2019.

Column 3 reports the estimated trend movement over the 60-year sample in degrees Celsius. $p(nt)$ is the robust p -value for a Wald test of no trend. $p(ns)$ is the robust p -value for a Wald test of no seasonality. $p(nts)$ is the robust p -value for Wald a test of no trend in seasonality. p is the estimated autoregressive coefficient.

R^2 is the adjusted coefficient of determination.

Asterisks denote significance at the 5 percent level.

Table 5: DTR, conditional variance dynamics

Location	Province	$\Delta trend$	$p(nt)$	$p(ns)$	$p(nts)$	p	R^2
South							
Victoria	BC	0.47*	0.00	0.00	0.00	0.03*	0.04
Vancouver	BC	-0.21*	0.00	0.00	0.00	0.07*	0.02
Calgary	AB	-0.05*	0.04	0.00	0.07	0.04*	0.02
Edmonton	AB	0.47*	0.00	0.00	0.17	0.07*	0.02
Winnipeg	MB	-0.04	0.44	0.00	0.65	0.03*	0.02
Toronto	ON	-0.13*	0.02	0.00	0.30	0.04*	0.02
Ottawa	ON	-0.48*	0.00	0.00	0.18	0.05*	0.02
Montreal	QC	-0.20*	0.00	0.00	0.01	0.05*	0.02
Quebec	QC	0.09*	0.00	0.00	0.00	0.05*	0.03
Moncton	NB	-0.36*	0.00	0.00	0.65	0.02*	0.02
St John's	NL	0.05	0.70	0.00	0.90	0.01*	0.02
MGE		-0.04	0.70				
North							
Yellowknife	NT	-0.15*	0.01	0.00	0.00	0.05*	0.05
Eureka	NU	-0.11*	0.02	0.00	0.01	0.05*	0.05

Note:

All results are based on daily data, 1960–2019.

Column 3 reports the estimated trend movement over the 60-year sample in degrees Celsius. $p(nt)$ is the robust p -value for a Wald test of no trend. $p(ns)$ is the robust p -value for a Wald test of no seasonality. $p(nts)$ is the robust p -value for Wald a test of no trend in seasonality. p is the estimated autoregressive coefficient.

R^2 is the adjusted coefficient of determination.

Asterisks denote significance at the 5 percent level.

Table 6: Trends in the number of days of extreme temperatures

Location	Province	Summer days	Hot days	Hot nights	Frost days	Ice days
South						
Victoria	BC	9.88*	1.34*	NA	-23.67*	-1.65
Vancouver	BC	6.05*	-0.33	0.07	-20.85*	-3.20
Calgary	AB	0.83	-0.10	NA	-7.42	-10.37
Edmonton	AB	7.65	0.13	NA	9.16	-12.46*
Winnipeg	MB	8.21	-1.31	0.11	-7.39	-11.99*
Toronto	ON	18.48*	7.06*	2.20	-22.32*	-15.98*
Ottawa	ON	19.99*	5.94*	0.99*	-14.01*	-18.30*
Montreal	QC	20.34*	5.73*	2.90*	-20.57*	-16.54*
Quebec	QC	10.62*	0.53	0.11	-15.47*	-15.71*
Moncton	NB	16.41*	4.54*	0.07	-18.24*	-12.99*
St John's	NL	8.49*	-0.07	NA	-17.20*	-8.27
MGE		11.40*	2.07*	0.92*	-15.14*	-10.76*
North						
Yellowknife	NT	6.14*	0.05	NA	-5.60	-10.00*
Eureka	NU	NA	NA	NA	-18.26*	-8.78

Table 7: Trends in the average duration of extreme temperatures

Location	Prov	Summer days	Hot days	Hot nights	Frost days	Ice days
South						
Victoria	BC	0.46*	0.43	NA	-1.38*	-0.86
Vancouver	BC	0.10	-0.31	0.07	-0.32	-1.42
Calgary	AB	0.00	-0.09	NA	0.20	-0.44
Edmonton	AB	0.10	0.21	NA	-0.89	-1.02
Winnipeg	MB	0.51	0.00	0.11	-0.82	-1.52
Toronto	ON	1.21*	0.32	-0.13	-1.65	-0.82*
Ottawa	ON	0.42	0.00	0.68*	0.19	-1.62*
Montreal	QC	0.49	-0.06	0.49	-1.98*	-1.11*
Quebec	QC	0.55*	-0.13	0.11	0.33	-0.68
Moncton	NB	0.50*	0.51	0.07	-0.40	-0.31
St John's	NL	0.40*	-0.07	NA	-2.31*	-0.44*
MGE		0.43*	0.10	0.20*	-0.87*	-0.92*
North						
Yellowknife	NT	0.33	0.08	NA	-2.34	1.14
Eureka	NU	NA	NA	NA	11.58	6.18

Table 8: Temperature indices

Conditions	Indices	Definition
Warm	Summer days	Days with $t_{max} > 25^{\circ}\text{C}$
	Hot days	Days with $t_{max} > 30^{\circ}\text{C}$
	Hot nights	Days with $t_{min} > 22^{\circ}\text{C}$
Cold	Frost days	Days with $t_{min} \leq 0^{\circ}\text{C}$
	Ice days	Days with $t_{max} \leq 0^{\circ}\text{C}$