

Online Appendix
Does Getting Forecasts Earlier Matter? Evidence from
Winter Advisories and Vehicle Crashes

Vaibhav Anand

Online Appendix

Does Getting Forecasts Earlier Matter? Evidence from Winter Advisories and Vehicle Crashes

A Data	3
A.1 Weather Advisory Data	3
A.2 Forecast data from NDFD	4
A.3 Daily weather observations	5
A.4 Vehicle Crash Data	6
A.5 Primary sample construction	7
B Estimating the effect of lead time on crash risks	8
B.1 Effect of lead time: Using linear functional form	8
B.2 Effect of lead time: Using quadratic functional form	9
B.3 Effect of false alarms on crash risk	9
B.4 Economic benefits of longer lead times	10
C Effect of lead time on snowplow activity	12
D Effect of lead time on visits	14
D.1 Mobile phone location data	14
D.2 Relation between visits and crashes	15
D.3 Effect of advisory lead time on visits- Descriptive evidence	17
D.4 Interaction between visits and road treatment	17
E Robustness Checks	20
E.1 Using fatal crash data	20
E.2 Using single day advisories	21
E.3 Using all advisories with heterogenous effects	21
E.4 Assigning incremental lead time to multi-day advisores	22
E.5 Aggregating multi-day advisory days by adding outcomes	23
F Additional Figures and Tables	29

A Data

A.1 Weather Advisory Data

The historical data on weather advisories issued by the NWS are digitized and maintained by the Iowa Environmental Mesonet (IEM). This data set is a collection of geospatial format files that provide information on geographic coverage, timing, and other details for the warning message for all weather advisories issued since 1986. Data on winter weather advisories are fully available only after 2005. The digital archival of the NWS advisories relies on NWS's Valid Time Extent Code (VTEC) system, which allows for systematic parsing of the information in advisories. The VTEC system for winter advisories was operational by November 2005.

The archive database provides the time an advisory is issued by the local weather forecast office, the name of the issuing weather forecast office, the time the advisory goes in effect (i.e., the forecasted time of onset of the hazardous event), the time the advisory expires, the nature of the hazardous event, and the names of affected counties. Figure F.8 shows an example of a winter advisory issued by the Milwaukee WFO. The advisory was issued for 20 counties in the state of Wisconsin at 2:53 PM on February 4, 2019. The adverse weather event related to this advisory was forecast to occur between 6 PM on February 5 to 6 AM February 6. This information allows me to capture for each date and county whether a winter advisory is active in the county on that date, and the lead time, i.e., the hours between advisory issuance and predicted event start.

A.1.1 Timing of advisory issuance and predicted event onset

Figure F.11 Panel B below shows the distribution of advisories (gray bars) and their lead times (solid line) by the predicted hour of the onset of weather event onset. The figure shows that there is indeed a relation between event onset time and advisory issuance decision and its lead time. However, this pattern is less likely to be a result of any attribute of crash risk. The figure shows that most advisories are issued for events that are predicted to start at 0000 (midnight), 0600, 1200, or 1800 hours. These advisories also have the longest lead times on average. The second peak occurs around 0300, 0900, 1500, and 2100 hours. Fewer advisories are issued with event start time outside of these hours, and such advisories also have shorter lead times on average. This is likely due to the fact that when forecasters know in advance of an adverse weather, they “round” the event onset hour to the nearest 6-hour cut-off (or the three hour cut-off), possibly for convenience. Such advisories have longer lead times. Whereas

the advisory for an event that develops at shorter notice, is issued with a more precise prediction of the onset time. However, there seems to be a uniform number of longer lead advisories issued at 0000, 0600, 1200, and 1800. While it is possible that forecasters take into account the event onset time and associated crash risk, we do not see a clear pattern in the data.

Figure F.11 Panel A below shows the distribution of advisories (gray bars) and their lead times (solid line) by the hour of advisory issuance. The figure shows that there is a relation between advisory issuance hour and advisory issuance decision and its lead time. This pattern is a result of the schedule followed by the National Weather Service in issuing the routine weather forecasts and advisories. Routine forecasts are issued twice a day at around 4 am and 4 pm, with ad-hoc forecasts and advisories issued as and when needed. It is expected that ad-hoc forecasts have shorter lead times because they are for events that are likely to develop at shorter notice. This figure suggests that the decision about the hour of advisory issuance is primarily driven by a predetermined schedule rather than any unobserved attribute of crash risk.

A.2 Forecast data from NDFD

I obtain the historical daily forecast data from National Digital Forecast Database (NDFD). The NDFD data provide the gridded forecasts for snow generated by the Weather Forecast offices (WFOs) and the Weather Prediction Center (WPC). The forecasts are available with different lead times. For my purpose, I obtain the snow forecast for the next 0 to 24 hour period. However, The NDFD forecasts for snow are issued every 6 hours and are valid for the 6-hour periods. For example, a 12-hour lead time snow forecast issued at 0600 PM 01-Jan-2010 is the amount of snowfall forecasted during 0600 AM-1200 PM on 02-Jan-2010. I obtain these 6 hourly gridded data of snow forecasts in GRIB format from NDFD. In this section, I explain the process of converting the snow forecast over a 6-hour period to a forecast of snow over a desired 24-hour period.

I construct the 0 to 24-hour lead time forecast for the period 1200 UTC to 1159 UTC (i.e. 0700 AM to 0659 AM EST) by adding the four separate 6-hour forecasts in the following way (say, for the period 1200 UTC 01-Jan-2010 to 1159 UTC 02-Jan-2010):

Thus, the above estimated 0 to 24-hour lead time forecast of snow amount for the period 7am 01-Jan-2010 EST to 7am 02-Jan-2010 EST is what an individual would receive at 7am EST on 01-Jan-2010. (Since 1200 UTC is 7 AM EST same day). Once

Table A.1: Example of NDFD Forecast valid times

Forecast valid period	Forecast issued at
1200 UTC 01-Jan-2010 to 1800 UTC 01-Jan-2010	1200 UTC 01-Jan-2010
1800 UTC 01-Jan-2010 to 0000 UTC 02-Jan-2010	1200 UTC 01-Jan-2010
0000 UTC 02-Jan-2010 to 0600 UTC 02-Jan-2010	1200 UTC 01-Jan-2010
0600 UTC 02-Jan-2010 to 1200 UTC 02-Jan-2010	1200 UTC 01-Jan-2010

The table shows the construction of 0 to 24-hour lead-time forecast for snow using the NDFD forecasts. The forecast is what is available to an individual at 1200 UTC 01-Jan-2010 for the snow during the period 1200 UTC 01-Jan-2010 – 1159 UTC 01-Feb-2010. (1200 UTC is 7 AM EST)

I obtain this 0 to 24-hour snow forecast for each grid point, I aggregate it at each county-date level to estimate the 24-hour ahead snow forecast. To do this, I use the population centroid for each county provided by the U.S. Census Bureau, and calculate the inverse-distance of each grid point in the county from the population center. Then, I aggregate the gridded forecasts by taking the inverse-distance weighted-average from the population centroid of that county.

A.3 Daily weather observations

I obtain daily weather observations from the National Oceanic and Atmospheric Administration’s (NOAA) Global Historical Climate Network (GHCN) database. This database provides daily weather monitor readings for weather stations across the 50 US states and the District of Columbia. It is an aggregation of records from several agencies that in turn collect the monitor readings from their network stations.

The reported daily minimum and maximum temperature readings are the recorded temperatures at a specific time on a day. The reported daily snow and rain readings are the accumulated amounts for the last 24-hour period. Most agencies require their network stations to report monitor readings once a day at a fixed hour, typically around 7 AM local time. While most stations report within a few hours of the suggested reporting time, some stations may report several hours later. This variance in reporting time creates a potential problem for the estimation of snowfall and rainfall amount during a calendar day. To address this, while calculating the daily snow and rain amounts,

I consider monitor readings for only those weather stations that report between 5 am and 9 am.

For each county, I estimate the daily accumulated snowfall, accumulated rainfall, and minimum temperature recorded by aggregating the weather monitor readings. To do this, I use the population centroid for each county provided by the U.S. Census Bureau, and calculate the inverse-distance from county population centroid for each pair of county-station in my sample. Then, I aggregate the weather observations for each county by taking the inverse-distance weighted-average from the population centroid of that county. This approach is consistent with prior studies that aggregate temperature and air pollution monitor readings for a region as the inverse distance-weighted average of all available readings from the monitors located within a radius of the region centroid (e.g., Currie and Neidell (2005) and Heutel et al. (2021)).

A.4 Vehicle Crash Data

I assembled this data set through requests to a state’s DoT, the Highway Safety Information System (HSIS) database, or through internet downloads from the respective DoT’s website. I limit my analysis to those Midwest and Northeast states for which I was able to obtain data for the years between 2008 to 2018.³⁶ These states are Illinois, Iowa, Indiana, Maine, Massachusetts, Michigan, Minnesota, New Jersey, Ohio, South Dakota, and Wisconsin. Crash data for Illinois, Iowa, and Minnesota are available to me only for the years 2010-2018, 2009-2018, and 2010-2015, respectively. My results are robust to limiting my sample to eight states for which I have crash data for all the years during 2008-2018.

One limitation of using police accident reports is that they likely undercount the number of vehicle crashes (Bhargava and Pathania, 2013; Blincoe et al., 2015). States typically require a vehicle crash to be reported to the police if the crash results in injury or death of a person, or property damage in excess of a threshold dollar amount. The threshold for property damage varies by state and is typically between 500-1000 USD in the eight states in my sample. As a result, minor crashes that result in a small property damage or minor injury may go unreported.

Using the police accident reports may also lead to potential biases in my estimate

³⁶The earliest digitized winter weather advisories are available from 2006, and the earliest digitized forecast data is available from 2008.

of the effect of lead time on crash risk. First, there are differences across states in the criteria for reporting a crash to the police. Similarly, the reporting criteria may also vary over time in my study. However, my empirical strategy exploits the variation in lead time and crashes within the same county, year, and month. As long as the reporting criteria do not change in a county within a month, my results may not be affected by the variations in the criteria. Second, if crash reporting patterns are correlated with the lead time on winter advisories, then my estimate of the effect of lead time on crash risk may be biased. When winter weather is worse, people may be less inclined to wait on the road to file a crash report. Police may also find it difficult to respond and reach promptly to crash sites on days with worse weather conditions, which often experience higher crash rates due to weather. However, my empirical strategy controls for the realized weather conditions on a day. So, conditional on the realized weather, crash reporting patterns are less likely to be correlated with advisory lead time. Also, as I show in this study, longer lead times on advisories result in more proactive winter road maintenance activities. So, longer lead times are likely to result in less snow accumulation on roads, enabling police to respond earlier and record a crash report. If this is the case, then my estimates will be biased upward, i.e. toward finding a less negative or more positive effect of lead time on crash risk.

A.5 Primary sample construction

My sample includes data for the 11 states during 2008-2019. I drop county-dates with missing values for weather, forecast, and advisory related variables. The preliminary data includes 2,387,517 county-date observations from 734 counties. 104,431 county-dates receive a winter advisory, of which 70,396 are the first day, 31,846 are the second day, and 2189 days are the third or subsequent day of advisory. The average crashes per 100,000 people are 5.94 and 10.88 on days without and with advisory, respectively. The average advisory lead time is 17.06 hours.

To account for multi-day advisories, I consider the whole advisory window as a single observation unit. I focus on single-day advisories and multi-day advisories that span two days, which accounts for nearly 97.3% of advisories issued. So, for my primary sample, I drop 1893 advisories that span three or more days. For advisories spanning two dates, I aggregate the crashes, weather monitor readings, and snow forecasts over the advisory window. Specifically, for each two-day advisory in a county, I replace the two county-date observations with a single county-date observation that has the initial lead time, and the value of crashes, snow, rain, temperature recorded, and snow

forecast as the average over two days. In robustness checks, I also aggregate by taking the total value of crashes, snow, and rain recorded over two days instead of average. My primary sample includes 2,296,451 county-date observations from 734 counties. 70,041 county-dates receive a winter advisory, which includes single day as well as aggregated multi-day advisory observations. The average crashes per 100,000 people are 5.92 and 11.11 on days without and with advisory, respectively. The average advisory lead time is 16.7 hours.

B Estimating the effect of lead time on crash risks

B.1 Effect of lead time: Using linear functional form

My linear specification is

$$y_{ct} = \sum_{l=-m}^n \left[Adv_{c,t+l} + \beta_l Leadtime_{c,t+l} + \eta_l Err_{c,t+l} \right] + \sum_{l=-m}^n \left[\gamma_l W_{c,t+l} + \lambda_l X_{c,t+l} \right] + \Phi_{cym} + \epsilon_{ct} \quad (B.1)$$

where $Leadtime_{c,t}$, the key variable of interest, is the lead time of the advisory on date t in county c in hours. To account for the cumulative effects of lead time on crash risk, I estimate a distributed lag model following Shrader et al. (2023). I use 2 days of lead and lag around the day of advisory to estimate the cumulative effects of lead time over 5 days around advisory, i.e., $m = n = 2$. The coefficient β_l captures the effect of an additional hour of lead time on crashes that occur l days before the advisory day.³⁷ Table F.3 column 1 presents these estimates. The cumulative effect of an additional hour of lead time, as on q^{th} day since advisory, is measured as $\sum_{l=-2}^q \beta_l$.

Figures F.13 plots the cumulative effect of an additional hour of lead time (y-axis), beginning with two days prior and up to two days after the day of the advisory. The x-axis denotes the days since the first day of the advisory (negative values indicate days before and positive values indicate days after the advisory.) The shaded bands represent the 95% confidence interval based on standard errors clustered two-way at WFO and date level. The figure shows that an additional hour of lead time reduces crashes on the advisory day, but does not affect crashes materially on earlier or later days.

³⁷Alternatively, the coefficient β_l captures the effect of an additional hour of lead time on advisory issued for date $t + l$ on crashes that occur on date t .

B.2 Effect of lead time: Using quadratic functional form

While the linear functional form is easier to interpret and use for performing counterfactual calculations, the marginal effect of lead time may vary with forecast horizon length. To account for this non-linear nature of the marginal effects, I also estimate a quadratic functional-form of my main specification:

$$y_{ct} = \sum_{l=-m}^n \left[\beta_{l1} \text{Leadtime}_{c,t+l} + \beta_{l2} \text{Leadtimesq}_{c,t+l} + \eta_{l1} \text{Err}_{c,t+l} + \eta_{l2} \text{Errsq}_{c,t+l} \right] + \sum_{l=-m}^n \left[\text{Adv}_{c,t+l} + \gamma_l \mathbb{W}_{c,t+l} + \lambda_l \mathbb{X}_{c,t+l} \right] + \Phi_{cym} + \epsilon_{ct} \quad (\text{B.2})$$

where $\text{Leadtimesq}_{c,t}$ and $\text{Errsq}_{c,t}$ are squares of the lead time and snow forecast errors, respectively. The coefficients β_{l1} and β_{l2} capture the linear and quadratic effects of an additional hour of lead time on crashes that occur l days before the advisory day. Table F.3 column 2 presents these estimates. I use the estimates from this quadratic specification in my counterfactual analysis to estimate the economic benefits of longer lead times in section 2.4 and section B.4.

B.3 Effect of false alarms on crash risk

To examine the effect of false alarms on the response to longer lead times in my study, I compare the advisories that are preceded by ‘false alarms’ in the preceding 7 days to those that are not preceded by false alarms. Specifically, I re-estimated my main specification in equation 1 on two subsamples, created from my primary sample—one with county-dates that are preceded by a false alarm, and another with county-dates that are not preceded by a false alarm. I define an advisory as a false alarm if the realized snow on the first two days of the forecast event is less than an inch. For this analysis, I focus on advisories issued for four types of events that experience snow related hazard: blizzards, lake effect snow, snow storms, and severe winter weather. Together, these account for nearly 94% of all advisory days. Figure F.17 plots the effect of lead time on crash risk for advisories that are preceded by a false alarm (Panel A) and that are not preceded by a false alarm (Panel B) in the last 7 days. At short to medium horizons, the effect of lead time doesn’t vary significantly with false alarm experience. However at longer horizons, the cumulative effect of advisories preceded by false alarms is smaller compared to that of advisories not preceded by a false alarm. In fact, for the longest horizon, an experience of false alarm completely removes the

cumulative effect of lead time on crash risk. This suggests that experiencing a false alarm recently may cause people and institutions to pay less attention to longer horizon advisories, and engage in ad-hoc adaptation measures that likely shift crash risk from the current day to another day.

B.4 Economic benefits of longer lead times

In this section, I discuss the steps to estimate the economic savings due to the reduction in crashes. I measure the benefits due to advisories that arrive with some positive lead time relative to a hypothetical scenario when all advisories arrive with zero lead time.

First, I estimate the total number of crashes avoided as a result of advisories with positive lead times relative to the hypothetical scenario of zero lead time on all advisories. To do this, I start with the estimated regression model in Equation B.2. This provides the model predicted number of vehicle crashes per 100,000 people in county c on date t for the observed lead time on advisories. I denote the predicted crashes by $\hat{crash}_{ct}^{obs}$. Specifically, I use the estimates from the quadratic specification in equation B.2 to calculate the predicted crashes $\hat{crash}_{ct}^{obs}$, i.e.

$$\begin{aligned} \hat{crash}_{ct}^{obs} = & \sum_{l=-m}^n \left[\hat{\beta}_{l1} Leadtime_{c,t+l} + \hat{\beta}_{l2} Leadtimesq_{c,t+l} + \hat{\eta}_{l1} Err_{c,t+l} + \hat{\eta}_{l2} Errsq_{c,t+l} \right] + \\ & \sum_{l=-m}^n \left[\hat{Adv}_{c,t+l} + \hat{\gamma}_l W_{c,t+l} + \hat{\lambda}_l X_{c,t+l} \right] + \hat{\Phi}_{cym} \end{aligned} \quad (B.3)$$

Replacing the actual lead times with zero in Equation B.3 provides the estimated crashes in county c on date t , i.e. $\hat{crash}_{ct}^{zero}$, under a scenario where the advisory comes with zero lead time. Thus, the predicted number of avoided crashes per 100,000 people due to the positive lead time on advisory in county c on date t is

$$\begin{aligned} \delta_{ct} = & \hat{crash}_{ct}^{zero} - \hat{crash}_{ct}^{obs} \\ = & - \sum_{l=-m}^n \left[\hat{\beta}_{l1} Leadtime_{c,t+l} + \hat{\beta}_{l2} Leadtimesq_{c,t+l} \right] \end{aligned} \quad (B.4)$$

Next, I estimate the economic savings from the predicted number of avoided crashes using the economic cost estimates of vehicle crashes from Blincoe et al. (2015), a study conducted by the National Highway Traffic Safety Administration of the Department of Transportation. Blincoe et al. (2015) estimate the economic cost of vehicle crashes in the US for the year 2010. They account for both the direct costs (value

of life, medical costs, legal, emergency service, insurance administration, and property damage costs) and the indirect economic costs (lost productivity, workplace losses, and congestion costs) of crashes. They estimate that the average economic cost in 2010 dollars of a property damage only (PDO) crash is 6,000 USD per-damaged-vehicle, of an injury crash is 21,000 USD per person,³⁸ and of a fatal crash is 1.4 million USD. I calculate the average cost of a single crash as the weighted average of the costs of three types of crashes, where the weights are the proportions of PDO, injury, and fatal crashes in the US in 2018.³⁹ These calculations estimate the average cost of a single crash to be around 20,000 USD in 2018 dollars.⁴⁰

Using these estimates, the total economic value of winter advisory lead times in reducing vehicle crashes in county c in calendar year yr is given by

$$\sum_{t \in yr} \delta_{ct} \times population_{ct} \times 20,000 \quad (B.5)$$

where $population_{ct}$ is the population (in 100,000 people) of county c in year yr .

B.4.1 Allowing lead time and forecast errors to move jointly

The above estimates of economic savings assume that forecast errors remain fixed as we change the lead times to zero in the counterfactual. However, these estimates may change when we allow both the advisory lead time and underlying forecast errors to vary together. Figures F.9(c) and F.9(d) show that as lead time decreases, both positive and negative errors decrease in magnitude. In my sample, the average positive and negative errors for zero lead time advisories are 1.4 cm and -2.1 cm, respectively. The corresponding means for advisories with more than 24 hours of lead time are 4 and -2.5, respectively. To account for the co-movement of forecast error and lead time, I re-estimate the number of avoided crashes by replacing the positive and negative forecast errors for all advisories by the corresponding mean error for zero lead time advisories. To do this, I assume that, in the counterfactual scenario, all advisories arrive with zero lead time and all snow forecast errors on those days are equal to the average forecast

³⁸The cost of injury varies from 4,000 USD for minor injury to 1 million USD for serious injuries. I estimate the average cost of an injury crash using the proportions of different injury levels in 2010. My estimations are limited by the assumption that the mix of injuries remains the same over time and geography.

³⁹For the purposes of this study, the proportions of three types of crashes have not changed much over time. The proportion of PDO, injury, and fatal crashes in 2010 and 2018 are (70%, 29.5%, 0.5%) and (71.4%, 28.1%, 0.5%) respectively.

⁴⁰17,300 USD in 2010 dollars.

error on zero lead time days in the sample. However, I allow independent movement of positive and negative forecast errors. To account for the co-movement of forecast error and lead time, I re-estimate the number of avoided crashes by replacing the positive and negative forecast errors for all advisories by the corresponding mean error for zero lead time advisories (in addition to replacing the positive lead times with zero). Specifically, I assume that a forecast error err_{ct} on an advisory day t in county c is replaced by the counterfactual error $e\hat{r}_{ct}$, such that

$$e\hat{r}_{ct} = \begin{cases} E[err|leadtime = 0 \text{ and } err < 0] & \text{if } err_{ct} < 0 \\ E[err|leadtime = 0 \text{ and } err > 0] & \text{if } err_{ct} > 0 \end{cases}$$

where $E[err|leadtime = 0 \text{ and } err < 0]$ is the average of all negative forecast errors on days that receive an advisory with zero lead time.

The black dashed line in figure 5 shows the best-fit trend line for the updated estimations. The figure shows that allowing forecast errors to move with lead times increases the annual economic savings by nearly USD 20 millions. This result is consistent with how forecast errors affect crashes, as shown in Figure F.18(b). The solid line plots the effect on crashes on the advisory day. The dashed line plots the cumulative effect of forecast error on crash risk over five days around the advisory day. Allowing the errors to move with lead time is equivalent to shifting the positive errors from marker D to C and the negative errors from marker A to B. This would result in higher crashes in the hypothetical scenario, resulting in greater economic savings.

C Effect of lead time on snowplow activity

To examine whether road crews perform a greater level of road maintenance activities when advisories arrive with longer lead time, I use the Snow Plow Truck Location data for the state of Iowa maintained by the DoT for the period October 2014 to April 2019.⁴¹ This dataset provides detailed location and operation level data for snowplows collected using the Automated Vehicle Location (AVL) system. The AVL system continuously tracks and records with a time stamp the snow plow ID, its location coordinates, speed, direction of movement, plow position, and the application rate of solid (pounds per second), liquid (gallons per second), and prewet material (pounds per second). These data are collected by the AVL system every few seconds. To estimate

⁴¹The historical data is available as geodatabase at <https://data.iowadot.gov/documents/historic-snow-plow-truck-location-avl/explore>

the amount of material applied, I first calculate the average rate of application between two consecutive time stamps. Then, I multiply the average rate of application with the time difference between those time stamps. To calculate the distance travelled between two time stamps, I add the distance between the coordinates recorded at corresponding time stamps.

For my analysis, I aggregate the high frequency observations at the county-date level to calculate the following measures of road maintenance activity: (1) whether a snowplow is active on a day, i.e., yes or no, (2) total solid material applied, (3) total liquid material applied, and (4) total prewet material (salt mixed with liquid) applied. I also calculate three additional measures: (5) total distance travelled in miles by snow plow, (6) total duration in hours that snow plows are active, (7) total distance in miles plowed (when plow is engaged). I focus on the road treatment activities that involve the application of anti-icing and de-icing materials for two reasons. First, although, I can observe whether a snowplow is dispensing a road treatment material or plowing the road, I cannot differentiate between distance traveled by a snowplow for 'commuting' (i.e., traveling from the shed to the section of road being treated) and for 'road treatment' (actually doing road maintenance). Second, road treatment measures such as the application of anti-icing liquid or prewet material are likely performed before the storm to reduce snow accumulation and crash risk during a snow storm.

Below, I discuss the effect of lead time on distance traveled, duration of operations, and miles plowed. The effect on other activities are discussed in the main text. Figure X Panel A and B show that, relative zero lead time, both distance traveled and duration of operations fall sharply at first, then increase with lead time. One potential reason for the drop in distance travelled (and duration) relative to no lead time may be the pre-positioning of snowplows closer to their operating locations when advisories arrive with enough lead time. When there is no lead time, snowplows may have to travel from their base locations to the operating location on the same day. Figure X panel C shows that the effect of lead time on miles plowed is mixed—relative to zero lead time, the miles plowed increase for moderate lead times, but decrease for longer lead times. Since, plowing usually begins after a substantial snow accumulation, it is likely that longer lead times may not affect snow plowing. In fact, increased anti-icing and de-icing road treatment due to longer lead times may even reduce the need for plowing after a storm.

D Effect of lead time on visits

D.1 Mobile phone location data

To test whether longer lead times result in fewer visits by people outside their homes, I use the mobile phone location data collected by SafeGraph. Using the latitude and longitude location of a smartphone, SafeGraph determines the number of daily visits by unique visitors to various points of interests (POI). For each county-date, I aggregate the number of visits by five categories based on North American Industry Classification System (NAICS) classification: retail, leisure (includes restaurants), commercial, educational, and healthcare.

To estimate the effect of lead time on visits, I use the variation in lead time and visits within the same county-month-year. There are limitations to working with mobile phone location data. SafeGraph calculates the number of visits to each point of interest (PoI) by tracking the GPS location of a sample of smartphone devices. The number of devices tracked (coverage) may vary across places and over time. So, the variation in the number of visits to a place (e.g., a retail store) may come from both the true variation in the number of visits and the variation in the coverage (i.e., the sample of devices tracked).⁴² This may be a potential source of concern for the estimation strategy. For example, counties that receive advisories with shorter lead time, may also have lower coverage (fewer devices tracked), and so fewer visits recorded. Similarly, severe weather may affect the coverage differently at different places (say, fewer devices tracked in rural counties compared to urban counties). However, such variation are likely not a concern for my design since my estimates are identified off the variation within the same county-month-year.

Another potential concern is that coverage may vary with advisory lead time within the same county. For example, large storms may receive longer lead time advisories and reduce the coverage due to severe weather conditions. However, in all my specifications, I control for weather severity and snow forecasts at the county-date level. So, it is less likely that the variation in coverage within the same county-month-year, after controlling for observed weather and snow forecast, is correlated with advisory lead time.

I normalize the number of visits at the county-date level by calculating the visits per 100,000 devices tracked in the county on that day.⁴³ Since, I observe the number

⁴²In the SafeGraph data, I observe the number of visits to a PoI on each date. However, I observe the number of devices tracked daily at the state level only.

⁴³This approach is similar to the one used in Monte et al. (2023) that normalizes monthly visits

of devices tracked at the state-date level only, I first estimate the county coverage by multiplying the state coverage by the ratio of the county population to the state population in that year. So, the normalized visits in county c on date t are

$$visits_{ct} = \frac{rawvisits_{ct}}{devices_{st} \times \frac{population_c}{population_s}} \times 100,000$$

where $devices_{st}$ are the number of devices observed in state s on date t , and $population_c$ and $population_s$ are county and state populations in that year. I also check the robustness of my results by simply normalizing the daily visits in a county using the county population, i.e.

$$visits_{ct} = \frac{rawvisits_{ct}}{population_c} \times 100,000$$

D.2 Relation between visits and crashes

Longer lead time may reduce the number of people visiting several places. This may happen because either people change their travel plans or places such as businesses and schools decide to close or reduce the hours of operation when they receive a weather advisory in advance. In both cases, the resulting effect is to reduce the number of people commuting to visit places outside their homes. This in turn is likely to reduce the number of vehicles on the road and the number of crashes. Figure D.1 provides descriptive evidence for the relationship between visits and crashes for county-dates that receive no winter advisory. The figure shows a binned scatter plot of crashes per 100,000 people (x-axis) by average visits per 100,000 devices (y-axis) within each of the ten decile bins for the visits. The figure shows a positive correlation between crashes and visits.


To further examine the effect of visits on crashes, I estimate the following fixed-effect specification using county-date observations that do not receive any advisory.

$$\begin{aligned} \log(crashes)_{cd} = & \beta \log(visits)_{cd} \\ & + \gamma_{d-1} W_{c,d-1} + \gamma_d W_{cd} + \gamma_{d+1} W_{c,d+1} + \lambda X_{cd} + \Phi_{cym} + \epsilon_{cd} \end{aligned} \quad (D.6)$$

where $\log(crashes)_{cd}$ and $\log(visits)_{cd}$ are the log of crashes and visits in county c on date d . The rest of the specification is the same as discussed in equation 1. $W_{c,d-1}$, W_{cd} , and $W_{c,d+1}$ are the controls for realized weather for the previous, current, and the

using census tract level sampling rate, which are only available monthly.

Figure D.1: Crashes by visits on days with no weather advisory



Notes: The figure shows the binned scatter plot of the average crashes per 100,000 people (x-axis) for the average visits per 100,000 devices (y-axis) within each of the ten decile bins for visits. The sample county-date observations that did not receive an advisory. The markers plot the average crashes per 100,000 people.

next day, respectively. $\mathbb{X}_{cd}$ includes variables *DayofWeek*, *Workday*, and *Weeknum* to control for the effects of day of week, holidays, and seasonality in traffic, respectively. Φ_{cym} are county-year-month fixed effects that allow me to use the within county-year-month variation in visits and crashes.

Table D.2: Linear effect of advisory lead time on visits

Dependent Variable:	log(crashes)
<i>Variables</i>	
$\log(visits)$	0.311 (0.037)
County-Year-Month Fixed Effects	Yes
Observations	319,934
Within R ²	0.14

Notes: The table shows the results from estimating the regression models in Equation D.6. The sample includes all county-date observations during Jan 2018-Dec 2019 that do not receive a winter advisory. The dependent variable is log of crashes per 100,000 people. $\log(visits)$, the key variables of interest, is log of total visits to places away from home per 100,000 people based on the mobile phone location data from SageGraph. The specification includes controls for current, previous, and next day's weather, day of week, workday, and week of year effects. Additionally, it includes county-year-month fixed effects. The regressions are weighted by county population. Standard-errors clustered at wfo-date level are in parentheses. Significance Level: ***, 0.01, **, 0.05, *, 0.1

The coefficient β on the primary variable of interest $\log(visits)_{cd}$ measures the percentage change in crashes when visits increase by 1%. Table D.2 presents the regression estimates for equation D.6. The coefficient on $\log(visits)$ is 0.22 and statistically sig-

nificant at 1% level. This shows that a one percent increase in visits increases crashes by 0.22%. Thus, days with higher visits to places away from home also have higher crash rate. Next, I examine whether longer lead time results in fewer visits to POIs.

D.3 Effect of advisory lead time on visits- Descriptive evidence

Figure D.2 plots the mean total daily visits to all POIs by average snow for different lead-time buckets on days that receive an active advisory. The figure is a binned scatter plot of the average visits per 100,000 people (y-axis) to all POIs by average realized snow in inches (x-axis) within each of the six snow bins as mentioned in section 1.2. Each line in the plot corresponds to county-days that receive advisory with lead time falling in one of the four bins of advisory lead times in hours: $\{0, (0,24], (24-48], >48\}$. The solid black line shows the average visits by realized snow for those county-dates which receive winter advisory with zero lead time. The gray long-dashed line shows the average visits by realized snow for those county-dates which receive winter advisory with lead time of more than zero hours but less than or equal to 24 hours, and so on.

The figure shows that on days when an advisory arrives with zero lead time, visits are high for low amount of snow, but falls sharply as snow increases. However, on days when an advisory arrives with some lead time, the visits fall with lead time for nearly all levels of snow.

D.4 Interaction between visits and road treatment

D.4.1 Effect of lead time on visits: More vs less likelihood of snowplow treatment

In this section, I examine whether the impact of lead time on visits is lower on days that are more likely to receive road treatment compared to days that are less likely to be treated. Specifically, I compare days that receive less than 2 inches of snow with those that receive more than 2 inches of snow. I estimated my main specification on two subsamples separately—county-dates that receive more than 0.01 inches but less than 2 inches snow, and county-dates that receive more than 2 inches of snow. If treated roads result in higher traffic, then we should expect to see a smaller negative impact of longer lead times when snowfall is more than 2 inches, i.e., when snow plows are more likely to treat roads. Figure F.19 shows the effect of lead time on total visits on days with less than 2 inches of snow (panel A) and on days with more 2 inches of snow (panel B). The figures show that cumulative visits fall sharply with lead time on days that are less likely to be treated by snowplows. Figures F.19(c)–(f) plot the cumulative




Figure D.2: Visits per 100,000 devices by snow for different advisory lead times

Notes: The figure shows the binned scatter plot of the average visits per 100,000 people (x-axis) for the realized average snow in inches (y-axis) within each of the six snow bins in inches, i.e., $\{<0.01, 0.01-0.5, 0.5-1, 1-2, 2-3, 3-5, >5\}$. The markers with whiskers plot the average visits per 100,000 people and the associated 95% confidence interval. Each line in the plot corresponds to county-days that receive advisory with lead time falling in one of the four bins of advisory lead times in hours: $\{0, 0-24, 24-48, >48\}$. The solid black line joins the markers that plot the average visits by realized snow for county-dates that receive winter advisory with zero lead time. The dark gray long-dashed line corresponds to the average crashes for county-dates that receive winter advisory with lead time of more than zero hours but less than or equal to 24 hours, and so on.

effect of an additional hour of lead time on total, discretionary, and non-discretionary visits, respectively. The solid (dashed) line plot the estimates on days that receive less than (more than) 2 inches of snow. The figures show that when snowfall is above two inches, visits fall on the day of the advisory, but bounce back by the next or the day after. However, when snowfall is below 2 inches, visits fall more sharply on the advisory day and continue to fall in the subsequent days. These results suggest that the two mechanisms may interact, such that some of the decline in visits due to longer lead time might be offset by increased traffic on roads that are timely treated by snowplows.

D.4.2 Effect of lead time on visits: By distance to the treated roads

In this section, I use detailed granular data from the state of Iowa to examine whether places closer to the treated roads experience higher visits. Fig F.20 illustrates this analysis using the map of the Linn county, the second most populous county in Iowa with a population of nearly quarter million people. Panels A and B show the sections of the road treated by snow plows on Jan 2 and Jan 3 in the year 2018, respectively.

Panels C and D show the location of different POI, based on Safegraph data, and whether they are within 500 meters of a treated road (blue if within 500 meters, else red) on the two dates. First, I compare the visits to the same place between two dates. The effect on visits are estimated off the within place variation in lead time and snow plow activity. Specifically, I estimate the below version of the specification in equation 1 with place-month-year fixed effects using the county-date observations that receive at least an inch of snow and some road maintenance.

$$y_{it} = \sum_{l=-1}^1 \left[\beta_l \text{Leadtime}_{c,t+l} + \text{PlowAct}_{i,t} + \beta_l^* \text{Leadtime}_{c,t+l} \times \text{PlowAct}_{i,t} \right] + \sum_{l=-1}^1 \left[\text{Adv}_{c,t+l} + \eta_l \text{Err}_{c,t+l} + \gamma_l \mathbb{W}_{c,t+l} + \lambda_l \mathbb{X}_{c,t+l} \right] + \Phi_{iym} + \epsilon_{it} \quad (\text{D.7})$$

where y_{it} are total visits to the place i today (t) and tomorrow ($t+1$), and PlowAct_{it} is an indicator variable that is equal to 1 if the place i is within 500 meters to a road that received anti-icing treatment yesterday ($t-1$) or today (t). Φ_{iym} are place-year-month fixed effects. I also interact the *Leadtime* variable with an indicator *PlowAct*. Table F.4 column 1 and 3 present the estimation effect on same day visits for discretionary and non-discretionary places, respectively.

Second, I compare places which are closer to the treated road with places which are farther within the same county-date-industry. The effect on visits are estimated off the variation in road treatment within a county-date. Specifically, I estimated a version of equation D.7 but with county-date-naic fixed effects. The sample includes the county-date observations that receive at least an inch of snow and have at least one snowplow active. The regression estimates:

$$y_{it} = \sum_{l=-1}^1 \left[\text{PlowAct}_{i,t} + \beta_l^* \text{Leadtime}_{c,t+l} \times \text{PlowAct}_{i,t} \right] + \Phi_{ctn} + \epsilon_{it} \quad (\text{D.8})$$

where Φ_{ctn} are county-date-naic (industry) fixed effects that ensure a comparison between the places within the same county, same date, and same industry classification (retail, leisure, etc.). Since, our variation is at the county-date-naic level, all county-date measures (except those interacted with *PlowAct*) drop out. Table F.4 column 2 and 4 present the estimation effect on same day visits for discretionary and non-discretionary places, respectively.

The coefficient on *PlowAct*, though not statistically significant, is always negative.

This may suggest that visits are lower to places closer to the roads that are treated. The coefficient on the interaction term ($PlowAct \times Leadtime$) is small and not statistically significant. The interaction coefficient is negative for discretionary visits in both specifications, but has mixed sign for non-discretionary visits (however, the standard error on the positive coefficient is relatively larger). The results suggest that visits fall with lead time (the coefficient on $Leadtime$ in column 2 and 4 is negative), and they seem to fall more for places that receive road treatment (based on the coefficient on the interaction term). The latter finding, though not unambiguously unidirectional, is contrary to our expectation that early road treatment may result in more visits, conditional on lead time. It is possible places that are most affected by travel decisions are also the places that are prioritized for road treatment. Retail shops or restaurants in remote locations may neither see significant variation in visits due to lead time nor receive snow plow treatment. Shops in busy locations are likely to experience both the variation in visits and road treatment. The coefficient on $leadtime$ alone is close to zero for discretionary visits, which supports this hypothesis. In such a case, my specifications in column 1 and 3 may not be able to tease out the two effects.

E Robustness Checks

E.1 Using fatal crash data

Fatal crash data is available at the national level since 1975 under the Fatality Analysis Reporting System (FARS), provided by the National Highway Traffic Safety Administration (NHTSA). Although, the FARS data cover all US states and are measured robustly, the FARS data may not be the best option to estimate the benefits of longer lead times of winter advisories. First, fatal crashes account for nearly 19% of the total economic costs of crashes in the US (Blincoe et al, 2010), which includes both the direct and indirect costs (including the value of life). So, examining total crashes based on police accident reports is likely to help us better account for the benefits of longer lead times.

Second, fatal crashes are relatively rare—there are roughly 110 fatal crashes per day nationally during 2008-2019. So, there may not be enough fatal crashes at the county-date level to provide variation needed for a robust econometric analysis in my setting. Other studies that use FARS data often aggregate the outcome variable (i.e. crashes) at some level. For example, Anderson and Davis (2023) aggregates fatal crashes at census tract-month level to examine the effect of ridesharing on traffic fatalities; and Smith

(2016) uses daily crashes at the national level to examine the effect of daylight saving on fatal crashes. Similarly, McClelland and Iselin (2019) aggregates crashes at state-year level, and French and Gumus (2014) aggregates at county-week level. During 2008-2019, about 97.3% of county-date observations have zero fatal crashes. This contrasts with 28.4% of county-dates with zero crashes of all types in my sample based on police accident reports. So, using police accident reports can provide useful variation since my estimation strategy depends on using within county variation.

Third, a related point is that using FARS data will likely identify the effect of lead time off the variation in larger counties, which are likely to have enough fatal crash instances. However, I appreciate that these estimates, even if based off the variation in larger counties, could be useful in examining the effect of lead times.

However, using the FARS data can provide a robustness check for two things. First, whether the effect of lead time observed for all crashes is consistent when measured using only fatal crashes. Second, to check whether my results are valid when estimated using a national level (sub)sample of crashes. Using the FARS data for the period 2008-2019, I re-estimate my main specifications. Figure F.16 plots the estimates of lead time effect on fatal crashes using FARS data. Figure F.16(a) shows the results of estimating my main specification on FARS data. Figure F.16(b) replicates one of the robustness checks from the paper that uses only those county-dates that receive a winter advisory. Results are consistent with those obtained using police accident reports.


E.2 Using single day advisories

In this robustness check, I re-estimate the main specification in equation 1 using county-dates with single-day advisories or no advisory at all. Figure E.3 shows an illustrative sample, which would include dates with a green check, i.e., the dates with no advisories as well as those with single-day advisories. I drop dates with an active multi-day advisory, i.e., dates marked with red cross.

E.3 Using all advisories with heterogenous effects

In the second test, I re-estimate the main specification using all observations that include observations with single-day and multi-day advisories, as well as with no advisory. But to account for multiday advisories, I estimate the effect of lead time separately for the initial day and subsequent day of advisories. Specifically, I interact the $leadtime_{c,t}$ variable with an indicator variable $Initialday_{c,t}$ that is 1 if the date t is the initial day

Figure E.3: Sample with single-day advisories or no advisory at all



Notes: Illustration of single-day and multi-day advisories. Advisories A1 and A2 are single-day advisories that remain in effect only for a day. Advisory A2 is a multi-day advisory that remains in effect for two days. 12/13 is the initial day (1st day) and 12/14 is the subsequent day (2nd day) of the same advisory

of advisory, else it is 0. This allows heterogeneous effects of lead time on different days of a multi-day advisory. My specification is

$$\begin{aligned}
 y_{ct} = & \sum_{l=-m}^n \left[Adv_{c,t+l} + \sum_j \beta_{l,j} Leadtime_{c,t+l,j} \times Initialday_{c,t+l} \right] + \\
 & \sum_{l=-m}^n \left[\sum_k \eta_{l,k} Err_{c,t+l,k} \right] + \\
 & \sum_{l=-m}^n \left[\gamma_l \mathbb{W}_{c,t+l} + \lambda_l \mathbb{X}_{c,t+l} \right] + \Phi_{cym} + \epsilon_{ct}
 \end{aligned} \tag{E.9}$$

The coefficient $\beta_{l,j}$ captures the effect of the initial day's j^{th} bin lead time on crashes that occur l days after the initial day of advisory. For example, for advisories shown in Figure E.3, $\beta_{l,j}$ captures the effect of j^{th} bin lead time for the advisory days 12/11, 12/13, and 12/17, which are the initial days of A1, A2, and A3 advisories, respectively.

However, estimating the effect of lead time by pooling all advisories, single-day as well as multi-day, using the specification in equation 1 assumes homogenous effect of leadtime on different days of a multi-day advisory.

E.4 Assigning incremental lead time to multi-day advisores

An alternative measure of lead time is to assign incrementally higher lead time to each subsequent day based on the number of hours between the advisory issuance and the

start of the day. For example, consider the advisory A2 in Figure E.3, which is issued at 4 pm on Jan 12 for a snow storm beginning at 7 am on Jan 13 and valid until 12 pm on Jan 14. The measure used in column 3 and 4 assigns a lead time of 15 hours for both Jan 13 and Jan 14. The alternative measure assigns a lead time of 15 hours for Jan 13 and 32 hours for Jan 14.


Table 1 column 5 and 6 present the estimates using this alternative measure of lead time. Column 5 shows the cumulative effect of an advisory’s initial day lead time (assuming heterogeneous effects). These estimates are consistent with the main results. However, column 6, which plots the cumulative estimates assuming homogenous effects, shows smaller marginal returns of lead time at longer horizons. This is likely due to the way we assign incrementally higher lead times to subsequent days of multi-day advisories. Figures E.7(f)-(j) break down the cumulative effect and plot the effect on five different days around advisory, beginning two days before and up to two days after the advisory day. The figure shows that crashes fall with lead time on the current day, but increases with lead time at longer horizons on the previous day, which results in decreasing returns to lead time. However, this pattern is not present in specifications that assign equal lead time to all the days of an advisory. Figure E.4 breaks down the cumulative effect for the main results shown in Figure 3. Similarly, Figure E.5 breaks down the cumulative effect using single-day advisories presented in Table 1 column 2. Figure E.6 breaks down the cumulative effects in Table 1 column 3 and 4, which use all advisories but assigning the same lead time to initial and subsequent days. When the subsequent days are assigned incrementally higher lead time, the decreasing returns to lead time in cumulative effect is likely driven by ‘mechanically inflated’ lead times on subsequent days of multi-day advisories. These subsequent days, which are preceded by severe weather on the previous day as shown in Figure F.10, may be picking up the mechanical association between today’s longer lead time and previous day’s higher crash risk.

E.5 Aggregating multi-day advisory days by adding outcomes

In my primary sample, I consider the whole advisory window as a single observation unit. Specifically, for a two-day advisory in a county, I replace the two county-date observations with a single county-date observation that has the initial lead time, and the value of crashes, snow, rain, temperature recorded, and snow forecast as the average over two days. In this robustness check, I aggregate the days of multi-day advisories by taking the total value of crashes, snow, and rain recorded over two days instead of


average. Figure F.15 plots the marginal effects of lead time by lead time horizon by estimating the specification in equation 1. The figure shows that results are consistent with my main estimates in figure 3.

Figure E.4: The effect of lead time on crashes- Primary Estimations



Notes: The figure plots the break down of the cumulative effect of advisory lead time on crashes per 100,000 people by estimating the flexible functional form of lead time in equation 1. The cumulative effect corresponds to the estimates in Table 1 Column 1. The estimates are based on the primary sample that considers the whole advisory window of multi-day advisories as a single observation unit. Panel (a) plots the effect of lead time on crashes that occur on date $t - 2$, i.e., 2 days before the advisory day, and so on. Panel (b) plots the effect of lead time on crashes that occur one day before the advisory day, and so on. The dependent variable is crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48], >48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in crashes per 100,000 people. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.


Figure E.5: The effect of lead time on crashes: Using single-day advisories




Notes: The figure plots the break down of the cumulative effect of advisory lead time on crashes per 100,000 people by estimating the flexible functional form of lead time in equation 1. The cumulative effect corresponds to the estimates in Table 1 Column 2. The estimates are based on a sample that includes only single-day or no advisory days as discussed in robustness check section 4.2. Panel (a) plots the effect of lead time on crashes that occur on date $t - 2$, i.e., 2 days before the advisory day, and so on. Panel (b) plots the effect of lead time on crashes that occur one day before the advisory day, and so on. The dependent variable is crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48] > 48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in crashes per 100,000 people. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.

Figure E.6: The effect of lead time on crashes- Heterogenous effects on subsequent days

Effect of initial day's lead time: Assuming heterogenous effects




Effect of lead time: Assuming homogenous effects




Notes: The figure plots the break down of the cumulative effect of lead time on crashes per 100,000 people. Panels (a)–(e) in row 1 estimate equation E.9 assuming heterogenous effects on different days of multi-day advisories. Panels (f)–(j) in row 2 estimate equation 1, assuming homogenous effects. The cumulative effect in row 1 and 2, corresponds to the estimates in Table 1 Column 3 and 4, respectively. Panel (a) and (f) plots the effect of lead time on crashes that occur on date $t - 2$, i.e., 2 days before the advisory day, and so on. The dependent variable is crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48] >48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in crashes per 100,000 people. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.

Figure E.7: The effect of lead time on crashes- Incremental lead time on subsequent days

Effect of initial day's lead time: Assuming heterogenous effects



Effect of lead time: Assuming homogenous effects



Notes: The figure plots the break down of the cumulative effect of lead time on crashes per 100,000 people. Panels (a)–(e) in row 1 estimate equation E.9 assuming heterogenous effects on different days of multi-day advisories. Panels (f)–(j) in row 2 estimate equation 1, assuming homogenous effects. The cumulative effect in row 1 and 2, corresponds to the estimates in Table 1 Column 5 and 6, respectively. In the sample, subsequent days of multi-day advisories are assigned incremental lead time. Panel (a) and (f) plots the effect of lead time on crashes that occur on date $t - 2$, i.e., 2 days before the advisory day, and so on. The dependent variable is crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48], >48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in crashes per 100,000 people. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.

F Additional Figures and Tables

Figure F.8: Sample winter weather advisory text message

```
635
WMUS43 KMKX 042053
WSWMKX

URGENT - WINTER WEATHER MESSAGE
National Weather Service Milwaukee/Sullivan WI
253 PM CST Mon Feb 4 2019

...FREEZING RAIN, SLEET AND SNOW EXPECTED TUESDAY NIGHT...

.A wintry mix will quickly overspread southern Wisconsin Tuesday
evening, tapering off or ending by sunrise on Wednesday.
Precipitation amounts won't be especially high, but it is the
likelihood of freezing rain and sleet that raises concern for a
hazardous impact to motorists Tuesday night, possibly stretching
into the Wednesday morning commute.

Ice accumulations of around 1/10 of an inch are expected, mainly
east and south of Madison. Temperatures will be in the 20s Tuesday
night.

WIZ046-047-051-052-056>060-062>072-050500-
/O.NEW.KMKX.WM.Y.0006.190206T0000Z-190206T1200Z/
Marquette-Green Lake-Fond Du Lac-Sheboygan-Sauk-Columbia-Dodge-
Washington-Ozaukee-Iowa-Dane-Jefferson-Waukesha-Milwaukee-
Lafayette-Green-Rock-Walworth-Racine-Kenosha-
Including the cities of Montello, Westfield, Oxford, Neshkoro,
Endeavor, Berlin, Princeton, Markesan, Fond Du Lac, Plymouth,
Sheboygan Falls, Howards Grove, Oostburg, Baraboo, Reedsburg,
Prairie Du Sac, Sauk City, Portage, Columbus, Lake Wisconsin,
Lodi, Beaver Dam, Waupun, Mayville, West Bend, Germantown,
Hartford, Mequon, Cedarburg, Grafton, Dodgeville, Mineral Point,
Barneveld, Madison, Watertown, Fort Atkinson, Jefferson,
Waukesha, Brookfield, New Berlin, Menomonee Falls, Muskego,
West Allis, Wauwatosa, Greenfield, Franklin, Oak Creek,
South Milwaukee, Cudahy, Darlington, Shullsburg, Benton, Belmont,
Argyle, Blanchardville, Monroe, Brodhead, Janesville, Beloit,
Whitewater, Delavan, Elkhorn, Lake Geneva, East Troy, Racine,
and Kenosha
253 PM CST Mon Feb 4 2019

...WINTER WEATHER ADVISORY IN EFFECT FROM 6 PM TUESDAY TO 6 AM
CST WEDNESDAY...

* WHAT...A mixture of freezing rain, sleet and snow is expected.
Sleet and snow is the main concern west of Madison, while
freezing rain and sleet expected to the east and south of
Madison. The freezing rain could result in ice accumulations of
around 1/10 of an inch, especially across far southeast
Wisconsin.

* WHERE...Portions of east central, south central and southeast
Wisconsin.

* WHEN...From 6 PM Tuesday to 6 AM CST Wednesday.

* ADDITIONAL DETAILS...Plan on slippery road conditions. The
hazardous conditions could impact the Tuesday evening and
Wednesday morning commute.


PRECAUTIONARY/PREPAREDNESS ACTIONS...

A Winter Weather Advisory means that periods of snow, sleet or
freezing rain will cause travel difficulties. Expect slippery
roads and limited visibilities, and use caution while driving.

The latest road conditions for the state you are calling from can
be obtained by calling 5 1 1.
```


Notes: The figure shows a winter weather advisory issued by the Milwaukee weather forecast office on February 4, 2019. The advisory shows the time of issuance, the nature and severity of the forecasted winter event, a description of the potential risk to people, the timing of the winter event, and a list of counties and cities the advisory is issued for.

Figure F.9: Observed rain, temperature, and snow forecast error by advisory lead time




Notes: The figures show the distribution of observed rail (Panel A), minimum temperature recorded (Panel B), and the distribution of forecast error when forecast overestimates (Panel C) and underestimates (Panel D) the snowfall at the county-date level by advisory lead time for the main sample. Box and whiskers plot the measured variable (y-axis) for county-dates that receive the winter advisory with lead time in one of the six 12-hour bins (x-axis). The lower hinge, mid-line, and upper hinge of boxes show the 25th, the 50th, and the 75th percentile. Whiskers stretch from the 5th percentile to the 95th percentile. The dotted red line plots the mean.

Figure F.10: Level of snow on initial and subsequent days of advisory




Notes: The figure plots the average snow recorded (y-axis) on the initial days (solid line) and subsequent days (dashed line) of multi-day advisories, by advisory lead time as of the initial day.

Figure F.11: Timing of advisory issuance and predicted event onset




Notes: The figures plot the distribution of the number of advisories (gray bars) and the lead time in hours (solid black line) by the hour of advisory issuance (Panel A) and the predicted hour of event onset (Panel B).

Figure F.12: Advisory lead time, performance, and snow forecasts errors over time




Notes: Panel (a) plots the average lead time of winter advisories in my sample by year (x-axis). Panel (b) plots the root mean square error of snow forecasts in my sample by year (x-axis). Panel (c) plots false alarm ratio in red (number of false alarm divided by total advisories) and hit ratio in black (number of days with more than 1 inch of snow that also receive advisory divided by the number of days that receive more than 1 inch of snow). False alarm is measured as a day that receives an advisory but receives less than one inch of total snow on the same day and following day. The solid and dashed lines are the best fit trend lines for the corresponding measures.

Figure F.13: The effect of advisory lead time on crashes per 100,000 people




Notes: The figure plots the estimated effect of advisory lead time on crashes per 100,000 people by estimating the linear functional form of lead time in equation B.1, beginning with three days prior and up to three days after the day of the advisory. The dependent variable is crashes per 100,000 people. The key variable of interest is lead time in hours. The sample includes all county-date observations that receive as well as that do not receive a winter advisory. The x-axis denotes the days since the first day of the advisory (negative values indicate days before and positive values indicate days after the advisory.) The y-axis shows the change in crashes per 100,000 people. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.

Figure F.14: The effect of advisory lead time on crashes per 100,000 people




Notes: The figure plots the estimated effect of advisory lead time on crashes per 100,000 people by estimating the flexible functional form of lead time in equation 1. The sample includes all county-date observations that receive as well as that do not receive a winter advisory. The dependent variable is crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48], >48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in crashes per 100,000 people. The solid black line plots the same day effect of the lead time at different horizons. The dashed, long-dash, dot-dashed, and dotted lines plot the cumulative effects over 3, 5, 7, and 9 days around the day of the advisory, respectively. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.

Figure F.15: The effect of advisory lead time on crashes per 100,000 people: aggregating multi-day advisory days by adding the outcomes




Notes: The figure plots the estimated effect of advisory lead time on crashes per 100,000 people by estimating the flexible functional form of lead time in equation 1. The sample aggregates multi-day advisory days by adding the total crashes as discussed in section E.5, instead of averaging them. The dependent variable is crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48], >48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in crashes per 100,000 people. The solid black line plots the same day effect of the lead time at different horizons. The dashed, long-dash, dot-dash, and dotted lines plot the cumulative effects over 3, 5, 7, and 9 days around the day of the advisory, respectively. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.

Figure F.16: The effect of lead time on fatal crashes




Notes: The figure plots the estimated effect of advisory lead time on fatal crashes per 100,000 people by estimating the flexible functional form of lead time in equation 1. Panel (a) shows the estimates using observations that receive or do not receive a winter advisory. Panel (b) shows the estimates using only those days that receive an advisory. The dependent variable is fatal crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48], >48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in fatal crashes per 100,000 people. The solid black line plots the same day effect of the lead time at different horizons. The dashed and dot-dashed lines plot the cumulative effects over 3 and 5 days around the day of the advisory, respectively. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.

Figure F.17: The impact of experiencing false alarms on the lead time effect on crash risk



Notes: The figure plots the estimated effect of advisory lead time on crashes per 100,000 people by estimating the flexible functional form of lead time in equation 1. Panel (a) shows the estimates using only those observations from the primary sample that experience a false alarm in the last 7 days. Panel (b) shows the estimates using only those observations from the primary sample that do not experience a false alarm in the last 7 days. The dependent variable is crashes per 100,000 people. The key variables of interest are the five indicator variables that capture which lead time bin, $b \in \{(0,12], (12-24], (24-36], (36-48], >48\}$, the advisory lead time falls in. The x-axis shows the average lead time in the corresponding bin. The y-axis shows the change in crashes per 100,000 people. The solid black line plots the same day effect of the lead time at different horizons. The dashed and dot-dashed lines plot the cumulative effects over 3 and 5 days around the day of the advisory, respectively. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level.


Figure F.18: The effect of snow forecast error on crashes per 100,000 people



Notes: The figure plots the estimated effect of snow forecast error on crashes per 100,000 people using the primary sample. Panel (a) plots estimates based on the flexible functional form specification in the Equation 1. The key variables of interest are the six indicator variables that capture which forecast error bin, $\{<-2.0, (-2.0,-1.0), (-1.0,-0.5), [-0.5,0.5], (0.5,1.0], (1.0,2.0], >2.0\}$, the snow forecast error falls in. The x-axis shows the average error in the corresponding bin. The markers plot the estimated effect of forecast error (in cm on x-axis) on vehicle crashes per 100,000 people (y-axis). The solid black line plots the same day effect. The dashed and dot-dashed lines plot the cumulative effects over 3 and 5 days around the day of the advisory, respectively. The shaded gray ribbon represents 95% confidence interval around the same day estimates. Standard Errors are clustered two-way at WFO and date level. Panel (b) plots estimates based on the quadratic specification in the Equation B.2 for the same day effect (solid line) and 5 day cumulative effect (dashed line). The black markers (A and D) are average negative and positive errors on advisories with greater than zero lead time. The red markers (B and C) are average negative and positive errors on advisories with zero lead time.

Figure F.19: The effect of lead time on visits by snow level


Marginal effect of lead time on all visits by horizon length



(a) All visits- snow below 2 inches

(b) All visits- snow above 2 inches

Cumulative effect of an additional hour of lead time by type of visit



(c) All visits


(d) Retail and Leisure

(e) Office, Health, Education


Notes: The figures plot the estimated effect of advisory lead time on visits per 100,000 devices. Panels (a) and (b) plot the marginal effect of lead time using a specification similar to Equation 1 for the same day (solid line), over 3-days (dashed) and over 5-days (dot-dash) around advisory. Panel (a) plots the effect for days that receive below 2 inches of snow. Panel (b) plots the effect for days that receive more than 2 inches of snow. Gray ribbon is the 95% confidence interval for the same day effect. Standard Errors are clustered two-way at WFO and Date level. Panels (c), (d), and (e) plot the cumulative effect of an additional hour of lead time using a specification similar to Equation B.1 for all visits, discretionary visits (retail and leisure), and non-discretionary visits (commercial, health, and education). The solid lines (dashed lines) in Panels (c)–(e) plot the effect on days that receive below (above) 2 inches of snow. The confidence intervals are not shown for clarity. Gray ribbon is the 95% confidence interval. The regressions include county-year-month fixed effects. All regressions are weighted by county population. Standard Errors are clustered two-way at WFO and Date level.

Figure F.20: Plow activity and Point of Interests Locations

Road treatment by snow plow




(a) Road treatment on Jan, 02, 2018




(b) Road treatment on Jan, 03, 2018

Road treatment and point of interest locations



(c) PoIs on Jan, 02, 2018



(d) PoIs on Jan, 03, 2018

Notes: The figure shows the map of Linn county in the state of Iowa. Panels A and B show the sections of the road treated by snowplows on Jan 2 and Jan 3 in the year 2018, respectively. Panels C and D show the location of different POI, based on Safegraph data, and whether they are within 500 meters of a treated road (blue if within 500 meters, else red) on the two dates.

Table F.3: The effect of advisory lead time on crash risk

Dependent Variable:	crashes per 100,000 people	
	(1)	(2)
$Leadtime_{c,t-2}$	0.005 (0.006)	0.016 (0.011)
$Leadtime_{c,t-1}$	-0.023** (0.009)	-0.014 (0.015)
$Leadtime_{c,t}$	-0.059*** (0.012)	-0.045** (0.020)
$Leadtime_{c,t+1}$	0.0007 (0.008)	0.005 (0.018)
$Leadtime_{c,t+2}$	-0.001 (0.007)	0.002 (0.016)
$Leadtimesq_{c,t-2}$		-0.0002 (0.0002)
$Leadtimesq_{c,t-1}$		-0.0002 (0.0003)
$Leadtimesq_{c,t}$		-0.0003 (0.0004)
$Leadtimesq_{c,t+1}$		-0.0001 (0.0004)
$Leadtimesq_{c,t+2}$		-0.0001 (0.0004)
County-Year-Month FE	Yes	Yes
Mean Crashes (advisory days)		11.11
Observations		2,296,451

Notes: The table shows the results from estimating the regression models in Equation B.1 (Columns 1) and Equation B.2 (Columns 2) using the primary sample described in section 1.1.5. The dependent variable is crashes per 100,000 people. $Leadtime_{c,t}$, $Leadtime_{c,t-l}$, and $Leadtime_{c,t+l}$, the key variables of interest, are lead times (in hours) of winter advisories active on date t , $t-l$, and $t+l$, respectively. All specifications include controls for current, previous, and next day's weather and snow forecasts, advisory issuance type, advisory issuance time, day of week, workday, and week of year effects. All specifications include county-year-month fixed effects. All regressions are weighted by county population. Standard-errors clustered two-way at wfo and date level are in parentheses.

Table F.4: Effect of lead time on visits by snow plow activity

Visit Type:	Retail and Leisure		Commercial, Health, and Education	
Dependent Variable:	visits	demeaned visits (by place-mon-yr)	visits	demeaned visits (by place-mon-yr)
	(1)	(2)	(3)	(4)
$PlowAct_{c,t}$	-0.093 (0.287)	-0.054 (0.150)	-0.806 (0.479)	-0.142 (0.177)
$Leadtime_{c,t} \times PlowAct_{c,t}$	-0.007 (0.013)	-0.009 (0.007)	0.011 (0.013)	-0.005 (0.005)
$Leadtime_{c,t}$	-0.016 (0.010)		-0.053 (0.016)	
place-month-year FE	Yes	No	Yes	No
county-date-naic FE	No	Yes	No	Yes
Mean visits	10.26		11.76	
Observations	976,077	832,685	464,643	396,280

Notes: The table shows the estimated effect of snow plow treatment, advisory lead time, and their interaction on daily visits to places based on Safegraph granular data. Columns 1 and 3 show the results from estimating the regression models in Equation D.7 which controls for place-year-month fixed effects. Columns 2 and 4 show the results from estimating the regression models in Equation D.8 which controls for county-date-naic fixed effects. The sample includes all place-date observations from counties that receive at least an inch of snow and have at least one snow plow active on that day. The sample in Columns 1 and 2 includes ‘discretionary’ places (leisure and retail). The sample in columns 3 and 4 includes ‘non-discretionary’ place (commercial, health, and education). All regressions are weighted by county population. Standard-errors clustered twoway at wfo and date level are in parentheses.