PROJECTION BIAS IN THE CAR AND HOUSING MARKETS

Meghan R. Busse
Devin G. Pope
Jaren C. Pope
Jorge Silva-Risso*

February 2012

Preliminary: Please do not circulate

Abstract

Projection bias is the tendency to overpredict the degree to which our future tastes will resemble our current tastes. We test for evidence of projection bias in two of the largest and most important consumer markets – the car and housing markets. Using data for more than forty million car transactions and four million housing purchases, we explore the impact of the weather on purchasing decisions. We find that the choice to purchase a convertible or a 4-wheel-drive vehicle is highly dependent on the weather at the time of purchase in a way that is inconsistent with classical utility theory. Similarly, we find that the hedonic value that a swimming pool adds to a house is, on average, $1400 more when the house goes under contract in the summertime compared to the wintertime.

*Busse: Kellogg School of Management, Northwestern University and NBER; D. Pope: Booth School of Business, University of Chicago and NBER; J. Pope: Department of Economics, Brigham Young University; Silva-Risso: School of Business Administration, UC Riverside. We are grateful to Chris Bruegge and Ezra Karger for valuable research assistance. We also thank Stefano DellaVigna, Emir Kamenica, Ulrike Malmendier, Ted O'Donoghue, Mathew Rabin, Dick Thaler, and seminar participants at the University of Chicago, Cornell University, UC Berkeley, UC Los Angeles, and Yale University for helpful suggestions.

Many decisions that people make require predicting their future preferences. For example, choosing a job, deciding where to live, planning a vacation, deciding whether to have a baby, and purchasing a home are all important life decisions that require predicting future utility across a variety of choice dimensions.

The standard economic model assumes that the expected value of an individuals' predicted future utility will, on average, match realized utility. Evidence from psychology, however, suggests that individuals may be systematically biased when predicting future utility. A specific bias that has received considerable attention is projection bias (Loewenstein, O'Donoghue, and Rabin, 2003). Projection bias refers to the tendency of an individual to overpredict the degree to which their future tastes will resemble their current tastes. For example, the popular adage "never shop on an empty stomach" is explained by projection bias because consumers are likely to overpredict the degree to which their future selves will appreciate the purchases that their current selves crave.

While projection bias has intuitive appeal for situations like shopping while hungry, an open question is whether this bias can impact the type of important life decisions that were used in the motivation above, like deciding where to live and what job to choose. In this paper, we test for projection bias in two high-stakes environments: the purchases of cars and houses, the two biggest purchases that most consumers make.

Cars and houses are durable goods; therefore consumers must predict at the time of purchase the preferences that they will have in the future when consuming these goods. Projection bias suggests that consumers may mistakenly purchase a car or a house that has a high utility at the time of purchase, but whose utility will not always be so high given the other states of the world that the consumer will experience while owning the car or house. We test the extent to which seasonal and idiosyncratic weather variation at the time of purchase can cause consumers to overweigh the value that they place on certain car and housing characteristics. Projection bias predicts that consumers will overvalue warm-weather car types and housing characteristics (e.g. convertibles and swimming pools) when the weather is warm at the time of purchase and overvalue cold-weather car types (e.g. 4-wheel-drive vehicles) when the weather is cold and snowy at the time of purchase.

We begin by exploring these hypotheses in the car market using administrative data for more than forty million transactions of new and used cars from dealerships around the U.S. We find that the volume of convertible and 4-wheel-drive sales is highly influenced by idiosyncratic variation in temperature, cloud cover, and snow fall. We show that for convertibles, warmer weather and clearer skies in general leads to a higher number of sales. Controlling for seasonal sales patterns, our estimates suggest that an MSA that experiences a mean temperature that is 20 degrees higher than normal will experience a 0.22 percentage point increase in the percentage of total cars sold that are convertibles. Given a base rate of 2.6% of cars sold that are convertibles, this represents a 8.5% increase in the fraction of convertible cars sold. We find large and significant effects both in the spring and in the fall (e.g. an abnormally warm week in November increases the fraction of cars sold that are convertibles). Importantly, we also show that abnormally warm weather does not impact convertible sales when the temperature is already high (when average daily high temperature is already more than ~80 degrees Fahrenheit). Purchases of 4-wheel-drive vehicles are also very responsive to abnormal weather variation - particularly snow fall. Our results suggest that a snow storm of ~10 inches will increase the fraction of cars sold that have 4-wheel drive by ~2 percentage points over the next 2-3 weeks (a ~6% increase over the base rate of 33.5%). This effect is robust to using an event study design that uses large storms as events.

The data allow us to rule out several alternative explanations for these findings. For example, a dynamic analysis indicates that the increase in convertible sales and most of the increase in 4-wheel-drive sales due to abnormal weather cannot be explained by short-run substitutions in car purchases from one week to the next (a "harvesting effect"). We also present evidence that learning about a car during a test drive (which for a convertible may be easier to do on a warm day) is unlikely to explain the results we find. For example, cloud cover (which does not limit the ability to test drive a vehicle as temperature might) has a large impact on sales. Furthermore, individuals who previously owned a convertible and thus may be less likely to need a test drive are also affected by idiosyncratic weather conditions. Finally, we look at the impact of the weather at the time of vehicle purchase on the probability that a car is traded in quickly for a different car. This analysis, which uses unique car identifiers to follow cars over time in our data, suggests that a car is more likely to be returned quickly when purchased on a day with abnormally good weather variation - evidence in favor of projection bias.

The second part of the paper turns to identifying projection bias in the housing market using a repeat-sales methodology for over four million housing transactions. This methodology allows us to see the value that certain house characteristics (e.g. swimming pool) have at different times of year by looking at houses that sold once in the summertime and again in the wintertime while also controlling for variation in overall housing trends across time and space. We find evidence that a swimming pool adds more value to a house that goes under contract in the summertime than it adds to the same house that goes under contract in the wintertime. Specifically, a house with a swimming pool that goes under contract in the summertime sells for an average of 0.3 percentage points more than the same house when it goes under contract in the wintertime. Given the average value of homes with swimming pools in our dataset, this effect suggests a swing in value of approximately $1400 between summer and winter contract dates.

This result is robust to a variety of different specifications and subsamples of the data. Our within-house identification strategy helps us to rule out concerns about unobserved housing characteristics that are correlated with houses that have swimming pools or with the type of people who buy and sell houses with swimming pools. Our fixed-effects framework also allows us to control for seasonal patterns in houses overall in order to identify the interaction between seasonal weather and houses with swimming pools. We also discuss and rule out the possibility that a home with a swimming pool may be worth more due to immediate utility gains (during the season of purchase). Finally, we provide the results for three other housing characteristics whose value may fluctuate across seasons - central air, lot size, and fireplaces. We find no evidence that the hedonic value of these characteristics vary with seasonal temperature and discuss likely explanations for this finding.

Our findings are significant for several reasons. Firstly, the car and housing markets in and of themselves are large and important. Identifying, and potentially correcting, systematic errors in these markets can have valuable welfare implications. Perhaps more importantly, our results suggest that projection bias may be prevalent in other important decisions (getting married, choosing a job, etc.) that are similarly distinguished by having large stakes and inexperienced decision makers.

Our paper is also related to a growing literature that uses field data to test models from behavioral economics (see DellaVigna (2009) for a review). More specifically, our paper relates to a small literature that empirically explores projection bias in field settings (Read & van Leeuven, 1998; Conlin, O'Donoghue, & Vogelsang, 2007; Simonsohn, 2010).[1] Our paper is most similar to the work by Conlin, O'Donoghue, and Vogelsang (2007) who test for projection bias in catalog orders. They convincingly show that decisions to purchase cold-weather items are overinfluenced by the weather at the time of purchase. Specifically, they find that if the temperature at the time of a purchase is 30 degrees lower, consumers are 0.57 percentage points more likely to return the item (3.95%). Our paper complements this earlier work. We extend the existing research by providing evidence of

---

[1] In the psychology literature, the type of projection bias that we explore in this paper is most closely related to the work on hot/cold empathy gaps and visceral states (see for example, Nisbett and Kanouse (1968), Loewenstein (1996), Loewenstein, Nagin, & Paternoster (1997), Van Boven & Loewenstein (2003), Nordgren, van der Pligt, and van Harreveld (2006, 2007). Loewenstein and Schkade (1999) provide a useful review of the psychological evidence for projection bias.

projection bias in two markets of even greater economic importance with much higher stakes. The richness of our data also allow us to explore not only how projection bias impacts sales volume, but also how it can have an impact on prices (housing market).

The paper proceeds as follows: Section I provides a simple, conceptual framework for projection bias following Loewenstein, O'Donoghue, & Rabin (2003), Section II explores the data, empirical strategy, and results for the car market, Section III describes the data for our housing analysis, our empirical strategy, and the results we find in the housing market, and Section IV provides a conclusion along with a brief discussion of the broader implication of our findings.

I. Conceptual Framework

We are interested in understanding how projection bias may influence durable goods purchases following the framework of Loewenstein, O'Donoghue, & Rabin (2003). To begin, suppose that a person has state-dependent preferences such that her instantaneous utility of consumption, c, in state, s, can be represented as $u(c, s)$. Furthermore, consider an individual who is currently in state s' who is attempting to predict her future instantaneous utility of consumption, c, in state s: ($\tilde{u}(c, s|s')$). An accurate prediction would be represented by $\tilde{u}(c, s|s') = u(c, s)$.

Loewenstein, O'Donoghue, & Rabin (2003) argue that projection bias causes agents' predictions about future utility to be unduly influenced by the state they are in at the time of the prediction. Specifically, an individual exhibits projection bias if

(1) $$\tilde{u}(c, s|s') = (1-\alpha)u(c, s) + \alpha\left(u(c, s')\right),$$

where α is a number between 0 and 1. If $\alpha = 0$, then the individual accurately predicts her future preferences, whereas if $\alpha > 0$, an individual perceives her future utility to reflect a combination of her true future utility along with the utility that consumption c would provide in her current state s'.

This simple model of projection bias can be extended to an intertemporal-choice framework. Consider, for example, the instantaneous utility that a person receives in time t from purchasing a convertible in time t $(conv_t)$. Her true utility can be represented by

(2) $$U^t(conv_t,\ldots,conv_T) = \sum_{\tau=t}^{T} \delta^\tau u(conv_\tau, s_\tau),$$

where $\delta \leq 1$ is her standard discount factor. Once again, following Loewenstein, O'Donoghue, & Rabin (2003), a person with projection bias perceives her intertemporal preferences to be

(3) $$\widetilde{U}^t(conv_t,\ldots,conv_T|s_t) = \sum_{\tau=t}^{T} \delta^\tau \tilde{u}(conv_\tau, s_\tau|s_t),$$

where $\tilde{u}$ represents the perceived instantaneous utility described by Equation (1).

This framework illustrates that an individual's perceived intertemporal utility of purchasing a convertible at time t, $\widetilde{U}^t$, is overly influenced by s_t. Specifically, we would predict that when s_t is a very good state of the world for consuming a convertible (warm, sunny weather), an individual has a higher perceived utility of purchasing the convertible than when s_t is a very bad state of the world (cold, cloudy weather).

A challenge involved with empirically testing for projection bias is that the state at the time of purchase s_t, while unduly influential for agents with projection bias, also matters for agents that do not have projection bias (see Equation (2)). While this is less of a problem for durable goods like cars and houses, it may be troublesome if the number of periods is small.[2] For example, if an individual only plans to own a convertible for a few days, then the state of the weather on the day of purchase is quite important since a significant fraction of the consumption will occur in that state of the world. In fact, this logic suggests that we would expect a rational agent to be more willing to purchase a convertible in the Spring than in the Fall since it is likely that they will get to consume the convertible during a greater number of "good" states of the world if they purchase it in the Spring. Similarly, one might imagine that home buyers would be willing to pay slightly more for a home with a swimming pool when they are moving in at the beginning of the summer relative to the amount they would be willing to pay if they moved in at the end of the summer. Thus, simply finding that people are willing to pay more for a home with a swimming pool and a convertible when the weather is nice outside could be a response by agents who are accurately predicting their future utility and does not necessarily provide evidence of projection bias.

[2] As the number of periods gets infinitely large, the state in the current period becomes increasingly unimportant when the discount rate is close to one.

Our empirical strategies allow us to overcome this identification problem. In the housing market, we overcome this problem by using the fact that the purchase decision of a home (the date the home goes under contract) is made, on average, two months before the closing date. This lag between the decision and move-in dates allows us to distinguish between a rational response to the weather state at the time of purchase and a response by agents with projection bias. Specifically, we find evidence that swimming pools are most highly valued when homes go under contract in August (the hottest month of the year). While this fits a model of projection bias (since it is the state at the time of the decision that matters), it is not consistent with a more standard model of how people should value a swimming pool since the homebuyers will likely move into their homes in October or later (perhaps the worst time from a rational perspective to purchase a house with a swimming pool).

In the car market, we utilize idiosyncratic weather shocks to overcome this identification problem. Specifically, we control for the time of year when the car purchase is made and test for the impact of abnormally warm or cold weather on purchase decisions. By controlling for the time of year, this strategy eliminates all seasonal patterns in car purchases (e.g. the value to purchasing a convertible in the spring rather than the fall). Abnormally good or bad weather on the day/week of the car purchase will still contribute a very small amount to the overall level of utility from that purchase (even for agents without projection bias). However, given that most cars are owned for a considerable period of time, we argue that - absent projection bias - the impact of one or two days of good or bad weather should have a very small impact on the overall utility of economic agents with standard preferences.

One final note regarding our conceptual framework relates to whether or not individuals correctly anticipate the path of states ($s_t, \ldots s_T$). It is possible that individuals are more likely to predict a greater number of warm-weather states in the future when the current weather is warm relative to when the current weather is cold.[3] Loewenstein, O'Donoghue, & Rabin (2003) assume that individuals correctly anticipate the path of states, but err when predicting the utility that those

[3] Some psychological evidence suggests that being in a hot or cold state may make associated states of the world seem more likely in the future (see for example, Risen & Critcher (2011) and Li, Johnson, and Zaval (2011)).

states, combined with a given consumption, will generate. In practice, these two errors (projection bias of utility and projection bias of states) both lead to similar incorrect predictions of future utility and are thus very difficult to separate. It is not clear that it is important to separate these two mechanisms given that they both can be thought of as projection biases and that they both lead to predictable decision errors.[4] However, we argue that projection bias of one's utility is likely to be the underlying mechanism. We posit this due to the prevalence of weather information that is available to people during the time of our study. (It is hard to imagine that in today's information age looking out the window is the best way to gauge what the weather will be like over the next several years.) In addition, Conlin, O'Donoghue, & Vogelsang (2007) (who also grapple with this question) cite Krueger & Clement (1994) who find that students at Brown University did a reasonable job of estimating temperature levels in Providence for different days of the year. While this is not conclusive, it provides support for the mechanism of projection bias of future utility as opposed to the mechanism of making incorrect predictions about future weather patterns.

II. Car Market

Data and Empirical Strategy. The data used in our analysis contain information about automobile transactions from a sample of about 20% of all new car dealerships in the U.S. from January 1, 2001 to December 31, 2008. The data were collected by a major market research firm, and include every new and used car transaction that occurred at the dealers in the sample. For each transaction, we observe the exact date and location (MSA) of the purchase, information about the vehicle purchased, the price paid for the car, and Census-based demographic information on the customer.

The data indicate that car transactions occur all year round, but are most common during the summer months. Of primary interest in this paper is the seasonal trend in convertible and 4-wheel-drive purchases. In Panel A of Figure 1, we illustrate the percentage of total car transactions that

[4] At the same time, one might argue that the policy implications for these two mechanisms differ. One mechanism suggests that people need help to better understand seasonal weather patterns while the other mechanism suggests that people need help to appreciate how their utility changes.

were convertibles by month of the year. Overall, convertibles make up between 1.5 and 3% of total cars purchased. The data show a strong seasonal pattern in which the percentage of cars sold that are convertibles is highest in the early spring. For seven out of the eight years, the percentage of cars purchased that are convertibles peaks in April. While springtime is the most popular time to buy a convertible, the percentage of cars sold that are convertibles is still relatively large in the winter months. The annual winter troughs in percentage of cars sold that are convertibles are well over half the magnitude of the corresponding spring peaks. These seasonal differences in convertible purchases are consistent with the standard model of state-dependent preferences discussed in the conceptual framework section: consumers do seem to take into consideration the season of the year when making convertible purchases since those first few months of consumption in the warm-weather state will likely increase total discounted utility.

Similarly, Panel B of Figure 1 illustrates the percentage of total car transactions that were 4-wheel-drive vehicles by month of the year. Four-wheel-drive transactions range between 20% and 35% of total car transactions. Panel B shows a seasonal pattern in which 4-wheel-drive vehicles are particularly popular in the early winter months (purchases usually peak in December). As was the case for convertibles, this is not yet strong evidence for projection bias since a standard model of state-dependent preferences would predict that the discounted utility of a four-wheel-drive is highest at the beginning of the winter.

We expect there to be a large amount of heterogeneity in the seasonal differences shown in Figure 1 depending on the geographic location of the dealership. To illustrate this heterogeneity, we perform a simple cut of the data by dividing MSAs into two groups: above or below the median level of monthly temperature variation.[5] Figure 2, like Figure 1, displays month-to-month sales of convertibles (Panel A) and 4-wheel-drive vehicles (Panel B) as a percentage of total cars sold, but does so separately by the variable temperature areas (e.g. Chicago) and the non-variable temperature areas (e.g. Miami). Perhaps surprisingly, Panel A shows that the overall percentage of convertibles purchased in these two types of MSAs is not too different. However, it is clear that the amount of

[5] For each MSA, we calculate the variance of month-by-month average temperature data. MSAs are then classified as above the median if their temperature variance is larger than the median temperature variance in the sample.

seasonal variation is higher in the variable temperature MSAs. Panel B shows that there is a large level difference in the percentage of 4-wheel-drive vehicles purchased across the MSA types and once again the variable temperature areas appear to a more pronounced seasonal pattern.

Our identification strategy involves testing whether abnormal weather conditions (controlling for time of year in order to eliminate seasonal purchasing patterns) are correlated with abnormally high or low sales volume of convertible and 4-wheel-drive vehicles. To do this, we collapse the data to the MSA-week-year level.[6] After collapsing, we create variables that represent the percentage of total cars sold that were convertibles or 4-wheel-drive vehicles in each MSA-week. Weekly weather data at the MSA level is also merged in.[7] These data will allow us to test whether the abnormal weather leads to abnormally high or low levels of convertible and 4-wheel-drive purchases. We proceed by first presenting the results for convertibles followed by the results for 4-wheel-drive vehicles.

Baseline Convertible Results. We begin the analysis by using two MSAs (Chicago and Miami) as examples of the effects that we find. Panel A of Figure 3 plots the convertible percentage of all cars sold in Chicago for each week between 2001 and 2008. As expected given the temperature variation that exists in Chicago, we see a strong seasonal pattern in which convertible sales range from approximately 1.5% of total cars sold in the wintertime to 3-4% of total cars sold in the spring. In accordance with our empirical strategy outlined above, we want to obtain a measure of abnormal convertible sales. To do this, we regress the weekly convertible percentage of total cars sold in Chicago on year fixed effects and week-of-the-year fixed effects. The residuals from this regression, which range from approximately -0.75% to 1% are plotted in Panel B of Figure 3. A week with a 0.5% residual is a week in which the convertible percentage of total cars sold was 0.5 percentage

---

[6] Alternatively, the data can be collapsed to the day level. We choose to do most of the analysis at the week level for three primary reasons. First, while the data contain an exact day of purchase, the paperwork may be signed and dated later than the actual date the deal was made. Thus, using day-to-day level variation is noisier than variation at the week level. Second, week-level data largely eliminates the need to worry about weekday/weekend effects as well as holidays and other events that can cause abnormal sales volume. Third, many weather events (e.g. snow storms) occur across multiple days making a weekly analysis more appropriate.

[7] The MSA-level weather data was gathered by using wolframalpha.com. The nearest weather station to each of the metropolitan areas was used. If weather data was not available for 90% or more of the daily observations, then the second-nearest weather station was used. Approximately, 10% of the MSAs required using data from the second or third-nearest weather station to the MSA.

points higher than our regression predicted for that week of the year. Figure 4 illustrates how we obtain abnormal temperature by week in Chicago. Panel A of Figure 4 shows the average daily temperature for each week. Panel B, which once again nets out week-of-the-year effects, illustrates that any given week in Chicago may be up to 20 degrees Fahrenheit hotter or 20 degrees Fahrenheit colder than would be predicted by average seasonal patterns in the data.

To test for projection bias, we want to know whether the abnormal convertible sales illustrated in Panel B of Figure 3 are positively correlated with the abnormal temperature values in Panel B of Figure 4. We find that these residuals are indeed highly correlated (correlation coefficient = 0.35; t-stat = 7.62). The size of this correlation suggests that an increase in residual temperature value by 20 degrees results in a 0.36 percentage point increase in the convertible percentage of total cars sold (a 14.4% increase given the baseline of 2.5%).[8]

A natural question is whether abnormally high temperature is only effective in the early spring. In other words, people may buy a convertible as soon as it warms up in the spring time - but may not be impacted by abnormal temperature variation in the fall. Figure 5 provides the scatter plots for abnormal temperature and abnormal convertible sales in Chicago separately for each quarter of the year. The results suggest a strong and significant correlation in quarters 1, 2, and 4 (t-stats: 5.1, 3.7, and 4.7 respectively). We argue that the lack of correlation in quarter 3 (July, August, and September) likely reflects the fact that since the weather is already so warm during quarter 3, abnormally high temperature does not increase the instantaneous utility for owning a convertible - a necessary condition for projection bias to cause an increase in purchases. Particularly important, however, is the strong positive and significant correlation in quarter 4. Similar to springtime, an abnormally warm-weather week during November in Chicago results in a large increase in the percentage of convertibles sold.

The impact of abnormal weather variation on convertible sales that we find in Chicago may not generalize to all types of MSAs. We use Miami, FL as the second example for how weather impacts

---

[8] We use 20 degrees as a convenient way to think about the overall size of the effect. The extremes of the data are temperature residuals of approximately -20 and 20 degrees. Thus, 20 degrees can be thought of as an extreme temperature value in the data relative to average, or the difference between having a somewhat lower temperature value than average (-10 degree residual) compared to a somewhat higher temperature value than average (10 degree residual).

convertible sales. Figures 6, 7, and 8 replicate Figures 3, 4, and 5 using data from the Miami-Ft. Lauderdale MSA. Figure 6 illustrates a much weaker seasonal pattern in convertible sales in Miami than was found in Chicago. In addition, Figure 7 shows that the mean daily temperature is both warmer on average and less variant in Miami than in Chicago. Due to the warmer average temperature in Miami, we would predict that abnormally warm temperature in Miami does not increase the residual fraction of cars purchased that are convertibles nearly as much as abnormally warm weather in Chicago. To directly test this, we combine the data from Figures 6 and 7 into a scatter plot. We find that the overall correlation between residual convertible sales and residual temperature in Miami is actually negative (although not statistically significant; t-stat: -0.6). Figure 8 shows that the correlation is not significant for any of the 4 quarters of the year.

We combine the data for all MSAs to estimate the impact of temperature on convertible sales across our entire sample. Table 1 reports coefficients and standard errors from a regression of the convertible percentage of total cars sold on residual weather variables (temperature, rain, snow, slush, and cloud cover) once again controlling for MSA*Year fixed effects and MSA*Week-of-the-Year fixed effects. Given the varied size of the MSAs in our sample, we weight the regression based on the total number of cars sold. Column 1 indicates that when the temperature is 1 degree higher than expected in a given MSA, the MSA is expected to experience a 0.011 percentage point increase in the convertible fraction of total cars sold. Thus a 20-degree swing in temperature in any given week, can change the convertible percentage of cars sold by 0.22 percentage points (an 8.5% change relative to the weighted base rate of 2.6% of cars sold being convertibles). Liquid inches of rain, snow, and slush fall all have negative impacts on the convertible percentage of cars sold, although these effects are relatively small given the amount of variation in rain, snow, and slush fall that exists in the data. Cloud cover is also very important for convertible demand. As the sky goes from completely clear to completely cloudy, convertible sales decrease by 0.172 percentage points. Thus, a clear sky (relatively to completely overcast) increases convertible demand by the same amount as ~16 degrees higher temperature.

The next four columns in Table 1 break down the impact of residual temperature on convertible sales by quarter of the year. Consistent with the Chicago example shown in Figures 3-5, the effect of residual temperature is large and statistically significant in quarters 1, 2, and 4, but insignificant in quarter 3 (when baseline temperature is already quite warm in most areas). Cloud cover - which is arguably important no matter what time of year - is large and significant in all quarters (including quarter 3).

As our Chicago and Miami examples exemplify, the overall effects that we present in Table 1 are likely to mask important heterogeneity that exists in the data. To better understand this heterogeneity, we estimate the impact of residual temperature on convertible sales separately for MSA-weeks with different mean values for the daily high temperature. The mean value for the daily high temperature for each MSA-week was obtained by simply calculating the average of the daily high temperatures in a given MSA-week across the different years in our sample. We then group MSA-weeks into 5-degree buckets by average daily high temperature. We regress the convertible percentage of total cars sold on residual temperature for MSA-weeks in each bucket. Figure 9 plots the estimated coefficients of residual temperature with standard error bars for 5-degree buckets of expected temperature values. For example, the leftmost point plotted in the graph is the estimated coefficient for MSA-weeks whose average daily high temperature is less than 35 degrees. This figure illustrates that abnormally high and low temperature values have large and significant impacts on convertible sales when the baseline temperature for a given MSA-week is less than 80-85 degrees. The point estimates for these degree bins range from 0.010% to 0.019%. As the average daily high temperature rises above 80 degrees, however, we find that abnormal temperature variations have little effect on convertible sales. In fact, we find negative values at the very highest temperature ranges suggesting that an increase (decrease) in mean daily temperature over these hot baselines can have a negative (positive) impact on convertible purchases. These heterogeneous effects explain the zero-effect that we found for Miami since Miami's expected temperature is nearly always above 80 degrees.

Baseline 4-Wheel-Drive Results. While buying a convertible may seem especially attractive on a warm day, it is cold and snowy days that make four-wheel-drives seem like an especially good idea. Table 2 presents the impact of abnormal weather variation on the 4-wheel-drive percentage of total cars sold. As we expected, the results we find are the opposite of what we found for convertible sales. We find that colder temperature values lead to more 4-wheel-drive purchases. For example a 20-degree change in temperature can lead to a 1.0 percentage point change in the percentage of 4-wheel-drive vehicles purchased (a ~3% change relative to the weighted baseline of 33.5% of cars sold being 4-wheel drive). We also find a large, positive impact of snow and slush on abnormal 4-wheel-drive transactions. One inch of liquidized snow (~10 inches of snow) leads to a 1.02 percentage point increase in the percentage of total cars sold that have 4-wheel drive. The effects for snow fall are significant in quarters 1 and 4 (the standard errors for quarters 2 and 3 indicate that we do not have sufficient snow-fall variation to estimate effects in these quarters). The effect of snow fall is larger in quarter 4 than in quarter 1. However, the significant impact of snow fall in quarter 1 suggests that even a snow storm that occurs towards the end of the winter season can have a powerful impact on 4-wheel-drive purchase behavior.

The amount of snow fall each week has a very different distribution from the distribution of temperature. Snow fall is usually zero in most MSA-weeks, but can have very large values in other MSA-weeks. The nature of this variable suggests a modeling approach along the lines of an event-study design. What happens in the weeks leading up to and after a big snow storm? We present the results from an event-study design in Figure 10. We choose the events to be the largest snow storm of the year (measured in amount of snow fall) for the MSAs in our sample that have above median weather variation (this excludes places with no snow variation). We then plot the weighted residuals of the 4-wheel-drive percentage of total cars sold for the 12 weeks before and the 12 weeks after each of these events. As can be seen in Figure 10, we find limited evidence that individuals increase their 4-wheel-drive purchases leading up to a snow storm. We then see a large spike at the event date such that the percentage of cars sold that have 4-wheel-drive goes up by almost 1 percentage point.

This effect diminishes but continues to be significant for two more weeks before returning to baseline.

Our analysis uses the percentage of total cars sold that have 4-wheel drive as the outcome of interest. Thus, a change in this measure can be due to an increase in 4-wheel-drive purchases or a decrease in purchases of vehicles without 4-wheel drive. Analysis on the level amount of convertible and 4-wheel drive purchases made (in logs) confirm the finding that convertible purchases increase substantially during warm-weather weeks, but show that 4-wheel-drive purchases actually decrease during and after snow storms - but not by nearly as much as purchases of cars without 4-wheel drive. Thus, it is worth noting that the 4-wheel-drive results are driven in part by a drop in overall volume. After a snow storm, it is individuals who are very interested in purchasing a 4-wheel drive that are the most likely to go to the dealership.

Dynamic Analysis. The effects that we find for convertibles and 4-wheel-drive vehicles suggest that, due to projection bias, idiosyncratic weather differences from week to week can have a large impact on the types of cars that people choose to purchase. One concern with this story, however, is that abnormal weather may appear to be increasing the demand for certain types of cars, but is actually just causing short-run substitutions in car purchasing behavior. An example of this "harvesting" story is that a consumer may be interested in purchasing a convertible sometime in the next month and then actually makes her purchase whenever it happens to be a nice day outside.[9] In fact, our previously noted finding that abnormally warm weather in November can impact convertible purchases and a snow storm in February can impact 4-wheel-drive purchases casts doubt on harvesting as the sole cause of our results. However, these end-of-season purchases cannot rule out harvesting entirely as a contributing factor to our results.

In order to directly address short-run intertemporal substitution of purchases, we perform a dynamic analysis that adds to the weather residual variables during the week of purchase a one-week lead and 12 weeks of lagged weather residual variables. By including lag variables, we are able to test

[9] The fact that more convertibles are bought in spring than winter and the reverse for four-wheel-drives suggest that there is harvesting in respons to the overall seasonal pattern of the weather. However, this does not mean that harvesting happens in response to idiosyncratic weather variation.

whether having cold or hot weeks leading up to the week of purchase can have an impact on how the current weather affects behavior. For example, in the convertible scenario, negative coefficients on the lag variables would be evidence of harvesting; such coefficients would imply that unusually hot weather in the previous weeks will reduce convertible purchases in the current period, while unexpectedly cold weather in the previous weeks will increase convertible purchases in the current week. Complete harvesting would lead to negative coefficients on the lag variables whose sum entirely cancels out the positive coefficient on the current weather variable - resulting in a zero effect on the total fraction of cars sold across the year that were convertibles.[10]

Table 3 presents the results of this dynamic analysis for convertible purchases. The results once again show a large and significant effect of current weather on convertible purchases. The lag variables are all small relative to the current temperature coefficient, in most cases not statistically significant, and more often positive than negative. in the full year data (Column 1 of Table 3), there is no evidence that warmer than usual weather in the previous weeks affect the current week's sales. If anything, it appears that several weeks of warm weather in a row might lead to an even larger demand for convertibles. There is also no evidence that warm weather in the following week (the lead 1 variable) has a significant impact on current convertible sales, which serves as a nice placebo test.

Table 4 provides a similar analysis for 4-wheel-drive purchases. This analysis indicates that snow fall anytime in the last 3 weeks leads to an increase in the percentage of cars sold that are 4-wheel drive. There is, however, evidence of some short-run substitution in demand. The summation of the coefficients for lag 4 through lag 12 is -1.13%. Thus, approximately 47% of the positive effect of snow fall on 4-wheel-purchases (2.42%) could be considered due to harvesting. In other words, the increase in the percentage of 4-wheel drives purchased after a snow storm is smaller if there was a snow storm that occurred sometime in the previous two to three months. (Presumably, this is because some people purchased a 4-wheel drive in the wake of the earlier snow storm and no longer need to buy one.) Overall, this dynamic analysis suggests that the increase in demand for

[10] See Jacob, Lefgren, and Moretti (2007) for a similar analysis that tests for intertemporal substitution of crime using abnormal weather shocks.

convertibles and much of the increase in demand for 4-wheel-drive vehicles that we find due to abnormal weather variation cannot be explained by short-run intertemporal substitution in demand.

Test Drive Timing. One aspect of car purchasing that may lead to a correlation between weather and car purchase timing, particularly for convertibles, is the desire of most customers to test drive a car before buying. Suppose a customer is considering buying a convertible. Suppose that she does not suffer from projection bias, meaning that she has no problem accurately forecasting her utility from owning a convertible in various weather states. But suppose that, before she buys the convertible, she would like to be able to test out various features of the convertible: how convenient it is to put the top up and down, how much wind or road noise she experiences with the top down, etc. It is unpleasant to do such a test drive when the weather is cold, so she waits for a warm day to go to the dealership, test drive, and ultimately purchase the convertible. Alternatively, suppose that another customer suddenly needs a replacement car, perhaps because his current car has broken down and is no longer worth repairing. Suppose that a convertible is one of the vehicles he would consider purchasing, but the day he needs the new car is too cold to test drive a convertible. Unwilling to buy the convertible without being willing to test out the convertible features of the car, he buys a non-convertible instead.

The behavior of both of these types of customers would lead to a higher percentage of cars sold on warm days being convertibles relative to cold days, but not because of projection bias. We can get a sense of how important test drive timing might be for our results by considering the effect of cloud cover. There is no reason that a customer could not test drive a convertible on a day that is cloudy—as long as it is not cold or rainy. Thus, in our regressions, which control for temperature and rain, we should not see an effect of cloud cover if the reason for the correlation between temperature and convertible purchases is test drives. However, projection bias should lead to warm, sunny days being days on which people are particularly likely to buy convertibles, rather than warm cloudy days. Indeed, if we examine the results in Table 1, we find that unusually cloudy days has a significant negative effect on the percentage of cars sold that are convertibles, consistent with projection bias. It is particularly noteworthy that cloudy days have a negative effect in all four quarters, and its effect is largest in the third quarter, when days are generally warm. This third quarter effect is especially suggestive that people buy more convertibles on warm days not because it is more possible to test drive them, but because it seems more attractive on such days to own a convertible.

Car Buyers who Previously Owned a Convertible or 4-Wheel Drive. Another alternative hypothesis that would explain our findings is that customers need somewhat extreme weather days (warm, sunny ones or cold, snowy ones) in order to actually learn what their utility will be from owning either a convertible or a four-wheel drive in such weather conditions. Under this hypothesis, a warm, sunny day does not lead a customer to overestimate the utility she will get from owning a convertible; instead it enables her to learn for the first time how high her true utility will be from owning a convertible in such weather states. Before considering this as an alternative hypothesis, we note that this type of extreme learning story—in which car buyers can't quite imagine what it would be like to own this car in another state of the world even when they have experienced that state of the world many times—starts to mesh together with exactly what projection bias is; namely, the inability to appreciate the utility that one will experience when the state of the world changes.

Even though projection bias and learning might look similar at their extremes, our data allow us to investigate somewhat more direct evidence for learning as an explanation. In our data, we observe what trade-in, if any customers bring when they buy a car. This means we can observe car transactions by customers whom we know have already owned a convertible or have already owned a four-wheel drive. Previous convertible owners are less likely to need to "learn" about what it is like to own a convertible during a warm weather state, and similarly for previous four-wheel-drive owners and cold, snowy states, so evidence that abnormal weather impacts these buyers is particularly strong evidence for projection bias.

If we look within the subset of transactions that use a convertible as a trade-in, we find that approximately 25% of these buyers purchase another convertible while 75% purchase a hardtop. Column 1 of Table 5 reports the results of our baseline specification if we restrict the sample to buyers who are trading in a convertible. While the standard errors are much larger due to the sample restriction, we continue to find a positive impact of residual temperature at the time of purchase on convertible demand. In fact, the point estimate is about six times larger than the point estimate in the entire sample - although because the convertible purchase rate in this sample (25%) is so much higher the point estimate is much larger in percentage point terms and a bit smaller in percentage terms.[11] In Column 2 of Table 5, we estimate the effect of weather on buyers who are trading in a four-wheel-drive vehicle. Overall, 78% of people who trade in a 4-wheel drive vehicle purchase another 4-wheel drive. In Column 2 we continue to find strong and statistically significant effects of abnormal weather—including temperature, snow fall, slush fall, and cloud cover—on 4-wheel-drive purchases for buyers who traded in a 4-wheel-drive vehicle. The estimated effects are substantially smaller in percentage terms than in the full sample, in large part because the unconditional probability of buying a four-wheel-drive vehicles is so high in this sample.

The fact that we find effects of abnormal weather in precisely the subsample of buyers who would seem to have the *least* to learn about their utility from owning either a convertible or a four-wheel drive cast doubt on a learning story explaining the effects that we find.

Returning Vehicles. Projection bias suggests that people can make mistakes when purchasing a durable good and that people may realize the mistake when the state of the world changes. Conlin, O'Donoghue, and Vogelsang (2007) make this case and specifically test for mistakes by analyzing whether cold-weather items (boots, gloves, etc.) purchased by mail order were more likely to be returned if the purchase was made during a very cold state. In the car market, projection-bias mistakes might be identified by seeing vehicles that were purchased during abnormal weather weeks reappear in the market (either as trade-ins or as subsequent used car sales) more quickly than cars that were purchased during normal weather weeks. The quick return of a vehicle to the market could indicate that the owner was not happy with the purchase he or she made.

Unfortunately, there are at least two reasons why testing for early returns in the car market is much harder than for catalog orders. First, is simply a data limitation. Although our data is

[11] The full sample results indicate that a 20-degree increase in abnormal temperature increases the percentage of vehicles sold that are convertibles by 0.22 percentage points in the full sample, a 8.5% increase relative to a base percentage of 2.6%. In the "convertible trade-in" subsample, the effect is a 1.2 percentage point increase, a 4.8% increase relative to a base percentage of 25%.

impressive and represents a 20% sample of all new car dealerships in the U.S., we can only identify "returned" cars that happen to be traded in or sold as a used car at one of the dealerships we observe. Said another way, for any vehicle whose sale we observe at some point, we have only a 20% chance of seeing that car's subsequent return or resale if that transaction happens at a dealership, and no chance of seeing it if that transaction happens person-to-person. Second, and perhaps more importantly, car dealerships do not offer the kind of "no-hassle return" policies that are common for catalog retailers. A mistake that is made when buying gloves can be easily fixed with a few minutes and a little postage. However, an individual who realizes they made a mistake after buying a convertible cannot return it so easily. To switch the convertible for a hardtop will likely require that the individual sell the convertible at a loss and buy the hardtop at a premium. Thus, even if mistakes are being made, the mistakes may not be large enough to merit fixing.

Despite these two concerns, we test for the impact of abnormal weather at the time of purchase on how quickly the car reappears in the market. Of the ~40 million cars that are sold in our dataset, 2.37% of them reappear within 1 year as a trade in or subsequent sale, 5.03% within 2 years, and 7.16% within 3 years.[12] On average in the U.S., owners keep their cars for just over 5 years (Polk, 2010).

Our empirical strategy is to estimate whether convertibles that were purchased when the weather was abnormally warm and four-wheel-drives that were purchased when the weather was abnormally snowy are more likely to reappear in our data within a short time frame than vehicles purchased under more typical weather conditions. The columns of Table 6 report results for regressions in which the outcome variable is an indicator that equals one for a given transaction if we observe the transacted vehicle reappear in our data as a trade-in or in another sales transaction within, respectively, 1, 2, or 3 years. We control for year*week*msa fixed effects to eliminate seasonal and geographic differences in how quickly cars are returned. Table 6 shows that convertibles are 1.272 percentage points more likely to be returned within a year than other types of cars. 4-wheel drives are also more likely to be returned (by 0.285 percentage points) than other types

[12] Unique identification numbers corresponding to individual VIN numbers are used to track vehicles over time.

of cars. The positive signs of the coefficients estimated for the interaction of convertible and residual temperature variables are consistent with projection bias: convertibles are more likely to be returned quickly when they were purchased during abnormally warm weather weeks. However, this result is statistically significant only in column 2. The point estimates suggest that when the weather is 20 degrees warmer, convertibles are 0.34 percentage points more likely to be returned within 2 years than hardtops (a 6.8% change off the base return rate of 5.03%). The residual temperature interaction with 4-wheel-drive vehicles is more consistently statistically significant, and indicates that a four-wheel drive is more likely to be returned within 1, 2, or 3 years if it is purchased in an abnormally cold week. Overall, our results for the effect of abnormal weather on returning vehicles is less strong that our evidence for the effect on purchasing vehicles. An important contributor to this is simply that the number of vehicles we see sold and then see reappear within our data is not that high. As a consequence, we have limited ability to identify differences in the rates at which vehicles are returned under different circumstances.

III. Housing Market

Data and Empirical Strategy. Our analysis is based on a housing dataset of more than four million single-family residential properties that sold at least once across the United States between 1998 and 2008. We purchased the data from a commercial vendor who had assembled these data from assessor's offices in individual towns and counties.[13] Since larger metropolitan areas are more likely to electronically archive their assessor data and sell it to commercial vendors, urban counties are over-represented relative to rural counties.[14] The data include the transaction price of each house and the sale date, the previous transaction price and sale date (if a previous sale occurred during the time frame of the data), a physical address (from which we obtain county and state indicators), and a consistent set of structural characteristics, including swimming pool, central air, fireplace, lot size, year built, square feet of living area, number of bathrooms, and number of bedrooms.

---

[13] The commercial vendor is Dataquick which is the source of housing data for many of the papers in the literature.
[14] Certain states are overrepresented in the data. For example, 30.7% of sales were in California, 14.1% in Florida, 8.9% in Ohio, 6.9% in Washington, 4.8% in Massachusetts, and 4.2% in Nevada. Data also include observations from AL, CO, CT, GA, HI, IA KY, MI, MN, MO, NC, NE, NY, OK, OR, PA, RI, SC, TN, TX, and VA.

While we observe the closing date for each home, we do not observe the date that the home went under contract - which is the relevant date for testing a model of projection bias. Throughout the analysis we assume that homes go under contract two months prior to the closing date. We base this assumption off of a small dataset consisting of homes in the Chicago area for which both contract date and closing date were available.[15] In this dataset, the average time between contract and closing dates was 52 days. After including a few days for price negotiations, we assume that purchase decisions were made on average two months prior to the closing date for each house. Because we lack exact data for the day the home went under contract, our empirical strategy in this section of the paper is going to be restricted to seasonal patterns rather than precise idiosyncratic weather differences at the time of purchase.

We clean the data in a similar manner as previous work using the same data in order to eliminate outlying observations. Housing transactions are dropped if the sales price was less than $5,000 or more than $5,000,000, if the house was built before 1900, if the square footage is less than 250 square feet or more than 10,000, if the number of bathrooms is less than .5 or more than 10, and if the number of bedrooms is less than 1 or more than 10. We also drop all new construction (age less than 2 years old).[16] We also restrict the sample to houses located in counties that report whether a home has a swimming pool. As described in our empirical section below, in order to perform a repeat-sales analysis, we restrict the sample to homes in our data for which we have more than one transaction record.

Table 7 provides some basic summary statistics for the final set of cleaned housing transactions in our dataset (~4.2 million). Average sales prices over the entire time frame of our data were approximately $275,000, again reflecting the fact that urban areas (and California) are overrepresented in our housing data sample relative to the entire population of housing within the United States. The Table also shows that 12% of homes in our sample had swimming pools, 30%

[15] We thank Steve Levitt and Chad Syverson for sharing this information with us.
[16] We do not have an indicator in the dataset for when a home is being sold for the first time. One potential concern is that new homes (which sell for a premium) may be more likely to have a swimming pool and may also have a strong seasonal pattern (which could bias in favor of the results we find). Because we lack an indicator for new homes, we simply drop homes that may fall in this category.

had central air, and 45% had at least 1 fireplace.[17] The average home was built in 1968 on a 0.32 acre lot with approximately 3 bedrooms, 2 baths, and 2,000 square feet of living area.

Our goal is to test for the presence of projection bias using a very simple empirical strategy and to provide the results in a graphical fashion. Our primary specification is

$$Log(Sales\ Price)_{itc} = \gamma_i + \theta_{tc} + \varepsilon_{itc}, \qquad (4)$$

where $Log(Sales\ Price)_{itc}$ is the log sales price of house i in sample-month*year t in county c. gamma_i is a fixed effect for house *i,* estimable because we use only houses we observe being sold more than once. Theta_tc is a sample-month*county fixed effect. The residual from Equation (4) represents, for each house transaction, how much more or less the house sold for than would have been expected after considering how much that exact house sold for on other occasions, and how much other houses sold for in the same sample-month and county. We will analyse the residuals from this regression by month and house type to see whether there is evidence of projection bias.

Results. We begin by calculating for homes with swimming pools the average residual (obtained from Equation (4)) by the month-of-year the house is assumed to have gone under contract. In Panel A of Figure 11, we plot these average residuals along with their standard error bars.[18] It is worth noting that the weighted average of the residuals for homes with swimming pools must sum to zero across months (because house fixed effects are included in the regression). Similarly, the weighted average of the residuals for all homes must sum to zero within every sample-month (because sample-month*year fixed effects are included in the regression). Therefore, whenever we see a positive average residual in a given month-of-year for houses with swimming pools, we know that the average residual for houses without swimming pools must be negative (although the magnitude of the average negative residual is approximately 10% as large as the positive residual since houses with swimming pools only represent about 10% of the data).

[17] Certain counties in our dataset reported all homes as having no central air or a fireplace (which provides an indication that these data were not systematically collected in those counties). Given that all of our analysis has county fixed effects, we leave these homes in the dataset in order to provide more observations for other housing characteristics of interest (e.g. swimming pools).

[18] Providing average residual values by month of the year is equivalent to reporting regression estimates for the interaction between the dummy variable for swimming pool and month of the year indicators.

Panel A of Figure 11 provides the first evidence that swimming pools add more value to a home in the summertime than in the wintertime. Specifically, homes with swimming pools that go under contract in the three hottest months of the year (June, July, and August) sell for on average 0.2 percentage points more than otherwise expected (this effect is jointly statistically signficant and individually significant for June and August) while homes with swimming pools that go under contract in the coldest months of the year (December, January, and February) sell for on average 0.16 percentage points less than otherwise expected (this effect is individually significant for December). Given that the average transaction value of houses with swimming pools in our data is ~$398,000, this represents a ~$1400 swing in value for homes with swimming pools that go under contract in the summertime relative to the wintertime.

One concern with this simple analysis is that while it is clear that the residuals for homes with swimming pools are showing a seasonal trend, it may not be the swimming pool that is causing the seasonal trend, but rather something else about homes with swimming pools. For example, perhaps the seasonal differences are being driven by large homes, which may be more likely than small homes to have swimming pools. To assuage this concern, we obtain conditional residuals for each month of the year, regressing the residuals from Equation (4) on a dummy variable for whether the house has a swimming pool and also on the other housing characteristics in our dataset (lot size, central air, fireplaces, square footage, number of bedrooms, and number of bathrooms). Panel B presents the conditional residuals on the swimming pool dummy variable for each of the twelve regressions. Controlling for the seasonal pattern of the other housing characteristics in our dataset does not change the overall effect of swimming pools, which continues to show that the value of a swimming pool is higher in the summertime than in the wintertime.

A common procedure when running hedonic models involves trimming the data to eliminate extreme residual values. For example, if the data suggest that a house sold for $100,000 and then sold two years later for $800,000, it is reasonable to assume that there was a data mistake or that the house was changed in a major way. To remove these types of observations, we trim the data to eliminate the top and bottom 1% of residual values and the top and bottom 5% of residual values.

Removing the top and bottom 1% of residual values eliminates homes that sold for more than 61.6% more than would be expected, or less than 61.6% less than would be expected. Removing the top and bottom 5% of residual values eliminates homes that sold for more than 24.1% more than would be expected or less than 24.1% less than would be expected.[19]

Figure 12 reports the swimming pool results once again (while also conditioning on other housing characteristics as in Panel B of Figure 1) with the 1% trim (Panel A) and the 5% trim (Panel B). The seasonal pattern for the value of swimming pools remains when trimming the data in this manner. The major advantage to this trimming is that the standard errors become much smaller due to the elimination of these high-variance observations. Our preferred specification (with the 5% trim), provides precise month-to-month point estimates for the value of a swimming pool and shows consistently higher values for homes with swimming pools that sold in the summertime (especially August) when compared to those same homes that sold in the wintertime.

Overall, we a priori identified four housing characteristics in our data that could arguably have a seasonal component. Along with swimming pools, these variables include central air conditioning, fireplaces, and lot size. In Figure 13, we report conditional residuals (with a 5% trim) for these three variables in the same way that we reported the conditional residuals for swimming pools. Panel A illustrates the results for central air. There appears to be a slight seasonal pattern, where central air is worth less in the wintertime than in the summertime, but the results are smaller and less statistically significant than those found for swimming pools. Panel B and Panel C present the results for fireplaces and lot size, respectively. There is no discernible pattern in either of these results. Furthermore, only 2 of the 24 month effects are significantly different from zero - similar to what one would expect as a result of randomness.

Why do we find small or no results for these other three housing characteristics? There are several possibilities. It could be that the instantaneous consumption value that these other characteristics provide to homeowners does not vary with season as much as the consumption value

[19] The 5% and 1% cutoffs for trimming are symmetric (the same for more than and less than expected) because our data consists of exactly 2 observations for each house and we include house fixed effects in our regression. Therefore, every observation in our sample with a positive residual has an observation in the data with a residual of the same magnitude but of the opposite sign.

of swimming pools across seasons. For example, it is not entirely clear when a large lot size is most valuable. It could be in the spring or fall when yards are very beautiful, or maybe the summer. Another reason for the lack of seasonal effects is that these other variables may simply be serving as proxies for other housing characteristics.[20] When we run a cross-sectional regression (no house fixed effects) of log house price on the housing characteristics in our dataset (while still controlling for sample-month*county), we find implausibly large coefficients on central air and fireplaces. This hedonic regression indicates that a fireplace adds 12.1% of value to a home while central air adds 18.7% of value.[21] These effects are likely so large because homes with central air are likely to be homes that have other positive features that we are unable to control for in our regressions (recently remodeled, granite countertops, new carpet, etc.).[22] Therefore, when we test for a seasonal pattern in the value of central air, we may be simultaneously also testing for the seasonal value of having an updated kitchen (which likely does not vary with season, or may even be the most valuable in the wintertime). The correlations that central air and fireplaces may have with other important housing characteristics make identifying a seasonal trend in the value of these characteristics difficult.

In Figure 14, we report the seasonal value of the other housing characteristics in our data which are unlikely to have a strong seasonal component (number of bedrooms, number of bathrooms, and square footage). We find no significant seasonal pattern for these housing characteristics. The lack of seasonal variation in all of the other housing characteristics in our dataset (both in terms of statistical significance and effect size) lends credibility to the large and significant effects that we find for swimming pools.

Our hypothesis is that higher temperature levels at the time of the purchase decision cause swimming pools to have a larger positive effect on the sales price of the house when compared to

[20] Of course, swimming pools may also have this problem. Intuition suggests, however, that a housing characteristic like central air may be more correlated with other housing amenities than a housing characteristic like a swimming pool. Furthermore, if swimming pools are correlated with other housing characteristics (that aren't necessarily seasonal), then it suggests we may be underestimating the true seasonal difference in the value of a swimming pool.

[21] This cross sectional regression suggests that a swimming pool increases the value of a home by 13.4%.

[22] For example, Levitt & Syverson (2008) run a cross-sectional hedonic regression using Multiple Listing Services data for houses in the Chicago area. These data include many additional housing characteristics (e.g. granite countertops). When controlling for this expanded set of housing characteristics, they find that central air only adds 6.8% of value to a home (approximately one third the size of our effect).

purchase decisions made during colder parts of the year. Up to this point, however, we haven't used exact temperature, but rather have been using month of the year as a proxy for temperature. Given the variation in weather that exists across the U.S. and across different years in our sample, month of the year is clearly not a perfect proxy for temperature. We remedy this by merging in weather data for every sample-month*county, which allows us to know the average daily high temperature for the month and location in which each house in our dataset went under contract.[23]

The underlying model for how weather and housing characteristics such as a swimming pool interact to impact housing sales is not obvious. For example, it could be that a swimming pool becomes more valuable for every 1 degree increase in temperature. Alternatively, the value of a swimming pool may be constant until the high temperature reaches some hot tipping point (e.g. 70, 80, or 90 degrees). Table 8 reports regression estimates of the impact of our housing characteristics interacted with the weather at the time of purchase for these different models. The first column in Table 8 indicates that for every 1 degree Fahrenheit increase in the average daily high temperature during the month in which the house went under contract, a swimming pool increases the sales price by 0.013%. This means that a house that sold when the average daily high temperature was, for example, 80 degrees sold for 0.65 percentage points more than the same house that sold when the average daily high temperature was 30 degrees. This effect is statistically significant and remains large and statistically significant when trimming the data to eliminate the top and bottom 5% of residual values (Column 2). The interaction effects of temperature and housing characteristics other than swimming pools are nearly all small and statistically insignificant (with the exception of central air being more valuable during high temperature months in Column 2). The next 3 sets of columns in Table 8 show the impact of the average daily high temperature being above a threshold of 70, 80, or 90 degrees. Once again we find large and mostly statistically significant effects for the value of a swimming pool. For example, the final column in the table suggests that houses with swimming pools that went under contract in a month where the average daily high temperature was more than

[23] The temperature information comes from the PRISM Climate Group based at Oregon State University, which provides consistent weather information all across the United States. More information on the weather data we use can be found at http://www.prism.oregonstate.edu/. We accessed the data on 3/12/2011.

90 degrees sold for 0.37% more than these same houses with swimming pools that did not reach that average temperature threshold in the month they went under contract.

IV. Conclusion

Many of the most important decisions that we make in life involve predicting our future preferences. This paper provides evidence that projection bias may limit our ability to accurately make these predictions. We show that projection bias causes consumers in the housing and car market to make decisions that are overly influenced by the weather at the time of the decision. We argue that our results imply that projection bias can have important implications for large-stakes markets and that this psychological bias merits additional study and attention. One specific extension that would be particularly useful would be to study projection bias for various other state variables - not just weather. For example, emotional states and states of dependency, are likely to influence important decisions like having a baby, whether to get married, and whether to accept a given job offer.

From a policy perspective, our results suggest that laws designed to help consumers better evaluate their decisions may be beneficial. For example, laws that allow consumers a "cooling-off period" for durable goods or extended contracts may provide significant benefits to consumers and may provide incentives for sellers to help buyers be in a "cool" state before an important transaction or contract is made.[24] The Federal Trade Commission has an explicit cooling-off rule that applies to situations when you buy an item in your home or at a location that is not the seller's permanent place of business. This law was made specifically to deal with high-pressure sale situations like with door-to-door salesman. However, this rule exempts real estate and automobiles although there still can be high-pressure sale situations for these important durable goods. In the case of cars, a law allowing car buyers to return a vehicle soon after the original purchase may help those consumers that were highly influenced by the weather at the time of their decision, while not imposing significant costs on more rational car buyers.

[24] See Camerer et al. (2003) for an extended discussion about cooling-off periods and their potential applications in settings where people make suboptimal choices.

References

Camerer, C., S. Issacharoff, G. Loewenstein, T. O'Donoghue, and M. Rabin. 2003. "Regulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism'" *University of Pennsylvania Law Review*, 151(3): 1211-1254.

Conlin, M., T. O'Donoghue, and T. Vogelsang. 2007. "Projection Bias in Catalog Orders." *American Economic Review*, 97(4): 1217-1249.

DellaVigna, S. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature,* 47, 315-372.

Jacob, B., L. Lefgren, and E. Moretti. 2007. "The Dynamics of Criminal Behavior: Evidence from Weather Shocks." *The Journal of Human Resources,* 42(3), 489-527.

Levitt, S. and C. Syverson. 2008. "Market Distortions When Agents Are Better Informed: The Value of Information in Real Estate Transactions." *The Review of Economics and Statistics,* 90(4), 599-611.

Li, Y., E. Johnson, and L. Zaval. 2011. "Local Warming: Daily Temperature Change Influences Belief in Global Warming." *Psychological Science,* 22(4), 454-459.

Loewenstein, G. 1996. "Out of Control: Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes,* 65, 272-292.

Loewenstein, G., D. Nagin, and R. Paternoster. 1997. "The Effect of Sexual Arousal on Predictions of Sexual Forcefulness." *Journal of Crime and Delinquency,* 32, 443-473.

Loewenstein, G., T. O'Donoghue, and M. Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics*, 118: 1209-1248.

Loewenstein, G. and D. Schkade. 1999. "Wouldn't It Be Nice? Predicting Future Feelings." in *Well-Being: The Foundations of Hedonic Psychology*, ed. Daniel Kahneman, Edward Diener, and Norbert Schwarz, 85-105. New York: Russell Sage Foundation.

Nisbett, R. and D. Kanouse. 1968. "Obesity, Hunger, and Supermarket Shopping Behavior." Paper presented at the 76th Annual Convention of the American Psychological Association, San Francisco, CA.

Nordgren, L., J. Van Der Plight, and F. Van Harreveld. 2006. "Visceral Drives in Retrospect: Explanations About the Inaccessible Past." *Psychological Science,* 17, 635-640.

Nordgren, L., J. Van Der Plight, and F. Van Harreveld. 2007. "Evaluating Eve: Visceral States Influence the Evaluation of Impulsive Behavior." *Journal of Personality and Social Psychology,* 93(1), 75-84.

Read, D., and B. van Leeuwen. 1998. "Predicting Hunger: The Effects of Appetite and Delay on Choice." *Organizational Behavioral and Human Decision Processes*, 76, 189-2005.

Risen, J. and C. Critcher. 2011. "isceral Fit: While in a Visceral State, Associated States of the World Seem More Likely." *Journal of Personality and Social Psychology,* 100(5), 777-793.

Simonsohn, U. 2010. "Weather to Go to College." *Economic Journal,* 120(543), 270-280.

Van Boven, L. and G. Loewenstein. 2003. "Social Projection of Transient Drive States." *Personality and Social Psychology Bulletin,* 29, 1159-1168.

Figure 1 - Seasonal Trends in Car Purchases. This figure illustrates the percentage of total cars that were sold in each month between 2001 and 2008 that were convertibles (Panel A) and 4-wheel drive vehicles (Panel B).

Panel A. Convertible Percentage of Cars Sold

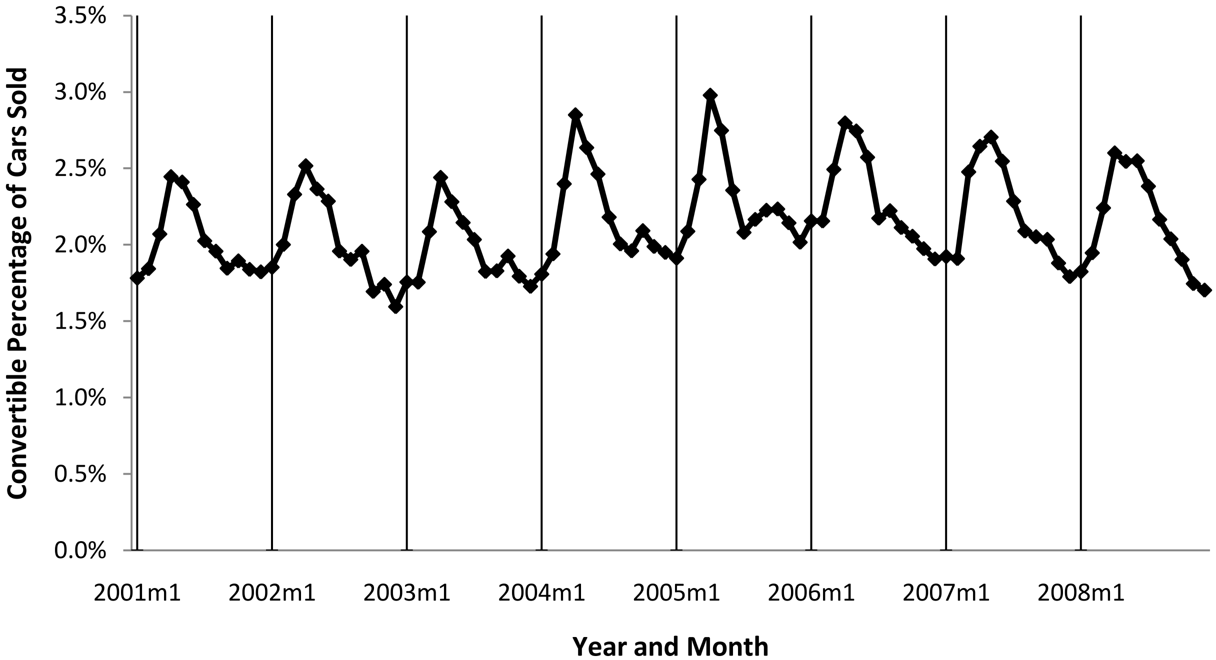

Panel B. 4-Wheel Drive Percentage of Cars Sold

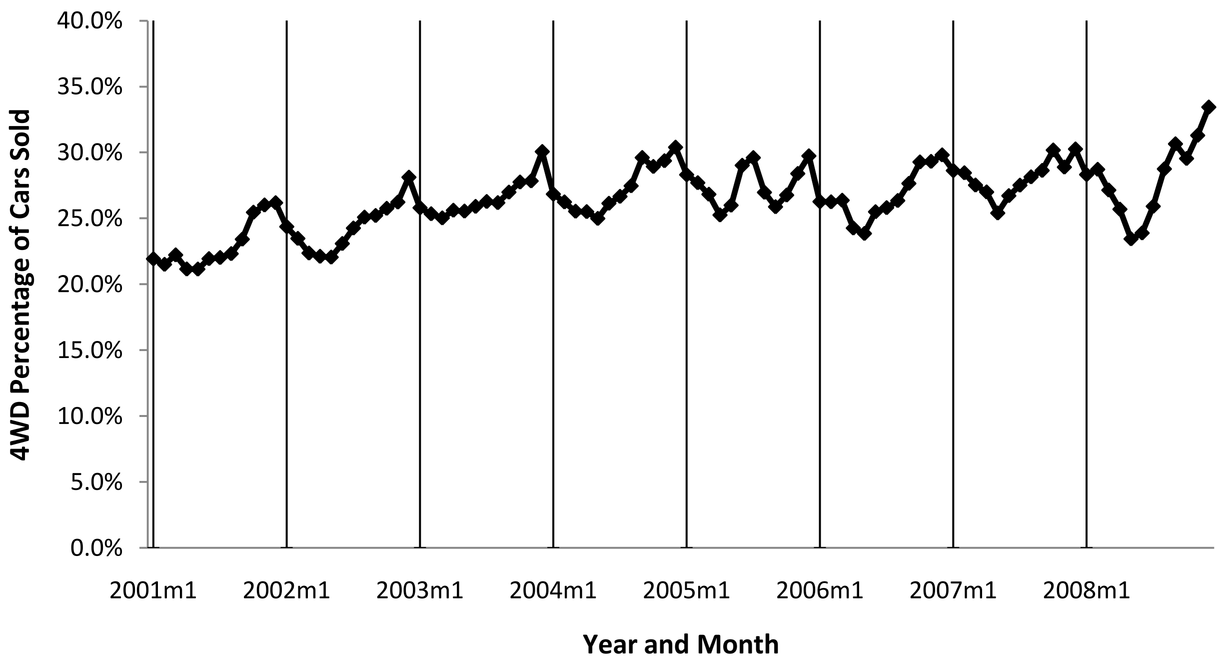

Figure 2 - Seasonal Trends in Car Purchases by Temperature Variation. This figure illustrates the percentage of total cars sold between 2001 and 2008 that were convertibles (Panel A) and 4-wheel drive vehicles (Panel B) separately for MSAs above and below the median level of monthly MSA temepraturе variation.

Panel A. Convertible Percentage of Cars Sold by Temperature Variation

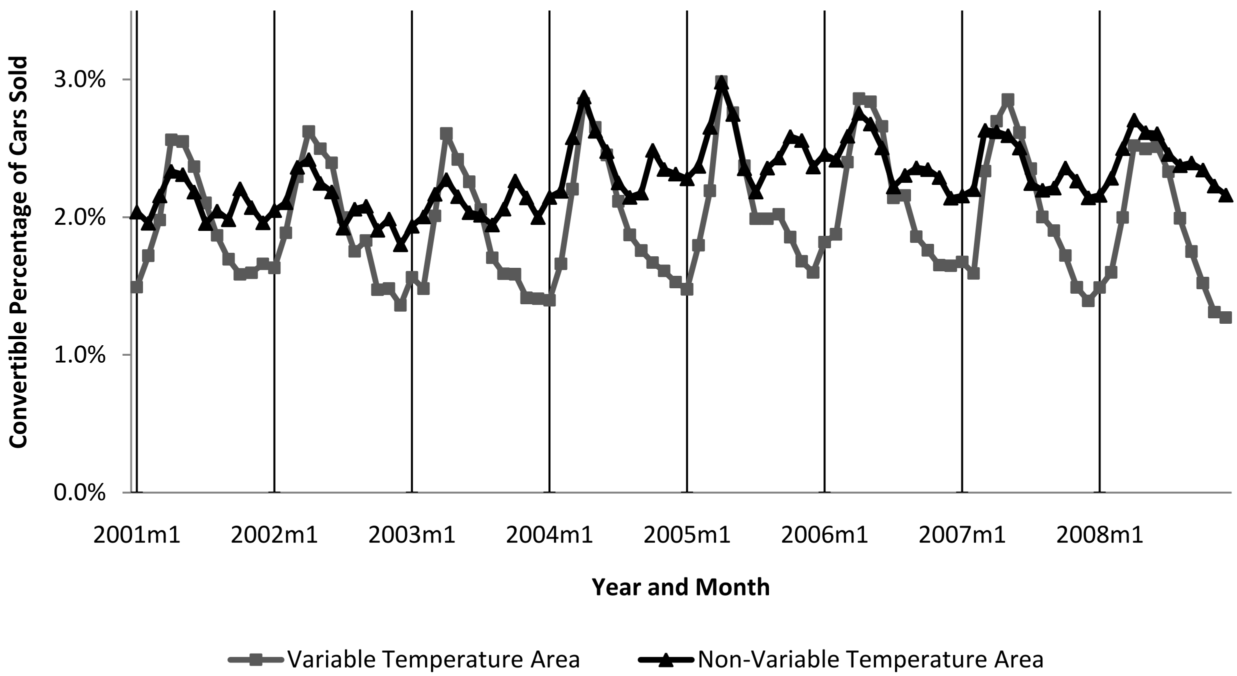

Panel B. 4-Wheel Drive Percentage of Cars Sold by Temperature Variation

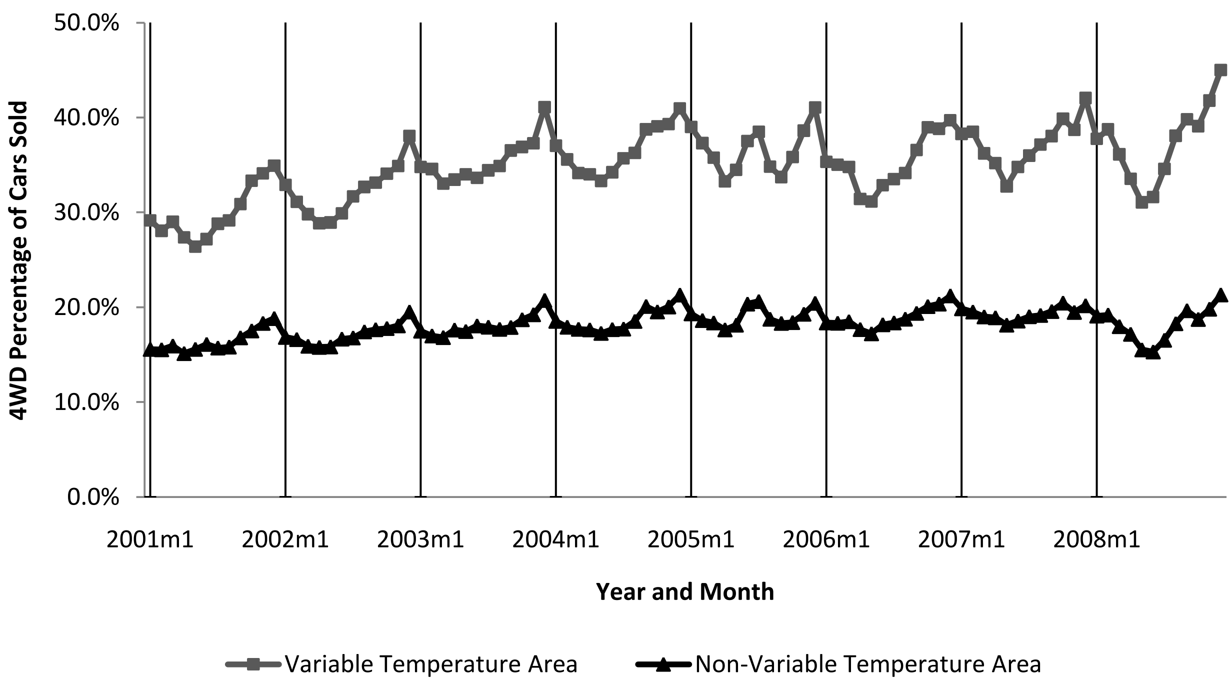

Figure 3. Residual Convertible Sales - Chicago. Panel A illustrates the percentage of cars sold in Chicago for each of the 52 weeks in a year that were convertibles. Panel B plots the residual convertible percentage of cars sold in each week.

Panel A. Convertible Percentage of Cars Sold - Chicago

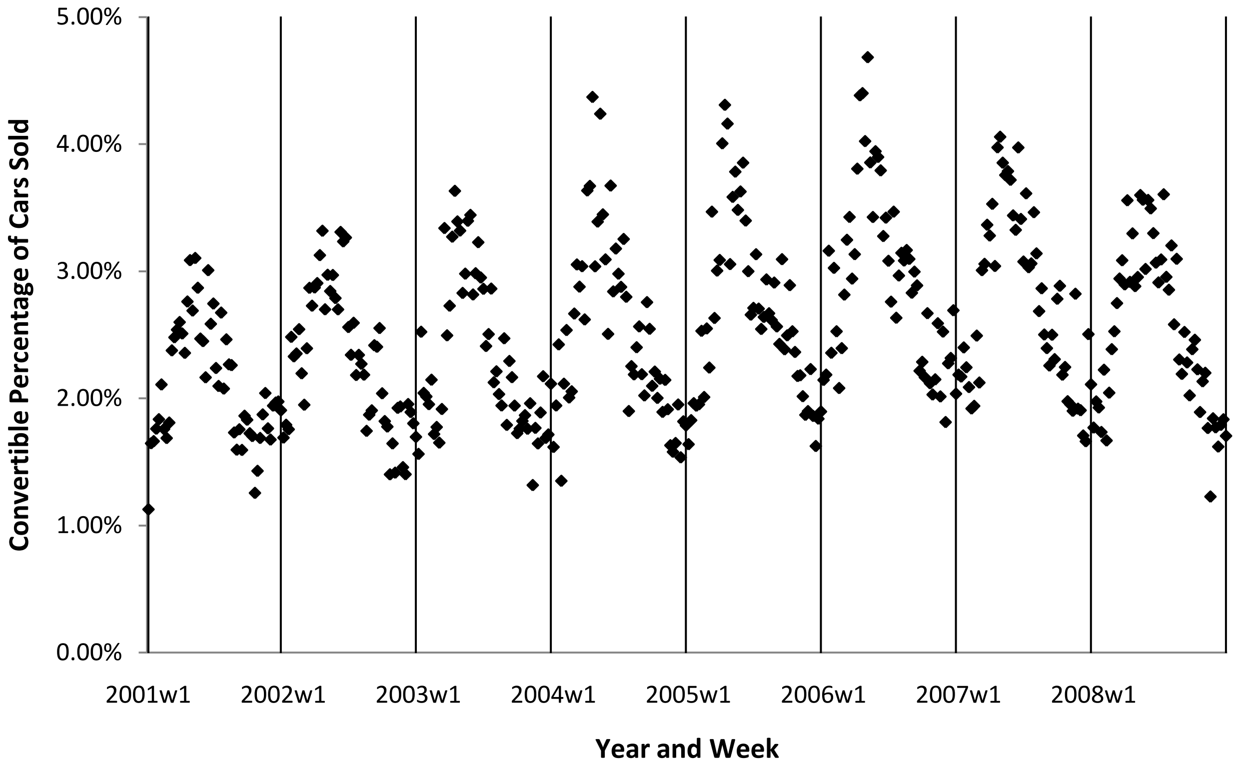

Panel B. Residual Convertible Percentage of Cars Sold - Chicago

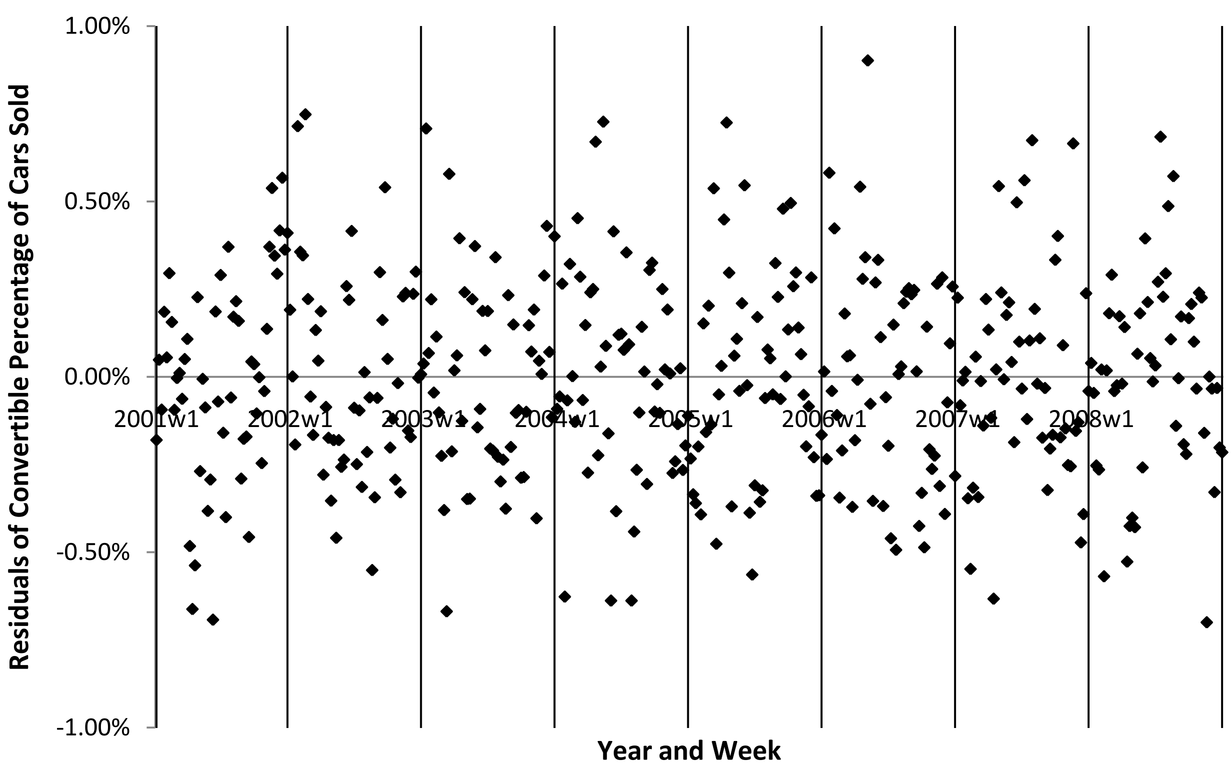

Figure 4. Residual Temperature - Chicago. Panel A illustrates the average daily high temperature in Chicago for each of the 52 weeks in a year. Panel B plots the residual average daily high temperature in each week.

Panel A. Mean High Temperature (Fahrenheit) - Chicago

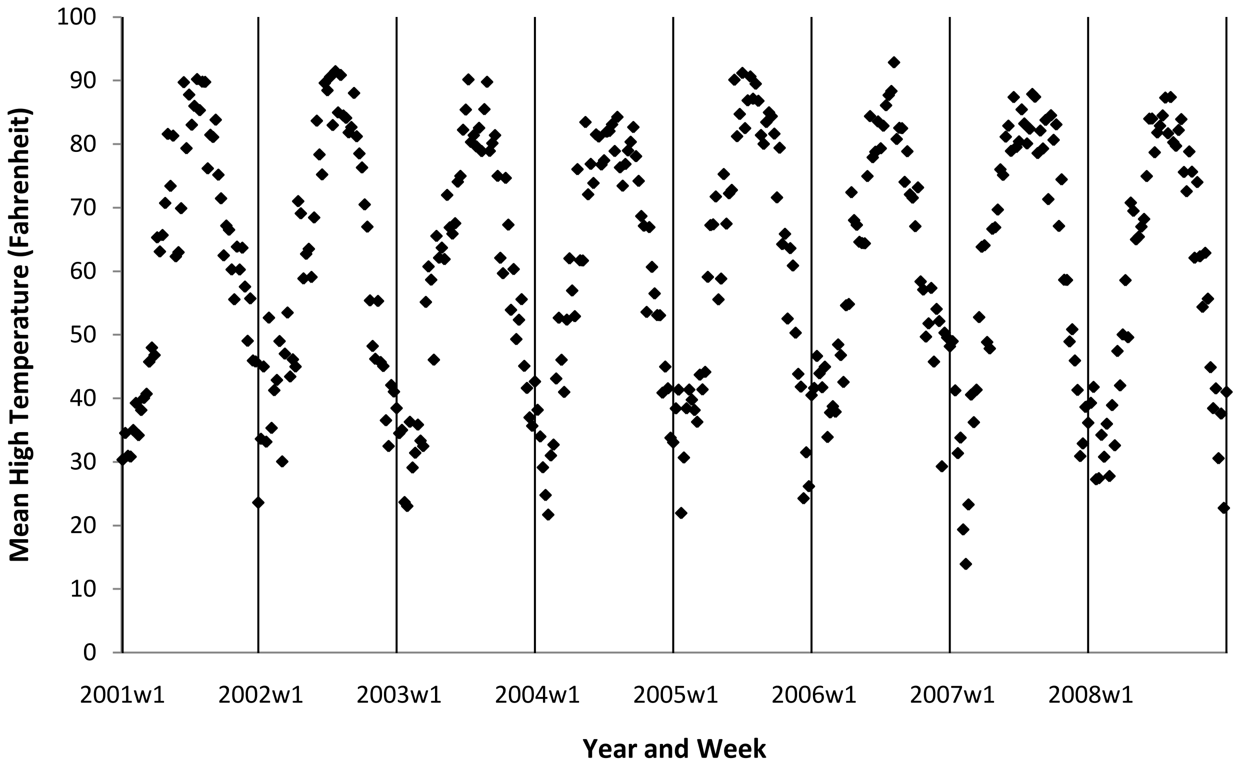

Panel B. Residual Mean High Temperature (Fahrenheit) - Chicago

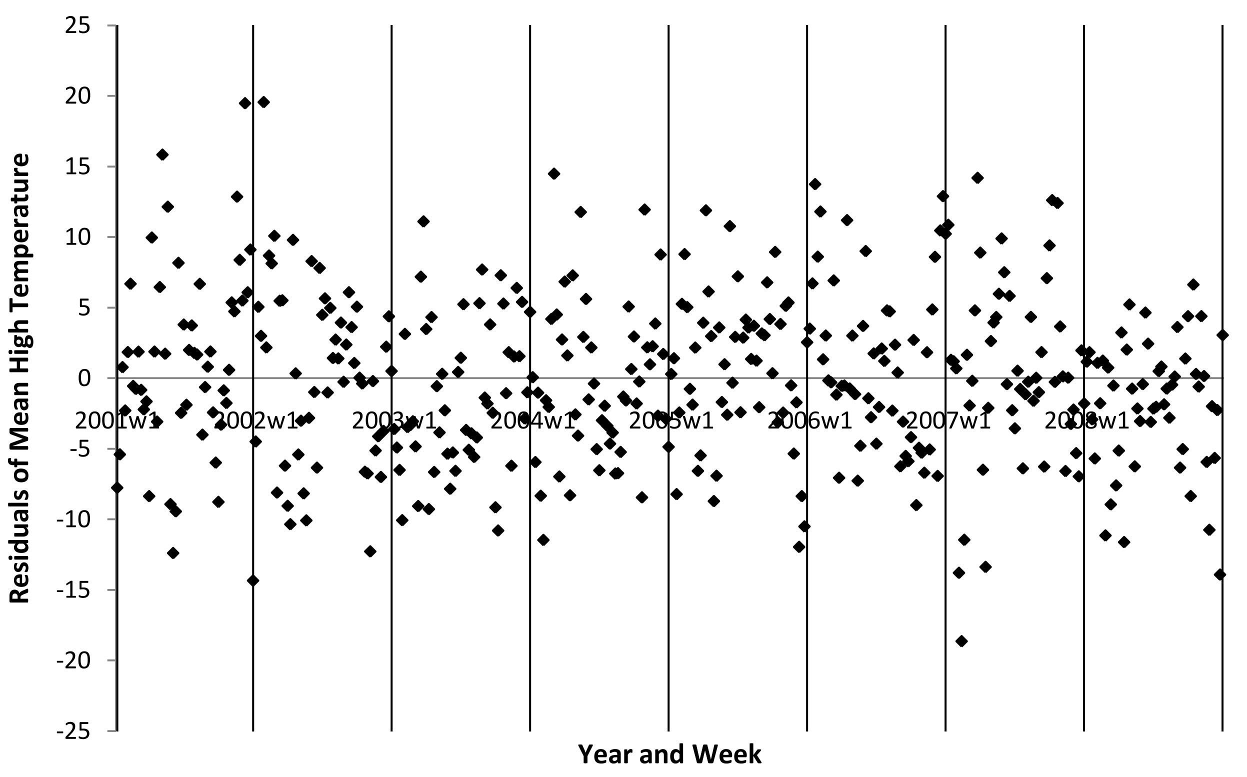

Figure 5. Temperature-Convertible Residuals - Chicago. This Figure provides scatter plots for the residuals of convertible percentage of cars sold (Panel B of Figure 3) and residuals of mean high temperature (Panel B of Figure 4) separately for each quarter of the year.

Panel A. Quarter 1

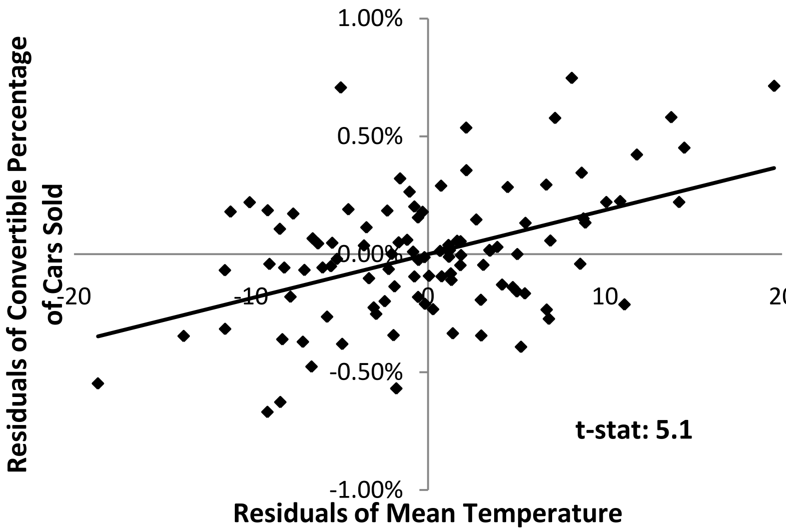

Panel B. Quarter 2

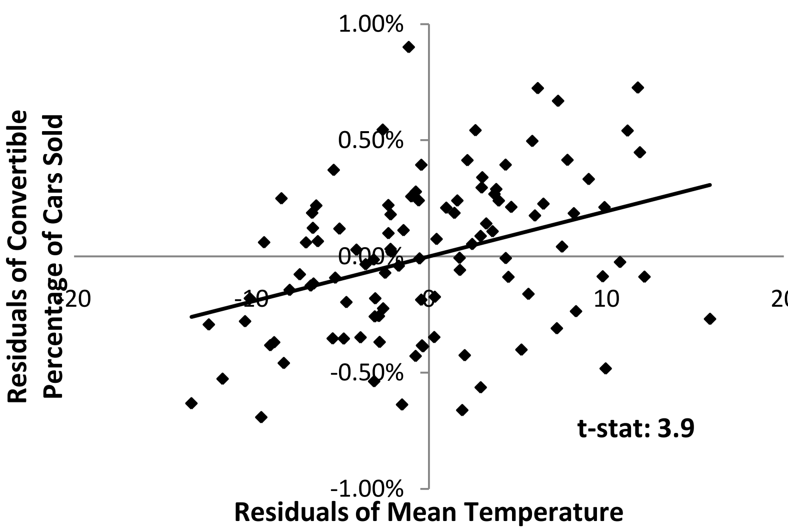

Panel C. Quarter 3

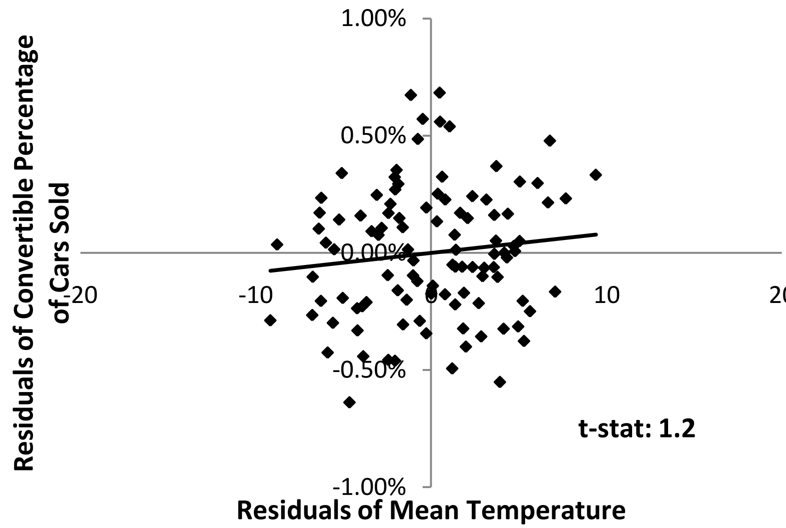

Panel D. Quarter 4

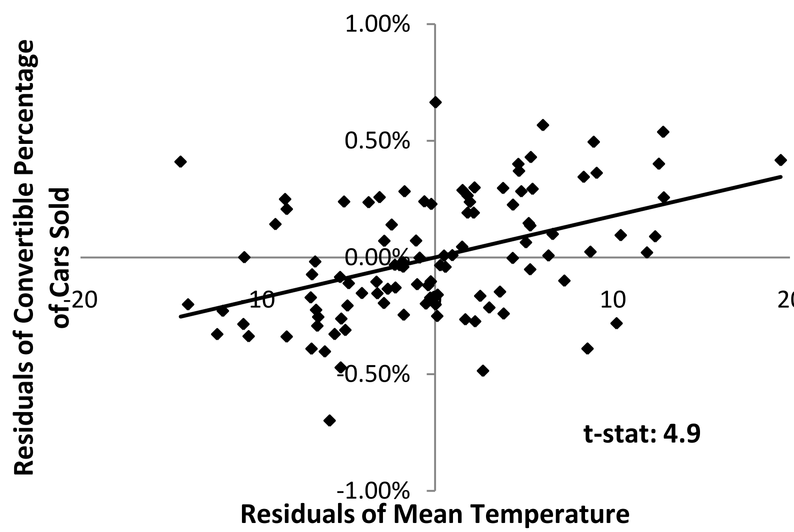

Figure 6. Residual Convertible Sales - Miami. Panel A illustrates the percentage of cars sold in Miami-Ft. Lauderdale for each of the 52 weeks in a year that were convertibles. Panel B plots the residual convertible percentage of cars sold in each week.

Panel A. Convertible Percentage of Cars Sold - Miami

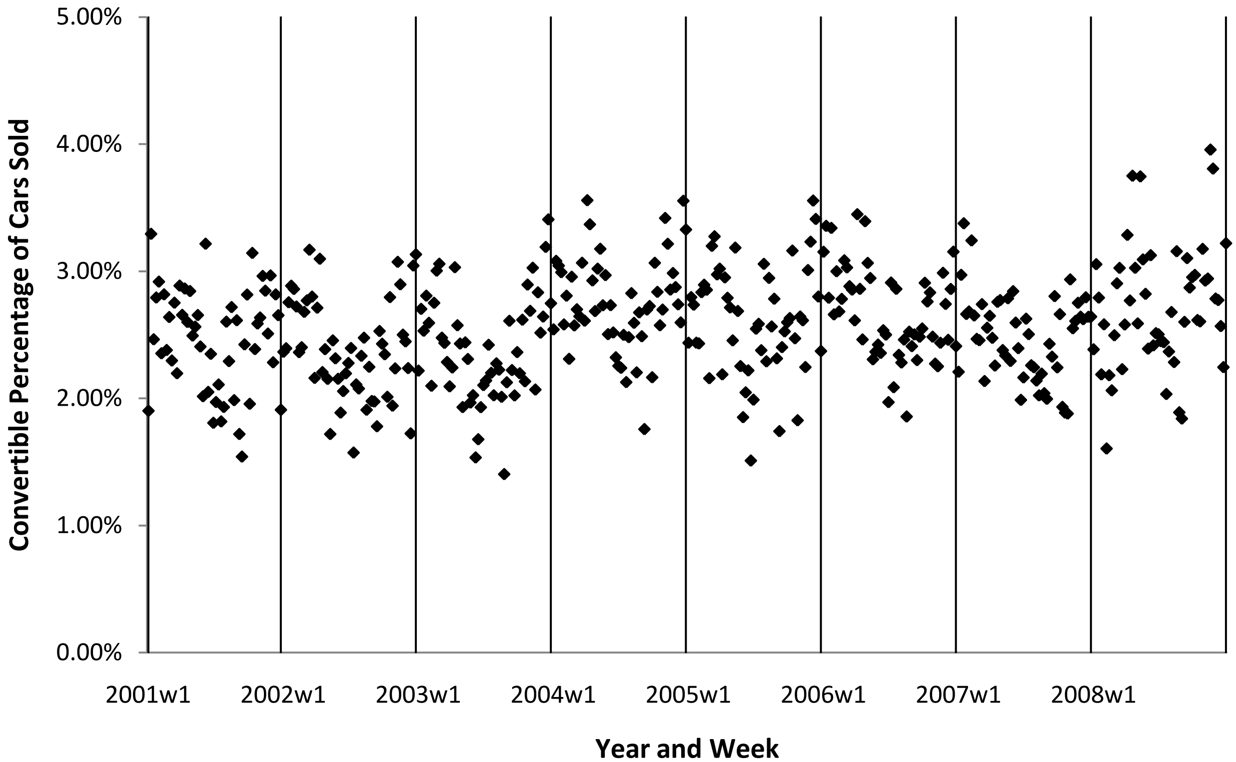

Panel B. Residual Convertible Percentage of Cars Sold - Miami

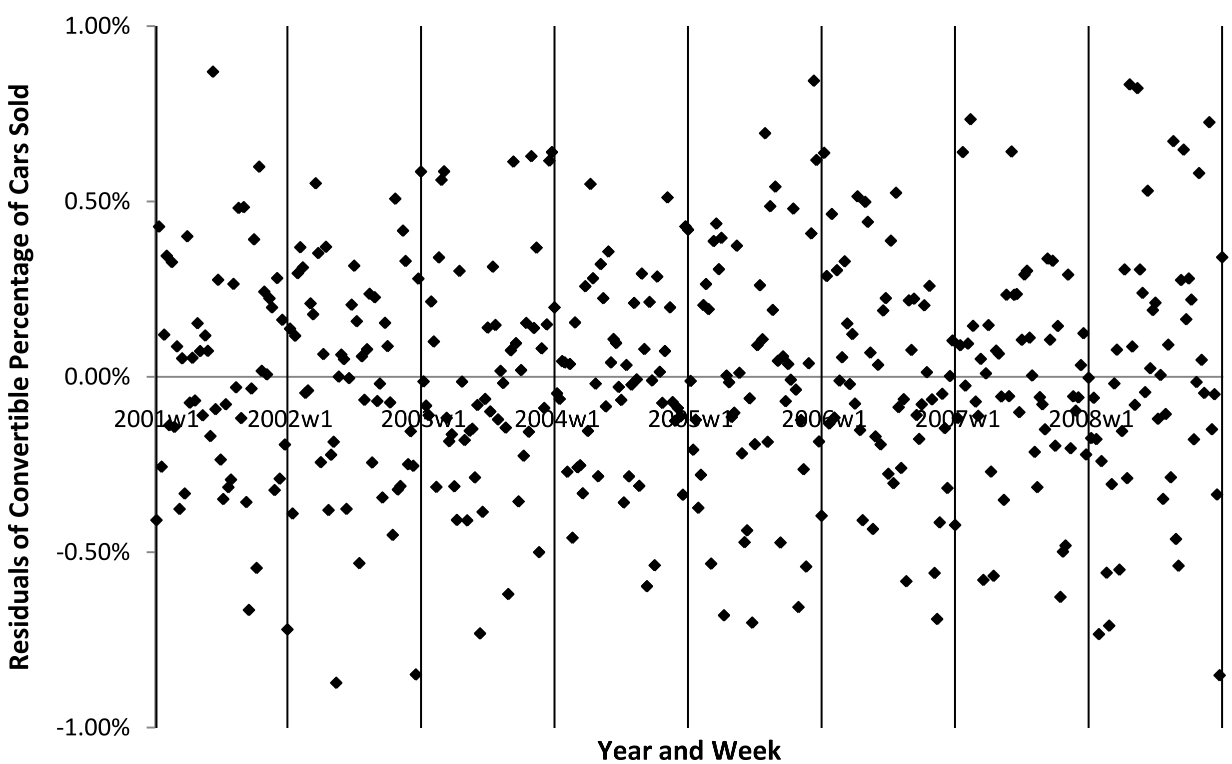

Figure 7. Residual Temperature - Miami. Panel A illustrates the average daily high temperature in Miami-Ft. Lauderdale for each of the 52 weeks in a year. Panel B plots the residual average daily high temperature in each week.

Panel A. Mean High Temperature (Fahrenheit) - Miami

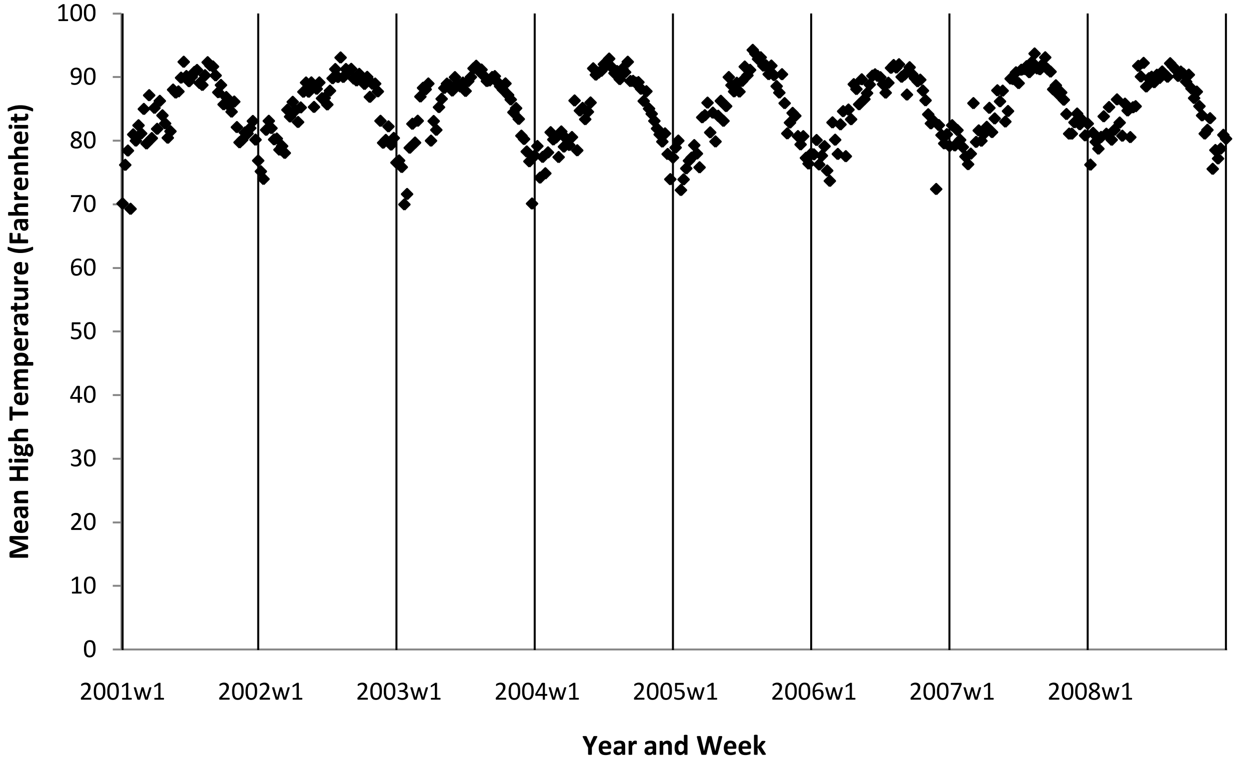

Panel B. Residual Mean High Temperature (Fahrenheit) - Miami

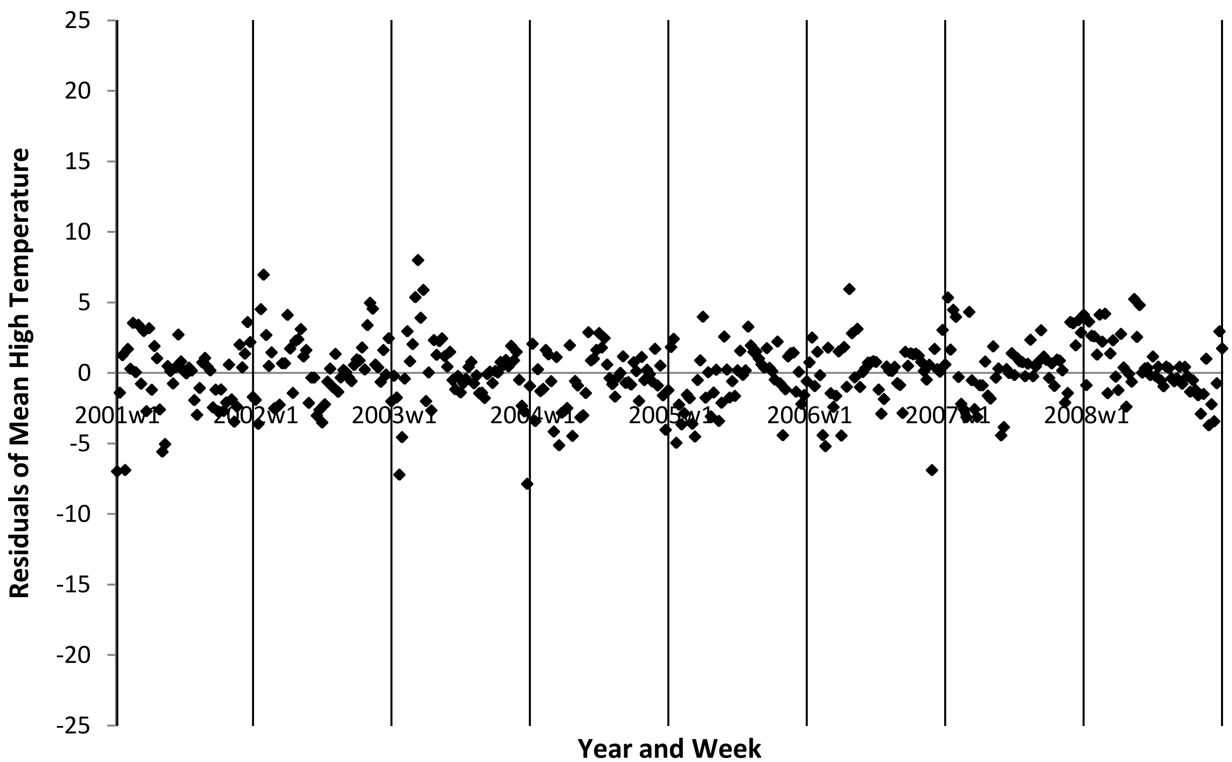

Figure 8. Temperature-Convertible Residuals - Miami. This Figure provides scatter plots for the residuals of convertible percentage of cars sold (Panel B of Figure 6) and residuals of mean temperature (Panel B of Figure 7) separately for each quarter of the year.

Panel A. Quarter 1

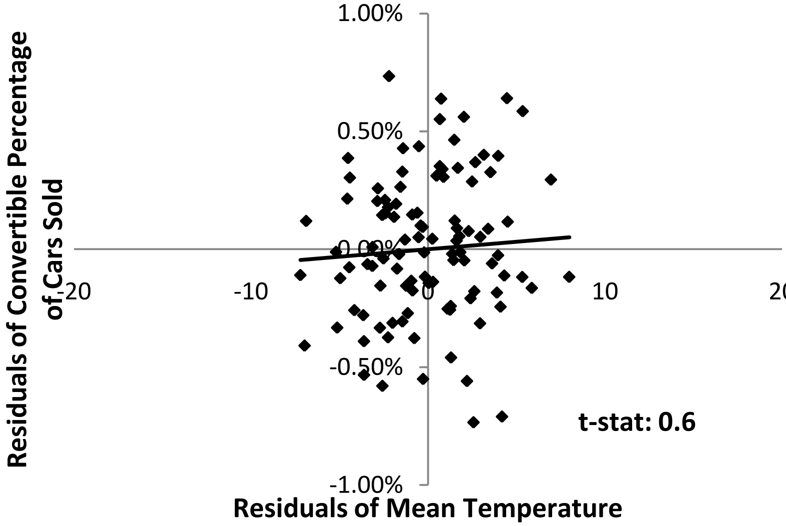

Panel B. Quarter 2

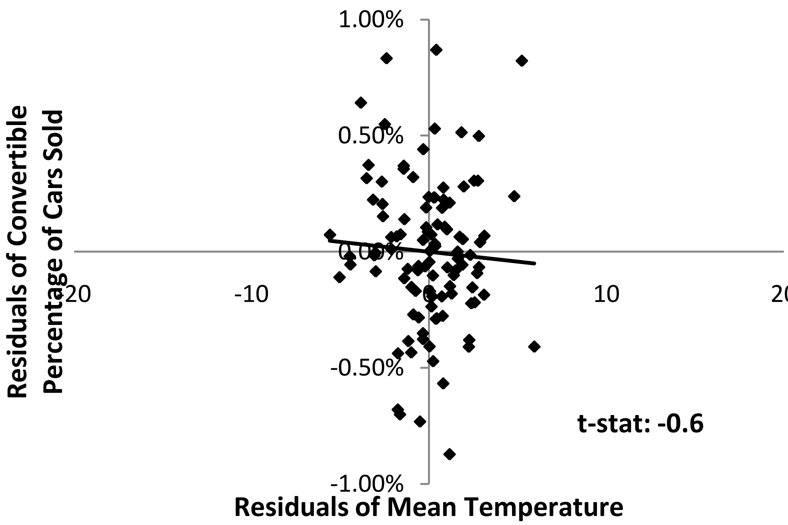

Panel C. Quarter 3

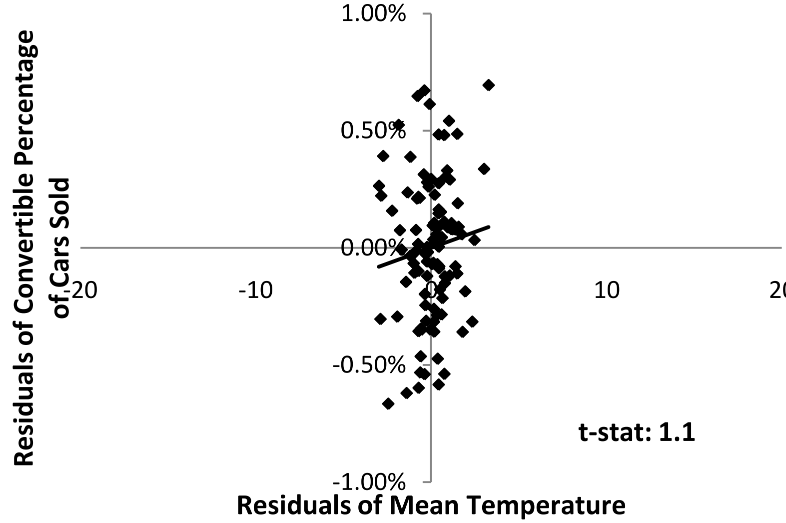

Panel D. Quarter 4

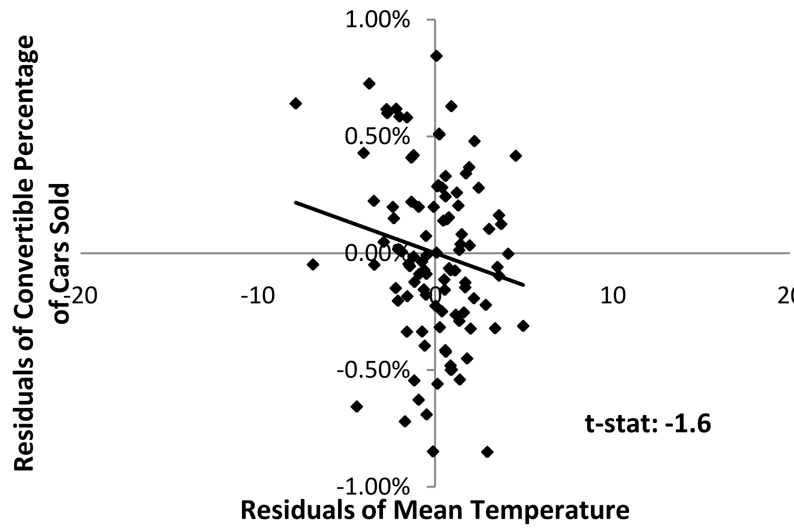

Figure 9. Abnormal Temperature and Convertible Sales by Usual Temperature. This Figure provides the coefficient values and 95% confidence intervals for the impact of residual mean daily high temperature on convertible percentage of total cars sold (the estimate in Column 1 of Table 1) separately by the typical mean daily high temperature for various MSA-weeks. For example, the dot furthest to the left represents the impact of abnormal temperature for MSA-weeks whose high temperature on average across the years in our sample was less than 35 degrees Fahrenheit.

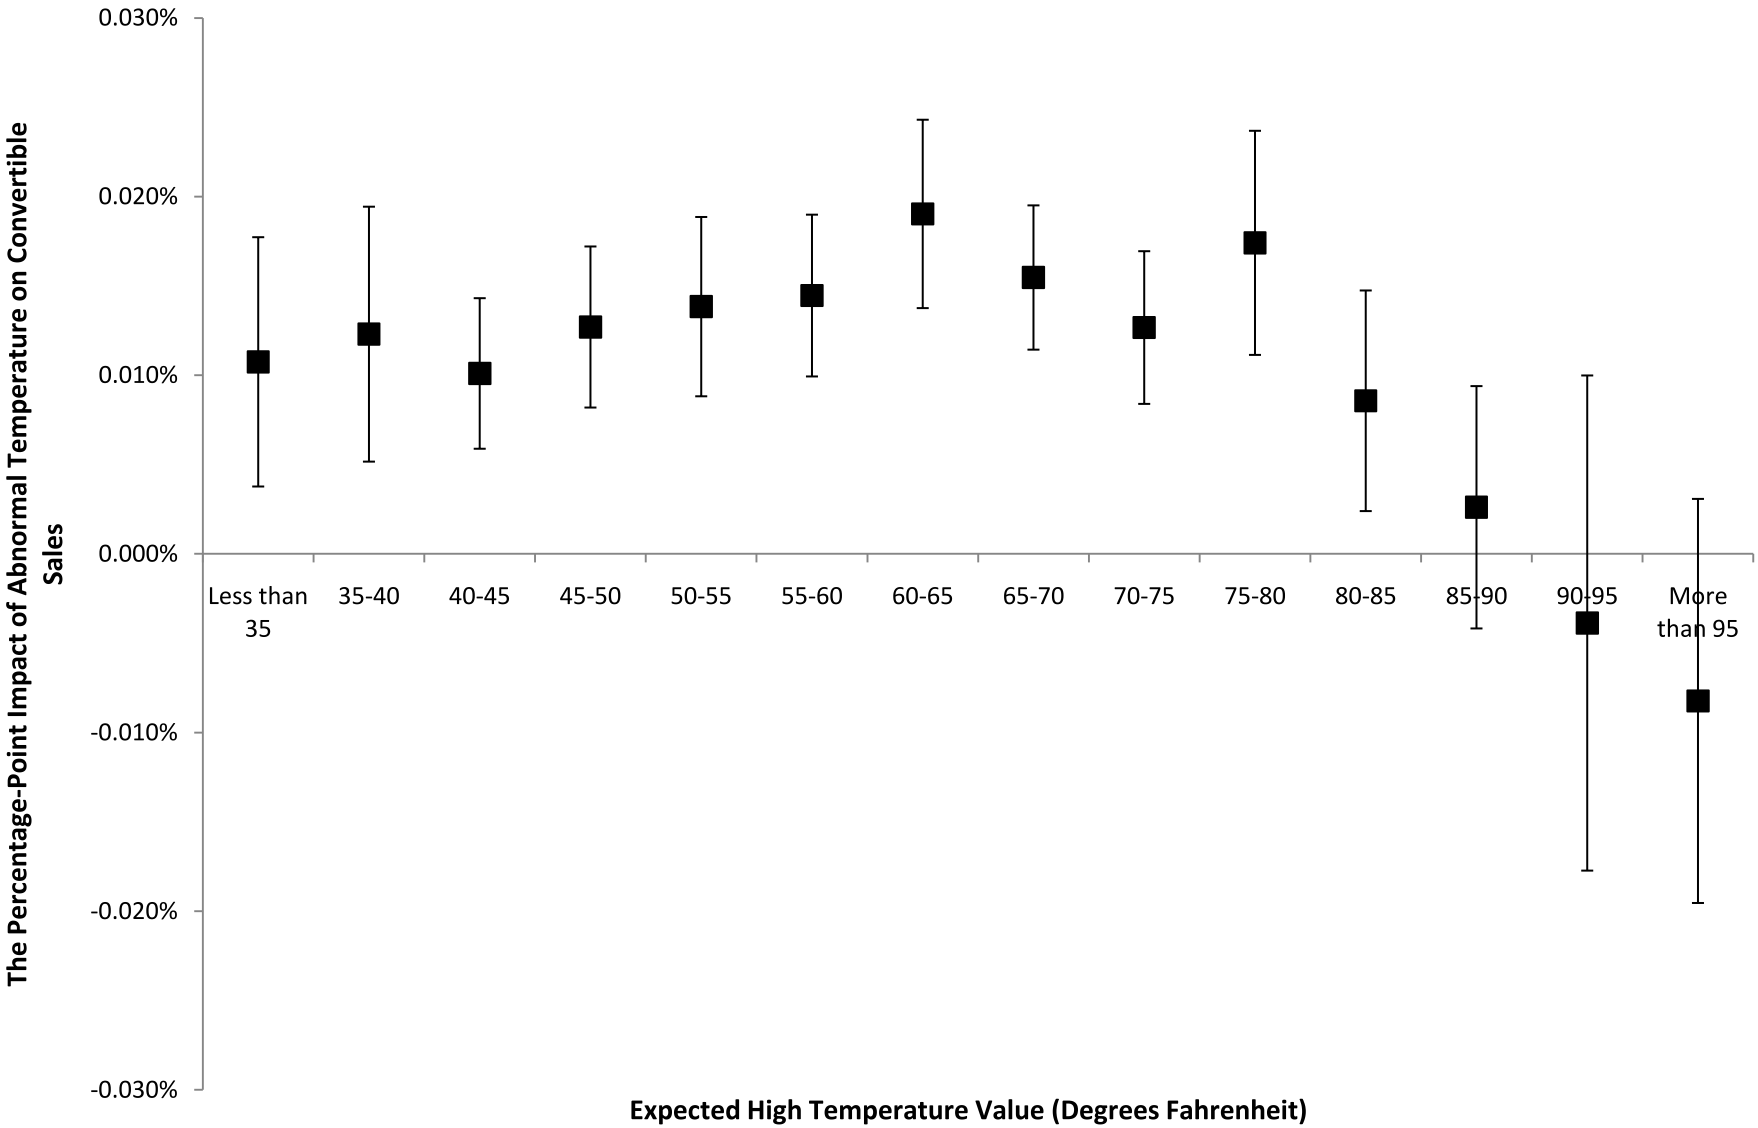

Figure 10. Snowfall and 4-Wheel-Drive Sales - Event-Study Design. This Figure plots the weighted average and 95% confidence intervals for the residuals of the 4-wheel-drive percentage of total cars sold for the twelve weeks leading up to and the twelve weeks after a snow storm event (week 0). The events were chosen to be the largest snow storm week of the year for MSAs that are above median in weather variation.

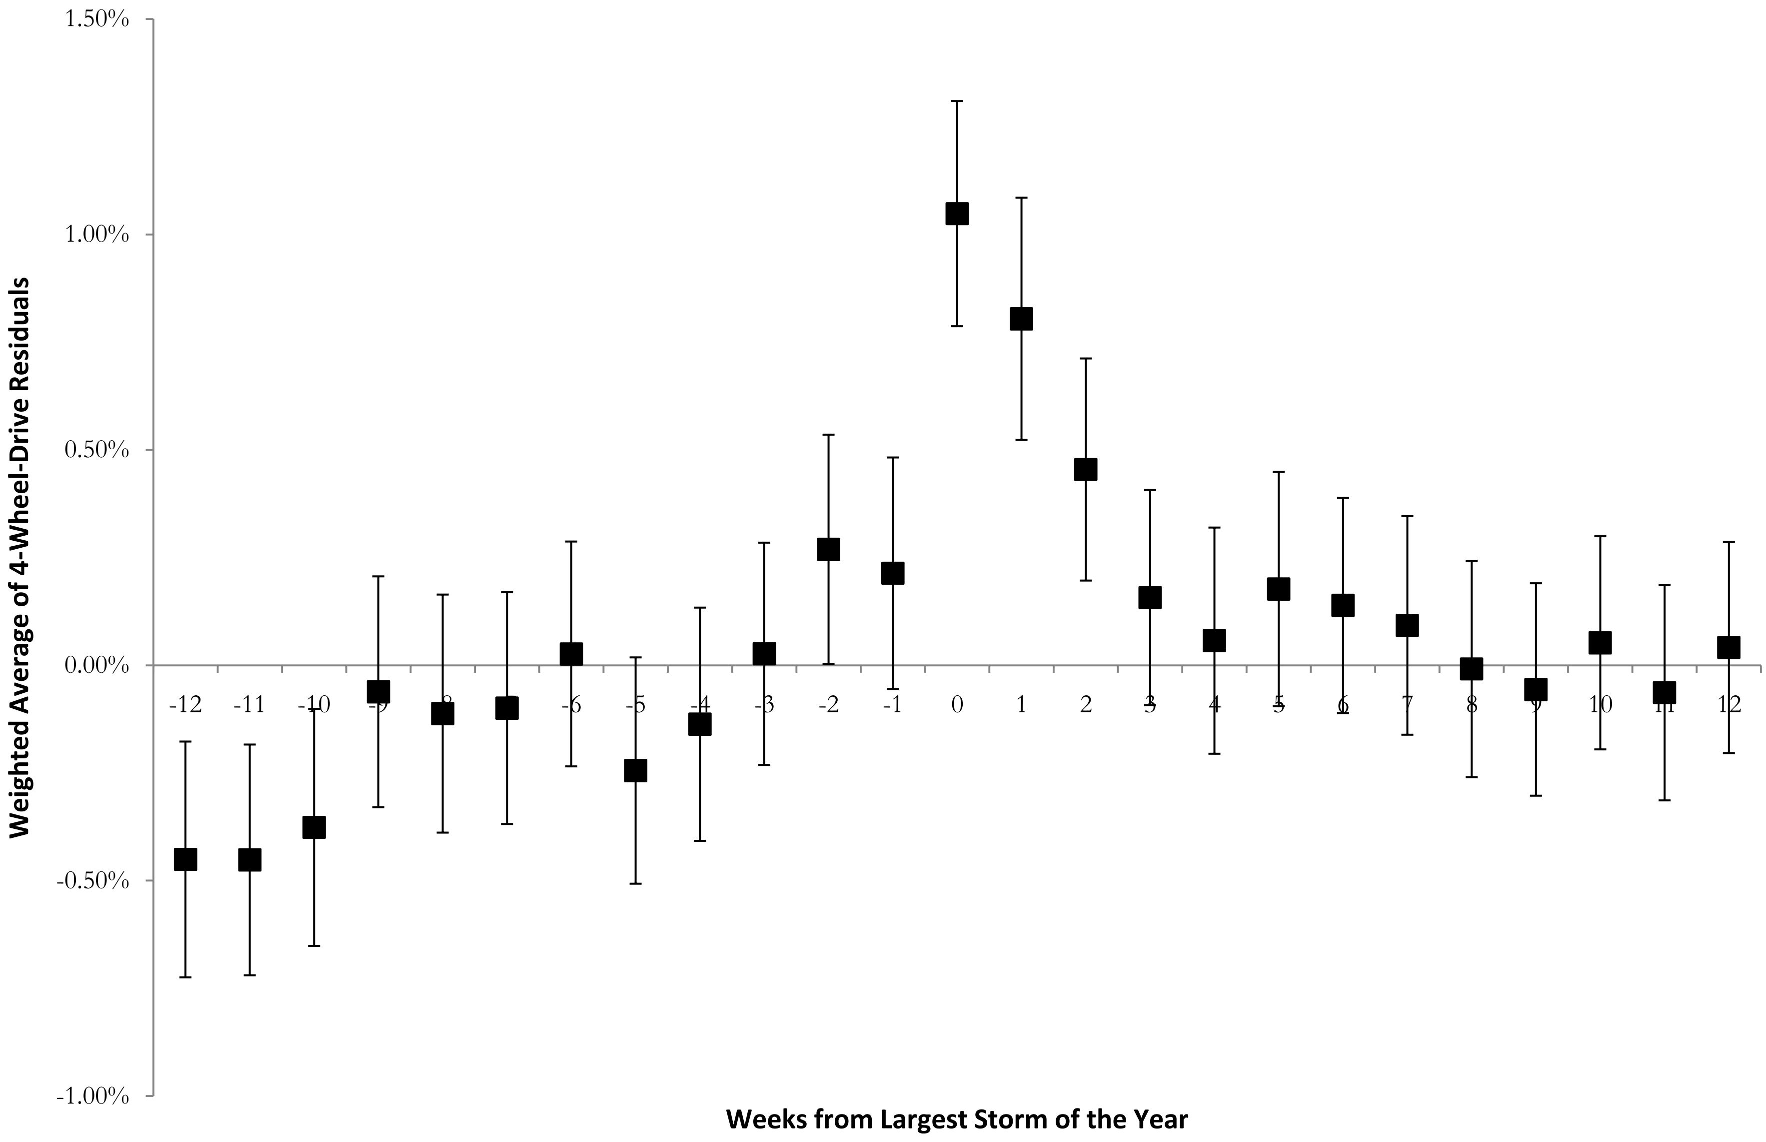

Figure 11 - Seasonal Value of a Swimming Pool. Panel A provides the average residual values for homes with swimming pools that go under contract during each month of the year. Panel B provides these residuals once again after controlling for other housing characteristics - regressing the residuals on a swimming pool dummy and all other housing characteristics for each month of the year. 95% confidence intervals are also presented.

Panel A. Residuals by Month

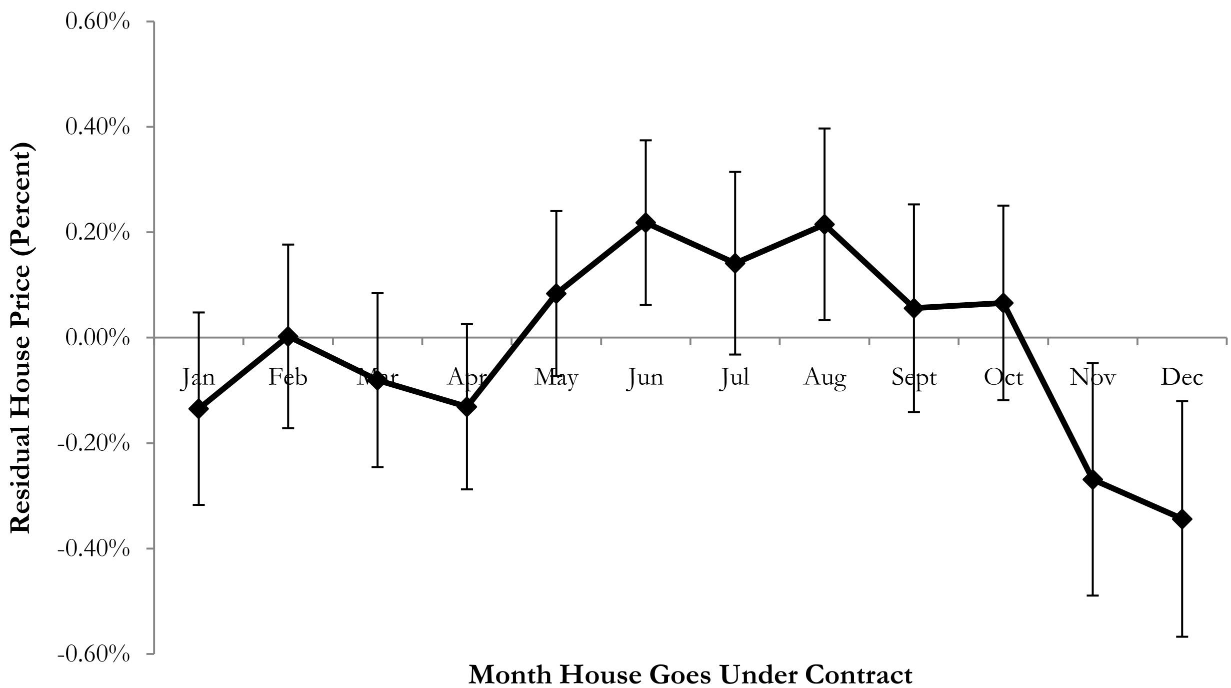

Panel B. Conditional Residuals by Month

Figure 12 - Seasonal Value of a Swimming Pool - Trimming. Panel A provides the average conditional residual values for homes with swimming pools that go under contract during each month of the year after eliminating residuals in the top and bottom 1%. Panel B provides these residuals once again after eliminating residuals in the top and bottom 5%. 95% confidence intervals are also presented.

Panel A. Conditional Residuals by Month - 1% Trim

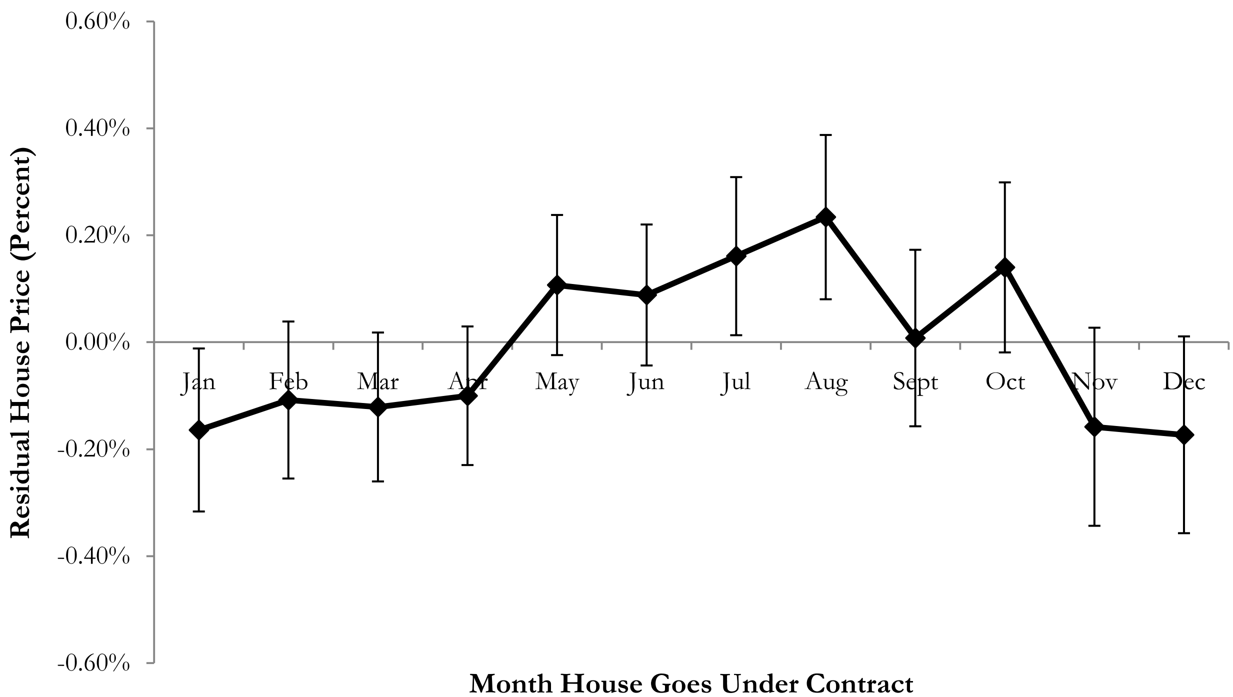

Panel B. Conditional Residuals by Month - 5% Trim

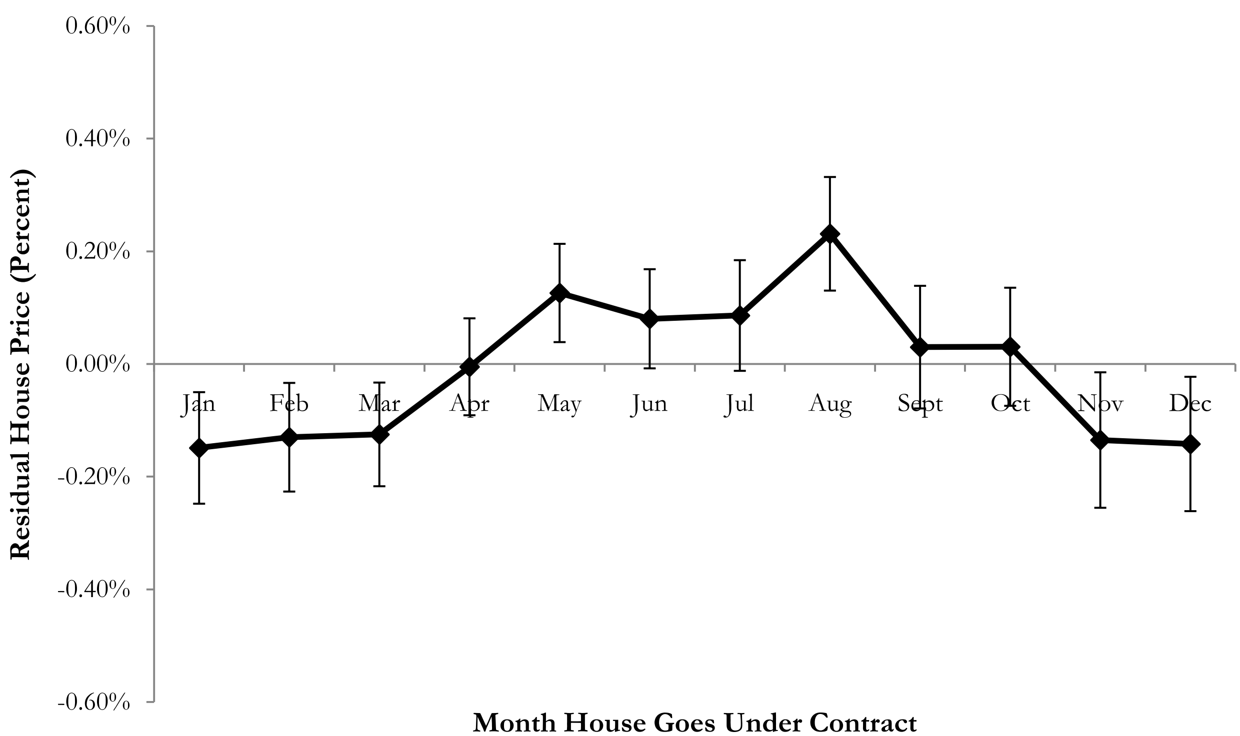

Panel A. Central-Air Residuals by Month

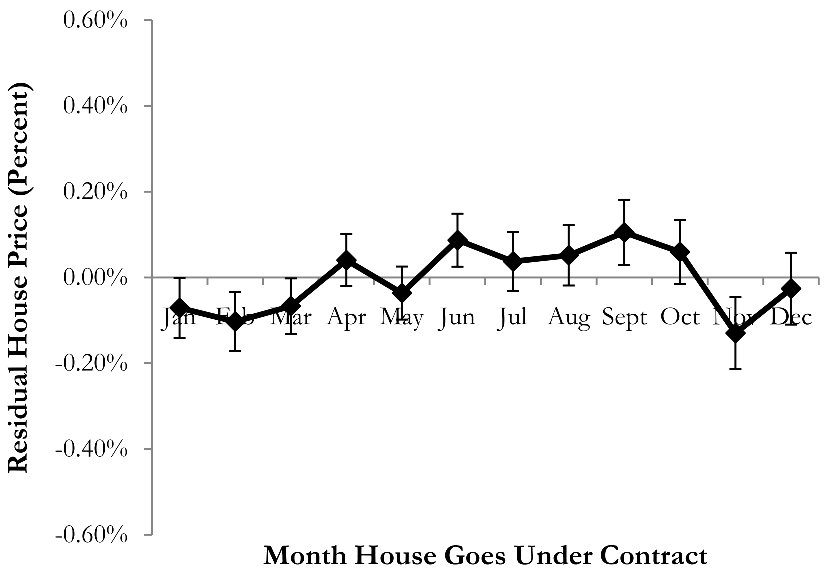

Panel C. Lot-Size Residuals by Month

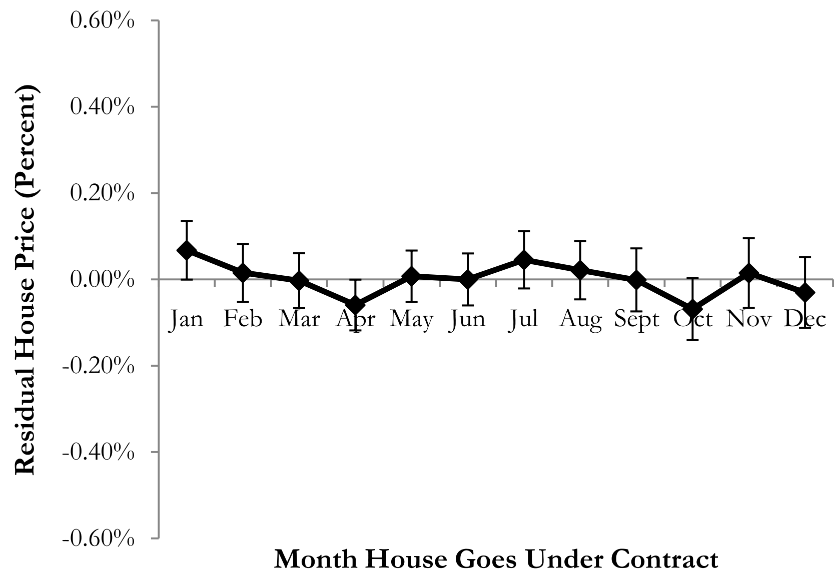

Panel B. Fireplace Residuals by Month

Figure 13 - Seasonal Value of Other Seasonal Housing Characteristics. This figure provides the average conditional residual values for homes with central air (Panel A), a fireplace (Panel B), and lot size in acres (Panel C) that go under contract during each month of the year. The top and bottom 5% of residuals are removed. 95% confidence intervals are also presented.

Panel A. Number-of-Bedroom Residuals by Month

Panel C. Square-Footage Residuals by Month

Panel B. Number-of-Bathroom Residuals by Month

Figure 14 - Seasonal Value of Non-Seasonal Housing Characteristics. This figure provides the average conditional residual values for homes with bedrooms (Panel A), bathrooms (Panel B), and square footage measured in thousands of square feet (Panel C) that go under contract during each month of the year. The top and bottom 5% of residuals are removed. 95% confidence intervals are also presented.

Table 1. Impact of Abnormal Weather on Convertible Purchases

	Dep. Var.: Convertible Percentage of Total Cars Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Residual Temperature	.011**	.014**	.010**	.002	.011**
	(.000)	(.001)	(.002)	(.002)	(.001)
Residual Rain Fall	-.005**	-.017**	-.006	-.000	-.003
	(.002)	(.004)	(.003)	(.003)	(.003)
Residual Snow Fall	-.022	-.006	-.082	-	-.034
	(.024)	(.032)	(.106)	-	(.034)
Residual Slush Fall	-.028**	-.020	-.028	-.026	-.033
	(.009)	(.014)	(.020)	(.026)	(.018)
Residual Cloud Cover	-.172**	-.125*	-.342**	-.171**	-.108*
	(.027)	(.053)	(.057)	(.052)	(.044)
MSA*Year F.E.s	X	X	X	X	X
MSA*Week-of-the-Year F.E.s	X	X	X	X	X
R-Squared	0.778	0.837	0.780	0.813	0.860
Observations	49,499	11,637	13,123	12,798	11,941

Notes: Coefficient values and standard errors are presented from OLS regressions of the convertible percentage of total cars sold on residual weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is an MSA*Week and is weighted by the total number of cars sold. Fixed effects are included for MSA*Year and for MSA*Week-of-the-Year (Week 1 - Week 52). The first column uses all the data while the next four columns present results separately for the four quarters of the year.
* significant at 5%; ** significant at 1%

Table 2. Impact of Abnormal Weather on 4-Wheel-Drive Purchases

	Dep. Var.: 4-Wheel-Drive Percentage of Total Cars Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Residual Temperature	-.050**	-.069**	-.024**	-.039**	-.063**
	(.002)	(.003)	(.004)	(.007)	(.004)
Residual Rain Fall	.014*	.021	.029**	.031**	-.003
	(.006)	(.014)	(.010)	(.010)	(.012)
Residual Snow Fall	1.02**	.73**	-.18	-8.11	1.18**
	(.05)	(.07)	(.20)	(25.3)	(.08)
Residual Slush Fall	.24**	.24**	.12*	-.14	.45**
	(.02)	(.04)	(.04)	(.09)	(.05)
Residual Cloud Cover	.378**	.351*	1.030**	.265	.405*
	(.082)	(.155)	(.158)	(.186)	(.150)
MSA*Year F.E.s	X	X	X	X	X
MSA*Week-of-the-Year F.E.s	X	X	X	X	X
R-Squared	0.964	0.972	0.971	0.970	0.972
Observations	68,431	16,517	17,101	17,320	17,493

Notes: Coefficient values and standard errors are presented from OLS regressions of the 4-wheel-drive percentage of total cars sold on residual weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is an MSA*Week and is weighted by the total number of cars sold. Fixed effects are included for MSA*Year and for MSA*Week-of-the-Year (Week 1 - Week 52). The first column uses all the data while the next four columns present results separately for the four quarters of the year.

* significant at 5%; ** significant at 1%

Table 3. Impact of Abnormal Weather on Convertible Purchases - Dynamic Analysis

	Dep. Var.: Convertible Percentage of Total Cars Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Residual Temperature Lead 1	.001	.003	.001	.005*	-.000
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature	.011**	.015**	.006**	.004	.010**
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 1	.001	.005**	-.005*	.001	.003
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 2	.003	.007**	-.001	.003	.000
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 3	.001	.002	-.001	.007**	-.001
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 4	.001	.002	-.003	.005	.001
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 5	-.001	.000	-.003	.003	-.001
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 6	.002*	-.001	.001	.010**	.004
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 7	.002*	.004**	.001	.002	-.003
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 8	.004**	.000	.006**	.008**	.003
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 9	.003**	.002	.000	.004	.004
	(.001)	(.002)	(.002)	(.002)	(.002)
Residual Temperature Lag 10	-.000	.001	-.002	.004*	-.003
	(.001)	(.002)	(.001)	(.002)	(.003)
Residual Temperature Lag 11	.000	-.002	-.000	.004	.007*
	(.001)	(.002)	(.001)	(.002)	(.003)
Residual Temperature Lag 12	.000	.001	-.004**	.005*	.000
	(.001)	(.002)	(.001)	(.002)	(.003)
MSA*Year F.E.s	X	X	X	X	X
MSA*Week-of-the-Year F.E.s	X	X	X	X	X
Residual Rain Fall (with Lead and Lags)	X	X	X	X	X
Residual Snow Fall (with Lead and Lags)	X	X	X	X	X
Residual Slush Fall (with Lead and Lags)	X	X	X	X	X
Residual Cloud Cover (with Lead and Lags)	X	X	X	X	X
R-Squared	0.809	0.875	0.790	0.791	0.873
Observations	36,873	8,068	9,696	9,908	9,201

Notes: Coefficient values and standard errors are presented from OLS regressions of the convertible percentage of total cars sold on residual weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Both the current residual weather as well as the lead and 12 lag residual weather variables are included in each regression. The coefficient values for rain, snow, slush, and cloud cover are omitted due to space constraints. Each observation is an MSA*Week and is weighted by the total number of cars sold. Fixed effects are included for MSA*Year and for MSA*Week-of-the-Year (Week 1 - Week 52). The first column uses all the data while the next four columns present results separately for the four quarters of the year.

* significant at 5%; ** significant at 1%

Table 4. Impact of Abnormal Weather on 4-Wheel-Drive Purchases - Dynamic Analysis

	Dep. Var.: 4-Wheel-Drive Percentage of Total Cars Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Residual Snow Fall Lead 1	.18**	-.01	-.90	3.22**	.27**
	(.06)	(.06)	(.29)	(1.2)	(.09)
Residual Snow Fall	1.01**	.71**	-.19	22.9	1.22**
	(.06)	(.09)	(.24)	(81.3)	(.10)
Residual Snow Fall Lag 1	.85**	.60**	.04	21.6	1.29**
	(.06)	(.08)	(.23)	(121.5)	(.11)
Residual Snow Fall Lag 2	.26**	.19*	.37	-247.8	.74**
	(.06)	(.07)	(.21)	(193.4)	(.13)
Residual Snow Fall Lag 3	.12	.26**	-.04	147.4	.40**
	(.06)	(.08)	(.15)	(432.5)	(.14)
Residual Snow Fall Lag 4	-.09	.01	-.39**	-250.9*	.26
	(.07)	(.08)	(.14)	(117.2)	(.17)
Residual Snow Fall Lag 5	-.14*	-.07	-.35**	-156.3	.05
	(.07)	(.08)	(.12)	(116.2)	(.22)
Residual Snow Fall Lag 6	-.26**	-.17	-.31**	-320.3**	-.07
	(.06)	(.08)	(.10)	(105.6)	(.39)
Residual Snow Fall Lag 7	-.15*	-.15	-.05	-71.1*	-.59
	(.06)	(.08)	(.09)	(34.8)	(.52)
Residual Snow Fall Lag 8	-.09	-.02	-.12	-39.4	1.58*
	(.06)	(.09)	(.09)	(32.6)	(.65)
Residual Snow Fall Lag 9	-.09	-.05	-.11	-1.4	1.65*
	(.06)	(.09)	(.08)	(1.9)	(.74)
Residual Snow Fall Lag 10	-.09	-.18	-.12	1.4	1.23
	(.06)	(.09)	(.08)	(.87)	(.77)
Residual Snow Fall Lag 11	-.09	.05	-.25**	.01	-.33
	(.06)	(.10)	(.07)	(.75)	(.76)
Residual Snow Fall Lag 12	-.13*	-.01	-.30**	.41	-1.68
	(.06)	(.09)	(.08)	(.29)	(1.2)
MSA*Year F.E.s	X	X	X	X	X
MSA*Week-of-the-Year F.E.s	X	X	X	X	X
Residual Temperature (with Lead and	X	X	X	X	X
Residual Rain Fall (with Lead and Lags)	X	X	X	X	X
Residual Slush Fall (with Lead and Lags)	X	X	X	X	X
Residual Cloud Cover (with Lead and	X	X	X	X	X
R-Squared	0.970	0.979	0.975	0.973	0.977
Observations	46,452	10,258	11,433	12,356	12,405

Notes: Coefficient values and standard errors are presented from OLS regressions of the 4-wheel-drive percentage of total cars sold on residual weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Both the current residual weather as well as the lead and 12 lag residual weather variables are included in each regression. The coefficient values for temperature, rain, slush, and cloud cover are omitted due to space constraints. Each observation is an MSA*Week and is weighted by the total number of cars sold. Fixed effects are included for MSA*Year and for MSA*Week-of-the-Year (Week 1 - Week 52). The first column uses all the data while the next four columns present results separately for the four quarters of the year.

* significant at 5%; ** significant at 1%

Table 5. Impact of Abnormal Weather on Convertible and 4-Wheel-Drive Purchases for Consumers Trading in a Convertible or 4-Wheel-Drive Vehicle, Respectively

	Dep. Var.: Convertible or 4-Wheel Drive Percentage of Total Cars Sold	
	Convertibles	**4-Wheel Drives**
Residual Temperature	.060**	-.044**
	(.020)	(.004)
Residual Rain Fall	.007	.003
	(.042)	(.011)
Residual Snow Fall	-.23	.61**
	(.60)	(.09)
Residual Slush Fall	-.45	.19**
	(.24)	(.04)
Residual Cloud Cover	-1.26	.89**
	(.679)	(.15)
MSA*Year F.E.s	X	X
MSA*Week-of-the-Year F.E.s	X	X
R-Squared	0.675	0.815
Observations	23,529	65,356

Notes: Coefficient values and standard errors are presented from OLS regressions of the convertible percentage of total cars sold (Column 1) and the 4-wheel-drive percentage of total cars sold (Column 2) on residual weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is an MSA*Week and is weighted by the total number of cars sold. Fixed effects are included for MSA*Year and for MSA*Week-of-the-Year (Week 1 - Week 52). The sample is restricted to people who were purchasing a vehicle while trading in a convertible (Column 1) or a 4-wheel drive (Column 2).
* significant at 5%; ** significant at 1%

Table 6. Impact of Abnormal Weather on Quickly Trading In a Vehicle

	Dep. Var.: Dummy Variable if Returned Within 1-3 Years		
	1 Year	**2 Years**	**3 Years**
Mean of Dependent Variable	2.37%	5.03%	7.16%
Convertible	1.272%**	2.302%**	2.905%**
	(.019%)	(.030%)	(.042%)
Convertible Interacted with:			
Residual Temperature	.006%	.017%**	.006%
	(.004%)	(.007%)	(.009%)
Residual Rain Fall	.008%	.002%	-.018%
	(.009%)	(.015%)	(.021%)
Residual Snow Fall	.181%	-.041%	-.142%
	(.131%)	(.222%)	(.289%)
Residual Slush Fall	.063%	.028%	-.116%
	(.053%)	(.094%)	(.131%)
Residual Cloud Cover	-.197%	-.036%	.332%
	(.138%)	(.228%)	(.312%)
4-Wheel Drive	.285%**	.929%**	1.634%**
	(.006%)	(.006%)	(.014%)
4-Wheel Drive Interacted With:			
Residual Temperature	-.003%*	-.005%*	-.013%**
	(.001%)	(.002%)	(.003%)
Residual Rain Fall	-.005%	-.005%	.001%
	(.003%)	(.006%)	(.008%)
Residual Snow Fall	.000%	.063%	.004%
	(.035%)	(.058%)	(.076%)
Residual Slush Fall	.002%	-.019%	-.048%
	(.016%)	(.028%)	(.038%)
Residual Cloud Cover	.006%	-.109%	-.124%
	(.047%)	(.078%)	(.106%)
Year*Week*MSA Fixed Effects	X	X	X
R-Squared	0.004	0.006	0.007
Observations	35,102,062	29,665,047	23,827,418

Notes: Coefficient values and standard errors are presented from OLS regressions of a dummy variable for whether the vehicle shows up in our dataset (as a trade-in car or another car sale) within 1, 2, or 3 years from the date of purchase on a convertible and 4-wheel drive dummy variable and an interaction between these car types and residual weather variables at the time of purchase - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is at the car level and MSA*Year*Week-of-the-Year fixed effects are included. The dataset is also restricted so as to eliminate all truncation (Columns 1-3 eliminate the last 1-3 years of car sales in the sample, respectively).
* significant at 5%; ** significant at 1%

Table 7. Housing Summary Statistics

	Mean	Standard Deviation	Minimum	Maximum
Sales Price	273,925	263,681	5,001	5,000,000
Swimming Pool	0.119	0.321	0	1
Central Air	0.304	0.460	0	1
Fireplace	0.455	0.337	0	1
Lot Size (Acres)	0.320	0.487	0	5
Year Built	1968	24	1900	2006
Square Footage	1679	734	250	10000
Bathrooms	2.06	0.84	0.5	10
Bedrooms	3.12	0.81	1	10
Observations	4,206,314	4,206,314	4,206,314	4,206,314

Table 8. The Impact of Temperature and Housing Characteristics on Residual Sales Prices

	Dependent Variable: Residual Housing Prices							
	Linear Temperature		Temperatre > 70° F		Temperatre > 80° F		Temperatre > 90° F	
Interaction of Temperature and:								
Swimming Pool	0.013%**	0.010%**	0.23%**	0.18%**	0.27%**	0.13%**	0.41%	0.37%**
	(.004%)	(.002%)	0.06%	0.03%	0.09%	0.04%	0.30%	0.14%
Fire Place	0.0006%	0.0013%	-0.01%	0.02%	0.06%	0.02%	0.18%	-0.12%
	(.0024%)	(.0011%)	0.05%	0.02%	0.07%	0.04%	0.24%	0.11%
Lot Acre	0.0006%	0.0006%	-0.06%	-0.05%*	0.00%	-0.02%	-0.68%*	-0.28%
	(.0024%)	(.0012%)	0.05%	0.02%	0.09%	0.04%	0.29%	0.15%
Central Air	0.0002%	0.0060%**	0.02%	0.10%**	0.02%	0.03%	-0.13%	-0.23%
	(.0025%)	(.0012%)	0.05%	0.02%	0.07%	0.04%	0.25%	0.12%
Square Footage (1,000s)	0.0043%	0.0004%	0.10%*	0.04%	0.09%	-0.03%	0.42%	0.22%
	(.0025%)	(.0012%)	0.05%	0.02%	0.08%	0.04%	0.27%	0.13%
Number of Baths	-0.0034%	-0.0008%	-0.03%	-0.03%	-0.05%	0.03%	0.07%	0.09%
	(.0020%)	(.0010%)	0.04%	0.02%	0.07%	0.03%	0.25%	0.12%
Number of Bedrooms	0.0004%	-0.0009%	0.01%	-0.02%	0.02%	0.00%	-0.34%	-0.09%
	(.0018%)	(.0009%)	0.03%	0.02%	0.06%	0.03%	0.17%	0.08%
Levels of All Variables	X	X	X	X	X	X	X	X
Trim 5%		X		X		X		X
Observations	4,145,410	3,731,014	4,145,410	3,731,014	4,145,410	3,731,014	4,145,410	3,731,014

Notes: The first two columns of this table present coefficients and standard errors from the regression of residual housing prices (from Equation (4) in the text) on the interaction between housing characteristics and linear temperature (average high daily temperature during the month the house goes under contract). The next three sets of columns report the interaction between housing characteristics and dummy variables for the average daily high temperature in the month of housing contract being above 70, 80, or 90 degrees Fahrenheit. The second column in each set restricts the sample to house sales whose residuals were not in the top or bottom 5%. The level effects of all variables (not just the interactions) are also included in all of the regressions.

* significant at 5%; ** significant at 1%