The effect of winter weather on U.S. economic activity

Justin Bloesch and François Gourio

Introduction and summary

The unusually cold and snowy 2013–14 winter substantially disrupted the routines of people across the United States, leading commentators and policymakers to ask if the weather affected economic activity as well. There were many media stories that supported this hypothesis. For instance, some employees were reported as unable to commute to work, and some projects, particularly in construction, were delayed due to equipment limitations or concerns about safety in the cold and snow. Supply chains were sometimes interrupted; for instance, steel production along the coast of Lake Michigan was affected because the boats delivering iron ore were unable to navigate the deeply frozen Great Lakes. Furthermore, retailers reported that households may have delayed shopping due to extreme weather. And finally, some expected that the higher heating costs and the expenses for home repairs (such as burst pipes) would hamper consumer spending.

Consistent with these anecdotes, economic indicators published early in 2014, such as industrial production, employment, and car sales, showed that economic activity had slowed substantially in December 2013 and January 2014. While the economic recovery following the Great Recession had appeared to accelerate in the fall of 2013, these statistics suggested a renewed slowdown. To illustrate these patterns, figure 1 depicts the evolution of several economic indicators:[1] the monthly change in nonfarm employment, the National Association of Purchasing Managers (NAPM) Index, lightweight vehicle sales, retail sales (excluding auto sales), manufacturing industrial production, and the Chicago Fed National Activity index (CFNAI), which itself summarizes a variety of indicators. These indicators are seasonally adjusted using statistical methods, which amounts to removing the effects of a "normal winter." In these figures, the three red dotted points correspond to December 2013, January 2014, and February 2014, respectively. The decline of these indicators during the December to February period is consistent with a slowdown in economic activity due to the weather, but could also have reflected other sources of weakness. Indeed, there was much controversy at the time on how much of the decline in the indicators was driven by the bad weather as opposed to other factors. The conventional wisdom was that a slowdown in economic activity due to weather would be very temporary; projects that had been delayed due to weather would eventually be finished and consumer shopping would likely resume.

Justin Bloesch is an associate economist and François Gourio is a senior economist in the Economic Research Department at the Federal Reserve Bank of Chicago.
The authors thank their colleagues at the Chicago Fed and Lisa Barrow, Olivier Deschênes, Charles Gilbert, Ben Herzon, Alejandro Justiniano, and Thomas Klier for useful discussions and feedback.

Economic Perspectives is published by the Economic Research Department of the Federal Reserve Bank of Chicago. The views expressed are the authors' and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

ISSN 0164-0682

FIGURE 1

Economic indicators

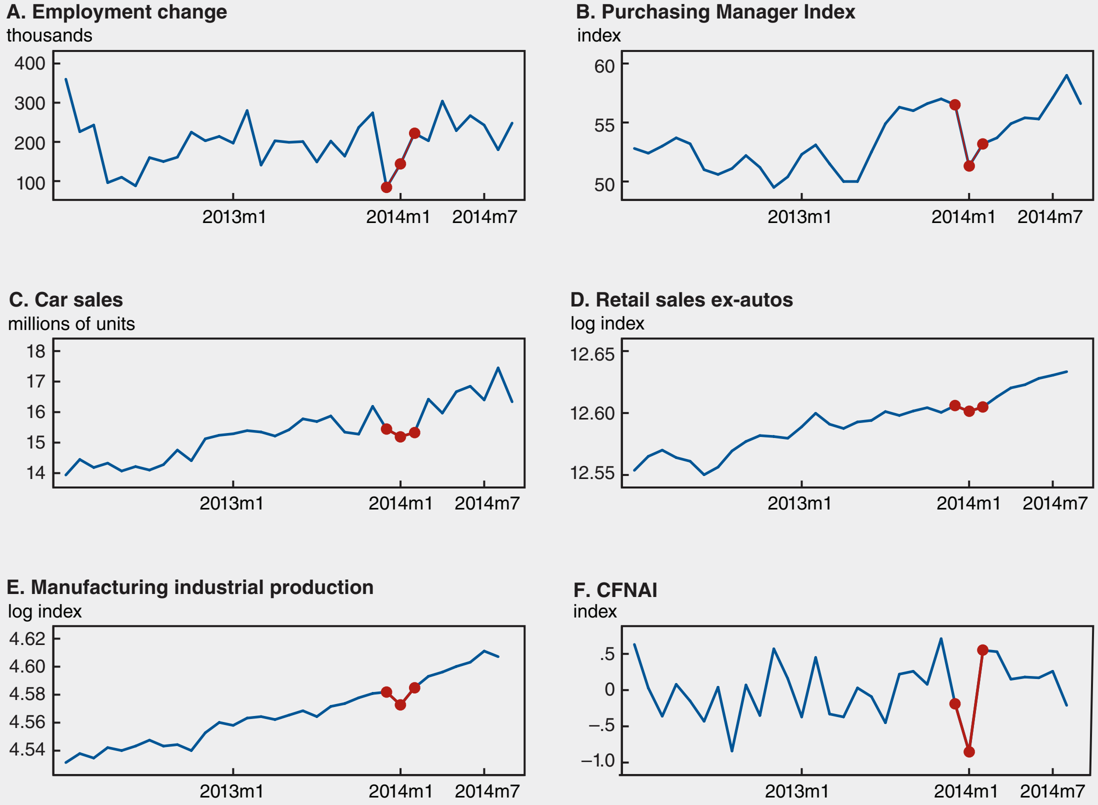

Source: Haver Analytics.

Whether the economic slowdown was due to winter weather or an underlying trend had implications for monetary policy. The Federal Reserve Open Market Committee (FOMC), which sets monetary policy for the United States, decided in its December 2013 meeting to start reducing the monthly volume of its asset purchases (from \$85 billion per month to \$75 billion per month). This "tapering" policy was motivated by improvements in the economy toward the Federal Reserve's inflation and employment targets during 2013, but it was explicitly made data dependent.[2] This means that if the economy was indeed becoming weaker persistently, the committee would likely continue its asset purchases at current levels rather than "taper" them. The weak economic data released early in 2014 made this a real possibility. However, if the disappointing data reflected only transitory weather effects, then the Federal Reserve would continue its gradual decline of asset purchases. The challenge for both for the committee and investors was to disentangle how much of the weakness of the economy was weather related. It is fairly unusual that the weather affects the economy in a very significant way, and hence there is little established knowledge, or even a good rule of thumb, that economists can rely on.[3] In part, it also reflects the fact that good-quality weather data were not readily available to allow economists to perform the statistical analyses they would need to estimate causal relationships. For instance, commonly used databases do not have data on aggregate snowfall for the United States, and temperature series are often area weighted rather than population weighted, which is probably better for physical science applications but less useful when one is trying to measure the economic impact of weather since it gives a large importance to some sparsely populated states. Reflecting this difficulty in measuring the precise effects of weather, the March 2014 FOMC statement noted simply that "growth ... slowed during the winter months, in part reflecting adverse weather conditions," with the qualifier "in part" hedging the statement but suggesting that staff work had not found weather to be the sole determinant of weakness.[4] In the related press conference, Federal Reserve Chair Janet Yellen stated that "we did spend a lot of time discussing weather and how it's affected businesses and households in various parts of the country—certainly weather has played an important role in weakening economic activity in [the first quarter]. It's not the only factor that is at work, and most projections for growth in the first quarter are reasonably weak."[5] Over time, however, economic data started to improve, as figure 1 shows, and most analysts came to attribute the winter weakness to weather. For instance, the April 30, 2014, FOMC statement noted that "growth ... has picked up recently, after having slowed sharply during the winter in part because of adverse weather conditions."[6] Chair Yellen reflected this in a speech on April 16, 2014, when she said that: "In recent months, some indicators have been notably weak ... [and] my FOMC colleagues and I generally believe that a significant part of the recent softness was weather related."[7]

Later, following a fairly strong increase in growth in the second quarter, it became folk wisdom that the weakness of growth in the first quarter was mostly weather related. For instance, Justin Wolfers wrote in the *New York Times* on September 26 that "Much of [the second quarter] growth is simply catching up from the first quarter when severe winter storms led the economy to contract. ... The snow, it seems, led spending to be deferred a quarter, rather than canceled."[8] Clearly, there is a tension between the initial assessment, which was highly uncertain regarding the effect of weather on economic activity, and the folk wisdom that emerged. The goal of this article is to resolve this tension by providing more robust statistical evidence regarding the effects of the weather on economic activity. To do so, we build on work started by some analysts at the Board of Governors of the Federal Reserve System (the Board) and at the private forecasting firm Macroeconomic Advisers (2014) and construct better data using records of individual weather stations across the entire continental United States. Our analysis improves over this previous work along two main dimensions: First, we use longer historical records, allowing us to increase significantly the length of data. Second, we use regional variation in economic activity to further increase the span of data available. We discuss later why having larger samples is especially useful in this context.

The rest of this article is organized as follows. We start by reviewing some related literature on the effects of weather on economic activity. We then present our measures of weather and discuss in particular the winter of 2013–14. Next, we present our empirical approach, which uses state-level measures of weather and economic activity to evaluate the effects of the weather on economic activity; then we show the results. We also present some results using national-level measures of weather and economic activity. Finally, we use our estimates to reassess how much of last winter's bad economic data was weather related. We finish with a note about the potential effects of climate change on our results.

Related research

The idea that weather is an important source of fluctuations of production is an old theme in economics. A century ago, when the economy was still in large part driven by agriculture, bad crops had a measurable effect on aggregate income. Thus, the Dust Bowl had a notable effect during the Great Depression. Weather may continue to have a significant economic impact in countries that are very reliant on agriculture, either because they are poor or because their exports are concentrated on a small number of crops. Even today, economists such as Jeffrey Sachs[9] attribute the limited economic development of some countries to their extreme climate. Extremes of temperature, dryness or humidity, and precipitation (rain or snow) make economic progress difficult for some countries in Africa and Asia. Closer to home, the Caribbean countries and Central America regularly experience hurricanes that destroy housing, infrastructure, and production capacity. Furthermore, the prospect of climate change raises the question of how the world economy will be affected by higher temperatures.

Dell, Jones, and Olken (2014) review the recent literature on weather and the economy, including their own study (2012), which shares many methodological similarities with our approach. Their focus is very different, however. They use annual country-level data on temperature (and precipitation) and gross domestic product (GDP) to estimate the effects of weather on GDP. They find a significant effect for poor countries, which appears to be largely, though not exclusively, driven by the impact on agriculture. An increase in the average annual temperature of 1 degree Celsius (that is, 1.8 degrees Fahrenheit) leads GDP to fall by 1.3 percent. Perhaps surprisingly, there seems to be little tendency for GDP to recover from its decline the following year, that is, little "bounceback." They find no effect on developed countries and no effect of precipitation. The contribution of their paper is to offer an identification of the effects of weather based on variation over time within countries, rather than on cross-country relationships. However, because they focus on annual data, they are unable to study the short-term movements in economic activity that may be due to weather in developed countries. In a related analysis, Deschênes and Greenstone (2007) measure the effects of weather on U.S. agricultural production using detailed geographic data. Here, too, measures of output are annual.

Most closely related to our study are three papers that were written contemporaneously with ours. Boldin and Wright (2015) calculate the effect of weather on national nonfarm payroll employment. One important conclusion they draw is that weather affects the seasonal adjustment. Colacito, Hoffman, and Phan (2014) and Deryugina and Hsiang (2014) both use cross-regional U.S. data to study the effects of weather on economic activity. An important difference is that they focus on annual measures of income or production rather than on the higher-frequency measures that we use. These papers also study the total annual weather effect, whereas we focus on the effect of unusual winter weather only.

Measuring the weather

Measuring the weather may seem to be a simple and straightforward exercise. However, exploring the details of the data reveals various challenges. First, we need to decide which measure of weather to study. Temperature alone does not fully capture the ways in which weather can affect economic activity; other factors may be important, such as precipitation, wind (direction and strength), and humidity, for example. Additionally, several variables may interact. Second, weather can be highly localized, and a snowstorm in southern Illinois is unlikely to have the same effect on employment as a snowstorm in Chicago. Because of this, the correct way to weight and aggregate our weather variables is not clear ahead of time. This section outlines our approach.

Our source of weather measurements is a data set called the U.S. Historical Climatology Network, which is part of the Global Historical Climatology Network (GHCN); these data were constructed by the National Climatic Data Center (NCDC), a part of the National Oceanic and Atmospheric Administration (NOAA).[10] This data set has daily measures of many weather variables, including temperature, snowfall, and total precipitation; in this article, we focus on temperature and snowfall. The data set reports conditions from about 1,200 weather stations throughout the United States. However, not all stations were in use in all years (that is, these data are an unbalanced panel). We use data from 1950 through 2014 for our estimation.[11]

There are a few potential issues with the quality of these data. First, changes in station design or practices sometimes introduce changes in measured temperatures. For instance, the station instrumentation may change; the station's neighborhood may change due to human activity or the station itself might be moved; the time at which observations are made during the day may change. These changes are especially important when we try to measure long-term changes in the mean temperature; some researchers have developed algorithms to take into account the changes. However, these adjustments are not available for our data.[12]

A possibly more important issue with the data is that stations are introduced partway through the data set, and some stations stop reporting measurements in the middle of the data set. This can cause problems with aggregation. To illustrate this, imagine constructing a state index for temperature in Illinois. Suppose that numerous stations are introduced in southern Illinois and some are discontinued in northern Illinois. Since the southern part of the state is on average substantially warmer than the northern part, the data would show a large increase in the statewide temperature even though the actual temperature never changed. Failing to account for the evolution of active stations over time could create artificial changes in measured weather conditions. When constructing our state weather indexes, we resolve this problem by aggregating deviations from local long-run averages. For example, suppose the temperature in Illinois is uniformly 2 degrees above average across the entire state. Suppose it is 52 degrees in southern Illinois but 42 degrees in Chicago, with a state average of 47 degrees. The normal temperatures for a given day are 50, 40, and 45 degrees, respectively. If the Chicago station drops out of the data set, then the state average temperature will suddenly jump to 52 degrees. It then appears that the state average is 7 degrees above normal, rather than the actual 2 degrees. However, if one were to average the deviations from normal, the observed average temperature would still be only 2 degrees above average.

This is precisely the process that we use to construct our weather indexes at the state and monthly level. We construct the index in six steps. We start from the daily temperature for a given weather station[13] and first calculate the average monthly temperature for each month and each year. Mathematically, for a month that lasts 30 days, and denoting $T_{s,d,m,y}$ as the temperature in day d of month m of year y in station s, we define $\overline{T}_{s,m,y} = \frac{1}{30}\sum_d T_{s,d,m,y}$. Second, we define the "normal weather" for a station and a month as the monthly temperature averaged over all years from 1950 through 2014. Mathematically, $\overline{\overline{T}}_{s,m} = \frac{1}{65}\sum_y \overline{T}_{s,m,y}$,

FIGURE 2

Temperature in all stations in Illinois

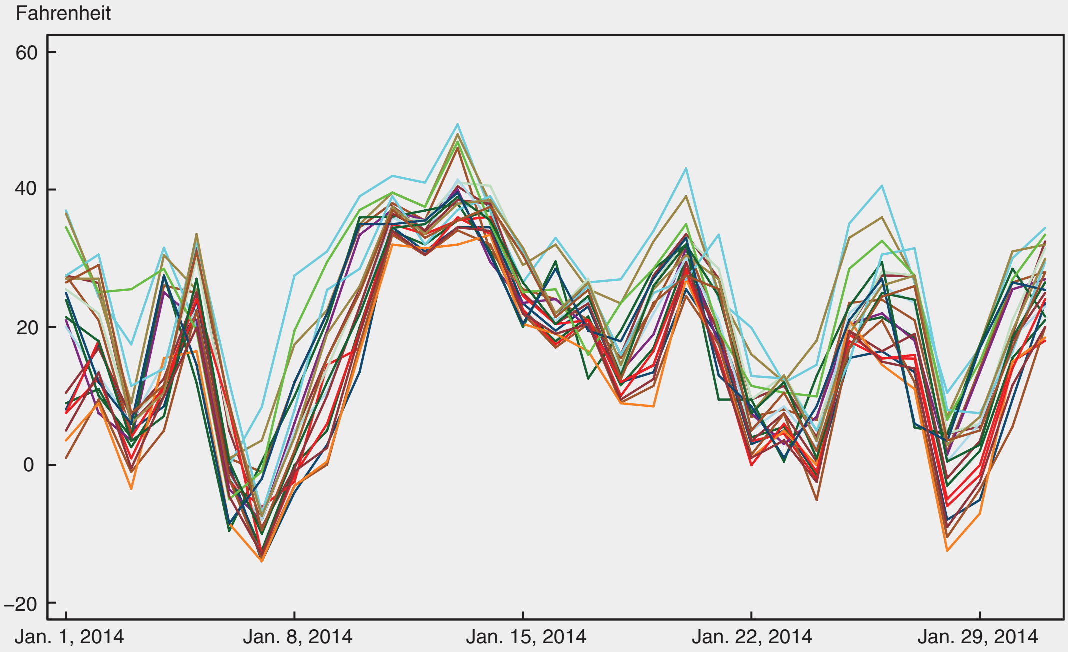

Source: National Climatic Data Center.

since we use 65 years of data. Third, we calculate the monthly deviation as the difference between the monthly average and the long-run normal, or $\hat{T}_{s,m,y} = \bar{T}_{s,m,y} - \bar{\bar{T}}_{s,m}$. This yields a measure of temperature deviation from its normal. It is important to note here that stations naturally experience different levels of variation: A day 20 degrees above or below normal will be more common in Minneapolis than in San Diego. Therefore, it is intuitive to normalize the monthly deviation by a measure of variability.[14] Hence, in the fourth step, we calculate a station- and month-specific measure of variability: the standard deviation across years of the monthly temperature $\bar{T}_{s,m,y}$, which we denote $\sigma^T_{s,m}$. The mathematical formula is

$$\sigma^T_{s,m} = \sqrt{\frac{1}{65}\sum_y \left(\bar{T}_{s,m,y} - \bar{\bar{T}}_{s,m}\right)^2}.$$

We then define the normalized deviation as the ratio of deviation to this standard deviation:

$$\tilde{T}_{s,m,y} = \frac{\hat{T}_{s,m,y}}{\sigma^T_{s,m}}.$$

The next step involves aggregating over all weather stations within a state. A refined approach would be to weight stations according to the population surrounding them, since economic activity is correlated with population. However, in the interests of simplicity, we calculate the simple average of $\tilde{T}_{s,m,y}$ across all stations in a state; this yields a temperature index that we denote $T_{i,m,y}$ (where i denotes the state):

$$\hat{T}_{i,m,y} = \frac{1}{N_{i,m,y}}\sum_{s \in i}\tilde{T}_{s,m,y}.$$

$N_{i,m,y}$ is the number of stations in state i in month m and year y, and the sum runs over all stations s in a state i. In a final step, we normalize this index so it has mean zero and standard deviation one:

$$T_{i,m,y} = \frac{\tilde{T}_{i,m,y} - E\left(\tilde{T}_{i,m,y}\right)}{\sigma\left(\tilde{T}_{i,m,y}\right)},$$

where E and σ denote the mean and standard deviation, calculated over the winter months.

One potential concern is that the simple average across all stations might be misleading if the weather is very different across the state. Figure 2 presents the temperatures of all stations in Illinois in January 2014;

even in this fairly large state, the co-movement of temperatures is striking. This suggests that the simple average may be good enough for our purposes.

We also construct regional (Midwest, West, Northeast, and South) and national weather indexes by weighting the state indexes according to their employment. Finally, in exactly the same way, we construct a snowfall index, replacing temperature T with snowfall data S. In this case, the lumpier nature of snowfall makes our simple averaging within a state less compelling, though figure 3 suggests that there is still some significant co-movement.

Figure 4 shows a scatter plot of our monthly state-level indexes of temperature and snow. As could be expected, the negative correlation is fairly strong (–0.49). Finally, figures 5 and 6 depict the correlograms of our temperature and snow indexes respectively; these figures provide a visual way to assess how long a good temperature (or snow) index lasts. While there is some significant correlation over a few days, we see that the correlation falls fairly quickly, especially for snowfall.

The 2013–14 winter in perspective

Much of the past winter's cold temperatures was caused by the "polar vortex," a low-pressure weather system that typically stays above the Arctic Circle during the winter, spinning in a tight bowl over high latitudes. It is held in place by the jet stream, a fast-moving, high-altitude wind that keeps cold air to the north and warmer air from the south from interacting. However, during the winter of 2013–14, the polar vortex slowed down, causing it to "wobble," much like a spinning top that loses momentum. This pushed the jet stream farther south than normal, bringing the cold arctic winds to lower latitudes.

The severity of the winter can be seen in figures 7 and 8, which show the deviation of first-quarter average temperatures and snowfall from the long-term averages. Clearly, this winter was cold and snowy, but the polar vortex did not impact the country evenly. Figures 9 and 10 (p. 10) show the weather deviations for each region. Temperatures were above average in the West of the United States. For the eastern half of the country, however, the winter was brutally cold. The first-quarter average temperature in the Midwest was about two standard deviations below the mean, making it the third coldest in our data and fairly similar to the two worst ones, 1979 and 1980. The Northeast similarly had its third-coldest first-quarter temperature in our data, and the South experienced its sixth coldest. As well as being cold, the first quarter of 2014 was also snowy in the Midwest, Northeast, and South.

Empirical approach using state-level data

To evaluate how weather affects economic activity, we use a commonly used statistical model known as regression analysis. The equation describing the model is

$$1)\quad \Delta \log Y_{i,m,y} = \alpha_i + \delta_{m,y} + \beta T_{i,m,y} + \gamma S_{i,m,y} + \varepsilon_{i,m,y},$$

where $\Delta \log Y_{i,m,y}$ is the change in the logarithm of a variable measuring seasonally adjusted economic activity (such as employment) in state i in month m of year y;[15] $T_{i,m,y}$ is our temperature index for state i in month m of year y; $S_{i,m,y}$ is our snow index. The factors α_i and $\delta_{m,y}$ are so-called fixed effects, that is, constants that depend solely on the state (α_i) or time ($\delta_{m,y}$). These factors serve to capture, respectively, the fact that some states grow faster on average and that all states tend to co-move, for instance, due to economic recessions. By removing this variation from the data, we obtain statistically more precise estimates of the weather effect. Finally, $\varepsilon_{i,m,y}$ is a so-called error term that captures factors other than temperature and snow, not constant across time or states, that affect economic activity.

The key assumption underlying this model is that these factors are uncorrelated with the temperature index $T_{i,m,y}$ and with the snow index $S_{i,m,y}$. This is plausible in our case since short-term variations in weather are unlikely to be caused by the factors thought to affect economic activity, such as productivity, interest rates, or consumer confidence.[16] This allows us to estimate the model using a simple technique known as OLS (ordinary least squares).[17] It is important to note that this model imposes several assumptions: First, the effect of our weather indexes on the growth rate of economic activity is linear, so that the effect of a one standard deviation increase in the temperature index is half the effect of a two standard deviation increase in the temperature index, and the exact opposite of a standard deviation increase in the temperature index. One might think that this is an unrealistic assumption. For example, in January 2014, the very low temperatures in Chicago had the extreme effect of leading many people not to commute to work, so perhaps the effect of very low temperatures is more than proportional. We performed some exploratory analysis and did not find support for such nonlinearities. For instance, one can create indexes to capture "extreme cold" or "extreme snow" by counting the number of days within a month with very low temperatures or very high snowfall. These indexes do not seem to convey important additional information relative to our simple index. However, this certainly deserves more study. A second important implicit assumption is that the effect of a high weather index is the same

FIGURE 3

Snowfall in all stations in Illinois

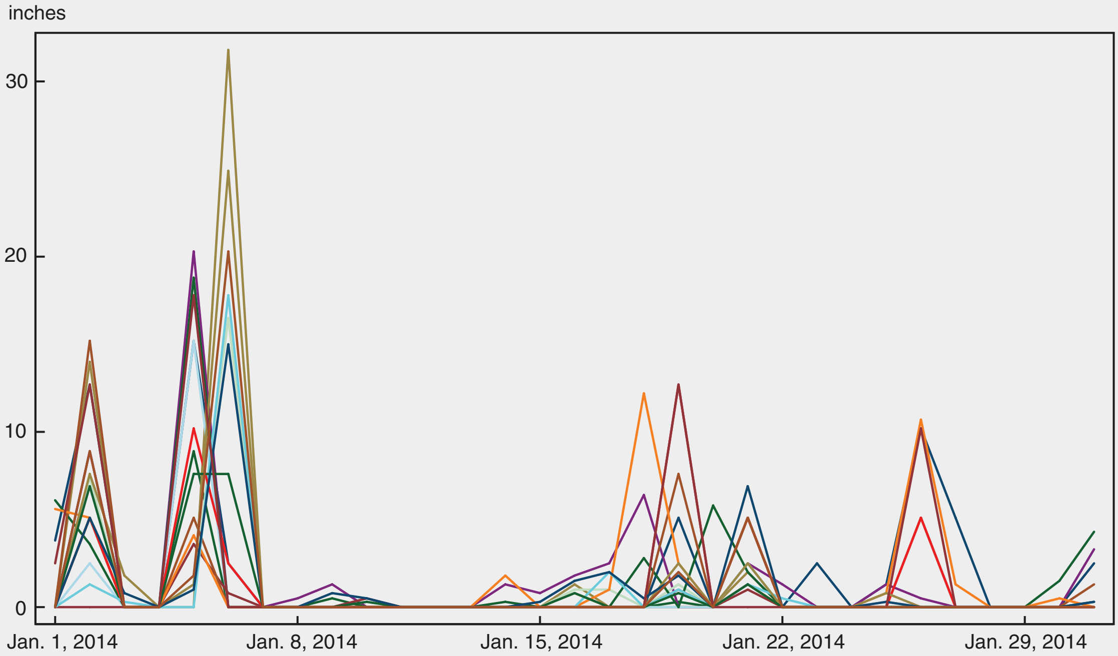

Source: National Climatic Data Center.

FIGURE 4

Correlation between temperature and snow (winter months)

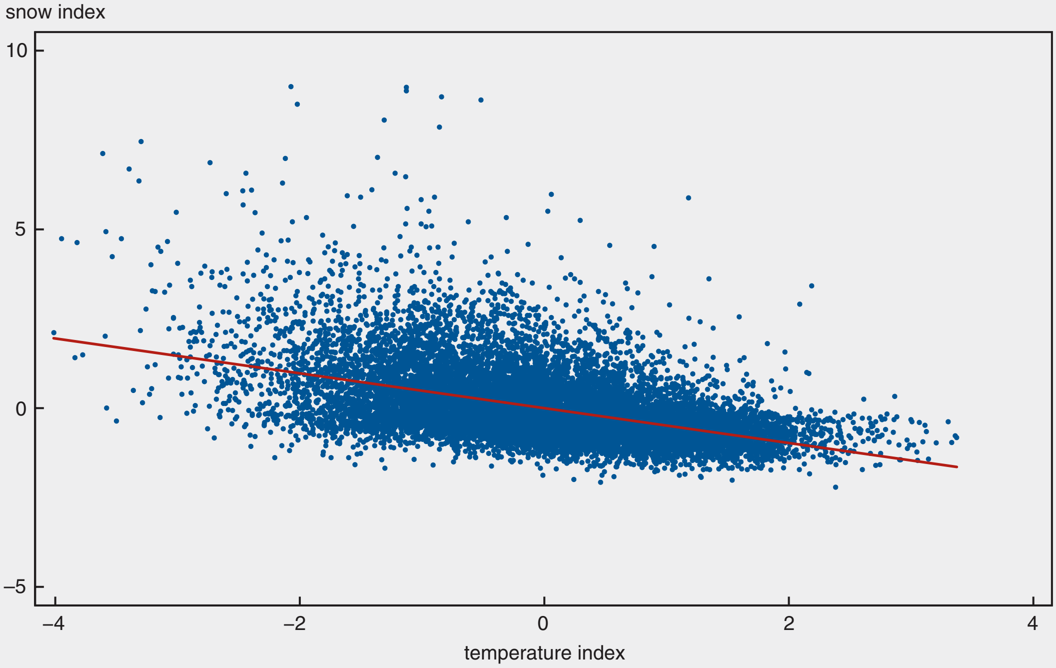

Source: Authors' calculations based on data from the National Climatic Data Center.

FIGURE 5

Correlogram of temperature index

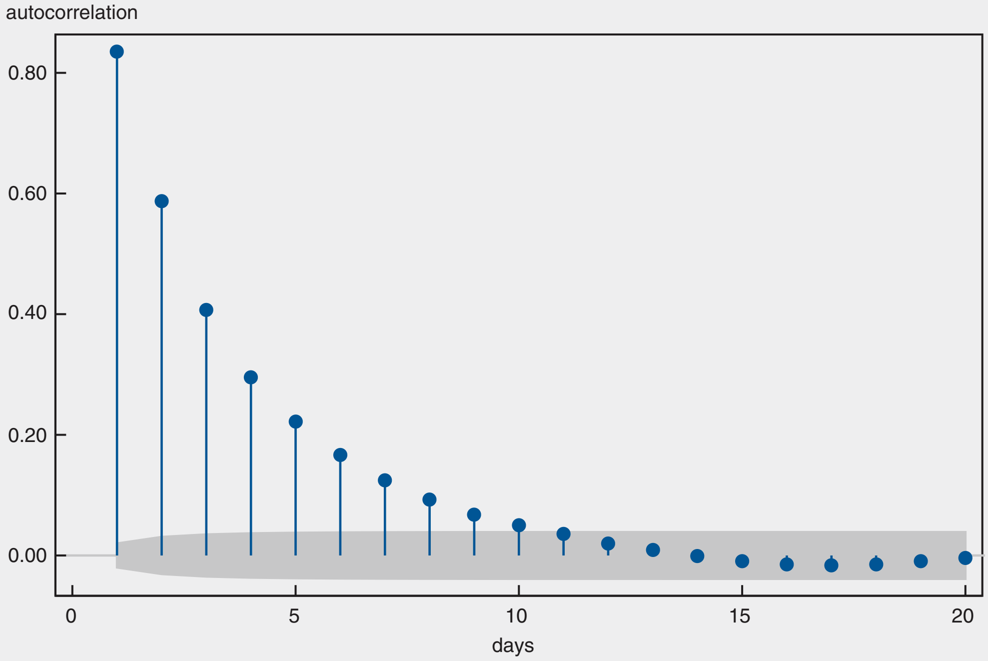

Source: Authors’ calculations based on data from the National Climatic Data Center.

FIGURE 6

Correlogram of snowfall index

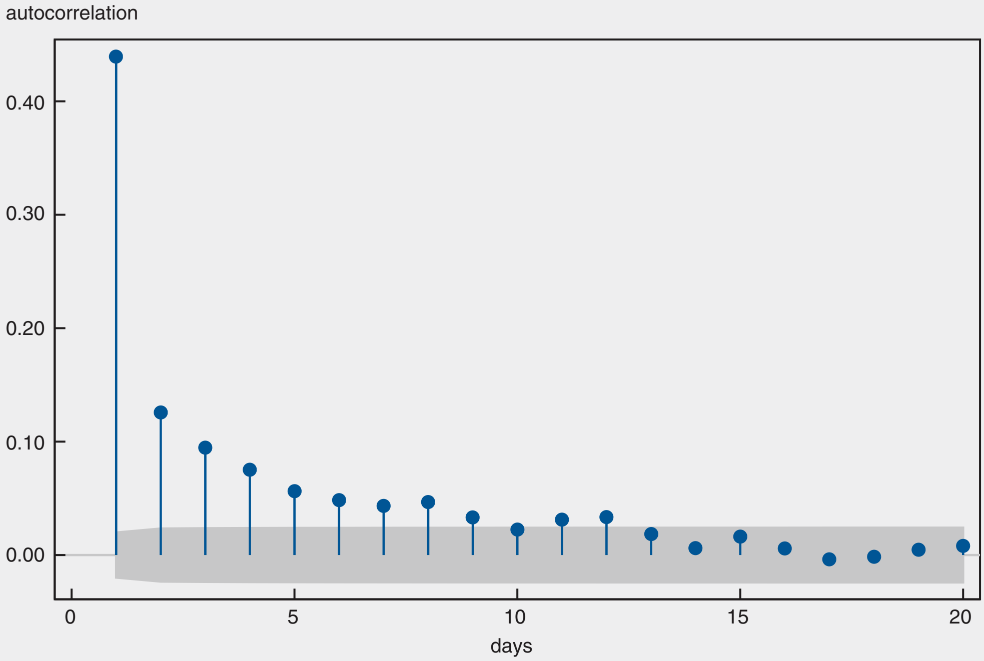

Source: Authors’ calculations based on data from the National Climatic Data Center.

FIGURE 7

National temperature deviation from 1950–2014 average

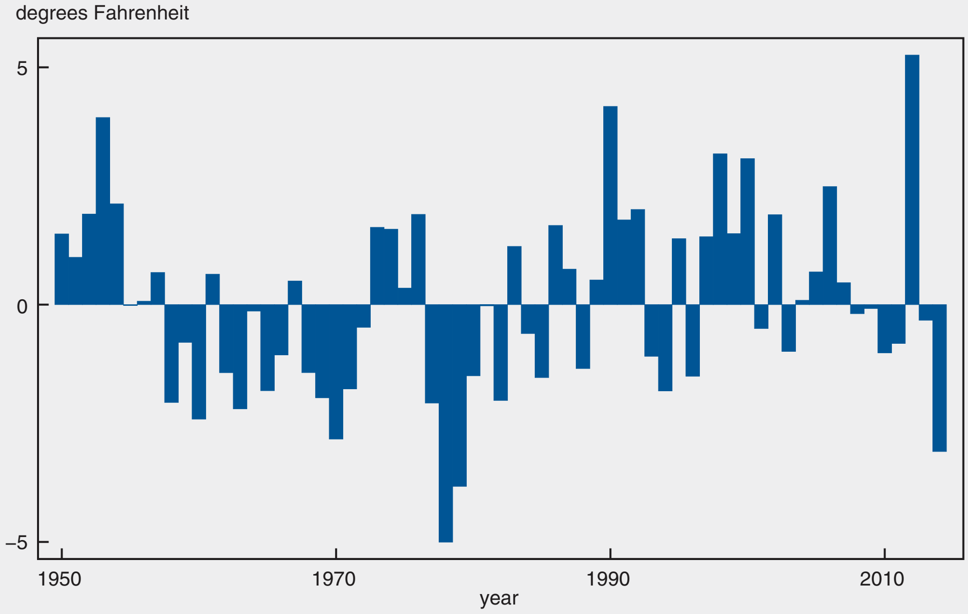

Source: Authors' calculations based on data from the National Climatic Data Center.

FIGURE 8

National snowfall deviation from 1950–2014 average

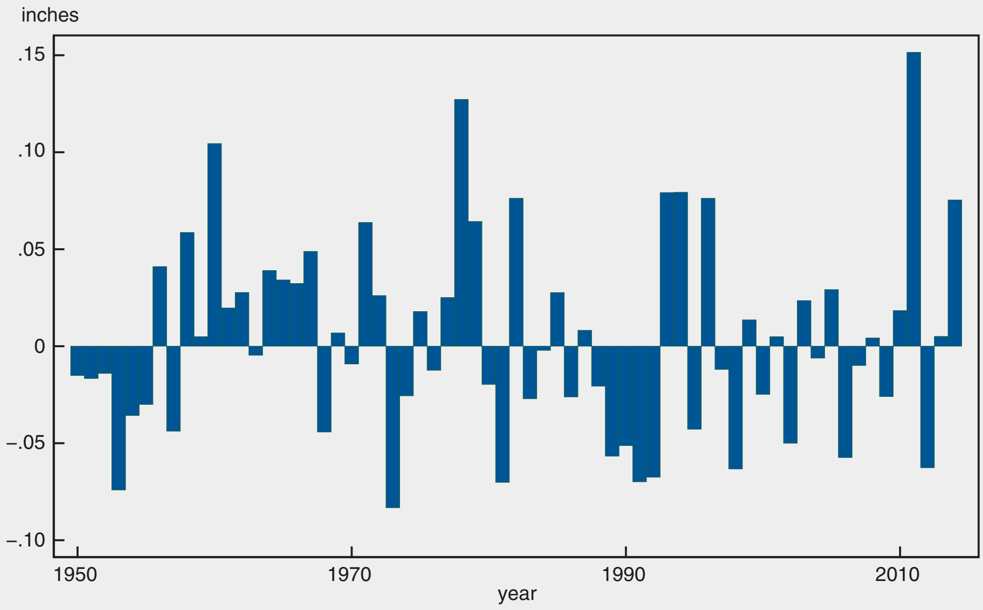

Source: Authors' calculations based on data from the National Climatic Data Center.

FIGURE 9

Regional temperature index for first quarter

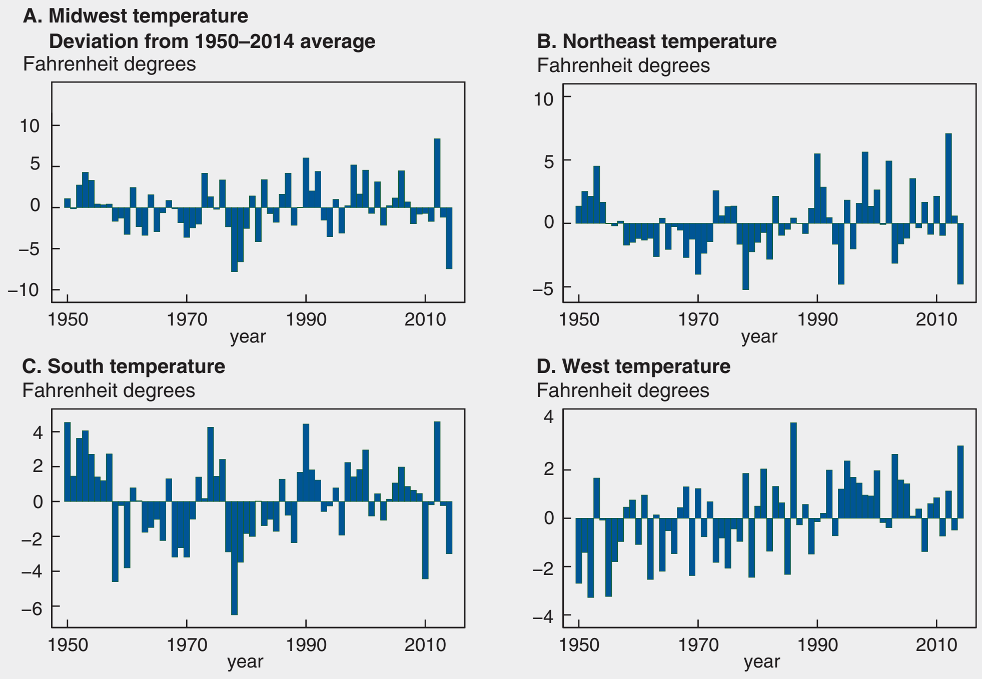

Source: Authors' calculations based on data from the National Climatic Data Center.

FIGURE 10

Regional snow index for first quarter

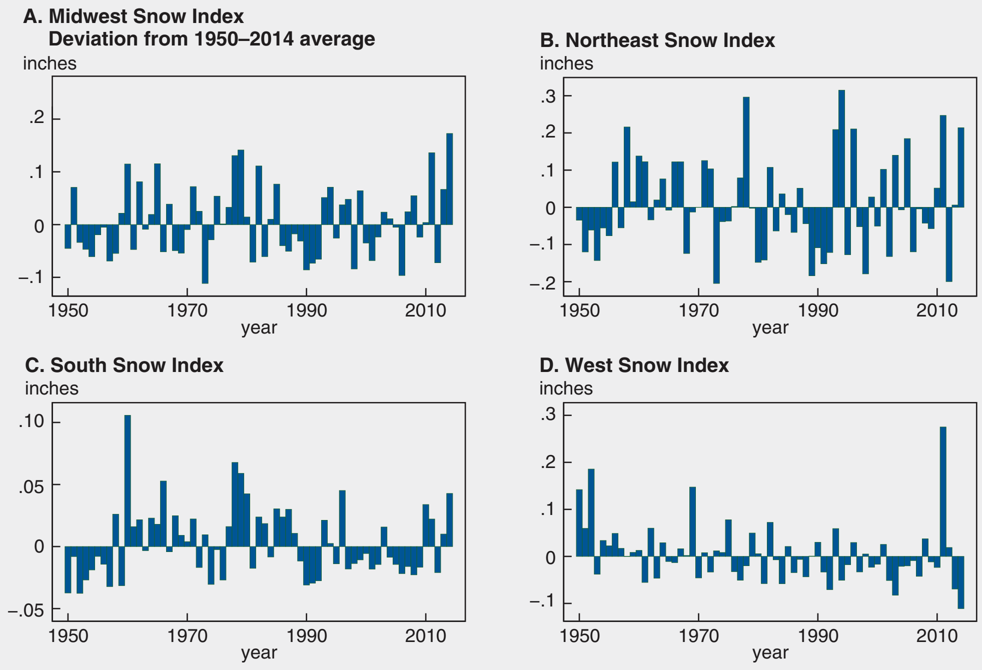

Source: Authors' calculations based on data from the National Climatic Data Center.

TABLE 1

Effect of temperature and snowfall indexes on state-level economic activity

	Nonfarm payrolls	Unemployment rate	U.I. new claims	Housing permits	Housing starts
Temperature	0.041***	−0.034	−0.967***	1.124*	2.431***
	(0.008)	(0.253)	(0.260)	(0.648)	(0.889)
Snowfall	−0.029***	0.100	0.861***	−2.091***	−1.905***
	(0.006)	(0.153)	(0.191)	(0.451)	(0.595)
Observations	37,154	22,517	25,084	19,900	25,660
R^2	0.262	0.451	0.148	0.118	0.098
Sample start	1950	1976	1971	1980	1970

Notes: Results from estimation of equation 1 using monthly data from November through March by ordinary least squares with state and time effects; standard errors in parentheses are two-way clustered by state and time. The left-hand-side variables are all in log changes, except the unemployment rate, which is in level change. Sample start date as shown; end date is 2014 for all series. U.I. indicates unemployment insurance. Temperature and snowfall indexes are normalized to have mean zero and standard deviation one for each state. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) level.
Source: Authors' calculations based on data from the National Climatic Data Center.

in all states. As we discussed above, our weather indexes are normalized to have the same standard deviation in all states, which makes this assumption more plausible, but it also deserves more study. A third assumption is that the average weather during the month is the relevant metric; for instance, we do not differentiate bad weather during the week from bad weather during the weekend.

Equation 1 estimates the effect of a given month's weather on the same month's economic variable Y. An important question is how long these effects last. To answer this question, we also estimate the same model, but allowing for lags in the weather:

$$2)\quad \Delta \log Y_{i,m,y} = \alpha_i + \delta_{m,y} + \sum_{k=0}^{K} \beta_k T_{i,m-k,y} + \sum_{k=0}^{K} \gamma_k S_{i,m-k,y} + \varepsilon_{i,m,y},$$

that is, the weather in the previous K months may affect Y. This specification allows us to evaluate the strength and speed of the bounceback from a bad weather spell.[18] The coefficient β_0 measures the effect in a given month of a one standard deviation change in the temperature index that month, the coefficient β_1 is the effect the following month, and so on.

Finally, note that we estimate this equation using the "cold season" only. That is, we define the temperature and snowfall indexes for November through March only (and allow the lags to work through until K months later). While weather affects the economy in the summer as well, the effects are likely to be different—for instance, high temperatures might have negative effects rather than positive effects and rainfall rather than snowfall might be relevant. This requires a different model.

Our analysis requires us to measure economic activity at the state level and at high frequency. Unfortunately, there is a relative paucity of economic data at this level of regional disaggregation and at this frequency. The main source of our data is the Bureau of Labor Statistics' *Current Establishment Survey* (CES), which surveys a large number of establishments each month regarding how many employees are on the payroll (during the pay period of the week including the 12th of the month). The headline national number ("nonfarm payrolls") is released the first Friday of the following month to considerable attention, but the same survey also produces monthly estimates of employment in each state and in each industry. These are our main sources of data. We also use monthly data on new unemployment insurance claims, housing starts, and housing permits.[19] Finally, we exclude Alaska and Hawaii from our analysis given their distinctive weather patterns and small population.

Results using state-level data

We first present the immediate effect of the weather, then discuss the bounceback. Finally, we study whether the economy has become less sensitive to weather over time.

Immediate effect

Table 1 presents the estimates of equation 1, which shows the effects of temperature and snowfall on total nonfarm employment, the unemployment rate, new unemployment insurance claims, housing permits, and housing starts. (Note that the economic data become available in different years, depending on the specific statistic.) A one standard deviation increase in the

TABLE 2

Effect of temperature and snowfall on the subcomponents of state-level nonfarm employment

	Temperature		Snowfall		Observations	R^2
Private nonfarm payrolls (A+B)	0.024***	(0.007)	−0.031***	(0.006)	14,143	0.426
A. Goods	0.080***	(0.016)	−0.056***	(0.014)	14,143	0.339
1. Mining	−0.016	(0.108)	−0.149*	(0.084)	12,067	0.058
2. Construction	0.185***	(0.045)	−0.181***	(0.032)	13,377	0.273
3. Manufacturing	0.036**	(0.016)	0.001	(0.011)	13,828	0.235
3a. Durable manufacturing	−0.077	(0.132)	−0.093	(0.096)	12,226	0.950
3b. Nondurable manufacturing	−0.003	(0.014)	0.004	(0.014)	13,257	0.125
B. Private services	0.009	(0.006)	−0.025***	(0.006)	14,143	0.353
B1. Trade, transportation, and utilities	0	(0.008)	−0.024***	(0.008)	14,143	0.322
B1a. Wholesale	0.011	(0.012)	−0.015	(0.011)	13,545	0.165
B1b. Retail	−0.003	(0.010)	−0.033***	(0.010)	14,143	0.259
B1c. Transportation and utilities	0.005	(0.018)	−0.003	(0.014)	13,848	0.184
B2. Information	0.013	(0.021)	0.021	(0.013)	12,679	0.126
B3. Financial services	0.017	(0.012)	0.001	(0.011)	14,143	0.142
B4. Professional and business services	0.003	(0.019)	−0.014	(0.016)	14,143	0.213
B5. Education and health	0.009	(0.008)	−0.022***	(0.008)	14,143	0.092
B6. Leisure and hospitality	0.041**	(0.016)	−0.067***	(0.014)	14,143	0.148
B6a. Accommodation and food	0.043***	(0.016)	−0.062***	(0.014)	13,839	0.146
B6b. Arts and leisure	0.054	(0.045)	−0.079	(0.051)	13,388	0.066
B7. Other services	−0.019	(0.013)	−0.043***	(0.012)	14,143	0.069
C. Total government	0.026**	(0.011)	−0.015	(0.010)	14,143	0.133
C1. Federal government	0.007	(0.030)	−0.010	(0.016)	13,393	0.624
C2. State government	−0.007	(0.016)	−0.030	(0.022)	13,819	0.027
C3. Local government	0.040***	(0.015)	−0.021	(0.014)	13,963	0.061

Notes: Results from estimation of equation 1 using monthly data from November through March by ordinary least squares with state and time effects; standard errors in parentheses are two-way clustered by state and time. All left-hand-side variables are in log changes. Sample is 1990–2014. Temperature and snowfall indexes are normalized to have mean zero and standard deviation one for each state. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) level.
Source: Authors' calculations based on data from the National Climatic Data Center.

temperature index during the winter leads nonfarm employment to grow by 0.04 percent, while a one standard deviation increase in the snowfall index leads to a decline of 0.03 percent. While these percentages are small, they can be important relative to the usual month-to-month fluctuations. For example, nonfarm payrolls are around 140 million nationwide, so the 0.04 percent estimated effect of temperature amounts to a difference of around 56,000 employees, which can make the difference between a "good" labor report and an "average" one. Importantly, the effects here are highly statistically significant, but the precision of the estimate is not extremely high. This may reflect the complexity of measuring the weather accurately.

The effect on the unemployment rate has the expected sign: A one standard deviation increase in the temperature index lowers the unemployment rate by 0.03 percentage points, and a one standard deviation increase in the snowfall index increases the unemployment rate by one-tenth of a percentage point (for instance, the unemployment rate would go from 5.8 to 5.9 percentage points). However, these effects are not statistically significant. The effects on new unemployment insurance claims, housing permits, and housing starts are all highly significant, of the expected sign, and fairly large: A one standard deviation shock to either snowfall or temperature moves new claims by about 1 percent and housing starts and permits by 1–2 percent. These are larger effects, but the underlying series are also more volatile, so that a percentage point difference is not extremely important.

Table 2 provides a breakdown of the employment effects by industry.[20] Some industries stand out as particularly affected, notably construction, hospitality, and, to a lesser extent, retail. Manufacturing is affected more by temperature than by snowfall. There is also an effect on education and health services, as well as government employment (the latter are very volatile, making the snowfall results not significant, despite the large negative point estimates). These education and government values may partly reflect school closures during bad weather. Overall, the results suggest that both the "supply" and "demand" channels are at work during the weather-related slowdown; that is, some sectors contract because it is impossible to produce, while some contract because there is less demand for their services.

Beyond showing the mechanics of the weather effect, these industry responses are useful because they allow us to identify episodes in which the weather may be an important driver of the economy. Thus, if a slowdown is associated with a large decline in construction, hospitality, and retail, it may in fact be weather related (even if the weather is not well measured).

Bounceback

To evaluate how long the effects of weather last, we estimate equation 2 with three lags ($K = 3$). These results are in table 3. The same-month impact effect of weather is obtained for each economic time series in the row $k = 0$ for temperature and snowfall, respectively. These are very similar to the effects found in table 1.[21] The novel result in this table is that last month's weather typically affects these economic time series with the opposite sign. For instance, a higher temperature index pushes nonfarm employment up by 0.043 percent the first month, but this recedes by 0.009 percent the next month and 0.025 percent the following month. This means that after two months, the effect of a higher temperature has largely receded. This pattern is very general across all time series, suggesting a strong bounceback, such that the level of economic activity returns roughly to where it was before the weather. A formal test of the hypothesis that a temperature shock has no effect on the level of economic activity three months later can be formulated as $\sum_{k=0}^{K}\beta_k = 0$; and for snowfall the test is $\sum_{k=0}^{K}\gamma_k = 0$. This amounts to a test of whether the weather only has transitory effects. For all of the series studied here, we cannot reject the hypothesis that weather has only transitory effects.[22] Overall, our results strongly support the intuitive notion of a bounceback.[23] The bounceback usually happens within a month or two, though in at least one case (nonfarm employment) there appears to be some remaining bounceback three months later.

Has the weather sensitivity declined?

The U.S. economy has changed in many ways since the 1960s and 1970s. New technologies for homebuilding have been developed, just-in-time inventory systems have been introduced, and there has been a shift away from industry and toward services. It is possible that as a result of these changes, the weather has less of an impact on the economy now than it once had. To evaluate this hypothesis, we estimate separately the effects of both temperature and snow on two subsamples: prior to 1990 and after 1990.[24] Table 4 reports these results. The temperature sensitivity of nonfarm employment, unemployment insurance claims, and housing permits and starts appears to have declined, though not all these changes are statistically significant. However, the snowfall sensitivity appears to have remained constant (for nonfarm employment) and may be even larger (for permits and starts). These last results could be due to structural changes in the homebuilding industry, such as the lengthening of the homebuilding season.

Methodology and results using national data

In this section, we present our results using national data. The main advantage of using national data is that there are many more economic data series available at a monthly frequency at the national rather than state level. The key disadvantage, which will become obvious as we proceed, is that by discarding regional variation in the weather, we have less data available, which does not allow as precise estimates of the weather effect. In part, this reflects the difficulty of disentangling the effects of snowfall and temperature, which are strongly negatively correlated (–0.60) in our national data.

Our methodology here is similar to our state-level work. We first construct a national temperature index and a national snowfall index by weighting the state indexes using nonfarm employment:

$$T_{m,y} = \sum_{i=1}^{48} \omega_{i,m,y} T_{i,m,y},$$

where $\omega_{i,m,y}$ is the share of national nonfarm employment in state i in month m of year y; and similarly for snowfall. We then run a simple time-series regression of an economic indicator on our national weather indexes:

$$3)\quad \Delta \log Y_{m,y} = \alpha + \beta T_{m,y} + \gamma S_{m,y} + \varepsilon_{m,y},$$

and later on allow for lags to capture the bounceback:

$$4)\quad \Delta \log Y_{m,y} = \alpha + \sum_{k=0}^{K} \beta_k T_{m-k,y} + \sum_{k=0}^{K} \gamma_k S_{m-k,y} + \varepsilon_{m,y}.$$

Table 5 presents the results. Overall, there are fewer statistically significant results, and often the temperature coefficient β has the "wrong" sign. For instance, the effect of snowfall on nonfarm employment is negative, but so is the temperature effect, so that this equation predicts that lower temperatures lead to higher employment. Both effects are statistically insignificant. This pattern is fairly general, though in the cases of industries or activities that are

TABLE 3

Effect of temperature and snowfall on state-level economic indicators with lags

Lag	Nonfarm employment	Unemployment rate	U.I. new claims	Housing permits	Housing starts
Temperature					
Current month (k = 0)	0.043***	– 0.011	–1.269***	1.568***	3.128***
	(0.007)	(0.247)	(0.280)	(0.683)	(0.915)
Last month (k = 1)	–0.009	–0.065	1.357***	–1.577**	–3.120***
	(0.007)	(0.262)	(0.347)	(0.724)	(0.888)
Two months ago (k = 2)	–0.025***	0.174	–0.053	–1.111***	–0.239
	(0.008)	(0.232)	(0.232)	(0.458)	(0.827)
Three months ago (k = 3)	0.006	0.104	0.109	0.033	–0.917
	(0.007)	(0.216)	(0.197)	(0.495)	(0.855)
Snowfall					
Current month (k = 0)	–0.027***	0.091	0.848***	–2.038***	–1.891***
	(0.004)	(0.152)	(0.189)	(0.445)	(0.598)
Last month (k = 1)	0.009	0.113	–0.275	0.944*	0.693
	(0.005)	(0.180)	(0.234)	(0.516)	(0.553)
Two months ago (k = 2)	0.011***	0.104	–0.457**	0.385	1.222*
	(0.006)	(0.149)	(0.190)	(0.548)	(0.665)
Three months ago (k = 3)	0.015***	–0.103	0.008	0.483	0.255
	(0.005)	(0.187)	(0.193)	(0.330)	(0.468)
Observations	36,964	22,466	25,036	19,852	25,612
R^2	0.294	0.449	0.152	0.120	0.101
Sample start	1950	1976	1971	1980	1970

Notes: Results from estimation of equation 2 using monthly data from November through March by ordinary least squares with state and time effects; standard errors in parentheses are two-way clustered by state and time. All left-hand-side variables are in log changes except the unemployment rate, which is in difference. U.I. indicates unemployment insurance. Temperature and snowfall indexes are normalized to have mean zero and standard deviation one for each state. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) level.
Source: Authors' calculations based on data from the National Climatic Data Center.

TABLE 4

Effect of temperature and snowfall on measures of state-level economic activity, pre- and post-1990

	Nonfarm employment	Unemployment rate	U.I. claims	Housing permits	Housing starts
Temperature, pre-1990	0.050***	–0.588	–1.369***	1.227	4.207***
	(0.011)	(0.373)	(0.359)	(1.371)	(1.350)
Temperature, post-1990	0.023***	0.374	–0.580*	1.049*	0.618
	(0.006)	(0.325)	(0.341)	(0.619)	(0.974)
Snowfall, pre-1990	–0.030***	0.056	0.663**	–0.987	–0.379
	(0.008)	(0.219)	(0.261)	(0.870)	(0.745)
Snowfall, post-1990	–0.026***	0.143	1.069***	–2.650***	–3.563***
	(0.006)	(0.214)	(0.261)	(0.632)	(1.062)
Observations	37,154	22,517	25,084	19,900	25,660
R^2	0.262	0.451	0.148	0.118	0.099

Notes: Results from estimation of equation 1 using monthly data from November through March by ordinary least squares with state and time effects; standard errors in parentheses are two-way clustered by state and time. The left-hand-side variables are all in log changes, except the unemployment rate, which is in level change. U.I. indicates unemployment insurance. Temperature and snowfall indexes are normalized to have mean zero and standard deviation one for each state and are interacted with two dummies, pre- and post- 1990. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) level.
Source: Authors' calculations based on data from the National Climatic Data Center.

TABLE 5

Effect of temperature and snowfall on national economic indicators

	Temperature		Snowfall		Observations	R^2
Nonfarm employment	−0.01	(0.028)	−0.009	(0.028)	854	0.002
Unemployment rate	−0.15	(0.151)	−0.16	(0.151)	791	0.000
Private nonfarm employment	−0.03	(0.031)	−0.014	(0.031)	854	0.001
Construction employment	0.163*	(0.098)	−0.191***	(0.098)	854	0.021
Retail sales (excluding cars)	0.03	(0.077)	−0.132*	(0.079)	571	0.011
Private average hours per worker	−0.05*	(0.028)	−0.181***	(0.029)	608	0.074
Industrial production (IP): Total	−0.061	(0.126)	−0.14	(0.126)	979	0.001
IP: Manufacturing	0.035	(0.137)	−0.13	(0.137)	979	0.002
IP: Utilities	−1.504***	(0.159)	−0.23	(0.167)	511	0.210
Lightweight vehicle sales	−0.97	(0.595)	−1.407**	(0.610)	572	0.009
CFNAI	−0.053	(0.085)	−0.12	(0.088)	570	0.003
New orders of core capital goods	−0.46	(0.442)	−1.573***	(0.450)	270	0.054
Shipments of core capital goods	−0.467*	(0.281)	−0.610**	(0.286)	271	0.017
Housing starts	0.92	(0.629)	−2.270***	(0.631)	667	0.055
Housing permits	0.58	(0.480)	−0.949**	(0.482)	655	0.022
Purchasing Managers Index	−0.069	(0.543)	0.45	(0.543)	791	0.002

Notes: Results from estimation of equation 3 using monthly data from November through March of all years by ordinary least squares (OLS). Standard errors in parentheses are simple OLS. All left-hand-side variables are in log changes, except for the unemployment rate (in change) and CFNAI (in level). Temperature and snowfall indexes are normalized to have mean zero and standard deviation one. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) level.
Source: Authors' calculations based on data from the National Climatic Data Center.

heavily affected by temperature or snowfall we do obtain clear and intuitive results. For instance, the coefficients on construction employment are similar to those obtained at the state level (0.163 on temperature and –0.191 on snowfall, compared with 0.185 and –0.181 for state-level data). Average hours worked, retail and car sales, housing starts and permits, and shipments and order of new capital goods are all affected significantly by snowfall. In the case of utilities production, the very strong negative effect of temperature is likely not an artifact but simply reflects the higher demand for heating. This shows that some sectors of the economy react positively to cold weather. Overall, the general message is that snowfall seems better at capturing the effect of weather on the economy, but these effects are more difficult to measure using aggregate data alone.

While it is reassuring that the magnitudes of the effects are similar (where available) in both exercises, this need not be the case. For instance, if there are spillovers across states such that bad weather in one state negatively affects economic activity in another state, and if the weather is positively correlated across the two states, our state-level regression would overestimate the effect of local weather. However, these spillovers are likely to be small.

Table 6 studies the bounceback in the national data by adding lags to equation 3. As in the state-level data, we find significant evidence of bounceback for the categories that are highly affected by temperature or snowfall. For instance, snowfall is estimated to reduce car sales by 1.3 percent on impact, but the bounceback is estimated to be 1.27 percent the next month. Similarly, average hours fall by 0.17 percent then rebound by 0.13 percent the next month. In many cases, however, the bounceback is estimated imprecisely, probably due to sparseness of data at the national level.

Revisiting the 2013–14 weather contribution

We are finally in a position to estimate the effect of the 2013–14 winter on economic activity. We present two sets of results—the first one based on the national model of the previous section and the second based on the state-level model. In all cases, we simply use the actual weather observed during the winter, together with the sensitivities estimated using historical data (that is, our estimates of β and γ), to obtain the effect of the observed weather on the growth rates of these economic indicators. Table 7 presents the results based on the state-level model (which is more precisely estimated), while table 8 (p. 18) shows the national results.[25] In the first row of table 7, we see that nonfarm employment displays the slowdown presented in figure 1 (p. 2): Employment growth rates of 0.06 percent in December and 0.1 percent in January were below the recent trend of about 0.16 percent (that is, 200,000 jobs created per month). The second row shows that, according to our estimates, the weather contributed negatively to the growth rate of nonfarm employment from November

TABLE 6

Effect of temperature and snowfall on national economic indicators with lags

	Temperature						Snowfall				
	Current month		Last month		Two months ago		Current month		Last month		Two months ago
Nonfarm employment	−0.002	(0.028)	−0.033	(0.028)	0.011	(0.027)	−0.011	(0.027)	0.011	(0.027)	0.027
Unemployment rate	−0.12	(0.152)	0.25	(0.152)	−0.21	(0.150)	−0.12	(0.150)	0.19	(0.151)	−0.21
Private nonfarm employment	0.004	(0.031)	−0.043	(0.031)	0.005	(0.030)	−0.016	(0.030)	0.002	(0.030)	0.032
Construction employment	0.193**	(0.092)	−0.252***	(0.092)	−0.082	(0.090)	−0.217**	(0.091)	−0.017	(0.091)	0.264***
Retail sales (excluding cars)	0.094	(0.078)	−0.10	(0.078)	−0.011	(0.077)	−0.12	(0.079)	0.194**	(0.080)	−0.015
Private average hours per worker	−0.014	(0.028)	−0.044	(0.028)	0.034	(0.027)	−0.170***	(0.028)	0.126***	(0.028)	0.008
Industrial production (IP): Total	−0.063	(0.120)	−0.061	(0.120)	0.037	(0.118)	−0.15	(0.118)	0.047	(0.119)	−0.003
IP: Manufacturing	0.037	(0.130)	−0.087	(0.130)	−0.006	(0.128)	−0.14	(0.128)	0.050	(0.129)	−0.004
IP: Utilities	−1.636***	(0.145)	0.906***	(0.146)	0.536***	(0.143)	−0.11	(0.152)	0.094	(0.152)	−0.24
Lightweight vehicle sales	−0.62	(0.605)	−0.23	(0.606)	0.001	(0.596)	−1.303**	(0.613)	1.274**	(0.621)	−0.051
New orders of core capital goods	−0.014	(0.087)	0.023	(0.087)	0.097	(0.086)	−0.11	(0.088)	0.179**	(0.089)	0.15
New orders of core capital goods	−0.34	(0.446)	0.19	(0.464)	−0.032	(0.462)	−1.621***	(0.467)	0.69	(0.467)	0.46
Shipments of core capital goods	−0.522*	(0.283)	0.11	(0.295)	0.608**	(0.294)	−0.659**	(0.297)	0.056	(0.296)	0.624**
Housing starts	1.345**	(0.620)	−1.783***	(0.621)	0.28	(0.610)	−2.375***	(0.616)	0.87	(0.620)	1.846***
Housing permits	0.70	(0.475)	−1.578***	(0.476)	1.239***	(0.470)	−1.114**	(0.472)	−0.18	(0.475)	2.004***
Purchasing Managers Index	−0.096	(0.552)	−0.036	(0.553)	−0.035	(0.543)	0.42	(0.544)	0.18	(0.548)	0.087

Notes: Results from estimation of equation 4 using monthly data from November through March of all years by ordinary least squares (OLS). Simple OLS standard errors are reported in parentheses. All left-hand-side variables are in log changes, except for the unemployment rate (in change) and CFNAI (in level). Temperature and snowfall indexes are normalized to have mean zero and standard deviation one. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) level.

Source: Authors' calculations based on data from the National Climatic Data Center.

TABLE 7

Estimated effect of 2013–14 winter using state model

		Nov.	Dec.	Jan.	Feb.	Mar.	Apr.	May
Nonfarm employment	Data	0.20	0.06	0.10	0.16	0.15	0.22	0.17
	Weather effect	−0.02	−0.02	−0.01	−0.04	0.00	0.03	0.03
Unemployment rate	Data	−0.20	−0.30	−0.10	0.10	0.00	−0.40	0.00
	Weather effect	−0.02	0.03	−0.05	0.06	0.03	−0.03	−0.27
New unemployment insurance claims	Data	−5.30	8.27	−9.27	5.43	−3.26	−2.29	−0.83
	Weather effect	0.24	0.22	−0.07	0.98	−0.31	−1.13	0.08
Housing permits	Data	−2.85	−1.46	−8.47	7.39	−1.09	5.73	−5.23
	Weather effect	−0.59	−0.32	0.95	−1.04	1.11	1.51	1.08
Housing starts	Data	16.60	−6.64	−14.21	3.40	2.34	11.24	−7.72
	Weather effect	−1.71	−0.27	0.85	−0.67	0.22	3.35	0.62

Notes: Based on state model with three lags. All results are in percentage growth rates, except for the unemployment rate, which is the change in percentage points.
Source: Authors' calculations based on data from the National Climatic Data Center.

through February, to the tune of 0.04 percent in February, or about 50,000 to 60,000 jobs. The weather effects are then reversed in April and May. However, the weather hardly accounts for the weak December and January employment numbers. Similarly, the unemployment rate grew by 0.06 percentage points in February due to weather, according to these estimates. Housing permits and starts were also affected in a significant way, but the estimated effects (about 1 percent) fall short of the observed magnitude of the decline in the data (14 percent for starts and 8 percent for permits in January, for instance).

In table 8, we see that the results with national data have the same flavor, but are less clear perhaps due to the imprecision of the estimation. For instance, the weather effect is now estimated to be positive for nonfarm employment during most months. However, this relies on an equation that was insignificant. More sensible results are obtained for construction employment, retail sales, average hours worked, and lightweight vehicle sales. For instance, hours fell 0.6 percent in February, of which 0.19 percent is attributed to the weather. Utility production grew 3.3 percent in January, of which 0.52 percent is attributed to the weather. The decline in starts and permits due to weather is about 3 percent. However, the timing does not fit the observed decline in indicators well. For instance, the CFNAI fell sharply in January, and our model attributes little of this to the weather; and housing starts and permits rebounded in February, contrary to our model's prediction. Overall, while some of the patterns observed in the data can be attributed in part to weather, this explanation is insufficient to explain the magnitude and timing of the slowdown.

How does climate change affect our results?

It is important to note the potential impact of climate change on our study. When we construct our weather index, we normalize by a base value, which we take to be simply the long-run average (1950–2014). However, it is conceivable that given rising global temperatures, the typical temperature in the United States increased during the period of observation. This would make, for example, a 25-degree day in November more anomalous in 2014 than in 1950. As noted above, our weather data are not adjusted for changes in instrumentation and other measurement issues. Without these adjustments, it is difficult to detect a trend in temperature.[26] However, in some cases it is possible to observe a positive trend starting in 1980, which is consistent with the evidence on climate change on the United States. To assess the effect of this potential trend on our results, we fitted a linear trend starting in 1980 to each weather index and reestimated our models. All of our results are nearly unaffected by this modification. This is not surprising since the effect of weather is intuitively identified using the short-run deviations of weather, which are much larger than the trend. Incorporating the trend has one significant consequence: It makes the 2013–14 winter look even harsher; that is, the weather deviation from normal is larger due to the positive trend. This implies that our estimated effect of that winter weather is larger, by about 20 percent, than we discussed in the previous section.

Conclusion

Our results overall support the view that weather has a significant, but short-lived, effect on economic activity. Except for a few industries, which are affected importantly (such as utilities, construction, hospitality,

TABLE 8

Estimated effect of 2013–14 winter using national model

		Nov.	Dec.	Jan.	Feb.	Mar.	Apr.	May
Nonfarm employment	Data	0.20	0.06	0.10	0.16	0.15	0.22	0.17
	Weather effect	0.01	0.02	−0.01	0.02	0.07	0.08	−0.02
Unemployment rate	Data	−0.20	−0.30	−0.10	0.10	0.00	−0.40	0.00
	Weather effect	0.16	−0.31	0.25	−0.19	0.31	−0.57	0.29
Retail sales (excluding cars)	Data	−0.38	0.55	−0.47	0.36	0.81	0.72	0.25
	Weather effect	−0.02	−0.11	−0.02	−0.07	0.35	0.09	0.02
Private average hours per worker	Data	0.30	−0.60	0.30	−0.60	0.89	0.00	0.00
	Weather effect	0.09	−0.10	−0.08	−0.19	0.30	0.03	−0.04
Industrial production (IP)	Data	0.59	0.20	−0.30	0.98	0.78	0.10	0.48
	Weather effect	0.12	−0.02	−0.06	−0.17	0.21	0.03	−0.05
IP: Manufacturing	Data	0.31	0.10	−0.93	1.23	0.81	0.30	0.40
	Weather effect	0.04	−0.04	−0.08	−0.20	0.14	0.10	0.01
IP: Utilities	Data	1.85	0.10	3.32	−0.28	−0.47	−5.30	0.20
	Weather effect	1.39	−0.14	0.52	0.15	0.88	−2.05	−0.65
Lightweight vehicle sales	Data	5.79	−4.74	−1.60	0.90	6.88	−2.81	4.29
	Weather effect	1.11	−0.79	0.13	−0.77	3.54	0.05	0.00
CFNAI	Data	0.71	−0.19	−0.85	0.55	0.53	0.15	0.18
	Weather effect	0.06	−0.15	−0.15	−0.03	0.41	0.15	−0.14
New orders on core capital goods	Data	5.72	−0.88	−1.90	0.10	4.58	−1.10	−1.41
	Weather effect	1.03	−1.15	−1.05	−2.18	2.21	0.60	−0.01
Shipments of core capital goods	Data	2.81	0.38	−1.91	0.83	2.16	−0.31	0.08
	Weather effect	0.73	−0.23	−0.97	−0.84	0.81	0.57	−0.83
Housing starts	Data	16.60	−6.64	−14.21	3.40	2.34	11.24	−7.72
	Weather effect	0.01	−0.72	−2.98	−2.77	3.10	5.50	−0.55
Housing permits	Data	−2.85	−1.46	−8.47	7.39	−1.09	5.73	−5.23
	Weather effect	−0.05	0.51	−2.87	−1.19	0.95	4.87	−1.78
Purchasing Managers Index	Data	57.00	56.50	51.30	53.20	53.70	54.90	55.40
	Weather effect	−0.12	0.20	0.53	1.14	0.56	0.22	0.03

Notes: Based on national model with two lags. All results are in percentage growth rates, except the unemployment rate, which is the change in percentage points.

Source: Authors' calculations based on data from the National Climatic Data Center.

and, to a lesser extent, retail), the effect is not very large, so that even the fairly bad weather during the 2013–14 winter cannot account entirely for the weak economy during that period. Other factors must have been at play. Indeed, the National Income and Product Accounts data suggest that an important share of the slowdown in the first quarter was driven by an inventory correction and the effect of foreign trade. Another simple hint that something more than the weather was at play is that the timing of the decline, measured in economic statistics in the period December through March, was uneven across indicators: Some declined in December and January, others in January and February, and so on, which seems inconsistent with a simple weather story. There are several directions in which it would be interesting to extend this work. First, better weather indexes could be constructed by weighting station data using very local employment. The importance of nonlinearities could also be studied in more detail, as could the differences across states in sensitivities to weather. Finally, local measurement of production and sales would enable us to extend this study and consider more outcomes.

NOTES

[1]These indicators come from a variety of data sources, including private or government surveys, trade associations, or administrative data. These statistics are followed closely by investors because they are released often and with little lag, and hence are more timely and less subject to revisions than the broader and more comprehensive measures such as gross domestic product.

[2]FOMC statements noted starting in December 2013 that "asset purchases are not on a preset course, and the Committee's decisions about their pace will remain contingent on the Committee's outlook for the labor market and inflation as well as its assessment of the likely efficacy and costs of such purchases." See www.federalreserve.gov/newsevents/press/monetary/20131218a.htm.

[3]Recently, there has been some renewed interest by economists in the question of how weather affects the economy, but this research was not relevant for the issues at hand, as we explain.

[4]The minutes from the March 2014 meeting provided more detail: "The information reviewed for the March 18–19 meeting indicated that economic growth slowed early this year, likely only in part because of the temporary effects of the unusually cold and snowy winter weather. ... The staff's assessment was that the unusually severe winter weather could account for some, but not all, of the recent unanticipated weakness in economic activity, and the staff lowered its projection for near-term output growth. ... Most participants noted that unusually severe winter weather had held down economic activity during the early months of the year. Business contacts in various parts of the country reported a number of weather-induced disruptions, including reduced manufacturing activity due to lost workdays, interruptions to supply chains of inputs and delivery of final products, and lower-than-expected retail sales. Participants expected economic activity to pick up as the weather-related disruptions to spending and production dissipated." See www.federalreserve.gov/monetarypolicy/fomcminutes20140319.htm.

[5] See www.federalreserve.gov/mediacenter/files/FOMCpresconf20140319.pdf.

[6]The minutes noted that "the information reviewed for the April 29–30 meeting indicated that growth in economic activity paused in the first quarter as a whole, but that activity stepped up late in the quarter; this pattern reflected, in part, the temporary effects of the unusually cold and snowy weather earlier in the quarter and the unwinding of those effects later in the quarter." See www.federalreserve.gov/monetarypolicy/fomcminutes20140430.htm.

[7] See www.federalreserve.gov/newsevents/speech/yellen20140416a.htm.

[8]Published on the *New York Times* website September 26, 2014; available at www.nytimes.com/2014/09/27/upshot/gdp-report-emphasizes-the-problem-of-conflicting-economic-signals.html.

[9]See, for instance, Gallup, Sachs, and Mellinger (1999).

[10]The data are available at ftp://ftp.ncdc.noaa.gov/pub/data/ghcn/daily. Our data set is version 3.12, retrieved in September 2014.

[11]One important data issue is that until recently, snowfall was often not reported unless it was snowing; that is, the data are reported as missing rather than zero. As we believe is standard practice, we attribute a zero snowfall to all missing observations (which may include some observations for which no data were actually observed).

[12]As a result of the lack of adjustments, our data do not exhibit very clear increases in average temperature. We believe the adjustments, while critical for the measurement of the trend in average temperature, are not important for the measurement of short-term weather. We discuss this in more detail in the last section.

[13]We calculate the daily temperature as the simple average of the minimum and maximum daily temperature, that is,

$$T = \frac{T_{\max} + T_{\min}}{2}.$$

[14]The underlying issue is whether normalizing helps capture the effect of unusual weather on economic activity. We hypothesize that economies in highly variable climates have adapted: For example, states with highly variable levels of snowfall may have the infrastructure in trucks and salt to deal with large snowfall events. This is largely an empirical question. In some explorations, we found that the precise normalization was not critical to our result, but this is an area that deserves future research.

[15]The log change approximates the percentage change in the variable Y, while reducing the effects of outliers and heteroskedasticity.

[16]However, it may be that economic activity in state i depends on weather in other states, for example, because of supply chains or because lower retail sales in one state affect production in another state. Because weather may be correlated across states, this could lead to a bias.

[17]The error terms $\varepsilon_{i,m,y}$ may be correlated across states and over time; we adjust the standard errors to take this into account using two-way clustering.

[18]Technically, this equation requires defining $T_{i,m-k,y} = T_{i,m-k+12,y-1}$ if $m - k \leq 0$.

[19]We are not aware of monthly data available on sales or production at the state level.

[20]The breakdown of employment by industry at the state level and at the monthly frequency is only available for the period 1990–2014, so we have fewer data and consequently fewer statistically significant results. The sample size varies further by industry because the Bureau of Labor Statistics' establishment survey does not report employment for some industries in some states.

[21]This is expected since our indexes of temperature and snowfall exhibit relatively little serial correlation; hence, adding lagged values to the regression does not affect the same-month impact estimates since current weather and lagged weather are roughly orthogonal.

[22]The only marginal case is the effect of temperature on housing permits, which is significant at the 7 percent level.

[23]Note, however, that the precision of the estimates does not permit us to rule out a small long-run effect.

[24]Technically, we interact both of our weather indexes with two dummies, pre- and post-1990, and run a single regression for each economic indicator.

[25]We construct the national implied weather effects from the state model by weighting the state-level predictions to adjust for the state size.

[26]There appears to be no trend in precipitation, even in "adjusted" data.

REFERENCES

Boldin, Michael, and Jonathan H. Wright, 2015, "Weather-adjusting employment data," Federal Reserve Bank of Philadelphia, working paper, No. 15-05, January.

Colacito, Riccardo, Bridget Hoffman, and Toan Phan, 2014, "Temperatures and growth: A panel analysis of the U.S.," University of North Carolina at Chapel Hill, working paper, December 18.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken, 2014, "What do we learn from the weather? The new climate-economy literature," *Journal of Economic Literature*, Vol. 52, No. 3, September, pp. 740–798.

__________, 2012, "Temperature shocks and economic growth: Evidence from the last half century," *American Economic Journal: Macroeconomics*, Vol. 4, No. 3, July, pp. 66–95.

Deryugina, Tatyana, and Solomon M. Hsiang, 2014, "Does the environment still matter? Daily temperature and income in the United States," National Bureau of Economic Research, working paper, No. 20750, December.

Deschênes, Olivier, and Michael Greenstone, 2007, "The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather," *American Economic Review*, Vol. 97, No. 1, March, pp. 354–385.

Gallup, John Luke, Jeffrey D. Sachs, and Andrew D. Mellinger, 1999, "Geography and economic development," *International Regional Science Review*, Vol. 22, No. 2, August, pp. 179–232.

Macroeconomic Advisers, 2014, "Elevated snowfall reduced Q1 GDP growth 1.4 percentage points," April 15, available at www.macroadvisers.com/2014/04/elevated-snowfall-reduced-q1-gdp-growth-1-4-percentage-points/.

Derivatives and collateral at U.S. life insurers

Kyal Berends and Thomas B. King

Introduction and summary

Insurance companies serve the important economic role of helping businesses and households to insulate themselves against risks. But these risks do not disappear from the economy—they remain on insurers' books, necessitating careful risk management among insurers themselves. Over the past two decades, one way that insurers have managed risk is through the use of derivative contracts,[1] which derive their value from the performance of an underlying entity. This underlying entity can be an asset, index, or interest rate. Some of the more common derivatives include forwards, futures, options, and swaps. Most derivatives, including interest rate swaps (IRS), have historically been traded over the counter (OTC) rather than on centralized exchanges.

The use of derivatives comes with its own set of costs related to the transaction, management, and collateralization of positions. With the implementation of the Dodd–Frank Act of 2010, those costs seem certain to rise. Among other provisions, the law requires the central clearing of certain types of OTC derivatives and mandates that those transactions must satisfy margin requirements that will in most cases require counterparties to post more collateral than was previously the case.[2] Forthcoming rules will impose additional collateral requirements on derivatives positions for which the central clearing mandate does not apply. Thus, the new rules for both cleared and noncleared derivatives could generate new costs for insurers or require changes in their business practices.

In this article, we review life insurers' use of OTC derivatives and discuss the possible implications of these new rules for their financial condition.[3] Although insurers represent a relatively small part of the derivatives markets, they are an interesting case study, in part because they report very detailed information about their derivatives positions and associated collateral in quarterly regulatory filings. We exploit these data to study how derivatives are used by insurers and to get a quantitative sense of what the new regulations are likely to imply for their business models.

The new regime poses several potential costs for insurers. For example, like many market participants, insurers will face a short-term fixed cost of adjusting operations and corporate structure to meet the new clearing and collateralization requirements, as well as ongoing expenses associated with trading, collateral optimization, and back-office functions; and insurers may also face some regulatory capital consequences.

Kyal Berends is a former senior associate economist and Thomas King is a senior financial economist in the Economic Research Department at the Federal Reserve Bank of Chicago. *The authors thank Anna Paulson, Rich Rosen, and other members of the Chicago Insurance Initiative for helpful feedback.*

ISSN 0164-0682

TABLE 1

Life insurers with the largest OTC derivatives portfolios

	Notional OTC derivatives	Statutory assets
	(- - - - *dollars in billions* - - - -)	
MetLife Inc.	188	603
Manulife Financial Corp.	151	267
Massachusetts Mutual Life Insurance Co.	137	202
New York Life Insurance Group	104	261
Nationwide Mutual Group	72	132
Voya Financial Inc.	71	193
Ameriprise Financial Inc.	65	110
AEGON	57	202
Lincoln National Corp.	52	222
Prudential Financial Inc.	44	545
Jackson National Life Group	29	186
Principal Financial Group Inc.	24	149
Allianz Group	21	116
Genworth Financial Inc.	19	70
AXA	16	166
Hartford Financial Services	12	179
American International Group	11	269
Aflac Inc.	7	111
Delaware Life Partners LLC	6	42
Sun Life Financial Inc.	5	19

Notes: OTC indicates over the counter. Includes interest rate swaps, caps, floors, collars, and swaptions; credit default swaps; total return swaps; and inflation-linked products. Data as of 2014:Q3.
Source: Statutory filings via SNL Financial.

In this article, however, we focus on one particular set of costs that has received attention, namely, costs related to reallocating insurers' portfolios to high-quality—and therefore low-yielding—assets in order to meet margin requirements.[4] We find that, overall, the requirements are unlikely to generate large costs for the industry as a whole through this channel—although there are some low-probability tail scenarios in which they could result in substantial forgone investment income for a few larger insurers. This finding is largely due to the fact that insurance companies already hold large amounts of high-quality unencumbered securities that could be pledged for this purpose, and indeed they may be natural collateral providers to other market participants.[5]

After reviewing insurers' use of derivatives and collateral in the following section, we develop a Monte Carlo exercise to attempt to quantify the amount of margin posted and revenue lost due to required margin under different scenarios for interest rates and insurer portfolio evolution. Then, we consider some ways that insurers may adjust their business practices in light of the new regulations. Two likely responses are to reduce the need for hedging by shifting more interest rate risk onto consumers or markets and to build up new sources of liquidity to cover cash needs. Depending on how these adjustments play out, they could expose insurers and their counterparties to new risks, especially in a crisis environment in which liquidity is constrained.

Life insurers' use of OTC derivatives

Insurers use a variety of types of derivatives for hedging different types of risks. Some of these derivatives, such as equity options and currency swaps, are typically exchange traded and are not affected by the Dodd–Frank rules. In this article, we focus on the interest rate and credit derivatives that are traded OTC, because those are the contracts to which the central clearing and collateralization requirements apply. Table 1 lists the 20 life insurance companies that participate most in the OTC derivatives market, as measured by the gross notional value of their positions in these instruments.[6] These companies include the largest insurers by assets, but derivatives usage is not perfectly correlated with firm size—it depends on a variety of factors, including lines of business and corporate structure. For example, some very large insurers, including TIAA-CREF and Northwestern Mutual, have OTC derivatives positions that are too small to be included in the table.

As a whole, the life insurance industry held \$1.1 trillion of notional value in OTC derivatives as of September 2014. For a sense of scale, we note the statutory assets of these companies totaled \$6.1 trillion. Relative to other market participants, such as commercial banks, the OTC derivatives portfolios of life insurers are relatively modest. The gross market value of their swaps positions was only about \$13.2 billion, and the net positions are likely smaller still.[7] However, derivatives portfolios are highly concentrated—over 50 percent of notional value of OTC derivatives in the industry is held by the four insurers with the largest swaps portfolios (MetLife, Manulife, Mass Mutual, and New York Life). The companies in the table collectively hold 97 percent of the industry's OTC derivatives.

Large insurance operating companies often reside within even larger, complex corporate structures. Thus, derivatives positions at the operating company may not give a complete picture of the derivatives activity at the whole firm. For companies that are publicly traded in the United States, it is possible to obtain some information on consolidated derivatives positions from SEC filings, although this information is not as detailed as what is available from regulatory reports. Table 2 shows the sum of interest rate and credit derivative positions for the largest companies for which such information is available. For these eight firms (which together hold about half of the positions listed in table 1), derivative exposures that are in subsidiaries other than life

TABLE 2

Selected operating company versus consolidated derivatives positions

	Operating company	Consolidated
	(- - - - - - - - - *dollars in millions* - - - - - - - - -)	
Prudential	50,179	316,283
MetLife	189,881	267,155
Manulife	169,550	209,486
AIG	17,685	90,446
Voya	69,773	73,614
Lincoln	60,009	56,864
Hartford	11,645	30,715
Principal	24,063	25,426

Notes: Data as of 2014:Q3. Includes all interest rate and credit derivatives. Operating company amounts may not match those in table 1 due to imperfect overlap between these categories and OTC derivatives.
Sources: Statutory filings and 10K reports via SNL Financial.

insurance operating companies constitute 55 percent of the holding companies' notional positions.

As shown in table 3, insurers' interest rate swaps positions at the industry level are roughly equally balanced between paying and receiving fixed rates.[8] This pattern also holds at the individual insurer level: A typical firm both receives and pays fixed rates. However, as we discuss later, the simultaneous positions in opposite directions reflect the hedging of different types of risk and, consequently, typically differ by maturity. Insurers also hold fairly large positions—about $327 billion in notional value—in other types of interest rate derivatives, especially caps and call swaptions, which hedge against rising rates. They also hold small amounts of total return and credit default swaps, which are used for asset replication purposes as well as hedging, and a smattering of miscellaneous products.

To understand the potential impact of the new collateral rules on the insurance industry, it is useful to review how these derivative positions function within insurers' business models. Insurers take very few directional positions using derivatives, relying on them almost entirely for hedging purposes. In particular, they hedge four broad types of risk.[9] First, they hedge the interest rate risk of their fixed-income portfolios. As of September 2014, the insurers in table 1 collectively held nearly $1 trillion in various types of bonds, exposing them to rising interest rates.[10] They use pay-fixed interest rate swaps and other interest rate derivatives to hedge against this risk. Statutory data on hedging purpose (not shown) indicate that about half of OTC derivatives positions serve this function.

Second, insurers hedge the risks of deposit-like liabilities, including funding agreements and guaranteed interest contracts (GICs). These may pay fixed or floating rates and span a spectrum of maturities, although they are typically much shorter than insurers' other liabilities. These contracts may also have option-like features that require more complex hedging strategies. Some of these strategies may involve relatively exotic derivatives for which central clearing is not available.

Third, insurers attempt to match the duration of their long-term insurance and annuity liabilities. For the simplest contracts, hedging these exposures involves receiving interest payments to match the payments that the firm is required to make. But, for most insurance and annuity products, cash flows are uncertain. Thus, unlike the security-specific hedging on the asset side, liability hedging can only be done imperfectly in an economic sense, since there is significant uncertainty about the timing and duration of future insurance claims. As shown in table 3, insurers are, on net, receivers of fixed payments in swaps, implying that on balance they are using swaps to add duration to their portfolios. This makes sense as many life insurance liabilities are very long duration—indeed, in some cases longer than can be achieved by buying fixed-income products in the cash market.

Finally, insurers hedge the optionality of their liabilities. This optionality can take a variety of implicit and explicit forms. For example, it is common for insurance companies to offer minimum-return guarantees on variable annuities, which they in turn hedge with a combination of OTC and exchange-traded derivatives. Furthermore, most annuities may be surrendered at the option of the beneficiary. Fixed-rate annuities are more likely to be surrendered when interest rates rise, precisely when they are most attractive from the issuer's point of view. Most of the caps, floors, and swaptions reported in the table are also used to hedge these types of risk, and interest rate swaps may be used as part of the strategy. Many of the nonoperating-company positions shown in table 2 are likely held by captive reinsurers, which also principally use them for this type of hedging.

It is important to recognize that derivatives portfolios reflect a mix of risk mitigation, accounting, and regulatory considerations. In particular, under FAS 133, insurers can receive hedge accounting treatment for derivatives positions that are deemed "effective hedges," and a similar treatment applies in statutory accounting. For example, insurers discount the value of future claims on insurance policies using an assumed maturity structure and discount rate, and they can receive hedge accounting treatment by entering into (usually long-dated) swaps that match these terms. Although long-term bonds might be able to match the duration

TABLE 3

Characteristics of insurer OTC derivatives portfolios

	Notional amount (millions of dollars)	Fair value	Maturity (% of notional) < 1 year	1–3 years	3–7 years	7–15 years	>15 years
Interest rate swaps	705,229	11,121	6	15	19	27	33
Receive-fixed	346,373	19,411	3	12	17	29	39
Pay-fixed	296,358	–9,570	9	20	22	23	25
Type not reported	62,498	1,279	7	10	16	35	32
Other rate products	326,961	1,696	27	30	27	12	3
Floors and puts	87,675	796	53	30	3	13	1
Caps and calls	184,603	716	16	33	39	11	1
Other	54,683	184	26	20	26	12	15
Credit default swaps	25,896	209	17	21	56	3	3
Bought protection	3,638	–33	33	27	31	2	7
Sold protection	16,871	183	19	19	55	4	2
Type not reported	5,386	59	0	24	73	2	0
Miscellaneous	32,957	–28	73	4	3	7	13

Notes: Includes data as of 2014:Q3 from the 20 life insurance operating companies with the largest OTC derivatives portfolios, as measured by notional value. Other rate products include interest rate collars and swaptions classified as "other." Miscellaneous includes total return swaps and inflation swaps.
Source: Statutory filings via SNL Financial.

of those same positions reasonably well and thus hedge them in an economic sense, such a strategy would not qualify for hedge accounting treatment. Insurers may have incentives to engage in offsetting swaps contracts to hedge both sides of the balance sheet to recognize accounting benefits.

One should also bear in mind that insurers' derivatives use takes place against a backdrop of regulatory controls. Some states require insurers to maintain a strict "derivatives use plan" that must meet with the approval of supervisors, and they also set limits on the quantity of derivatives activity. For example, New York prohibits swaps holdings with potential exposure in excess of 3 percent of admitted assets. ("Potential exposure" is a regulatory measure of the total amount of risk posed by an insurer's derivatives book.) Insurers must therefore choose carefully which risks to hedge and how best to use their limited derivatives capacity.

Margin requirements and Dodd–Frank

Because participants in derivatives contracts have risk exposures to their counterparties, they are typically required to post some form of collateral to each other. The Dodd–Frank Act standardizes these requirements for OTC derivatives transactions. Collateral requirements associated with derivatives trades are of two types. *Variation margin* captures the marked-to-market change in the value of positions on a daily or, in exceptionally volatile periods, intraday basis. This is meant to ensure that in the event of a default by one counterparty, the other counterparty can recover the fair value of the position. *Initial margin* is intended to cover possible losses incurred by the remaining counterparty *after* default, as it goes about liquidating or replacing the defaulted position. Thus, initial margin is typically calculated by assuming a certain amount of time for liquidation and using the data to estimate a worst-case scenario for the price moves of the position.

Even prior to the Dodd–Frank rules, it was standard for OTC derivatives counterparties to post some form of variation margin, and the exchange of initial margin was also common.[11] However, derivatives counterparties typically had a fair amount of leeway in how these requirements were satisfied. For example, they may have been able to post a variety of collateral types as margin or, depending on their bilateral agreements, post margin only when the change in the fair value of the position exceeded some threshold. Figure 1, panel A, shows margin posted by insurance companies in support of derivatives since 2013:Q1, when these data were first collected.[12] Figure 1, panel B, shows the collateral breakdown as of 2014:Q3. Note that, although variation margin constitutes the bulk of insurers' collateral positions, very little of this collateral consists of cash. This reflects the fact that most derivatives on insurers' books, if they require collateral at all, allow variation margin to be posted in the form of a range of securities.

Since June 10, 2013, new plain vanilla IRS and CDS index positions covered under Dodd–Frank have had to be cleared by a central counterparty (CCP) and collateralized accordingly. In particular, CCPs must

FIGURE 1

Fair value of collateral pledged by life insurance companies

A. Collateral posted over time

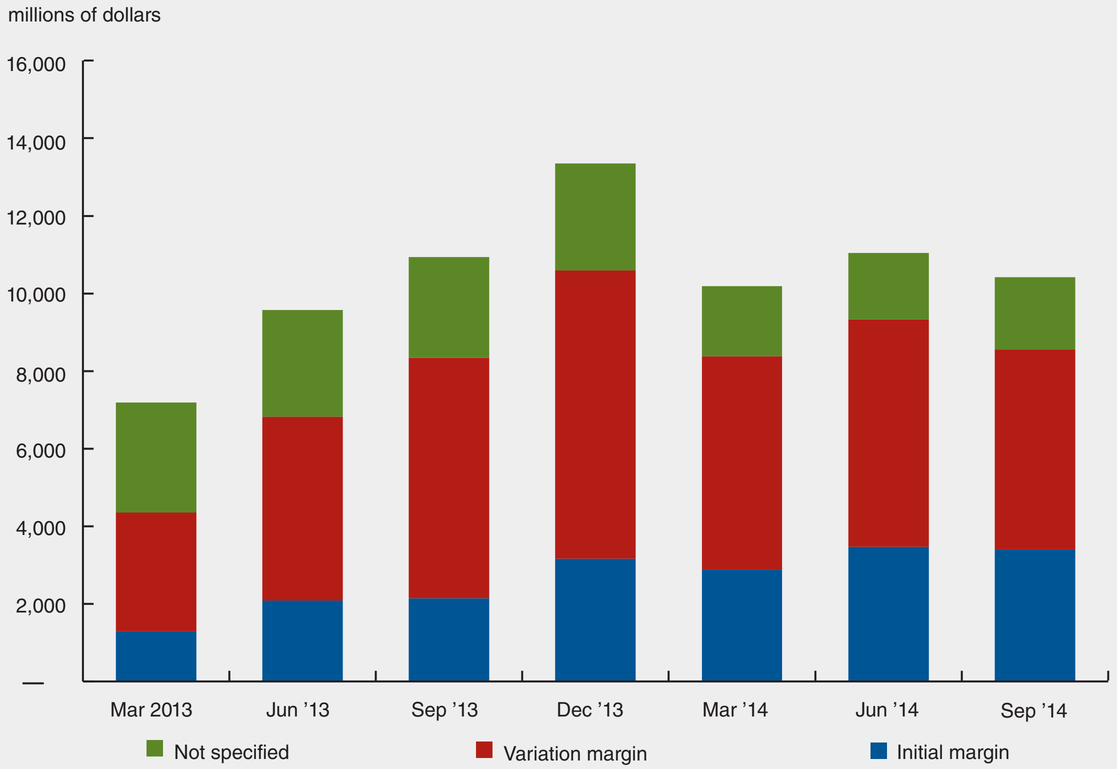

B. Securities posted as margin

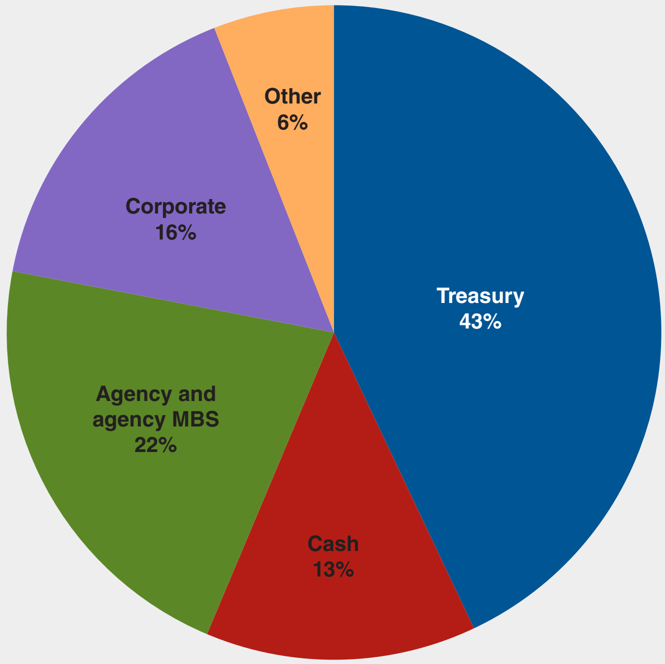

Notes: Data for 20 largest OTC derivatives users. Panel B as of September 2013. MBS indicates mortgage-backed securities.
Source: Statutory filings via SNL Financial.

require counterparties to post initial margin sufficient to cover a hypothetical five-day liquidation period with at least a 99 percent level of confidence and variation margin to cover daily fluctuations in the market value of positions. Forthcoming rules on uncleared trades are likely to impose a similar requirement for variation margin and a more stringent ten-day liquidation period for initial margin.

As shown in figure 1, margin posted by insurance companies to cover derivatives positions has indeed risen notably since the first quarter of 2013. In the 18 months surrounding the implementation date, insurers increased the collateral posted with derivatives counterparties by 45 percent, from $7.2 billion to $10.4 billion. Although both initial margin and variation margin have increased significantly, variation margin has fluctuated more. This is because variation margin is heavily influenced by external factors, such as interest rates. This volatility is suggestive of one type of risk that insurers now face—large movements in interest rates can require the transfer of large quantities of securities and, especially, cash into margin accounts. The following section discusses the scope of this risk in greater detail.

The types of collateral that can be posted to cover margin for cleared contracts, and the haircuts that apply, vary across CCPs. Initial margin is most often satisfied by high-quality securities, such as U.S. Treasury securities, although at least one major CCP has begun accepting investment-grade corporate bonds (within certain limits and subject to steep haircuts). In contrast, variation margin must be covered by cash. Moreover, the time frame within which clearing members must post variation margin after receiving a margin call is typically very short, often a matter of hours. (For uncleared trades, proposed rules would require most insurance companies—as "low-risk end-users"—only to update variation margin once per week and when the values involved rise above some *de minimus* amount.)

The burden of initial margin requirements is reduced to a degree by the possibility of netting potential moves in negatively correlated positions against each other. For example, if an insurer engages in a receive-fixed swap and a pay-fixed swap on similar terms with the same counterparty, that counterparty should expect price movements in the two contracts to offset exactly. Consequently, the margin needed to cover the position as a whole should be minimal, even though the margin needed to cover each swap individually might not be. For cleared trades, the extent to which such gains are available depends on the CCP's rules and models. For uncleared trades, the potential for margin offsets depends on the extent and terms of master netting agreements. In both cases, it also depends on the degree to which positions are concentrated at particular counterparties, since it is generally not possible to recognize portfolio-margining benefits from offsetting positions at different counterparties.[13]

The potential costs of the new collateral and clearing requirements span a variety of operational and economic considerations and are discussed more fully in a later section. It is clear, however, that the incidence of these costs—and, therefore, the nature of the industry's response—will depend greatly on the quantity of collateral that insurers end up having to post. We turn to this question next.

Collateral needs under alternative scenarios

In this section, we attempt to quantify the amount of collateral that may be necessary for life insurers to provide in support of cleared swaps positions in coming years. The results are essentially the product of three inputs: 1) a distribution for the possible path of interest rates; 2) calculations of how the value of each derivative contract type responds to the various interest rate configurations; and 3) an assumption regarding how insurers' derivative positions will evolve over time. Given institutional shifts in the industry and limited historical data, the last item is the most difficult of the three to ascertain. Therefore, we consider two different scenarios for the changes in the industry's derivatives mix that likely bracket the possibilities.

We summarize the methodology briefly here and describe it in detail in the appendix.

Model setup

For each of the 20 firms with the largest OTC derivatives holdings, as measured by notional value, we break down the interest rate swaps portfolio based on derivative type (pay- versus receive-fixed), maturity, and time since the contract was originated. We take the granularity of maturity buckets and contract ages to be annual, and we assume that the maximum maturity is 30 years.[14] For each type, we approximate each firm's notional holdings using a beta distribution over maturities, based on the 2014:Q3 data that were summarized at an aggregate level in table 1 (p. 22). Our assumptions about how this distribution evolves generate flows of derivatives originations and terminations in each year in our simulations for each firm in each type/maturity bin. Knowing the flows allows us to back out the distribution of contract ages for each swap bin. Since swaps valuation depends on the contract's remaining maturity and the fixed rate that applies to it, we can track the distribution of

TABLE 4

Five-day, 1 percentile fair-value changes for interest rate swaps of various terms

	Remaining maturity (years)				
Fixed rate (%)	1	3	7	15	30
2	–0.2	–0.9	–2.0	–3.2	–3.9
4	–0.2	–0.9	–2.2	–3.6	–5.1
6	–0.2	–0.9	–2.2	–4.1	–5.9
8	–0.2	–0.9	–2.5	–4.8	–7.6
10	–0.2	–0.9	–2.5	–4.9	–8.7

Notes: Considers rate changes from the average 2014:Q4 level of interest rates. Distribution of rate changes is based on daily data, April 1, 2004–September 30, 2014.

Source: Authors' calculations based on interest rate data provided by the Board of Governors of the Federal Reserve System.

swap rates within each bin, given a path of historical interest rates.

We assume that all swaps held by insurers are "plain vanilla" interest rate swaps (meaning standard contracts that exchange fixed and floating payments based on commonly used benchmarks and schedules). This assumption allows us to calculate the net present values of these contracts analytically, given an interest rate path. Furthermore, this assumption implies that the collateral requirements associated with central clearing apply to all of those contracts that are originated going forward.[15] We do not consider contracts originated prior to 2013:Q2 because, although many such contracts do involve margin agreements between the counterparties, the Dodd–Frank rules only require insurance companies to clear and post margin on plain vanilla swaps originated after June 2013.

Dodd–Frank mandates initial margin sufficient to cover a five-day liquidation period on cleared trades. To give a sense of the magnitudes involved, we calculate the range of initial margin values that could apply to swaps of various maturities and rates. Specifically, we calculate the distribution of five-day changes in value by drawing random five-day yield curve changes from the last ten years of data, and we apply these changes to rates that start at their 2014:Q1 level. Table 4 shows the resulting 99.7 percent quantiles, corresponding roughly to the levels of confidence used by CCPs.[16] In the absence of netting, the total initial margin required on a particular portfolio in the current interest rate environment would simply be given by the margin rates listed in the table, weighted by the amount of the portfolio in each corresponding bin. However, when calculating initial margin, CCPs and other counterparties generally allow for possible negative correlations between value changes for different derivatives positions in the same portfolio. This implies that one needs to evaluate the distribution of outcomes at the *portfolio* level. Our simulations of initial margin do this for each firm, at each date, for each simulated path of interest rates, based on our projections for how the distribution of swaps to which Dodd–Frank applies evolves over time.

Calculating variation margin, given a path of interest rates, is somewhat easier, since variation margin is simply equal to the net fair value of the swaps positions. Thus, for each firm, at each date, for each simulated path of rates, we calculate the net present value of swaps of each age in each type/maturity bin. Total variation margin is the sum of these values across swaps originated after 2013:Q2, weighted by the respective portfolio shares.

Interest rate simulations

We begin our computations in 2013:Q2. For the calculations through 2014:Q3, we use actual data on the yield curve to price the swaps portfolios. For projections beyond that date, we estimate a vector autoregression on Treasury forward rates, the Moody's Baa corporate bond yield, gross domestic product (GDP), and PCE inflation (based on the Personal Consumption Expenditures Price Index). We then take 10,000 draws from the estimated residual distribution and simulate forward ten years beginning with 2014 data. Any time a simulation results in a nonpositive rate in any quarter, we discard it and draw again. Figure 2 shows the distribution of simulated rate paths.

Portfolio scenarios

The amount of margin that will need to be posted against derivatives positions will depend crucially on how insurers adjust their derivatives portfolios going forward. The most natural assumption about this behavior may be simply that they keep the distribution across contract types and maturities unchanged at its current level, and this is indeed the first scenario that we consider. However, market participants generally anticipate that the net duration of the portfolio will lengthen going forward. This is also the situation in which margin is potentially greatest in the rising-interest-rate environment that we consider, and so it is worth modeling from a stress-testing point of view. Our second scenario is a variant of this outcome, in which insurers take new long positions by passively rolling over their maturing derivatives.

Specifically, in our "constant maturity distribution" scenario, we assume that the distribution of the *stock* of derivatives (that is, the percentage of the total in each type/maturity bin) is static. This means that the *flows*—that is, the amount of contracts originated or extinguished in each quarter—must generally be nonzero. We assume that the *gross* flows (the amount of notional value originated and canceled) are the

FIGURE 2

Distribution of simulated interest rate paths

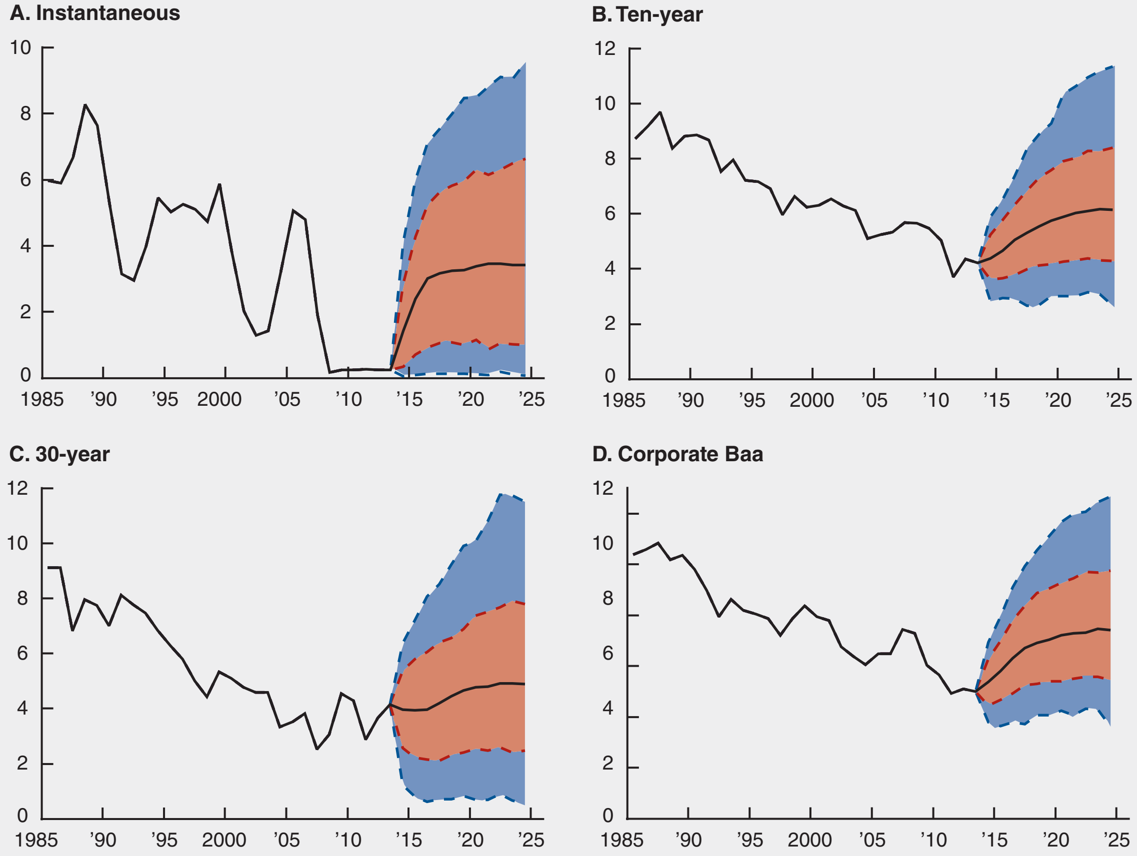

Notes: The figures show selected interest rate forecasts from the VAR model, with the zero lower bound imposed, based on 10,000 Monte Carlo draws. Solid black lines are means; shaded regions are 10–90% and 1–99% confidence bounds.
Sources: Board of Governors of the Federal Reserve System and authors' calculations.

minimum possible to achieve the *net* flow that keeps the stock distribution unchanged. In our "duration extension" scenario, we assume that insurers do not terminate any swaps going forward, but all contracts (either long or short) that mature are rolled into new 30-year receive-fixed swaps. Given the initial maturity distribution of swaps, this implies that by the end of the projection period, about 40 percent of the pay-fixed and 60 percent of the receive-fixed portfolio have rolled into new long swaps positions that are subject to Dodd–Frank.

Importantly, both scenarios assume that the overall *size* of insurers' derivatives portfolios stays constant. This assumption is simply for ease of comparison to current balance-sheet values. If, as seems nearly certain, the notional value of swaps positions continues to increase over time, the dollar values of posted margin—and the corresponding costs—will be proportionally higher.

Estimated margin

For each firm, we calculate the margin that would be required in each year under each scenario, given the distribution of interest rate paths shown in figure 2. As shown in figure 3, initial margin is forecast to rise steadily over the projection period in both scenarios. The smoothness and relative precision of the projected paths of initial margin reflect the fact that initial margin is largely driven by portfolio turnover, which is (by assumption) independent of the interest rate environment. However, the size of the increase depends crucially on the extent of the portfolio lengthening. In the constant-maturity case, it climbs about $2 billion between 2013:Q2 and 2014:Q3, reflecting portfolio changes that we have already observed, but then stays approximately constant for the remainder of the projection period. This outcome reflects the strong negative correlation between changes in receive- and pay-fixed values, which insulates the value of the overall

FIGURE 3

Required initial margin under alternative portfolio scenarios

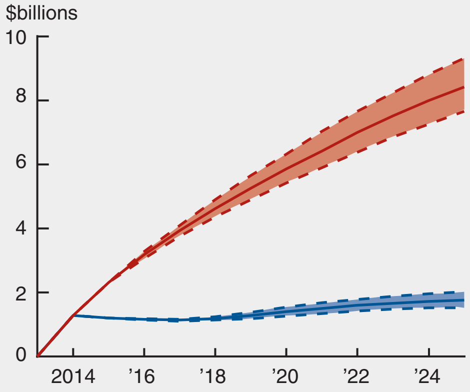

Notes: The figure shows the range of initial margin posted by the top 20 OTC derivatives users under the static-portfolio (blue) and duration-extension (red) scenarios. Solid lines represent means; shaded regions are 10–90% confidence bounds, based on the distribution of forecasted interest rates depicted in figure 2.
Source: Authors' calculations.

FIGURE 4

Required variation margin under alternative portfolio scenarios

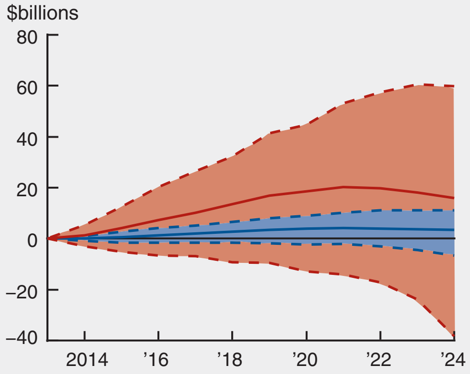

Notes: The figure shows the range of variation margin posted by the top 20 OTC derivatives users under the static-portfolio (blue) and duration-extension (red) scenarios. Solid lines represent means; shaded regions are 10–90% confidence bounds, based on the distribution of forecasted interest rates depicted in figure 2.
Source: Authors' calculations.

portfolio from interest rate shocks. In the duration-extension case, in which this offset gradually disappears, the required amount of initial margin climbs to about $8 billion by 2024.[17]

Under the constant-maturity scenario, the mean level of variation margin peaks at a level of about $4 billion after eight years, as shown in figure 4. Under the duration-extension scenario, this amount is considerably larger, at about $18 billion. Furthermore, the amount is very sensitive to the path of interest rates, with the 90 percent confidence interval in the duration-extension scenarios spanning a range of nearly $100 billion. Thus, the amount of variation margin that will be required from the industry in coming years is quite uncertain.

Potential costs of margin requirements

Initial margin

Table 5 reports sample firms' securities holdings that could, in principle, be used to meet initial margin requirements. In practice, two reasons that these total amounts of securities may not be able to be used for margin are that they are already pledged for some other purpose or that the CCP imposes a limit on how much may be used. As shown in column 2, encumbered assets generally represent a small portion of insurers' overall securities portfolios. To address the question of collateral limits imposed by CCPs, we apply the margining rules for cleared swaps adopted by the CME (Chicago Mercantile Exchange).[18] In particular, we assume that for margining purposes, each type of security is discounted by the amount shown in column 3, reflecting a typical haircut applied to that asset class by the CME. Furthermore, we apply the CME's rule that the sum of agency debt and agency mortgage-backed securities (MBS) used as collateral cannot exceed 40 percent of total collateral for any given customer and that corporate and foreign sovereign bonds cannot exceed the lesser of 40 percent of total collateral or $5 billion. The portfolio limits at the CME apply at the level of the futures commissions merchant (FCM), not the client, so an insurer effectively competes with the other clients of an FCM when trying to post corporate bonds. However, large insurers also have accounts at multiple FCMs; thus, it is not clear whether the effective limits on insurers should be considered to be greater or less than the limits imposed by the CCP. The table therefore considers both a case in which the CME rules are passed through one-for-one to insurers and a more conservative calculation in which insurers are not able to post any securities at all other than cash and Treasury securities. The results of these calculations, reflecting the approximate amount available for initial margin, are shown in columns 5 and 6, with the actual amount of margin (both initial and variation) currently posted shown for comparison in the final column. As the available securities exceed those being used by a factor of 6, even under the conservative assumptions, there is clearly a significant amount of spare capacity at present.

Furthermore, the amount of securities available to pledge is large compared with the amount of collateral

TABLE 5

Estimated collateral available for initial margin at top 20 swaps users

	Fair value of securities	Less: encumbered	= Available collateral	Assumed haircuts (%)	Potential collateralized positions: CME-like limits[a]	Potential collateralized positions: Cash and Treasury securities	Margin currently pledged
Cash and equivalents	16	0	16	2.5	15	15	1
Treasury securities	84	33	50	4.5	48	48	5
Agencies	25	4	22	6.0	20[b]	–	0
Agency MBS	85	22	63	11.0	56[b]	–	2
Foreign government	63	28	35	8.5	32[c]	–	0
Public corporates	700	117	583	20.0	61[c]	–	2
Total	972	204	768		232	63	10

[a] Uses portfolio limits for each insurer on each asset class based on those currently imposed on clearing members by the CME.
[b] Sum of agency debt and agency mortgage-backed securities (MBS) must be less than 40 percent of total portfolio.
[c] Sum of foreign government and corporate bonds must be less than 40 percent of total portfolio and $5 billion per insurer under CME-like limits.
Notes: Amounts in billions of dollars. Data as of 2014:Q3.
Sources: Statutory filings via SNL Financial and authors' calculations.

that was projected to be needed for initial margin in the previous section. Thus, it appears unlikely that collateral availability for initial margin will be a binding constraint for most insurers in the foreseeable future.[19] This is in contrast to the situation for many other types of derivatives market participants, which may have large OTC derivatives positions but do not necessarily hold large volumes of high-quality securities, giving rise to increased demand for collateral-transformation services.[20]

Although insurers have little incentive to engage in collateral transformation (apart, perhaps, from increased repo activity, as discussed later), the requirement to post initial margin will still involve some ongoing costs. CCPs typically charge fees of 10 to 25 basis points to service collateral (in addition to the other fees associated with central clearing). This is on top of any collateral administration fees charged by the insurer's clearing member.

Variation margin

The potentially costly scenario for insurers with respect to variation margin is one in which long-term interest rates rise significantly and spreads between the yields on their assets and overnight rates widen, even if these moves were to occur over a relatively long period. This is because such a scenario could involve having to sell bonds to meet variation margin on long-dated receive-fixed swaps; and the return on that margin would be low relative to that on bonds, representing an opportunity cost for the firm. Accordingly, we assume that variation margin on cleared swaps is posted in cash that is raised by selling corporate bonds and that it pays the effective federal funds rate.[21] Thus, the cost of variation margin is driven by the spread between the corporate bond rate and the fed funds rate in each of our Monte Carlo scenarios.[22] Figure 5 shows the corresponding distribution of losses (more precisely, forgone revenue), relative to what would obtain if there were no margin requirements.

For the constant-maturity scenario, the amount of projected variation margin was relatively small, and consequently the forgone revenue associated with variation margin is also small—$180 million per year by the end of the projection period in the mean case. While this amount is not trivial, it would not represent insurmountable costs for the industry. For example, profits at the firms in our sample were $24 billion in 2013,[23] so that even for extreme interest rate paths, margin-related costs would amount to less than 1 percent of earnings. In large part, this modest outcome has to do with the factors noted in the previous section that keep margin small when the distribution of swaps stays fixed.

Again, the scenario in which insurers extend the duration of their portfolios results in much larger median outcomes and a much wider range of possibilities. The mean cost of posting variation margin rises to about $760 million per year; and, for adverse interest rate outcomes (a steeply rising yield curve and a widening of the spread between corporate yields and the PAI), the cost could be over $2.5 billion per year.

Other costs

In addition to the opportunity cost of variation margin, there are other costs for insurers to consider. In particular, organizational and operational details may introduce

FIGURE 5

Forgone revenue from required margin under alternative scenarios

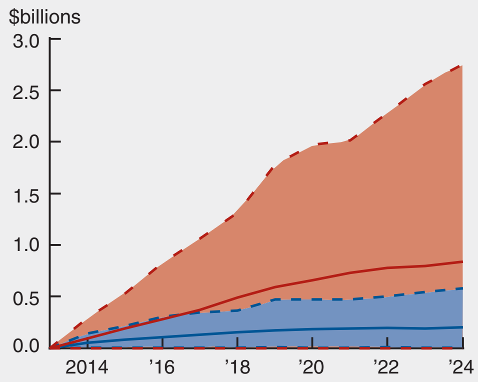

Notes: The figure shows the range of the opportunity cost of holding cash variation margin for the top 20 OTC derivatives users under the static-portfolio (blue) and duration-extension (red) scenarios. Solid lines represent means; shaded regions are 10–90% confidence bounds, based on the distribution of forecasted interest rates depicted in figure 2.

Source: Authors’ calculations.

complications, especially in the short run. For example, it may be that the subsidiaries of an insurance company that currently hold its swaps positions are not the same subsidiaries that hold its high-quality collateral. Insurers could respond by consolidating or rearranging the corporate structure or by transferring exposures and assets across entities.

Though relatively minor, there are also capital issues involved with collateral management. For example, collateral pledged for derivatives positions continues to be counted as an asset of the pledging insurance company, but it receives an additional risk-based capital charge, reflecting the risk that it may not be available to pay policyholder claims in the event of default.

The cost of derivatives trading may also increase. CCPs charge maintenance and transaction fees for swap clearing, although these are on the order of fractions of basis points. Perhaps more significantly, clearing members face significant new costs associated with account administration, default fund contributions to CCPs, and clearing. It is likely that they will pass on most of these costs to clients in the form of increased fees. The costs of trading uncleared derivatives are likely to increase by even more as liquidity deteriorates for such products.

Furthermore, in order to ensure that they can meet variation margin on an ongoing basis, insurers will have to maintain buffers of cash, highly liquid securities, or access to liquidity from other providers beyond the amount of margin that is required of them at each point in time. Using a similar calculation as we did earlier, we note that insurers would require an increment of about $2 billion in our constant-maturity scenario and $8 billion in our duration-extension scenario to keep cash on hand to satisfy, say, 99 percent of five-day movements in swaps positions—assuming that margin could be frictionlessly netted across all contracts and accounts. This compares with their current cash balances of about $16 billion.

Implications for the industry

Although we find it unlikely that the direct costs of posting margin will be unbearable for the life insurance industry, these costs could nonetheless amount to billions of dollars per year, and the bulk of this amount would fall disproportionately on a handful of larger firms. These firms thus have incentives to try to minimize their margin burden.

One obvious way to reduce collateral needs is simply to reduce derivatives positions. Since most insurer derivative use reflects hedging, rather than speculative activity, this could result in greater exposure to risk. However, much of insurers’ derivatives-based hedging activity reflects accounting and regulatory considerations, not necessarily economic ones. For example, some of insurers’ receive-fixed swaps are matched to specific bonds held on their balance sheets or otherwise qualify as highly effective hedges under GAAP. (This explains why they maintain large portfolios of both pay- and receive-fixed swaps.) Reducing hedging of this purely accounting sort would not necessarily increase overall risk. Indeed, the extent to which insurers are able to leave economic risks unhedged will be mitigated by regulatory pressure. It could mean an increase in GAAP earnings volatility or in regulatory capital requirements, but insurers would have to weigh those costs against the costs of holding margin.

On balance, however, insurers may move toward hedging strategies that require less collateral—particularly those that involve only cleared derivatives. As shown in table 3 (p. 24), insurers maintain sizable portfolios of caps, floors, and other derivatives that are not, for the moment, subject to central clearing. As noted earlier, many of these positions are intended to hedge the optionality embedded in various annuity and insurance products. If the cost of trading in these products rises significantly—or if liquidity deteriorates—insurers may find it advantageous to try to hedge some of these risks using swaps or exchange-traded products, which could introduce basis risk.

Another way for insurers to reduce the need for derivatives activity—or to cover the potentially higher costs of that activity—would be to shift some risk that is currently hedged using derivatives to other parties. For example, some companies may find it attractive to offer insurance or annuity products that offload some interest rate risk onto consumers. Indeed, insurance companies report that, as a result of the new rules, they are beginning to shift their mix of products by offering relatively less attractive pricing on products that provide long-term guaranteed payments and more aggressively marketing products with customer participation features, such as certain whole life policies. If insurers find it too expensive to hedge certain types of insurance products and pull back from offering them, significantly raise their prices, or modify them to pass through risks to customers, this could reduce the economic function they serve in providing risk-sharing services to the economy. The rules could also hasten the exit of insurance companies from variable annuities—which have proven expensive in the low-rate environment—as the costs of hedging guarantees on these products will increase. Many companies have already attempted to reduce their exposures to these products either by ceding them to captive reinsurers or by selling them outright.

Insurers will also increasingly need to maintain access to ready sources of liquidity to cover variation margin. Without such a liquidity buffer, insurers might have to make relatively large and rapid adjustments to variation margin during episodes of market volatility, perhaps contributing to fire-sale dynamics. One likely source of this liquidity is advances from Federal Home Loan Banks (FHLBs). Insurers maintain sizable portfolios of mortgage-related assets that qualify them for FHLB membership (Paulson et al., 2014). And, indeed, many insurers have begun to tap FHLBs for funds in recent years. Insurers may also turn to the broker-dealer sector to offer term repos against their securities portfolios or other collateral-transformation services. However, since term repos are not available to match the duration of long-dated swaps contracts, this strategy would be subject to rollover risk. Particularly for riskier collateral, insurers could find their liquidity sources evaporating during a crisis, perhaps at the same time that variation margin is rising due to volatile market conditions. Insurers could also look to sources of cash from elsewhere in their own corporate structures. Securities lending operations, for example, could potentially be scaled up to provide a source of cash for variation margin. Alternatively, firms may look for new ways to hold liquid assets without occupying balance-sheet space.[24]

With respect to the impact on capital, insurers may have an incentive to move derivatives activity outside of the insurance operating company, where they will not be subject to regulatory capital requirements. One way this could be done is through captive reinsurance. However, as noted earlier, captive reinsurers themselves may not maintain reserves of cash or high-quality securities adequate to meet margin requirements. Thus, insurers face conflicting incentives for corporate structure when it comes to swaps margin. On the one hand, they may wish to move derivatives transactions to nonoperating subsidiaries that face less-binding capital constraints. On the other hand, these subsidiaries themselves will be forced to hold high-quality collateral, reducing their profitability.

To the extent that insurers need to shift their assets into cash or liquid securities, they may look to offset the effect on returns by taking additional risks elsewhere. This activity could be similar to behavior that has been observed as insurers have faced weak investment returns in the persistently low-interest rate environment (Becker and Ivashina, 2013). While, in principle, larger cash positions and larger risky-asset positions may leave the aggregate risk of their assets unchanged, such a shift may well result in reduced liquidity for the industry, since the higher-quality assets would now be tied up as collateral.

Conclusion

Like other market participants, insurers that rely on OTC derivatives face challenges from the new Dodd–Frank regulations requiring the central clearing and collateralization of most of those positions. We have used Monte Carlo simulations to study the amount and type of collateral that insurers may have to hold against their interest rate swaps portfolios over the next decade. While we find that the industry-wide costs of collateralizing positions are likely to be modest, there are some low-probability tail scenarios in which they could be substantial for some large insurers, primarily because collateral must be posted in the form of low-yielding cash assets. We have discussed a variety of ways in which the industry might respond to these and the other costs associated with clearing and collateralizing derivatives positions.

NOTES

[1]Other risk-management techniques employed by insurance companies include insuring a large and diversified portfolio of risks (which reduces uncertainty), writing insurance on lines of business that act as natural hedges (for example, the mortality risk insurers face from life insurance contracts can partially offset the longevity risk associated with annuities), and sharing risk with other companies through reinsurance.

[2]The portion of the Dodd–Frank Act applying to most large life insurance companies took effect in June 2013. Title VII of Dodd–Frank mandates central clearing of certain types of swaps contracts, and in May 2013 the Commodity Futures Trading Commission (CFTC) finalized its rule indicating the specific classes of swaps for which central clearing will be required. These include all "plain vanilla" interest rate swaps, basis swaps, forward-rate agreements, and overnight index swaps (OIS) written in major currencies against the standard short-term interest rate benchmarks (the London interbank offered rate or LIBOR, the Euro interbank offered rate or EURIBOR, and, in the case of OIS, the fed funds rate). Credit default swaps are also covered under title VII, and the CFTC rule applies to CDS indexes on corporate debt. The Securities and Exchange Commission (SEC), which has yet to publish final rules, is responsible for single-name CDS contracts. The U.S. Department of the Treasury has determined that physically settled foreign exchange (FX) swaps are not subject to the Dodd–Frank central clearing requirements.

[3]Among insurance companies, the impact of the rules is only likely to be material for life insurers, not for property and casualty insurers, as the latter maintain substantially smaller OTC derivatives positions, both relative to their assets and in absolute terms. Unless otherwise stated, the terms "insurers" and "insurance companies" refer to life insurance companies in this article.

[4]See, for example, Festa (2013). Others have analyzed similar questions for other types of market participants. For example, Heller and Vause (2012) examine the collateral that swaps clearing requires from broker-dealers.

[5]While market commentary suggests that forgone revenue from investments likely represents one of the largest potential costs to the industry associated with Dodd–Frank OTC derivatives rules, our calculations do not include other possible costs associated with uncleared derivatives or operational and organizational costs that may result from the new clearing regime.

[6]These data come from quarterly statutory filings and cover only insurance-operating subsidiaries.

[7]We also note that although notional value is a convenient way of summarizing the size of a derivatives position, it is not a good measure of the potential loss or gain associated with that position, which is typically an order of magnitude smaller. For this reason, the importance of derivatives may be better captured by their "fair value," which reflects their economic worth based on current market conditions—see table 3 (p. 24).

[8]Interest rate swaps are an agreement between two parties in which one stream of future interest payments is exchanged for another, based on specific notional principal amounts. In a pay-fixed (or "receive-float") interest rate swap, a company makes fixed payments and in return receives a floating payment linked to an interest rate. In a pay-float (or "receive-fixed") interest rate swap, a company makes a floating payment linked to an interest rate and in return receives a fixed payment. In both cases, the fixed payment is agreed upon by both parties at the inception of the contract.

[9]Cummins, Phillips, and Smith (2001), Shiu (2007), and González, López, and Cunill (2011) investigate the factors that determine insurance companies' use of derivatives.

[10]See Berends et al. (2013) for a broader discussion of insurers' sensitivity to interest rates.

[11]For example, in the Bank of New York Mellon Corporation and *Insurance Risk's Collateral Management Survey 2013*, 7 percent of respondents (including a global sample of large insurers) indicated that they did not typically post variation margin, while 32 percent indicated that they did not post initial margin (available at https://www.bnymellon.com/_global-assets/pdf/solutions-index/collateral-management-survey-2013.pdf).

[12]Note that these data include collateral for both OTC and listed derivatives, but the amount associated with the latter is very small as insurers generally do not engage in much futures activity or write options.

[13]The move to central clearing could actually reduce netting opportunities in some situations by forcing insurers to clear some trades that could previously have been netted against other trades that will remain uncleared (and thus with non-CCP counterparties).

[14]Experiments using quarterly data did not yield substantially different results.

[15]Most fixed-to-floating swap contracts on insurers' books already satisfy the conditions for central clearing. Those that do not likely differ from clearing-eligible contracts in only relatively minor ways, such as the timing of interest payments or the day-count convention. As noted, we essentially assume away the other types of interest rate derivatives. Evaluating collateral that would have to be held against nonswap contracts would be a more challenging problem because of the diversity of such contracts and the complexity involved in computing their fair values. Most of these positions will not, at least initially, be centrally cleared. Margin requirements for uncleared derivatives have yet to be finalized but are certain to be more punitive than those for cleared positions. Given the harsher rules that will apply to these trades, insurers have an incentive to move away from such nonstandard contracts going forward, so that our assumption may not be much of an exaggeration. Furthermore, the framework developed by the Committee on Payment and Settlement Systems and the Technical Committee of the International Organization of Securities Commissions (CPSS-IOSCO, 2013) proposes exempting uncleared derivatives from initial margin requirements until 2019 for end-users with less than €3 trillion in notional value. Thus, initial margin on uncleared OTC contracts will likely not be collected from insurance companies until at least the middle of the projection period considered here. While most CDS contracts will be centrally cleared sooner and could in principle be incorporated into this analysis, those positions are a fairly small fraction of insurers' overall portfolios and do not seem likely to significantly affect the results.

[16]Dodd–Frank mandates a 99 percent level of confidence, but the CME (Chicago Mercantile Exchange), for example, uses 99.7 percent. As we do here, CCPs typically assess the distributions of derivative gains and losses for the purposes of calculating initial margin using a five- or ten-year look-back period. Indeed, the results are roughly in line with industry estimates, which have suggested that the initial margin requirements will amount to anywhere from 1 percent to 10 percent of the notional value of a single (one-way) swap contract. See, for example, Heller and Vause (2012).

[17]The calculations here assume that potential efficiencies from netting are completely exhausted—that is, that 100 percent of the fair-value gains in contracts is netted against the fair-value losses of contracts when determining potential portfolio losses for the purposes of calculating initial margin. In reality, these efficiencies may be smaller, either because contracts are cleared through multiple, separate accounts or because CCPs do not fully incorporate all netting possibilities into their initial margin calculations. The CME has recently begun offering cross-margining between futures and swaps positions, possibly allowing initial margin requirements to be reduced further for insurers with futures exposure.

[18]See http://www.cmegroup.com/clearing/financial-and-collateral-management/collateral-types-accepted-irs.html.

[19]This conclusion applies only to insurance operating companies. As noted earlier, several large insurance organizations have significant derivatives portfolios elsewhere in their corporate structure, and the legal entities that hold them do not necessarily maintain securities portfolios that would be adequate to cover initial margin under these assumptions.

[20]Indeed, the demand for high-quality collateral due to OTC derivatives requirements, bank regulatory requirements, and other sources could create potential opportunities for insurance companies themselves to expand their collateral-transformation services.

[21]Cash collateral for OTC variation margin receives price alignment interest (PAI) at a short-term interest rate in the corresponding currency. For dollar-denominated contracts, CME and LCH. Clearnet pay PAI at the federal funds rate.

[22]The federal funds rate is not a variable in our VAR, but it is nearly perfectly correlated with the instantaneous Treasury rate. In our simulations, we derive the fed funds rate path from the projections for this rate, as explained in the appendix.

[23]This amount reflects net income for domestic life insurance operating companies only. Earnings at the consolidated parent level are higher. For the ten firms in our sample that are publicly traded in the United States (and thus have easily available consolidated financial statements computed under generally accepted accounting principles [GAAP]), net income was $19.7 billion (relative to $16.46 billion in net income at the operating company level for these same firms).

[24]For example, although not explicitly tied to the Dodd–Frank rules, Prudential created an off-balance-sheet entity (a special-purpose vehicle or SPV) in November 2013 to hold Treasury securities. This structure enables the firm to source Treasury securities as “contingent liquidity” in exchange for notes issued to the SPV. The Treasury securities could be sold quickly to meet variation margin. See *Prudential Financial Inc. Annual Report, 2013* (p. 91).

APPENDIX: TECHNICAL DETAIL ON THE MONTE CARLO EXERCISE

We assume that all fixed-for-floating swaps are plain vanilla and that they are therefore subject to central clearing and initial margins reflecting a 99.7 percent confidence threshold for five-day losses. To calculate the change in the value of insurers' swaps positions under the simulated rate paths, slot each firm's portfolio into 60 buckets, reflecting receive-fixed versus pay-fixed positions and maturities of one through 30 years. We approximate the proportion of swaps in each of these buckets for each insurer by a beta distribution over the range 0–30 years with the parameters chosen to match the mean and standard deviation of each insurer's actual swaps portfolio, based on Schedule DB of their regulatory filings, as of September 2014.

For the "constant maturity distribution" scenario, we assume that the distribution of the stock of swaps held by each firm is fixed over time. This implies that the net flow in each maturity bucket must be nonzero in each quarter in order to keep the portfolio stable as contracts mature. In particular, the net amount of receive-fixed swaps originated by firm i at maturity m in each period must be

$$\Delta x_i\left(m, fixed\right) = \mathrm{B}[m,\alpha_i^{fixed},\beta_i^{fixed}] - \mathrm{B}[m+1,\alpha_i^{fixed},\beta_i^{fixed}],$$

where $B[\ldots]$ is the probability density function of the beta distribution, and α_i^{fixed} and β_i^{fixed} are the shape parameters for the receive-fixed swap distribution at firm i.[1] An analogous equation holds for the pay-fixed portfolio. In principle, this net amount could be obtained in a variety of ways. In particular, if the amount is positive, one could terminate y notional value each quarter and originate $\Delta x+y$ in new contracts, for any arbitrary number y. We assume that, within any type/maturity bin, a firm never terminates and originates contracts at the same time. Thus, if $x_i(m)$ is the desired notional value for the stock of swaps in bucket m and $x_i(m+1)$ swaps are rolling down into that bucket from maturity $m+1$, the firm will *either* (if $x_i(m+1) < x_i(m)$) originate swaps with $x_i(m) - x_i(m+1)$ notional value without terminating any of the existing ones *or* (if $x_i(m+1) > x_i(m)$) terminate swaps with $x_i(m+1) - x_i(m)$ notional value without originating any new ones. In the cases in which firms terminate swaps, we assume that they do so without regard to the contract's age or original maturity. New swaps are assumed to be originated at zero fair value.

For the "duration extension" scenario, we assume that the legacy swaps portfolio gradually matures over time. The amounts that mature are rolled into new 30-year receive-fixed swaps.

For pricing purposes, we assume that all swaps have quarterly payments, are indexed to the instantaneous risk-free rate, and are priced off of the same discount curve as Treasury bonds. The m-maturity swap rate at time t is given approximately by

$$R_t\left(m\right) \approx \frac{\delta_t\left(0\right)-\delta_t\left(m\right)}{\sum_{n=.0}^{m}\delta_t\left(n\right)},$$

where $\delta_t(n)$ is the time-t n-period discount rate. This formula is an approximation because the numerator is only strictly correct in continuous time and the denominator ignores intraquarter discounting. (If swap payments were made continuously, rather than quarterly, the formula would be exact.) The fair value (as a fraction of notional value) of a receive-fixed swap contract with remaining maturity m that was originated s periods ago, is given by the formula

$$FV_t\left(m,s\right) \approx \delta_t\left(m\right)-\delta_t\left(0\right)+R_{t-s}\left(m+s\right)\sum_{n=.0}^{m}\delta_t\left(n\right).$$

Consequently, to value the swaps portfolio, one must know both the distribution of remaining maturities and the distribution of origination dates *conditional on* the current remaining maturity.

To measure $\delta_t(n)$, we use zero-coupon Treasury rates through 2014:Q3 and projections for those rates from a vector autoregression (VAR) for subsequent dates. For the m-maturity yield $y_t(m)$, by definition,

$$\delta_t(m) = \exp\left[-my_t\left(m\right)\right].$$

The data are the zero- (instantaneous), one-, three-, seven-, 15-, and 30-year yields computed by Gürkaynak, Sack, and Wright (2007) over the period 1986:Q1–2014:Q3. The Moody's Baa corporate yield, real gross domestic product growth, and Personal Consumption Expenditures Price Index inflation are also included in the VAR. We begin the sample in 1986 because that is the first date at which 30-year yields become available.

Data are simulated from the VAR by drawing both from the distribution of parameter estimates and the distribution of error terms, assuming normality for both, and simulating forward 20 quarters from 2014:Q3. The zero lower bound is imposed by rejecting any draw for which any interest rate would be below zero at any time; in this case, the whole vector of shocks for that period is resampled. In addition, to reflect current forward guidance about the level of short-term rates (as well as current market expectations), we impose through rejection sampling that the fed funds rate cannot rise above 25 basis points until at least 2015:Q2.[2]

For each simulated value of the six Treasury yields that are included in the VAR, the entire yield curve is interpolated using a quadratic spline. This allows for the calculation of the swap rate associated with each of the 30 possible maturities at each point in time.

To calculate the opportunity cost of holding variation margin, we assume that variation margin must be posted in cash and is remunerated at the fed funds rate, consistent with current practice at the major clearinghouses. We approximate the federal funds rate in each simulation by the equation

$$ffr_t \approx 1.072y_t(0),$$

where the coefficient was estimated from an ordinary least squares regression, with an R^2 of 0.998. We assume that, under normal circumstances, the opportunity cost of holding cash is the Baa corporate bond yield. Since a significant portion of insurers' securities portfolios consist of bonds that are generally safer and thus typically pay lower yields than Baa corporates, this is a conservative assumption. However, occasionally in our simulations some Treasury rates (or the fed funds rate itself) may rise above the corporate bond rate, and in that case we use the higher rate. Specifically, the quarterly opportunity cost is then calculated as

$$VM_t \times (\max[r_t] - ffr_t)/4,$$

where r_t is the time-t vector of yields simulated from the VAR.

To calculate initial margin, we first estimate the covariance matrix of five-day changes in swap fair values between 2004:Q1 and 2014:Q3, across a 10 x 10 grid of maturities spanning zero to 30 years and swap rates spanning 0 percent to 10 percent. In each of the same 10,000 simulations used to calculate variation margin, we calculate the amount of each firm's receive- and pay-fixed portfolio that falls within each of the 100 bins in the grid. Multiplying these weights by the covariance matrix of swap value changes allows us to approximate the five-day variance of each portfolio.[3] The initial margin is assumed to be the 0.3 percent quantile of a normal distribution with this variance and a mean of zero.

NOTES

[1]The amount of 30-year swaps originated each period is simply equal to the stock of swaps maintained in the 30-year bin (that is, the normal PDF evaluated at that point).

[2]Since some parameter draws can imply nonstationary dynamics that lead to explosive behavior, we also impose restrictions to ensure that no projected rate exceeds its historical maximum. In addition, we impose that the spread of the corporate bond to the seven-year Treasury yield cannot be negative.

[3]This calculation assumes that the initial margin that applies to a *given* portfolio remains constant over time. In practice, central counterparty clearinghouses are likely to adjust margin requirements with the level of rates, as the conditional covariance matrix of swap values changes is not constant. Our calculation likely errs on the conservative side—estimating too much initial margin—because we forecast interest rates to rise, and the volatility of a given swap's value is generally decreasing in the level of rates.

Bank for International Settlements, Basel Committee on Banking Supervision, and Board of the International Organization of Securities Commissions, 2013, "Margin requirements for non-centrally cleared derivatives," Basel, Switzerland, September, available at http://www.bis.org/publ/bcbs261.pdf.

Becker, B., and V. Ivashina, 2013, "Reaching for yield in the bond market," working paper, October.

Berends, K., R. McMenamin, T. Plestis, and R. J. Rosen, 2013, "The sensitivity of life insurance firms to interest rate changes," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 37, Second Quarter, pp. 47–78, available at https://www.chicagofed.org/publications/economic-perspectives/2013/2q-berends-mcmenamin-plestis-rosen.

Cummins, J. D., R. D. Phillips, and S. D. Smith, 2001, "Derivatives and corporate risk management: Participation and volume decisions in the insurance industry," *Journal of Risk and Insuran*ce, Vol. 68, No. 1, March, pp. 51–91.

Festa, Elizabeth D., 2013, "New derivatives rules raise life insurers' collateral needs," *LifeHealthPro.com*, June 27, available at http://www.lifehealthpro.com/2013/06/27/new-derivatives-rules-raise-life-insurers-collater?page_all=1.

González, L. O., S. F. López, and O. M. Cunill, 2011, "Hedging with derivatives and value creation: An empirical examination in the insurance industry," *Frontiers in Economics and Finance*, Vol. 8, No. 1, pp. 17–42.

Gürkaynak, R. S., B. Sack, and J. H. Wright, 2007, "The U.S. Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, Vol. 54, No. 8, November, pp. 2291–2304.

Heller, D., and N. Vause, 2012, "Collateral requirements for mandatory central clearing of over-the-counter derivatives," Bank for International Settlements, working paper, No. 373, March.

Paulson, A., R. J. Rosen, K. Berends, and R. McMenamin, 2014, "Understanding the relationship between life insurers and the Federal Home Loan Banks," *Chicago Fed Letter*, Federal Reserve Bank of Chicago, No. 318, January, available at https://www.chicagofed.org/publications/chicago-fed-letter/2014/january-318.

Prudential Financial Inc., 2014, *Prudential Financial, Inc. Annual Report 2013*, Newark, NJ, available at http://www.prudential.com/media/managed/Prudential-AR2013.pdf.

Shiu, Y-M., 2007, "An empirical investigation on derivatives usage: Evidence from the United Kingdom general insurance industry," *Applied Economics Letters*, Vol. 14, No. 5, pp. 353–360.