PROJECTION BIAS IN THE CAR AND HOUSING MARKETS

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

June 2012

Abstract

Projection bias is the tendency to overpredict the degree to which one's future tastes will resemble one's 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 vehicle 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, a 4-wheel drive, or a vehicle that is black in color 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 and that central air add to a house is higher 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, Loren Pope, Mathew Rabin, Dick Thaler, and seminar participants at the Behavioral Economics Annual Meeting, Columbia University, Cornell University, Harvard Business School, MIT, Stanford University, UC Berkeley, UC Los Angeles, UC San Diego, UC Santa Barbara, University of Chicago, University of Illinois, University of Zurich, and Yale University for helpful suggestions.

Many decisions that people make require them to predict their future utility. For example, choosing a job, deciding where to live, planning a vacation, buying a car, 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 an individual's prediction of future utility will, on average, match his or her 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 individuals 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 a caution against projection bias: 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 such as shopping while hungry, an open question is whether this bias influences important life decisions for which people are likely to have strong motivations to make a good decision.

In this paper, we test for projection bias in two high-stakes environments: the purchases of vehicles and houses. Vehicles and houses are durable goods. When consumers purchase durable goods, they must predict at the time of purchase how much they will value consuming these goods in the future, including the enjoyment they will experience in a variety of future states of the world. Projection bias suggests that consumers may mistakenly purchase a vehicle or a house that has a high utility at the time of purchase, but whose utility will not be as high in other states of the world that the consumer will experience while owning the vehicle or house. We test the extent to which weather variation at the time of purchase can cause consumers to overweigh the value that they place on certain vehicle and housing characteristics. Projection bias predicts that consumers will overvalue warm-weather vehicle types and housing characteristics (e.g. convertibles and swimming pools) when the weather is warm at the time of purchase and overvalue cold-weather vehicle 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 transaction-level data for more than forty million transactions of new and used vehicles from dealerships around the U.S. We find that the sales of convertibles, 4-wheel drives, and vehicles that are black in color are highly influenced by idiosyncratic variation in temperature, cloud cover, and snowfall. We show that for convertibles, weather that is warmer and skies that are clearer than seasonal averages lead to a higher number of sales. Controlling for seasonal sales patterns, our estimates suggest that a location 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 vehicles sold that are convertibles. Given a base rate of 2.6% of vehicles sold that are convertibles, this represents an 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 vehicles 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 about 80 degrees Fahrenheit). Purchases of 4-wheel drive vehicles are also very responsive to abnormal weather variation—particularly snowfall. Our results suggest that a snow storm of approximately 10 inches will increase the fraction of vehicles sold that have 4-wheel drive by about 2 percentage points over the next 2-3 weeks (an approximately 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. Black vehicles are less likely to be purchased when the weather is warm and sunny. A 20-degree increase in temperature leads to a 2.1% (0.26 percentage point) reduction in the fraction of vehicles sold that are black, compared to a 12.6% baseline percentage. Moving from overcast to completely clear weather reduces the sales of black vehicles by 5.6% (0.71 percentage points).

The data allow us to rule out several alternative explanations for these findings. For example, a distributive-lag model 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 vehicle purchases from one week to the next (a "harvesting effect"). We also present evidence that learning about a vehicle 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. In particular, 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 have less to learn about their value for convertible attributes 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 vehicle is traded in quickly for a different vehicle. This analysis, which uses unique vehicle identifiers to follow vehicles over time in our data, suggests that a vehicle is more likely to be returned quickly when purchased on a day with abnormal weather—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 estimate the value that certain house characteristics (e.g. a swimming pool or central air) have at different times of the year by looking at two different sales for a single house, 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 0.4 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 $1600 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 also find evidence that the value of central air is higher when a home sells in the summertime. However, we find no evidence that the hedonic value of lot size or fireplaces vary with seasonal temperature and discuss likely explanations for this finding.

Our findings are significant for several reasons. First, 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, state-dependent utility, and low-frequency decision-making.

Our paper is 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

[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

of 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 projection bias in two markets of even greater economic importance. The richness of our data allows us to explore not only how projection bias impacts sales volume, but also whether it has an impact on prices.

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, empirical strategy, and results for the housing market. Section IV provides a conclusion along with a brief discussion of the broader implication of our findings.

I. Conceptual Framework

In this section we describe 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 utility 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

Harreveld (2006, 2007). Loewenstein and Schkade (1999) provide a useful review of the psychological evidence for projection bias.

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

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 ($conv_t$) and owning it until period 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 $0 \leq \delta \leq 1$ is her standard discount factor. Once again, following Loewenstein, O'Donoghue, & Rabin (2003), a person with projection bias perceives her intertemporal utility 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 bad state of the world for consuming a convertible (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)). If the number of periods is small, then it would be perfectly reasonable that the current state of the world has an important impact on the decision to buy. For example, in an extreme case, imagine an individual's decision to rent a car for a few days. It would be perfectly reasonable to be more likely to rent a convertible if the weather on the day of rental is nice since the consumption utility from the first period is a large part of the overall consumption utility. The advantage that we have in our paper is that we are focused on the purchases of very durable goods (vehicle and home purchases). For these purchases, we argue that idiosyncratically warm weather on the day of purchase should have a minimal impact on the probability of purchasing a particular type of vehicle or house since most vehicles and houses are owned for a considerable period of time. In fact, even very high initial discount rates consistent with present-biased preferences of the type described by Laibson (1997) and O'Donoghue and Rabin (1999) cannot easily explain the effect sizes that we find on vehicle purchases. It would be even harder to explain our housing purchase results using present-biased preferences because housing consumption does not occur at the time of the decision (houses go under contract many days before a house is sold).

It is also important to note that weather states are not uncorrelated. In general, there are many warm-weather states that occur sequentially in the summertime followed by many cold-weather states in the wintertime. We would expect an unbiased, rational agent to be more willing to purchase a convertible in the spring than in the fall since a consumer who buys in the spring is likely to experience a string of "good" states of the world starting immediately. A consumer who buys in the fall will have to wait months to experience a similar run of "good" states of the world. 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 (and can use the pool immediately) relative to the amount they would be willing to pay if they moved in after the end of the summer (and would have to wait until next summer to use the pool). Thus, simply finding that people are willing to pay more for a home with a swimming pool or are more likely to buy 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.

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 very 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 home buyers 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 vehicle 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 vehicle purchases (e.g. the value to purchasing a convertible in the spring rather than the fall).

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.[2] Loewenstein, O'Donoghue, & Rabin (2003) assume that individuals correctly anticipate the path of states, but err when predicting the utility that those 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. Thus, it is difficult to separate these two different types of projection bias and our analysis will not attempt to do so. However, there are several reasons to believe that projection bias of states is unlikely to be the underlying mechanism. The first is the prevalence of weather information that is available to people during the time of our study, including their own experience of local weather patterns. It is much harder to find information about future utility than it is to find information about future states. In addition, Conlin, O'Donoghue, & Vogelsang (2007) (who also comment on 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.

[2] 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)).

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 vehicle transaction that occurred at the dealers in the sample. For each transaction, we observe the date and location of the purchase, information about the vehicle purchased, and the price paid for the vehicle. Our locations are defined by Nielsen Designated Market Areas (DMAs), which divide the U.S. into approximately 200 areas. DMAs are defined to correspond to media markets, which means that DMAs corresponding to major cities will have higher populations than DMAs in more rural regions. Examples of DMAs in our data include Phoenix, Arizona; Tulsa, Oklahoma; Lansing, Michigan; and Billings, Montana.[3]

We will add to these data information about local weather. The weather data were collected by first using wolframalpha.com to find the weather station nearest to the principal city in each DMA. Weather data themselves were obtained for each weather station from Mathematica's WeatherData compilation.[4] Data were collected on temperature, precipitation, precipitation type, and cloud cover. Temperature is measured as the simple average of the seven daily high temperatures in the week, measured in degrees Farenheit. Precipitation is measured as the cumulative liquidized inches over the course of the week. If the only precipitation type reported in the week is rain, we classify the precipitation as rainfall (measured in inches). If the only precipitation type reported during the week is snow, we classify the precipitation as snowfall (measured in liquidized inches). If both rain and snow are reported during the week, we classify the precipitation as slushfall (measured in liquidized inches). Cloud cover is a simple average of the seven daily measures of the fraction of the sky covered by clouds.

The data indicate that vehicle 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-

[3] A list of all the DMAs covered by our data is available from the authors.

[4] If the weather station did not have weather data available for at least 90% of the 4745 daily observations between 1997 and 2010, data for the second- or third-closest weather station was used for that DMA. (There are 21 DMAs that use data from the second-closest station, and 6 that use data from the third-closest station.)

wheel drive purchases. In Panel A of Figure 1, we illustrate the percentage of total vehicle transactions that were convertibles by month of the year. Overall, convertibles make up between 1.5 and 3% of total vehicles purchased. The data show a strong seasonal pattern in which the percentage of vehicles sold that are convertibles is highest in the early spring. For seven out of the eight years, the percentage of vehicles purchased that are convertibles peaks in April. While springtime is the most popular time to buy a convertible, the percentage of vehicles sold that are convertibles is still relatively large in the winter months. The annual winter troughs in percentage of vehicles 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 for spring buyers relative to fall buyers.

Similarly, Panel B of Figure 1 illustrates the percentage of total vehicle transactions that were 4-wheel drive vehicles by month of the year. 4-wheel drive transactions range between 20% and 35% of total vehicle 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).[5] 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 4-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 DMAs into two groups: DMAs with above- and DMAs with below-median monthly temperature variation.[6] 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

[5] There is a mid-summer peak in 2005 which arose from record sales during GM, Chrysler, and Ford's employee discount pricing promotions. (Busse, Simester, and Zettelmeyer (2010) describe the effect of these promotions.)

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

vehicles 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 DMAs is not too different. However, it is clear that the amount of seasonal variation is higher in the variable-temperature DMAs. Panel B shows that there is a large level difference in the percentage of 4-wheel drive vehicles purchased in the two types of DMAs and once again the variable temperature areas appear to have 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 DMA-week level.[7] After collapsing, we create variables that represent the percentage of total vehicles sold in each DMA-week that were convertibles and that were 4-wheel drive vehicles. Weekly weather data at the DMA level are also merged in. These data will allow us to test whether abnormal weather leads to abnormally high or low levels of convertible and 4-wheel drive purchases. Note that our estimates will identify the effect of weather on the equilibrium sales of vehicles of different types. In other words, we will estimate not only the effect of weather on vehicle demand, but also the effect of any actions dealers take in response to their perception of increased demand for certain types of vehicles under particular weather conditions. Of course, if there is a supply effect, that is evidence that dealers believe buyers are influenced by projection bias, and respond accordingly. Our estimates identify the combined effect of changes in consumers' behavior and dealers' responses to those changes.

We proceed by first presenting the results for convertibles followed by the results for 4-wheel drive vehicles. Vehicles have other characteristics whose value might be weather related (sun roofs, air conditioning, snow tires, etc.). Many of these characteristics are either unobservable to us, or do

[7] 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.

not vary significantly in the data. However, in a later section we will consider the effect of weather on the sales of black vehicles.

Baseline Convertible Results. We begin the analysis by using two DMAs (Chicago and Miami) as examples of the effects that we find. Panel A of Figure 3 plots the percentage of all vehicles sold in Chicago that were convertibles 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 vehicles sold in the wintertime to 3-4% of total vehicles 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 vehicles 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 vehicles sold was 0.5 percentage points higher than our regression predicted for that week of the year. Figure 4 illustrates the seasonal pattern of 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 year and 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 positively and statistically significantly correlated (correlation coefficient = 0.36; t-stat = 7.9). 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 vehicles sold (a 14.4% increase given the baseline of 2.5%).[8]

[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).

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 positive correlation in quarters 1, 2, and 4 (t-stats: 5.2, 4.0, and 4.7 respectively). We argue that the lack of statistically significant 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, a week with abnormally warm weather 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 DMAs. We use Miami as the second example for how weather impacts convertible sales. Figures 6, 7, and 8 replicate Figures 3, 4, and 5 using data from the Miami-Ft. Lauderdale DMA. 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. Relatedly, the deviations of weather in Miami from weekly norms (Panel B of Figure 7) are smaller than in Chicago (Panel B of Figure 4). Due to the warmer average temperature in Miami, we would predict that abnormally warm weather in Miami does not increase the fraction of vehicles purchased that are convertibles by nearly as much as abnormally warm weather in Chicago. To test this directly, we calculate the correlation between the residual convertible sales in Panel B of Figure 6 and the residual temperature from Panel B of Figure 7. The overall correlation between residual convertible sales and residual temperature in Miami is actually negative (although not statistically significant; t-stat: -0.5). Figure 8 shows that the correlation is not statistically significant for any of the four quarters of the year.

We generalize from our Chicago versus Miami example by combining the data for all DMAs to estimate the impact of temperature on convertible sales across our entire sample. We do so by estimating the following specification.

(4) $$PercentConvertible_{rt} = \alpha_0 + \alpha_1 Weather_{rt} + \mu_{rt} + \tau_{rT} + \epsilon_{rt}$$

PercentConvertible measures the percentage of vehicles sold in DMA r during week t that were convertibles. *Weather* is a vector of weather variables for DMA r in week t—temperature, rainfall, snowfall, slushfall and cloud cover—defined previously in this section. (Summary statistics can be found in Table 9.) μ_{rt} are DMA*week-of-the-year fixed effects and τ_{rT} are DMA*year fixed effects. Given the varied size of the DMAs in our sample, we weight the regression based on the total number of vehicles sold in the DMA-week.

Table 1 reports the results of estimating Equation (4). Column 1 indicates that when the temperature is 1 degree higher than expected in a given DMA, the DMA experiences on average an increase of 0.011 percentage point in the convertible fraction of total vehicles sold. Thus a 20-degree swing in temperature in any given week, is predicted to change the convertible percentage of vehicles sold by 0.22 percentage points (an 8.5% change relative to the weighted base rate of 2.6% of vehicles sold being convertibles). Liquid inches of rain, snow, and slushfall all have negative impacts on the convertible percentage of vehicles sold, although these effects are relatively small given the amount of variation in rain, snow, and slushfall 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 (relative to completely overcast) increases convertible demand by the same amount as approximately 16 degrees higher temperature.

Another way to understand the size of the estimated effect would be to calculate the decrease in price that would be necessary in order to reduce sales by the same amount. Berry, Levinsohn, and Pakes (1995) estimated demand elasticities for 13 specific models of vehicles. Their estimates ranged from approximately -3 to -6. Assuming an average convertible price of $40,000 (the average in our data), price would have to fall by $1,133 (assuming an elasticity of demand of -3) to $567 (assuming an elasticity of demand of -6) in order for quantity demanded to fall by 8.5 percent, the predicted effect of a 20-degree change in temperature. This suggests that the size of the effects we find are much larger than the utility most people would get from owning a convertible during one week of particularly good weather.[9]

The next four columns in Table 1 break down the impact of temperature and other weather variables on convertible sales by quarter of the year. Consistent with the Chicago example shown in Figures 3 through 5, the effect of 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 illustrate, 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 temperature on convertible sales separately for DMA-weeks-of-the-year with different mean values for the daily high temperature. The mean value for the daily high temperature for each DMA-week-of-the-year was obtained by calculating the average of the daily high temperatures in a given DMA-week-of-the-year across the different years in our sample. We then group DMA-weeks into 5-degree bins by average daily high temperature for the corresponding DMA-week-of-the-year. We re-estimate Equation (4) for each bin. Figure 9 plots the temperature coefficients (with 95% confidence intervals) estimated for each 5-degree bin of average temperature values. For example, the leftmost point plotted in the graph is the estimated coefficient for DMA-weeks-of-the-year 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 DMA-week-of-the-year 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

[9] This is particularly true if one considers the possibility of renting a convertible in order to enjoy a week of unusually good weather, an option that would be less hassle and would not require buyers to bear the initial depreciation associated with buying a new vehicle.

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 may have a negative (positive) impact on convertible purchases. These heterogeneous effects explain the zero-effect of temperature on convertible sales 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 4-wheel drive vehicles seem like an especially good idea. Table 2 presents our estimate of the impact of weather variation on the 4-wheel drive percentage of total vehicles sold obtained by substituting *Percent4WheelDrive*, the percentage of total vehicles sold that are 4-wheel drive, on the left hand side of Equation (4). 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 leads 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 vehicles sold with 4-wheel drive). We also find a large, positive impact of snow and slush on 4-wheel drive transactions. One inch of liquidized snow (about 10 inches of snow) leads to a 1.02 percentage point increase in the percentage of total vehicles sold with 4-wheel drive. The effects for snowfall are statistically significant in quarters 1 and 4 (the standard errors for quarters 2 and 3 indicate that we do not have sufficient snowfall variation to estimate effects in these quarters). The effect of snowfall is larger in quarter 4 than in quarter 1. However, the significant impact of snowfall 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 snowfall each week has a very different distribution from the distribution of temperature. Snowfall is usually zero in most DMA-weeks, but can have very large values in a few DMA-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 (measured in amount of snowfall) that occurs in each July-to-June year in each DMA in our sample that has above median weather variation. (This excludes places with no snowfall.) We regress the 4-wheel drive percentage of total vehicles sold in each DMA-week on DMA*year and DMA*week-of-the-year, weighting by the total number of vehicles sold in each DMA-week. We obtain the residuals from this regression for each observation, and sort the residuals by the number of weeks before or after the largest snow storm of the year in the DMA where the observation occurred. Figure 10 plots the average of these residuals 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 vehicles 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 vehicles sold with 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 log number of convertible and 4-wheel drive purchases made 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 as much as purchases of vehicles 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, an individual who is going to purchase a 4-wheel drive vehicle appears to be more motivated go to the dealership than buyers of non-4-wheel drive vehicles.

Alternative Characteristics. We have estimated the effect of weather on the sales of vehicles with two characteristics whose utility is weather-related, convertible roofs and 4-wheel drive. One could imagine a variety of other characteristics whose value to a customer also varies with weather: air conditioning, sunroofs, snow tires, towing packages, etc. We cannot estimate the weather-related effects of all of these characteristics because some do not vary much in the data (air conditioning) and others we don't observe (sunroofs).

We briefly present the effect of one additional characteristic that we can observe and that varies in the data.[10] Light colors reflect solar radiation, while dark colors absorb it. This means that a black car can be oppressively hot and stuffy if it has been parked outside on a hot and sunny day. Car buyers seem to be familiar with this. Overall in our data, 12.6% of vehicles sold are black; however in Las Vegas, only 9.3% of vehicles sold are black, while in Phoenix the percentage is only 7.8. In Table 3, we report the results of regressing the percentage of vehicles purchased in a DMA-week that are black in color on weather variables, and on DMA*year and DMA*week-of-the-year fixed effects (the same specification as in Tables 1 and 2).

We find, in column 1, that the fraction of vehicles purchased that are black decreases by 0.013 percentage points for every degree increase in temperature. This means that a 20-degree increase in temperature would be associated with a 0.26 percentage point decrease in the sales of black vehicles, a 2.1% change relative to the baseline percentage of 12.6%. Sunshine matters, too. Going from an overcast week to a completely clear week lowers the percentage of black vehicles sold by 0.71 percentage points, or 5.6% relative to the baseline.

In columns 2 through 5 of Table 3, we split the estimates up by quarter. Hot weather and sunny weather reduce the sales of black vehicles in quarters 1, 2, and 4. In quarter 3, we find that sunshine matters even more than in other quarters, while temperature is estimated to matter less.[11]

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 vehicles 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 vehicles, but is actually just causing short-run intertemporal substitutions in vehicle 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

[10] We thank Loren Pope for suggesting this approach.

[11] In unreported results, we find evidence that cloud cover and temperature are strongly negatively correlated in quarter 3. Specifically, if we estimate the quarter 3 results without cloud cover, the estimated coefficient for temperature increases in magnitude to -0.023 and becomes statistically significant.

to be a nice day outside.[12] In fact, our previously noted finding that abnormally warm weather in November can affect convertible purchases and a snow storm in February can affect 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 estimate a distributive-lag model that adds to the weather variables during the week of purchase a one-week lead and 12 weeks of lagged weather variables. By including lag variables, we are able to test whether having cold or hot weeks leading up to the week of purchase influences how the current weather affects behavior. For example, in the convertible scenario, negative coefficients on the lag variables are interpreted as evidence of harvesting via the following argument. A negative coefficient on, say, the three week lag of temperature indicates that if the weather *three weeks ago* was hot, sales *this week* are lower by some amount than they otherwise would have been. This implies that if the weather *this week* is hot, sales *three weeks from now* will be lower by that same amount. We can thus use the lag coefficients to answer the question "If the weather is hot this week, how much lower will sales be in subsequent weeks?" The one week lag gives us an estimate for the effect of hot weather this week on sales one week from now, the two week lag estimates the effect of hot weather this week on sales two weeks from now, and so on. Thus, if we add up all our lag coefficients and find that they equal the negative of the current period coefficient, it suggests that any increased sales that occur due to hot weather this week are made up entirely of sales displaced from the twelve following weeks. More generally, the sum of the lag coefficients tells us how much of our estimated current period effect is due to intertemporal substitution.[13]

Table 4 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 coefficients on the lag variables are all small relative to the current temperature coefficient, in most

---

[12] The fact that more convertibles are bought in spring than winter and the reverse for 4-wheel drive vehicles suggests that there may be harvesting in response to the overall seasonal pattern of the weather. However, this does not mean that harvesting happens in response to idiosyncratic weather variation.

[13] See Jacob, Lefgren, and Moretti (2007) for a similar analysis that tests for intertemporal substitution of crime using abnormal weather shocks and Deschenes and Moretti (2009) who test for intertemporal substitution of mortality using abnormal weather shocks.

cases not statistically significant, and more often positive than negative. In the full data (Column 1 of Table 4), there is no evidence that warmer than usual weather in the previous weeks affects 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 5 provides a similar analysis for 4-wheel drive purchases. This analysis indicates that snowfall anytime in the last three weeks leads to an increase in the percentage of vehicles sold with 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 percentage points. Thus, approximately 47% of the positive effect of snowfall on 4-wheel drive purchases that occurred in the current, lead, and 3 lag weeks (2.42 percentage points) could be considered as arising from harvesting. In other words, the increase in the percentage of 4-wheel drive vehicles 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 vehicle 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 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 vehicle purchasing that may lead to a correlation between weather and vehicle purchase timing, particularly for convertibles, is the desire of most customers to test drive a vehicle before buying. Suppose a customer is considering buying a convertible, and 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. Now 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 vehicle, perhaps because his current vehicle has broken down and is no longer worth repairing. Suppose that a convertible is one of the vehicles he would consider purchasing, but on the day he needs the new vehicle it is too cold to test drive a convertible. Unwilling to buy the convertible without being able 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 vehicles sold on warm days being convertibles relative to cold days for reasons other than projection bias. The first type of customer that we outlined above would lead to harvesting (customers wait until a warm week to buy a convertible—so that they can test drive the vehicle). We already discussed and ruled out harvesting effects for convertibles in the previous section. However, the second customer type that we discuss above is not ruled out by our distributive lag model. Several pieces of evidence, however, argue against a test-drive learning story. For example, Figure 9 indicates that an extra degree of warm weather results in more convertible purchases even when the baseline temperature is in the 60-80 degree range. This is a range of temperature for which it is clearly possible for someone to test out the various car features comfortably. Our results thus suggest that it is more than simply testing the features of a car that cause warm weather to result in a higher number of convertibles being sold. We can also 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 have a significant negative effect on the percentage of vehicles sold that are convertibles, consistent with projection bias. It is particularly noteworthy that cloudy days have a negative effect in all four quarters, and the effect of cloudy days is largest in the third quarter, when days are generally warm. This third quarter effect is especially suggestive of the fact that people buy more convertibles on warm days not because it is more possible to test drive them, but because it seems more attractive to own a convertible on such days.

Vehicle Buyers who Previously Owned a Convertible or 4-Wheel Drive Vehicle. Another alternative hypothesis that would explain our findings is that customers need to test drive a vehicle on 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 4-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 vehicle buyers can't quite imagine what it would be like to own this vehicle 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 vehicle. This means we can observe vehicle transactions by customers whom we know have already owned a convertible or have already owned a 4-wheel drive vehicle. 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 4-wheel drive owners and cold or 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 non-convertible vehicle. Column 1 of Table 6 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 temperature at the time of purchase on convertible demand. The point estimate is about six times larger than the point estimate in the entire sample—although the larger estimate in percentage *point* terms is smaller in percentage terms because the convertible purchase rate in this sample (25%) is so much higher.[14] In Column 2 of Table 6, we estimate the effect of weather on buyers who are trading in a 4-wheel drive vehicle. Overall, 78% of people who trade in a 4-wheel drive vehicle purchase another 4-wheel drive vehicle. In Column 2 we continue to find strong and statistically significant effects of abnormal weather—including temperature, snowfall, slushfall, 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 4-wheel drive vehicle 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 4-wheel drive vehicle cast doubt on a learning story explaining the effects that we find.

Expensive Vehicles. One might worry that our finding that warm weather leads to higher convertible sales is simply spurious correlation of the following sort. Suppose that good weather puts people in a generally good mood, and that when people are in a good mood they spend money more freely. If that were so, then we might see good weather associated with higher convertible sales simply because convertibles are more expensive on average. (In our data, the average price of a vehicle that is not a convertible is $20,542, while the average price of a convertible is $30,845.)

We investigate this hypothesis by re-estimating Equation (4) (the specification reported in Column 1 of Table 1) replacing the dependent variable "percentage of vehicles sold that are

[14] 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, an 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%.

convertibles" with "percentage of vehicles sold whose price is greater than *X*." We estimate four variants of this alternative specification, with *X* equal to $20,000; $30,000; $40,000; and $50,000.[15]

Our original specification found that when the temperature rises by one degree (all else equal), the percentage of vehicles sold that are convertibles goes up by 0.011 percentage points (t-stat = 14.4). In our alternative specification, we find that when temperature goes up by one degree, the percentage of vehicles sold whose price is more than $20,000 is unchanged (coefficient estimate is 0.000, t-stat = 0.28). The estimated coefficient for a threshold price of $30,000 is 0.002 (t-stat = 0.93); for a threshold price of $40,000 the estimated coefficient is 0.001 (t-stat = 1.5); and for a threshold price of $50,000 the estimated coefficient is 0.000 (t-stat = 0.39). These results give no evidence that our original effect is driven by buyers buying more expensive vehicles in good weather.

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 vehicles 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 vehicle market is much harder than for catalog orders. First, is simply a data limitation. Although our data are impressive and represent a 20% sample of all new car dealerships in the U.S., we can only identify "returned" vehicles that happen to be traded in or sold as a used vehicle at one of the dealerships we observe. Said another way, for any vehicle whose sale we observe at some point, we have roughly a 20% chance of seeing that vehicle's subsequent return or resale if that transaction happens at a

---

[15] We run this analysis on all non-convertible vehicles. We eliminate convertibles from this analysis so that our results are not affected by our finding that higher temperatures are associated with increased convertible sales. For this analysis, we want to know whether higher temperatures increase the sales of high-priced vehicles absent a convertible effect.

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 that he or she has 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 (likely at a loss if the vehicle is new because of the rapid initial depreciation of new vehicles) and buy the hardtop at full price. 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 vehicle reappears in the market. Of the roughly 40 million vehicles that are transacted 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.[16] On average in the U.S., owners keep their vehicles 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 4-wheel drive vehicles that were purchased when the weather was abnormally cold 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 7 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 DMA*week fixed effects to eliminate seasonal and geographic differences in how quickly vehicles are returned. Table 7 shows that convertibles are, overall, 1.272 percentage points more likely to be returned within a year than other types of vehicles; 4-wheel drive vehicles are also more likely to be returned (by 0.285 percentage points) than other types of vehicles. The positive signs of the coefficients estimated for the interaction of convertible and 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

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

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 4.6% change relative to the baseline convertible return rate of 7.332%). The temperature interaction with 4-wheel drive vehicles is more consistently statistically significant, and indicates that a 4-wheel drive vehicle 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, while clearly suggestive, is less strong than 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.

Price Effects. We have shown in the previous sections that the percentage of vehicles sold that are convertibles is higher in weeks with warm and sunny weather, while the percentage of vehicles sold that are 4-wheel drives are higher during and just after weeks with cold, snowy weather. We argued that this is evidence of projection bias—that individuals are over-influenced by the current weather when they are making vehicle purchase decisions.

At a market level, one could describe this effect as an increase in demand associated with unusually warm and sunny, or unusually cold and snowy, weather. Thinking of the phenomenon this way, one might wonder whether there is an effect of projection bias not only on the quantities of vehicles people buy, but on the prices they pay. In a simple demand model, if the demand curve shifts out while an (upward-sloping) supply curve stays fixed, one would expect to see both higher prices and higher sales quantities.

There are several ways in which this simple model is not an ideal fit for the car industry. First, from a dealer's perspective, the supply of vehicles is not upward-sloping. Dealers can order vehicles from manufacturers as a fixed per unit invoice price in whatever quantity they wish. This corresponds to horizontal marginal cost curve for the dealer. If the dealer is selling vehicles in competitive market, the effect of an increase in demand should be increased sales, with essentially no increase in price.[17]

Second, a competitive price-taking market is not a very good description of the retail car industry. Individual buyers negotiate a price for a specific vehicle with the dealer. Whether the incremental buyers who are buying as a consequence of projection bias pay higher prices or lower prices than other buyers depends on the reservation prices and bargaining characteristics of projection bias buyers relative to other buyers. One might argue that projection bias buyers must have higher reservations prices than buyers on average, because they are being strongly swayed by temporary weather conditions. Similarly, one might argue that buyers who can buy "on impulse" must have high liquidity, and therefore likely higher incomes and higher reservation prices, than average buyers. Alternatively, one might argue that projection bias buyers are buyers who would not be buying this vehicle on another day, and that the influence of the weather has nudged them just above their point of indifference about buying. In this case, they might well have lower reservation prices than average buyers. Similarly, if dealers recognize which buyers are projection bias buyers, they may realize that they must offer a good price today, or lose the sale forever, since in another few days the weather will change and these buyers will no longer be in the market.[18]

Overall, we conclude that it is an empirical question whether prices for convertibles and 4-wheel drives will be higher in the same weeks that warm and sunny weather or cold and snowy weather leads to increased sales of these types of vehicles. We estimate the effect of weather on the prices of convertibles and 4-wheel drives using the following specification.

(5)
$$Price_{ijrt} = \beta_0 + \beta_1 Weather_{rt} + \beta_2 PurchaseTiming_{it} + f(Odometer_i, \beta_3) + \mu_{rt} + \tau_{rT} + \phi_i + \epsilon_{ijrt}$$

Price measures the price paid in transaction *i* for vehicle *j* that occurred during week *t* in DMA *r*. (In order to make our measure of price represent a customer's total wealth outlay for the vehicle, we

[17] Dealers place orders for vehicles months in advance, so over a horizon of several months, a dealer's supply of vehicles is predetermined. However, dealers can sell more or fewer vehicles on any given day, meaning daily vehicle supply is not fixed. (For more on how dealer supply and inventory affects prices, see Zettelmeyer, Scott Morton, and Silva-Risso (2007).

[18] We thank Glenn Ellison for suggesting this point.

define price as the contract price for the vehicle agreed upon by the buyer and the dealer, minus any manufacturer rebate the buyer received, plus any loss (minus any gain) the buyer received in negotiating a price for his or her trade-in.) *Weather* is a vector containing the temperature, rainfall, snowfall, slushfall, and cloudcover for week t in DMA r. *PurchaseTiming* is a vector containing indicators for whether transaction i occurred during the weekend, or at the end of the month, times in which salespeople may be willing to sell vehicles at a discount in order to hit sales volume targets. The specification also includes DMA*year (τ_{rT}), DMA*week-of-year (μ_{rt}), and "vehicle type" (ϕ_i) fixed effects. (A vehicle type is defined by the interaction of make, model, model year, trim level, doors, body type, displacement, cylinders, and transmission.) We estimate Equation (5) separately for new convertibles, used convertibles, new 4-wheel drives, and used 4-wheel drives. The specifications that estimate the effect of weather on used vehicle prices also include a linear spline in the vehicle's odometer, which allows vehicle prices to depreciate over time in a reasonably flexible way. (See Busse, Knittel, and Zettelmeyer (forthcoming) for use of a similar specification to estimate price effects in similar data.)

Table 8 reports the results of estimating equation (5). Generally speaking, we find that the effect of weather on prices is fairly small, even when it is statistically significant. In Column 1, which estimates the effect of weather on new convertible prices, none of the weather variables have statistically significant effects. In Column 2, which estimates the effect of weather on used convertible prices, an increase in temperature of 20 degrees would be estimated to increase the average price of convertibles sold during that week by $79.60, a very small amount compared to an average transaction price of $22,222 for used convertibles. In addition, one inch of liquidized snow (about 10 inches of snow on the ground) is associated with transaction prices that are lower by $114.48. For 4-wheel drives, the results are similarly small, and the directions of the effects are mostly counter-intuitive. A 20-degree drop in temperature would be predicted to decrease the average transaction price of a new 4-wheel drive by $16.60 and of a used 4-wheel drive by $40.60. These are very small effects relative to an average transaction price of $31,845 for new 4-wheel drives and $19,132 for used 4-wheel drives. In addition, for used 4-wheel drives, a 10-inch snowfall is predicted to decrease the average transaction price by $23.80, while going from a sunny to an overcast week is predicted to decrease the average transaction price by $54.06. The only coefficient that is statistically significant in the expected direction is the estimate of the effect of cloud cover on new 4-wheel drive prices; going from a completely sunny week to a completely overcast week is predicted to increase the price of a new 4-wheel drive by $36.44.

III. Housing Market

Data and Empirical Strategy. Our analysis is based on a housing dataset of more than four million single-family residential properties across the United States that sold at least twice between 1998 and 2008. We purchased the data from a commercial vendor who had assembled these data from assessors' offices in individual towns and counties.[19] Since larger metropolitan areas are more likely to archive their assessor data electronically and sell it to commercial vendors, urban counties are over-represented relative to rural counties.[20] The data include the transaction price and the sale date of each house, the previous transaction price and sale date, 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. We observe transactions at two different dates for a single house, but only a single set of characteristics for a given house—the characteristics that existed at the time of the second transaction.[21]

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

[19] The commercial vendor is Dataquick which is the source of housing data for many papers in the literature.
[20] 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.
[21] This is because we are essentially creating two transactions from a single observation—an observation which records the current sale price and sale date of a particular house, and the sale price and sale date of the most recent previous transaction for that same house if the previous transaction occurred during our data window.

contract date and closing date were available.[22] 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 will be restricted to examining the effect of seasonal weather patterns rather than precise idiosyncratic weather differences at the time of purchase.

We clean the data in a similar manner as previous work that has used this type of housing 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 0.5 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).[23] 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, our sample contains only transaction for houses that had a previous sale within our sample period.

Table 10 provides some basic summary statistics for the final set of housing transactions in our dataset (roughly 4.2 million observations). The average sales price over the entire time frame of our data was 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% had central air, and 46% had at least 1 fireplace.[24] The average home was built in 1968 on a 0.32 acre lot with approximately 3 bedrooms, 2 baths, and 1,700 square feet of living area.

[22] We thank Steve Levitt and Chad Syverson for sharing this information with us.

[23] 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.

[24] Certain counties in our dataset reported no homes as having either 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).

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

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

where $Log(Sales\ Price)_{itc}$ is the log sales price of house i in sample-month t in county c. γ_i is a fixed effect for house *i,* which we can estimate because we use only houses we observe being sold more than once. θ_{tc} is a county*sample-month fixed effect. The residual from Equation (6), e_{itc}, represents, for each house transaction, how much more or less the house sold for than would have been expected after considering how much that very house sold for on another occasion, and how much other houses in the same county sold for during the same sample-month. We will analyze the residuals from this regression by month-of-the-year and house type to see whether there is evidence of projection bias.

Results. We begin by calculating the average residual (obtained from Equation (6)) for homes with swimming pools by the month-of-the-year the house is assumed to have gone under contract (i.e. two months before the sale date). In Panel A of Figure 11, we plot these average residuals along with their 95% confidence intervals. It is worth noting that these average residuals by month-of-the-year for homes with swimming pools (if weighted by the number of transactions in each month-of-the-year) would sum to zero across months because house fixed effects are included in the regression. Similarly, the residuals for all homes sold within a single sample-month must sum to zero because sample-month fixed effects are included in the regression. Therefore, whenever we see a positive average residual in a given month-of-the-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 for houses without swimming pools is on the order of one-tenth the size of the average positive residual for houses with swimming pools since houses with swimming pools only represent about 12% of the data).

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 0.22

percentage points more on average than otherwise expected (this effect is jointly statistically significant and individually significant for June and August), while homes with swimming pools that go under contract in the three coldest months of the year (December, January, and February) sell for on average 0.18 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 about $398,000, this represents a roughly $1600 swing in value for homes with swimming pools that go under contract in the summertime relative to the wintertime.

Our finding that transaction prices are higher for houses with swimming pools that went under contract in the summer (especially in August) argues strongly against standard discounting or present-biased preferences as the reason for our results. The houses that we identify as selling in August are houses that will close in October, meaning that the buyers of those houses will move in just at the point in the year in which swimming pool season is the farthest away.

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 regress the residuals from Equation (6), e_{itc}, on all the house characteristics we observe, plus interactions for months-of-the-year. To be precise, we estimate:

$$(7) \qquad e_{itc} = \theta_0 + \theta_{1,t}\,\mu_t\,SwimmingPool_i + \theta_{2,t}\,\mu_t\,CentralAir_i + \theta_{3,t}\,\mu_t\,Fireplace_i + \theta_{4,t}\,\mu_t\,LotSize_i + \theta_{5,t}\,\mu_t\,Bedrooms_i + \theta_{6,t}\,\mu_t\,Bathrooms_i + \theta_{7,t}\,\mu_t\,SquareFootage_i + v_{itc}$$

e_{itc} is the estimated residual from Equation (6), and represents how much the log price observed for the sale of house *i* in county *c* in month *t* differs from what would be predicted from other sales of that house and from the sales of other houses in the same county and month. *SwimmingPool, CentralAir,* and *Fireplace* are indicator variables recording whether house *i* has the corresponding feature. *LotSize* measures the size of the lot in acres. *Bedrooms* and *Bathrooms* count the number of rooms of each type. *SquareFootage* measures the size of the house in square feet. μ_t is an indicator for the month-of-the-year in which the transaction occurs. The coefficients can be interpreted as follows: $\theta_{1,1}$ estimates how large on average the residual of log price (net of house and county*sample-month fixed effects) is for houses with swimming pools that sell in January, conditional on all the other house attributes we observe.

Panel B of Figure 11 presents the twelve swimming pool coefficients ($\theta_{1,1}$ through $\theta_{1,12}$) from Equation (7). Controlling for the seasonal effect of the other housing characteristics on the residual log price does not substantially change our estimates of the effect of swimming pools. Our results continue to show that the value of a swimming pool is higher in the summertime than in the wintertime, although with somewhat reduced statistical significance.

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. Because sales price in Equation (6) is measured as log price, the residual values are also measured in logs. Removing the top and bottom 1% of residual values eliminates homes whose sales price was about 60% more than or 60% less than what would be predicted Equation (6). Removing the top and bottom 5% of residual values eliminates homes that sold for about 25% more or less than Equation (6) would predict.[25]

Figure 12 displays the twelve swimming pool coefficients ($\theta_{1,1}$ through $\theta_{1,12}$) obtained by estimating Equation (7) for the 1% trim sample (Panel A) and the 5% trim sample (Panel B). The same general 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 confidence intervals become much tighter due to the elimination of these high-variance observations. Our preferred specification (with

[25] The 5% and 1% cutoffs for trimming are symmetric 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.

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 throughout the wintertime (November through March).

Along with swimming pools, we observe three additional housing characteristics in our data that we believe could have a seasonal component: central air conditioning, fireplaces, and lot size. In Figure 13, we report the estimated coefficients from Equation (7) associated with each of these characteristics, estimated on the 5% trim sample. Panel A shows the estimated coefficients for central air ($\theta_{2,1}$ through $\theta_{2,12}$). There appears to be a seasonal pattern in which central air is worth more in the summertime (especially June and September) and less in the wintertime (November and January-March). The results are smaller in magnitude than those found for swimming pools. Panel B and Panel C present the results for fireplaces and lot size, respectively. The results are smaller in magnitude and less statistically significant than the central air results in panel A. There is little evidence of a discernible pattern in either of these results.

Why do we find small or no results for fireplaces and lot size? 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 of swimming pools across seasons. People may enjoy using fireplaces from fall straight through to spring, and the value of having a large lot may be high both in the spring or fall when yards are very beautiful, and in the summer when people spend a lot of time outdoors.

In Figure 14, we report the seasonal value of 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 little evidence of a statistically or economically significant seasonal pattern for these housing characteristics. The lack of seasonal variation in the value of these characteristics (both in terms of statistical significance and effect size) lends credibility to the effects that we find for swimming pools and central air.

Our hypothesis is that higher temperature levels at the time of the purchase decision lead to higher sales prices for houses with swimming pools and central air when compared to purchase decisions made during colder parts of the year. Up to this point, however, we have not 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 county*sample-month, 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.[26]

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). In light of this, we estimate the following specification, whose results are reported in Table 11.

(8) $$e_{itc} = \delta_0 + \delta_1\, Temp_{tc}\, X_i + \delta_2 Temp_{tc} + \delta_3 X_i + \xi_{itc}$$

e_{itc} is the residual of the log sales price (net of house and county*sample-month fixed effects) obtained from Equation (6). X_i is a vector of the housing attributes we observe (swimming pool, central air, fireplace, lot size, bedrooms, bathrooms, and square footage). We measure $Temp_{tc}$ in four different ways. In Columns 1 and 2, $Temp_{tc}$ is the average daily high temperature in county *c* in month *t*, the month in which the house is inferred to go under contract. In the next three pairs of columns, $Temp_{tc}$ is an indicator variable that corresponds to whether the average daily high temperature in county *c* in month *t* is at least 70, 80, or 90 degrees, respectively. The first column in each temperature pairing in Table 11 reports results for the full sample while the second column reports results for the 5% trim sample. We multiplied all coefficients in Table 11 by 100 for ease of

[26] 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.

reporting. We can therefore interpret—as we do in the next paragraph—the coefficients as approximate percentage point changes.

The first column in Table 11 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 percentage points. 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). In Column 2, central air is also estimated to be more valuable during high temperature months. The interaction effects of temperature and the remaining housing characteristics in these two columns are nearly all small and statistically insignificant.

The next three pairs of columns in Table 11 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 90 degrees sold for 0.37 percentage points more than when these same houses went under contract in a month whose average temperature did not reach 90 degrees.

Although our housing results suggest that projection bias is at work in this market much as we found in the car market, our analysis would be even more compelling if we could see if houses with swimming pools that went under contract in the summertime were more likely to "fall through" and not actually close. This would be analogous in some ways to our results on returning vehicles and to Conlin, O'Donoghue, and Vogelsang's (2007) results on returning cold weather catalog items. Unfortunately our data preclude us from doing such an analysis since we don't have information on homes that went under contract but then did not close. However, this would be an interesting extension if one were able to acquire the relevant housing information to perform this test.

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 make these predictions accurately. We show that projection bias causes consumers in the car and housing markets 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.

From a policy perspective, our results suggest that consumers would benefit from laws designed to help them better evaluate their decisions. For example, laws that allow consumers a "cooling-off period" for durable goods or goods for which consumers sign extended contracts may provide significant benefits to consumers. Such laws could also provide incentives for sellers to help buyers be in a "cool" state before an important transaction or contract is made.[27] 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."[28] This rule was made specifically to deal with high-pressure sale situations such as door-to-door sales. The Federal Trade Commission's cooling-off rule does not apply to real estate and automobile sales even though there clearly can be high-pressure sale situations for these important durable goods. While our results suggest that some consumers might benefit from an opportunity to reverse a decision once they have "cooled-off," applying a cooling-off rule to vehicle purchases would provide other consumers an opportunity to game the system by "buying" a new convertible at the beginning of a holiday weekend and returning it after a few days, claiming to have had a change of heart.

Despite showing that projection bias can impact important consumer markets, there are many questions about projection bias that are left unanswered and that future research may be able to address. For example, it is unclear how easy it is to "de-bias" consumers. It is possible that simply

[27] See Camerer et al. (2003) for an extended discussion about cooling-off periods and their potential applications in settings where people make suboptimal choices.
[28] More information on the Federal Trade Commission's "Cooling-Off Rule" can be found on their webpage at: http://www.ftc.gov/bcp/edu/pubs/consumer/products/pro03.shtm .

providing consumers with information about projection bias or asking them to imagine how they will feel about their purchase in a different state of the world could lead to improved decision making. Another extension of our research 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.

References

Berry, S., J. Levinsohn, and A. Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841-890.

Busse, M., C. Knittel, and F. Zettelmeyer. Forthcoming. "Are Consumers Myopic?: Evidence from New and Used Car Purchases." *American Economic Review.*

Busse, M., D. Simester, and F. Zettelmeyer. 2010. " `The Best Price You'll Ever Get': The 2005 Employee Discount Pricing Promotions in The U.S. Automobile Industry." *Marketing Science,* 29(2): 268-290.

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.

Deschenes, O. and E. Moretti. 2009. "Extreme Weather Events, Mortality, and Migration." *Review of Economics and Statistics, 91(4),* 659-681.

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.

Laibson, D. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics, 112(2),* 443-77.

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.

O'Donoghue, T. and M. Rabin. 1999. "Doing It Now or Later." *American Economic Review,* 89(1): 103-24.

Polk. 2010. "Consumers Continuing to Hold onto Vehicles Longer." November 3, 2010. https://www.polk.com/company/news/consumers_continuing_to_hold_onto_vehicles_longer_according_to_polk

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. "Visceral 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.

Zettelmeyer, F., F. Scott Morton, J. Silva-Risso. 2007. "Scarcity Rents in Car Retailing: Evidence from Inventory Fluctuations at Dealerships." Mimeo, Northwestern University.

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

Panel A. Convertible Percentage of Vehicles Sold

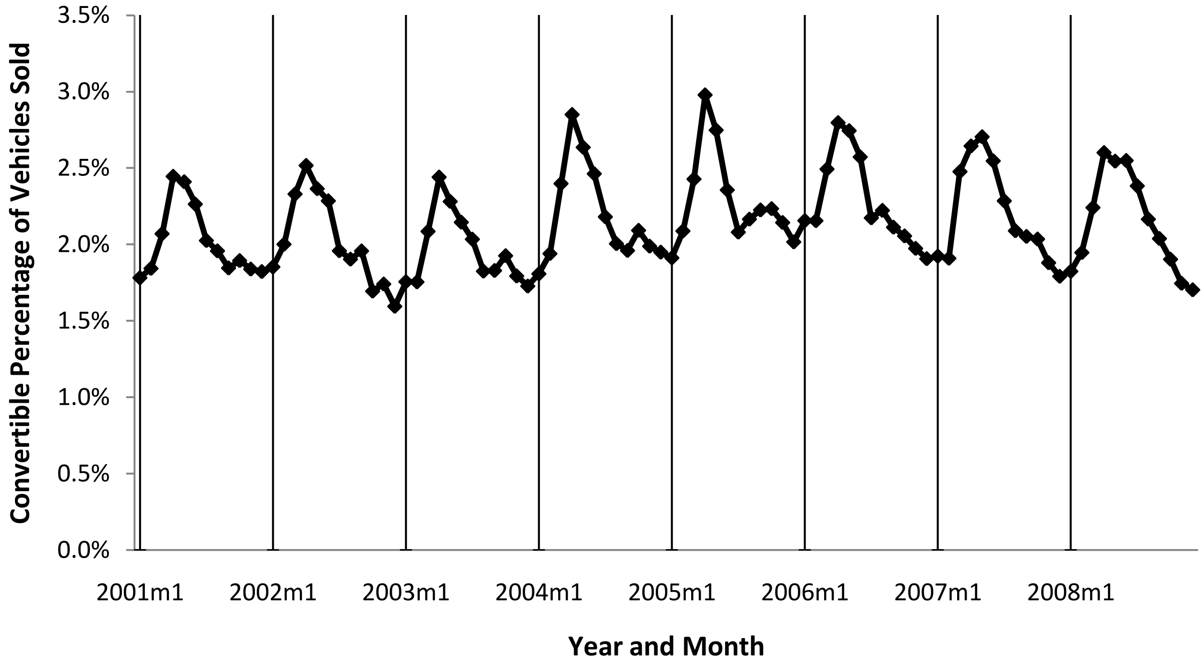

Panel B. 4-Wheel Drive Percentage of Vehicles Sold

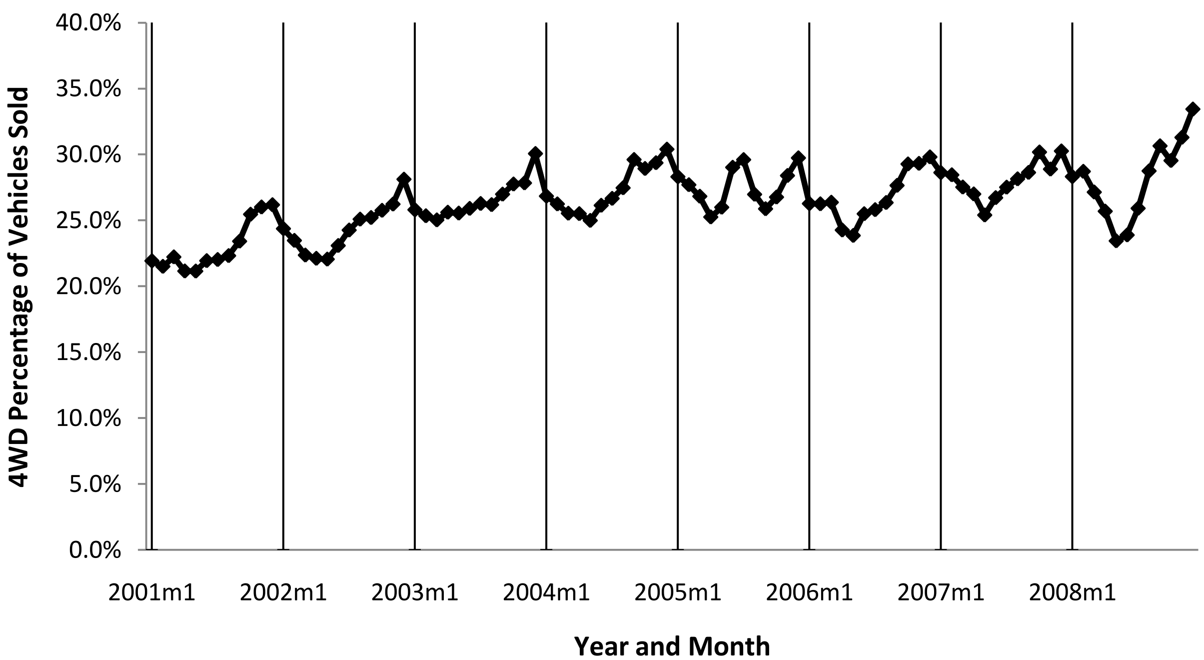

Figure 2 - Seasonal Trends in Vehicle Purchases by Temperature Variation. This figure illustrates the percentage of total vehicles sold between 2001 and 2008 that were convertibles (Panel A) and 4-wheel drives (Panel B) for DMAs with above- and below-median level of monthly DMA temeprature variation.

Panel A. Convertible Percentage of Vehicles Sold by Temperature Variation

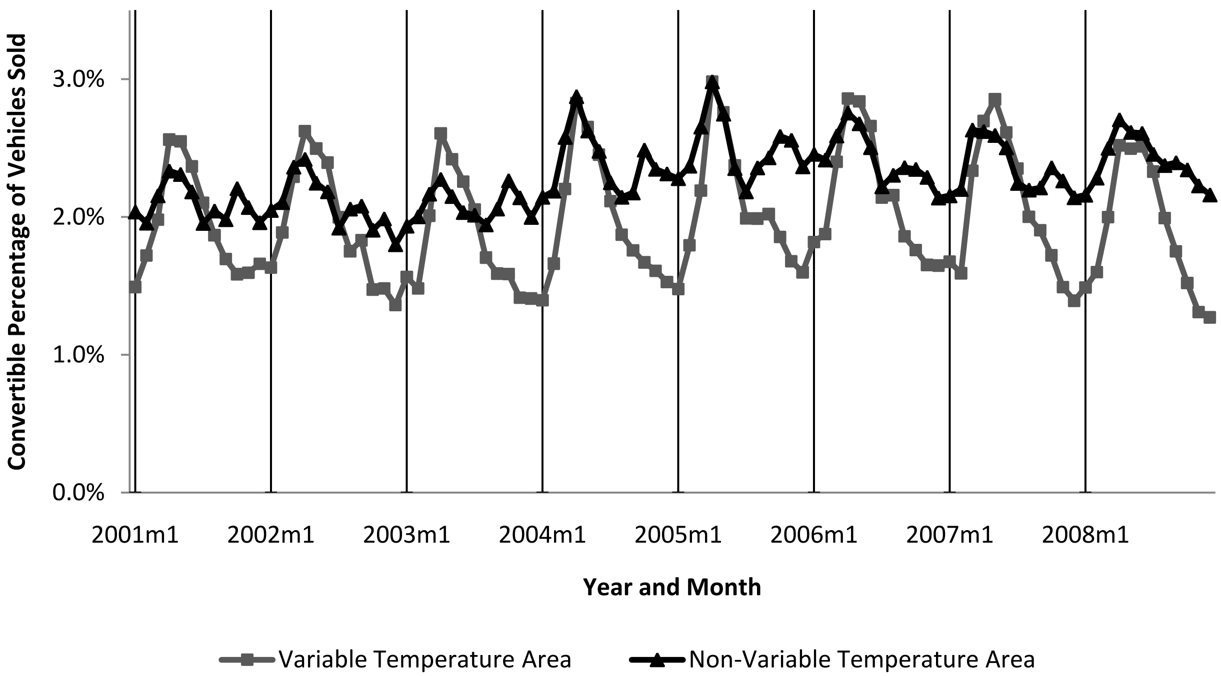

Panel B. 4-Wheel Drive Percentage of Vehicles Sold by Temperature Variation

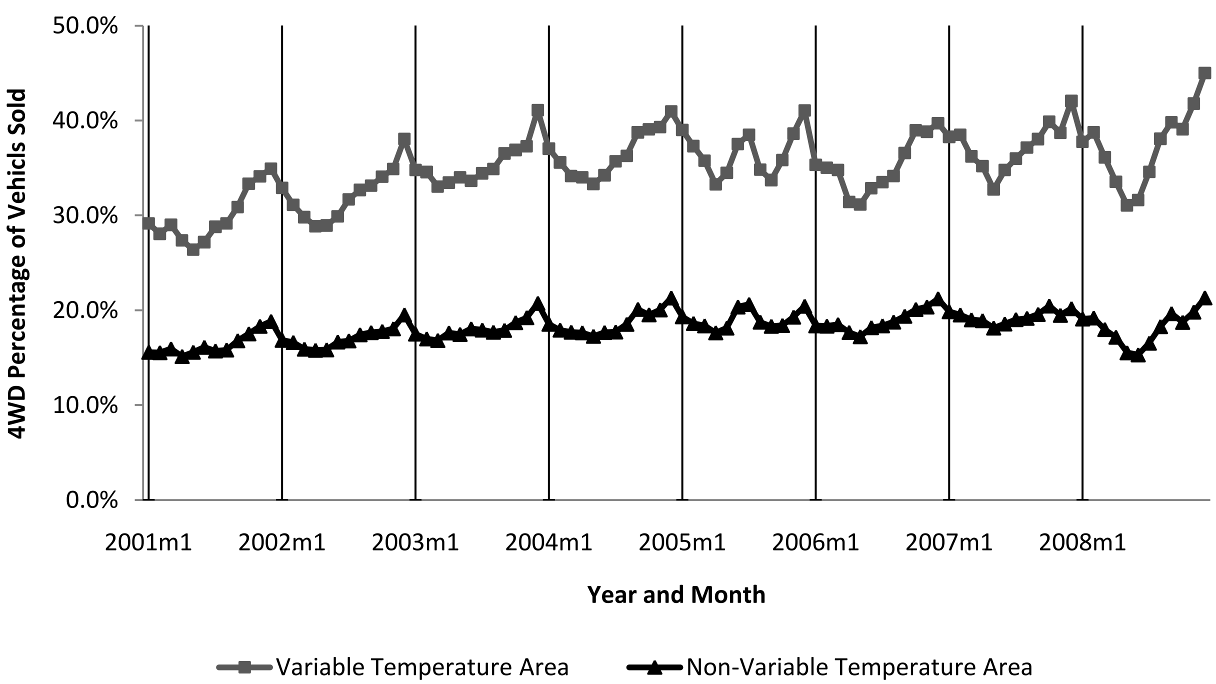

Figure 3. Convertible Sales - Chicago. Panel A illustrates the percentage of vehicles sold in Chicago for each of the 52 weeks in a year that were convertibles. Panel B plots the residual convertible percentage of vehicles sold in each week. (Residual is net of year and week-of-the-year fixed effects.)

Panel A. Convertible Percentage of Vehicles Sold - Chicago

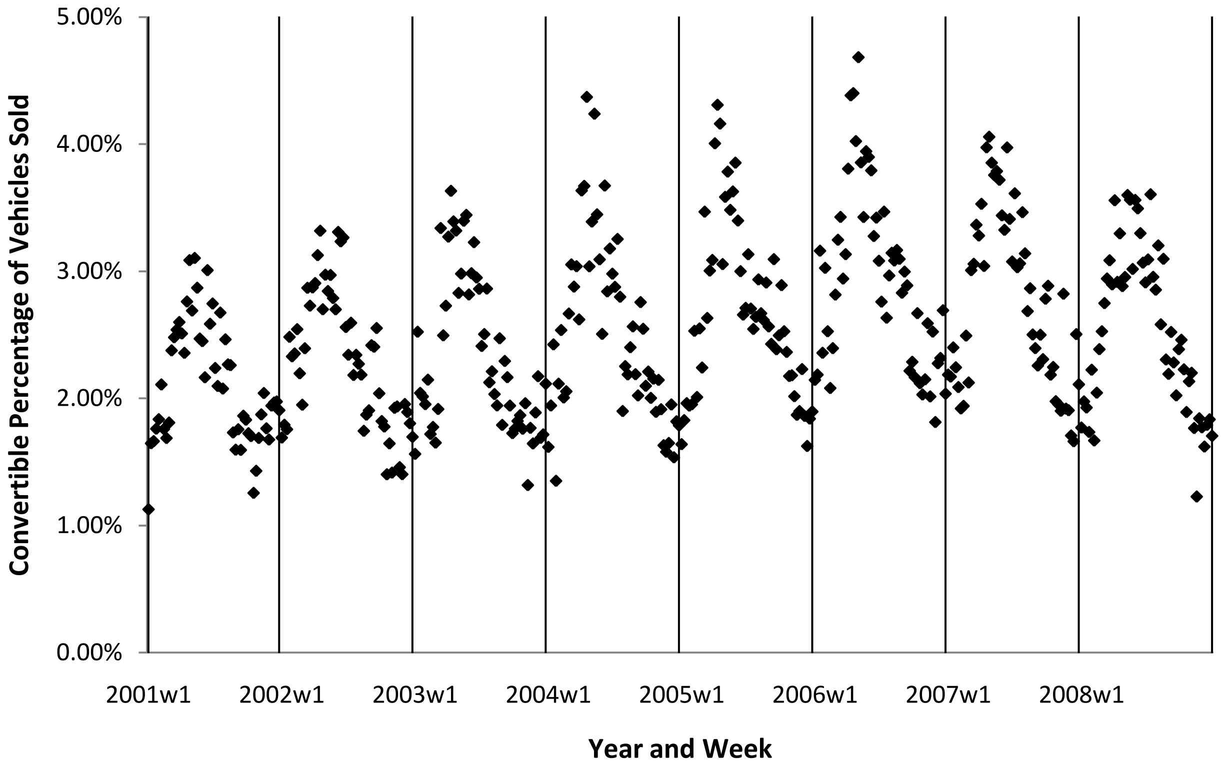

Panel B. Residual Convertible Percentage of Vehicles Sold - Chicago

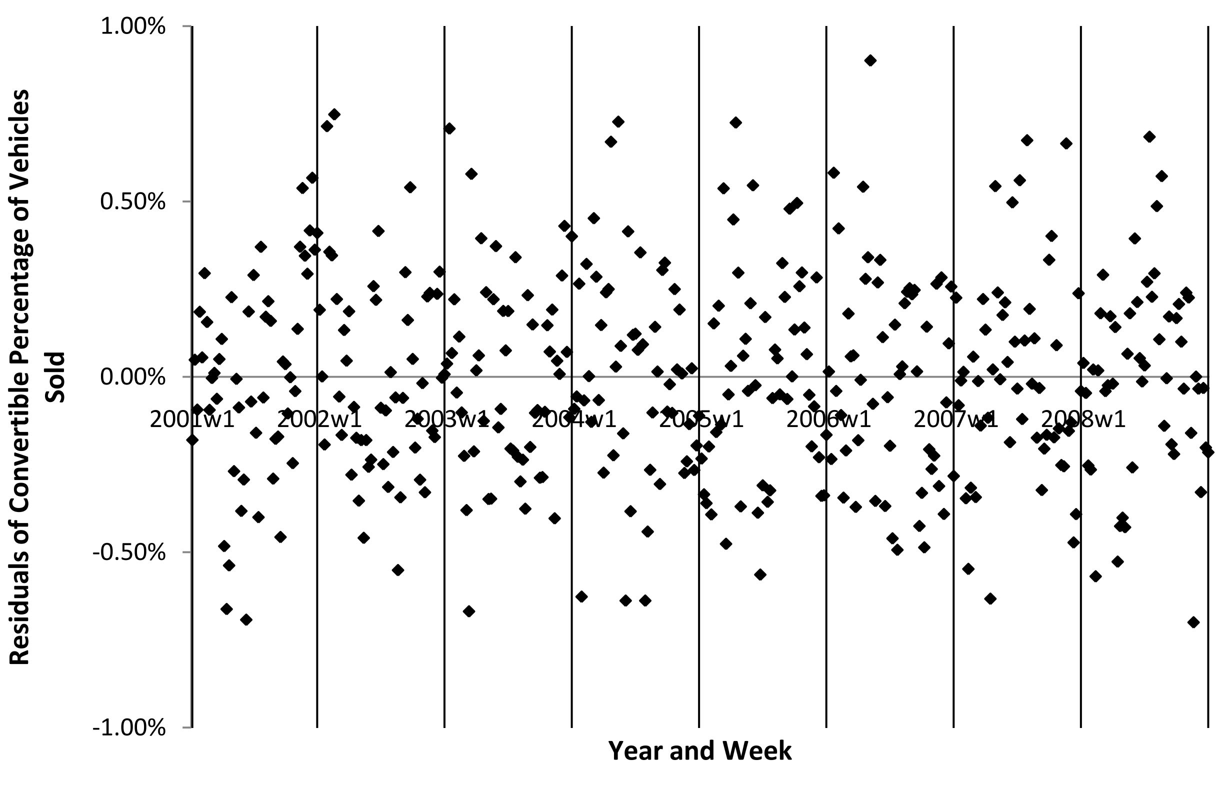

Figure 4. Temperature - Chicago. Panel A illustrates the average daily high temperature in Chicago for each of the 52 weeks in a year. Panel B plots the residual average daily high temperature in each week. (Residual is net of year and week-of-the-year fixed effects.)

Panel A. Mean High Temperature (Fahrenheit) - Chicago

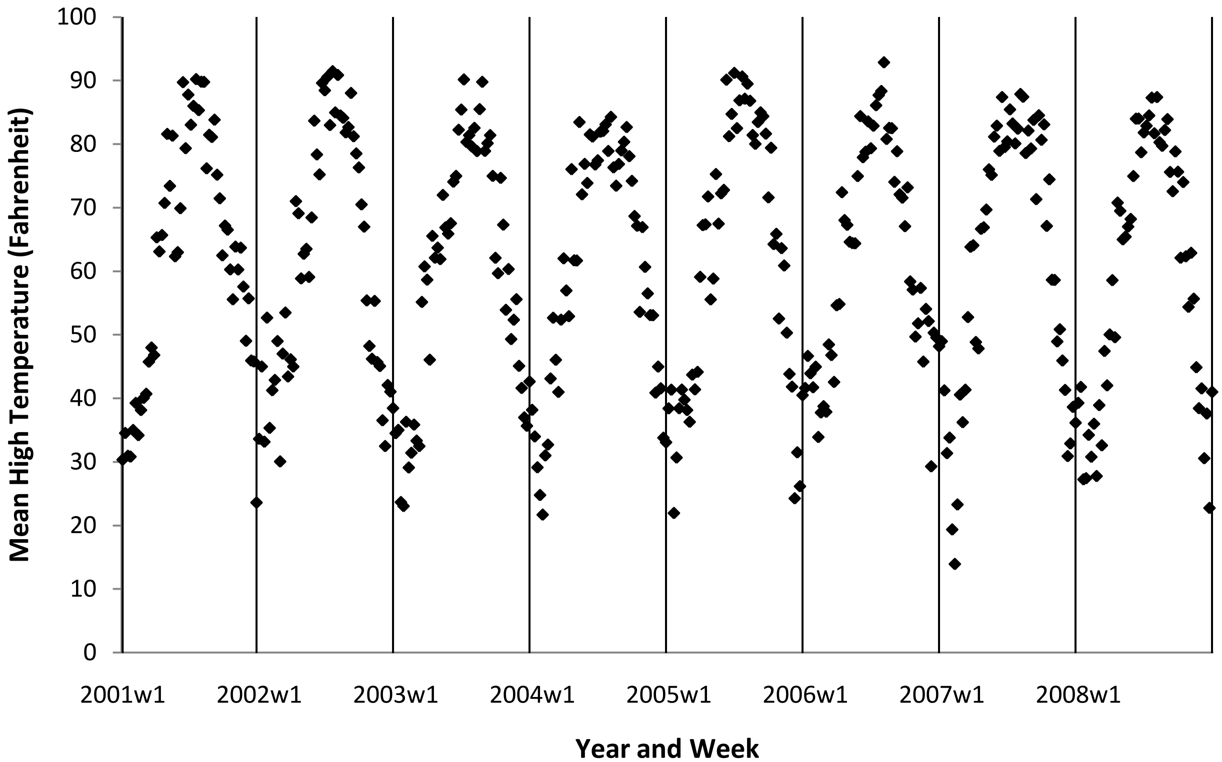

Panel B. Residual Mean High Temperature (Fahrenheit) - Chicago

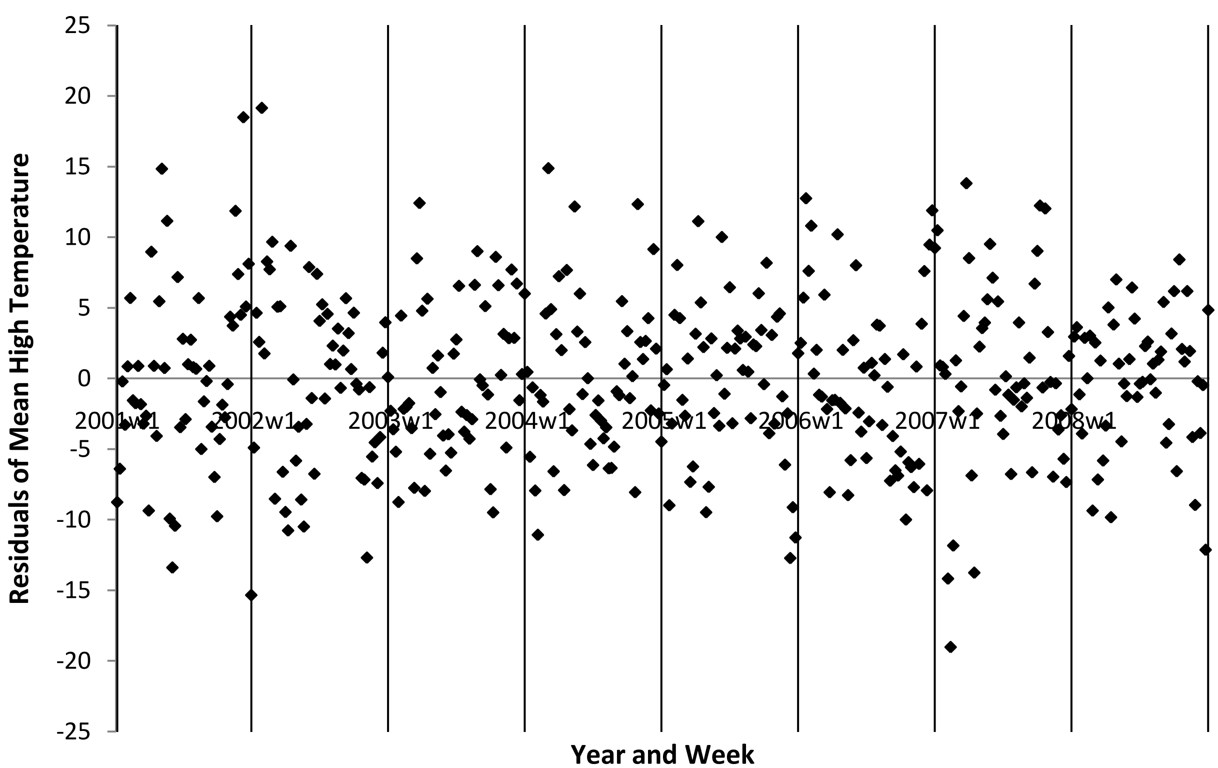

Figure 5. Temperature-Convertible Residuals - Chicago. This Figure provides scatter plots for the residuals of convertible percentage of vehicles sold (Panel B of Figure 3) and residuals of mean high temperature (Panel B of Figure 4) separately for each quarter of the year.

Panel A. Quarter 1

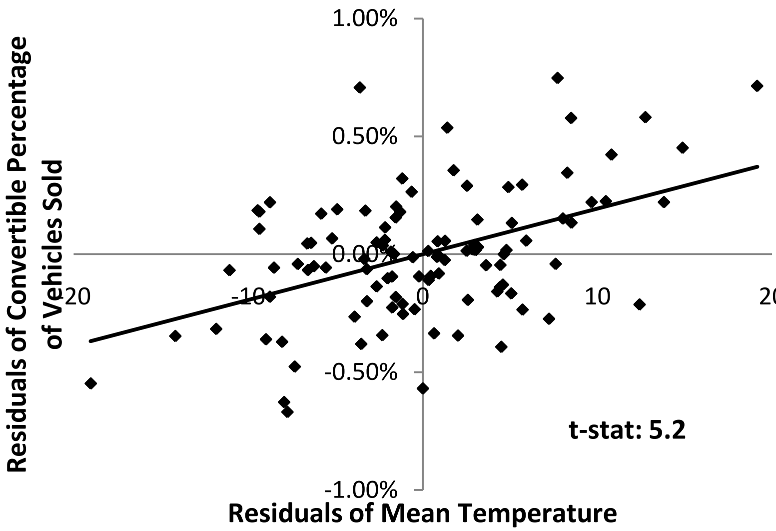

Panel C. Quarter 3

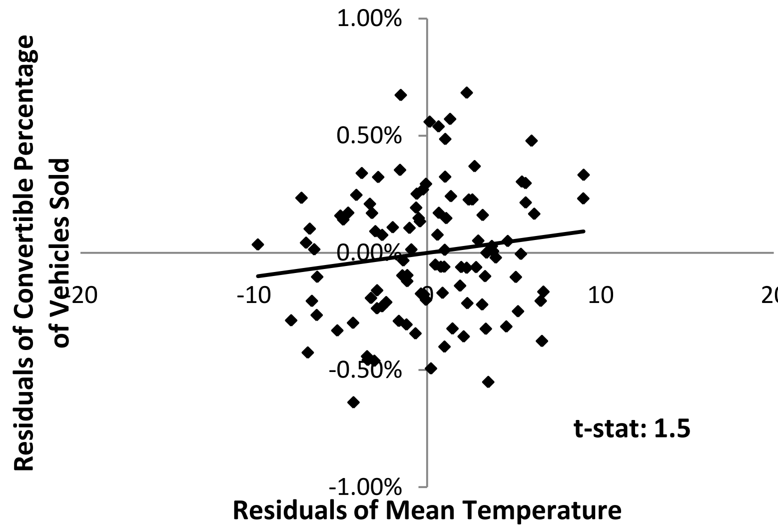

Panel B. Quarter 2

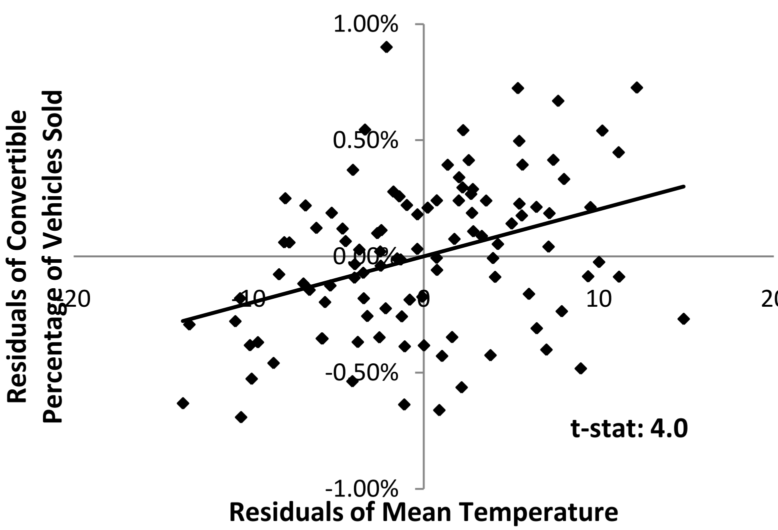

Panel D. Quarter 4

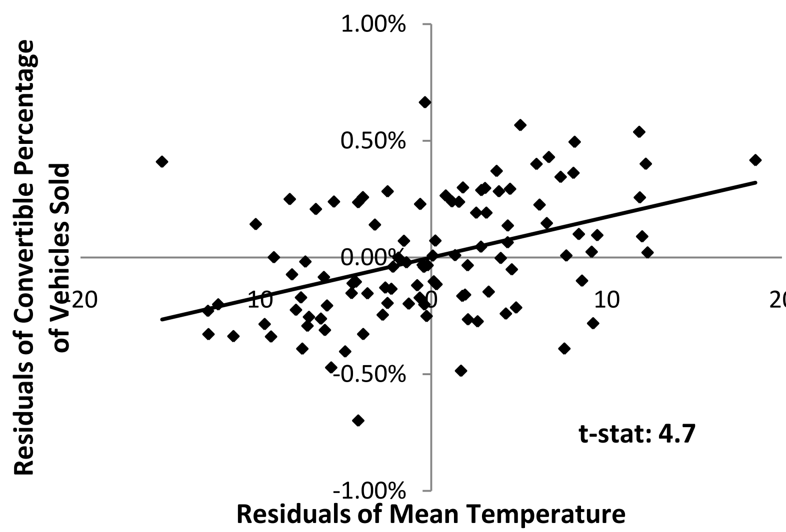

Figure 6. Convertible Sales - Miami. Panel A illustrates the percentage of vehicles sold in Miami-Ft. Lauderdale for each of the 52 weeks in a year that were convertibles. Panel B plots the residual convertible percentage of vehicles sold in each week. (Residual is net of year and week-of-the-year fixed effects.)

Panel A. Convertible Percentage of Vehicles Sold - Miami

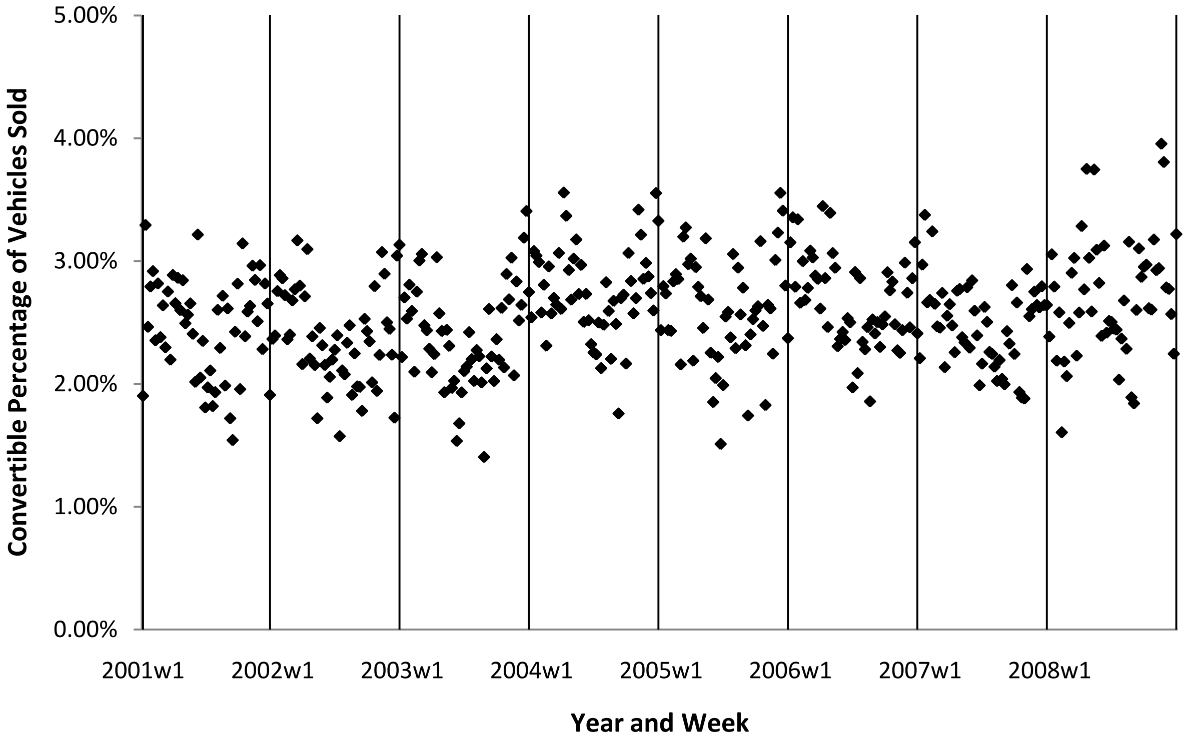

Panel B. Residual Convertible Percentage of Vehicles Sold - Miami

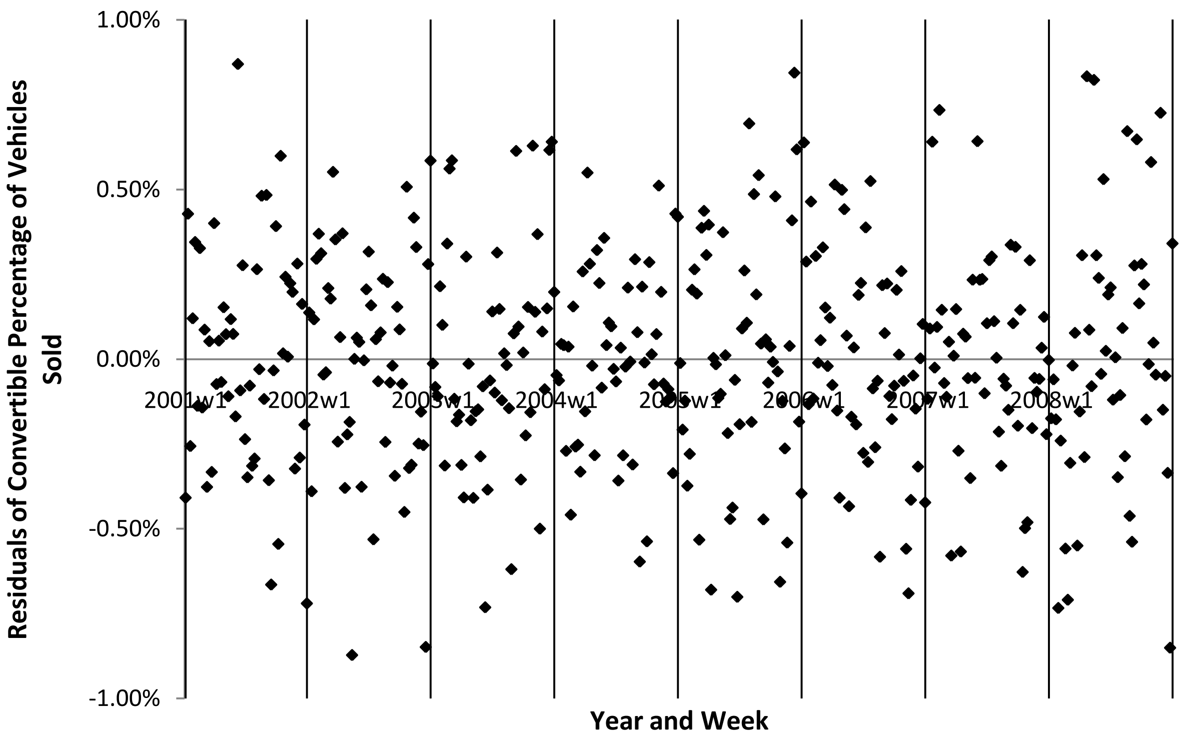

Figure 7. Temperature - Miami. Panel A illustrates the average daily high temperature in Miami-Ft. Lauderdale for each of the 52 weeks in a year. Panel B plots the residual average daily high temperature in each week. (Residual is net of year and week-of-the-year fixed effects.)

Panel A. Mean High Temperature (Fahrenheit) - Miami

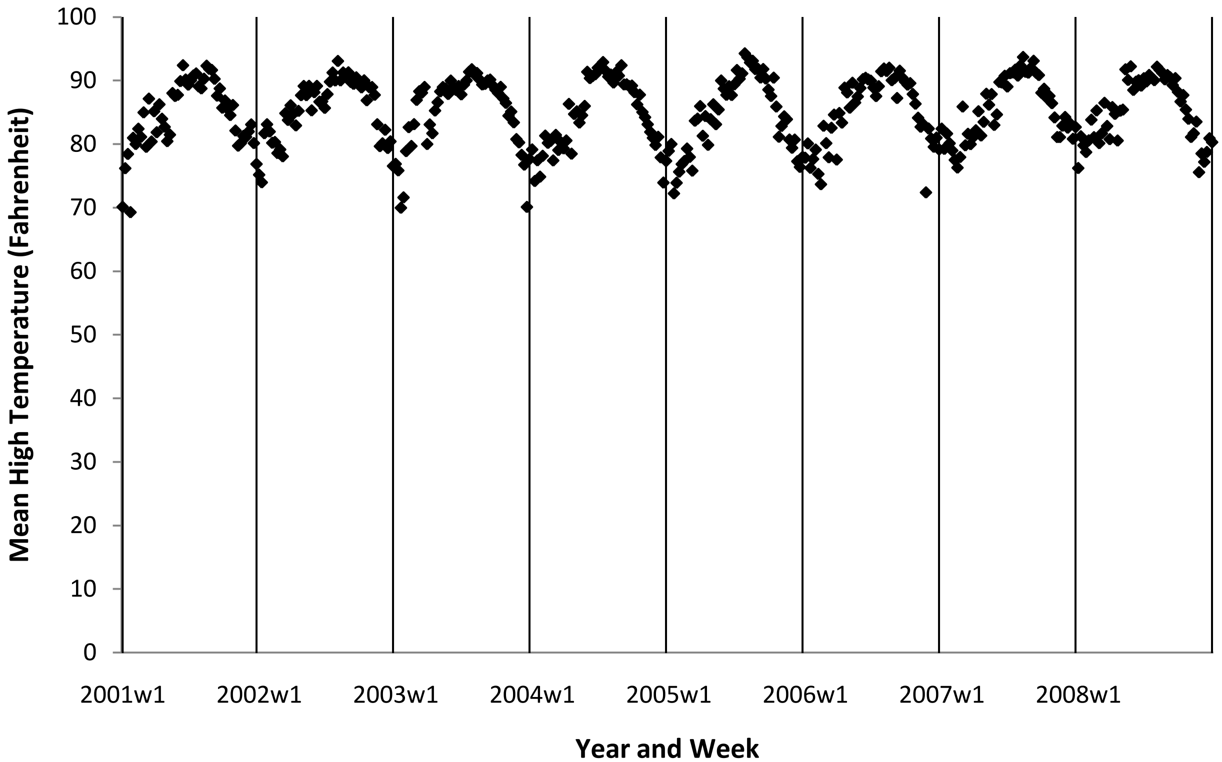

Panel B. Residual Mean High Temperature (Fahrenheit) - Miami

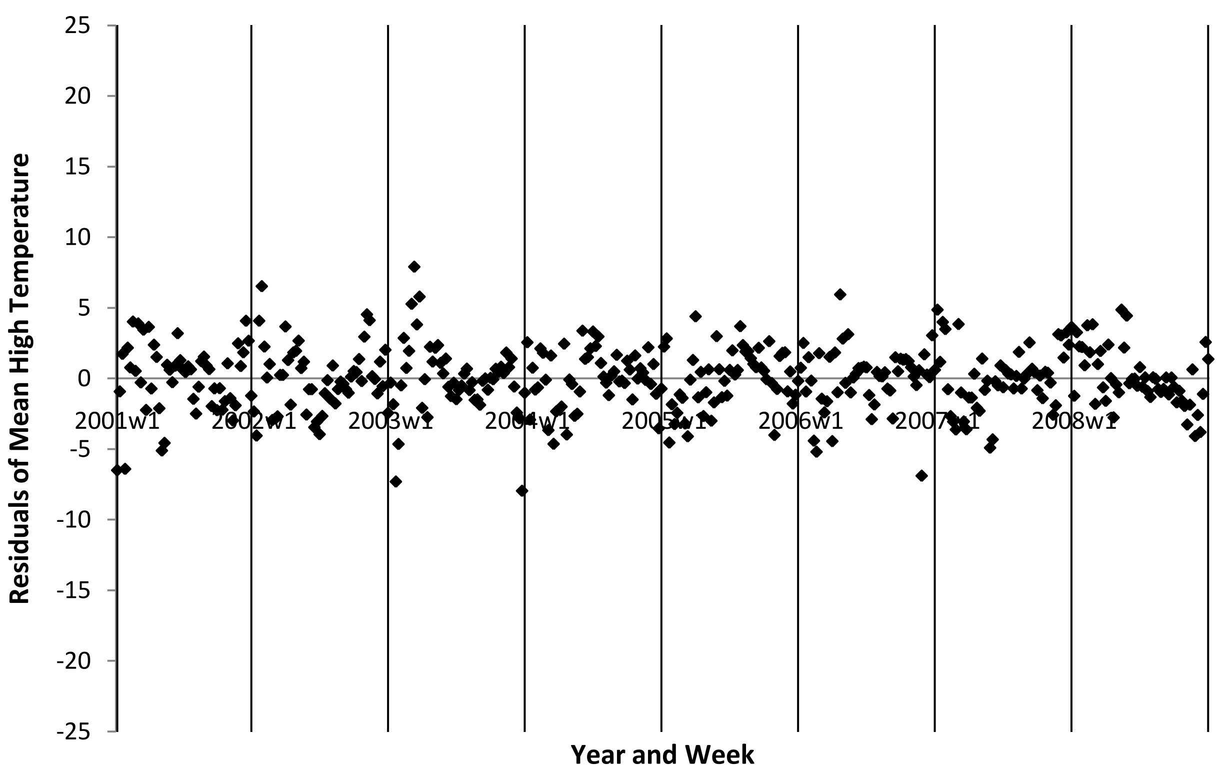

Figure 8. Temperature-Convertible Residuals - Miami. This Figure provides scatter plots for the residuals of convertible percentage of vehicles sold (Panel B of Figure 6) and residuals of mean temperature (Panel B of Figure 7) separately for each quarter of the year.

Panel A. Quarter 1

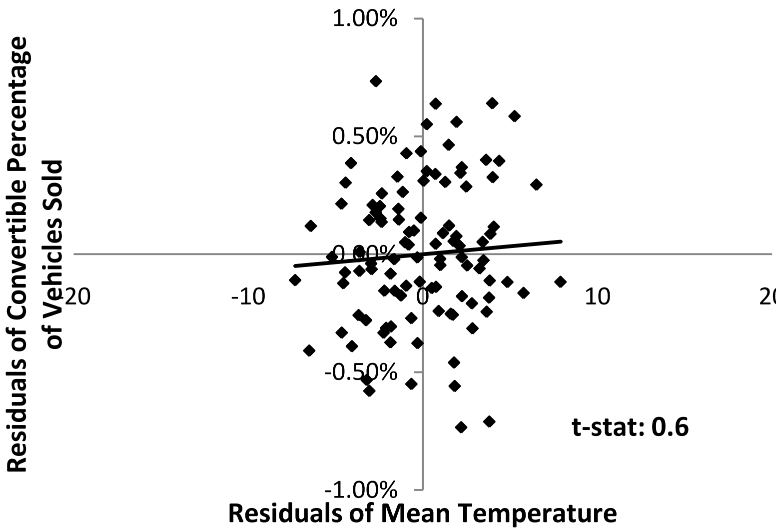

Panel C. Quarter 3

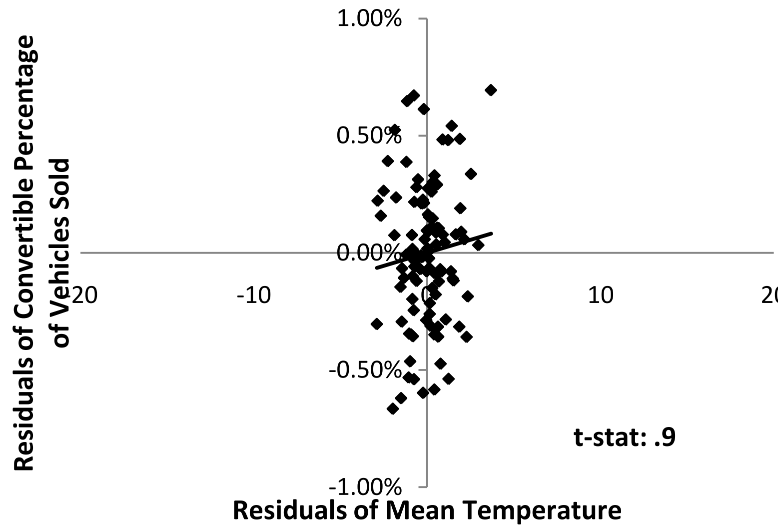

Panel B. Quarter 2

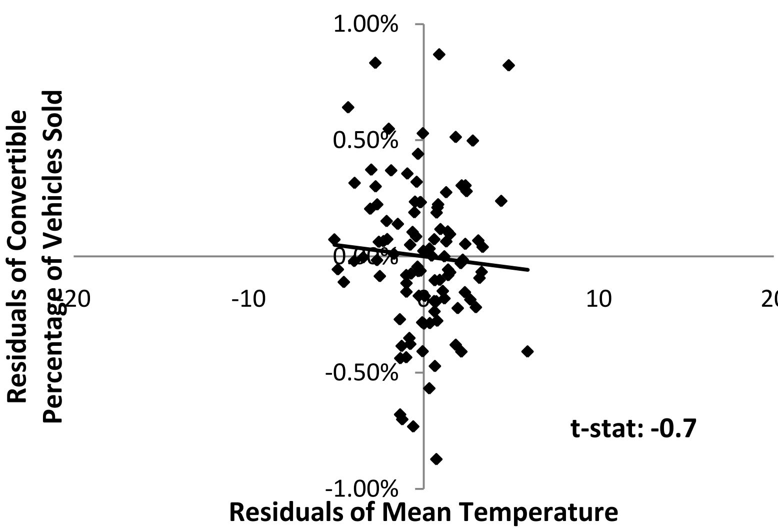

Panel D. Quarter 4

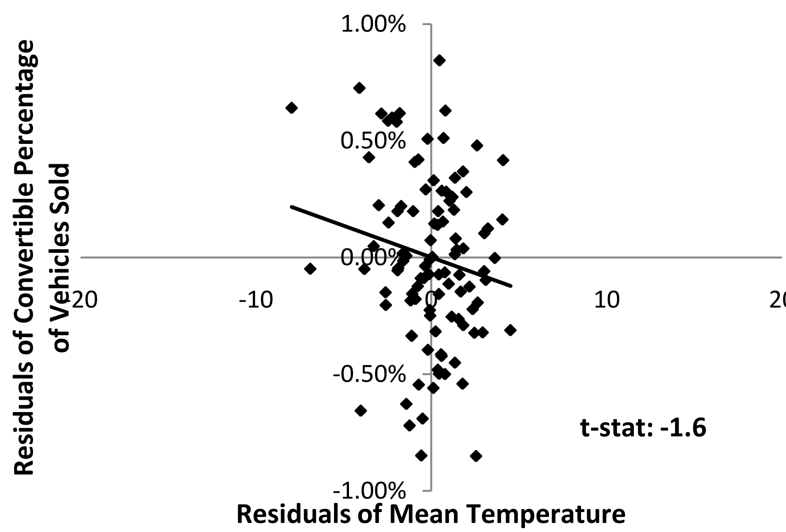

Figure 9. Temperature and Convertible Sales by Usual Temperature. This Figure provides the coefficient values and 95% confidence intervals for the impact of mean daily high temperature on convertible percentage of total vehicles sold (the estimate in Column 1 of Table 1) when the effect is estimated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

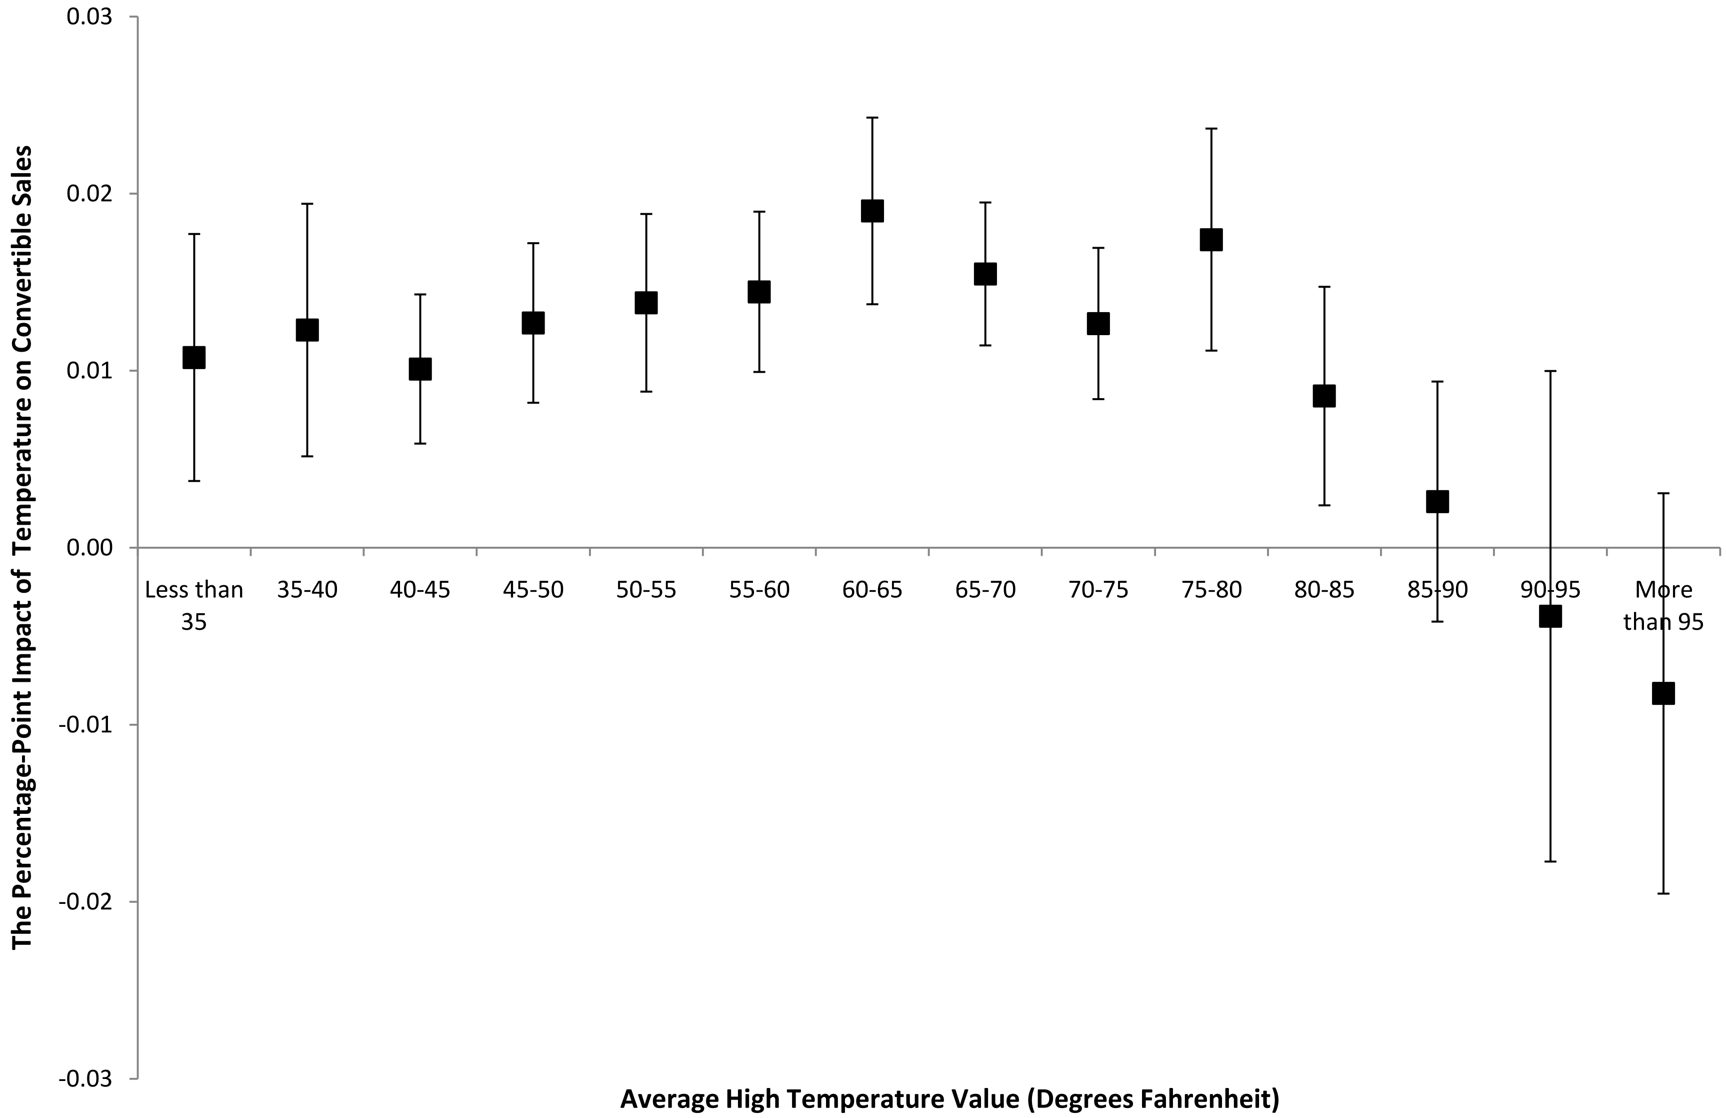

Figure 10. Snowfall and 4-Wheel Drive Sales - Event-Study Design. This Figure plots the weighted average and 95% confidence intervals for the residuals of the 4-wheel drive percentage of total vehicles sold for the twelve weeks leading up to and the twelve weeks after a snow storm event (week 0). The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

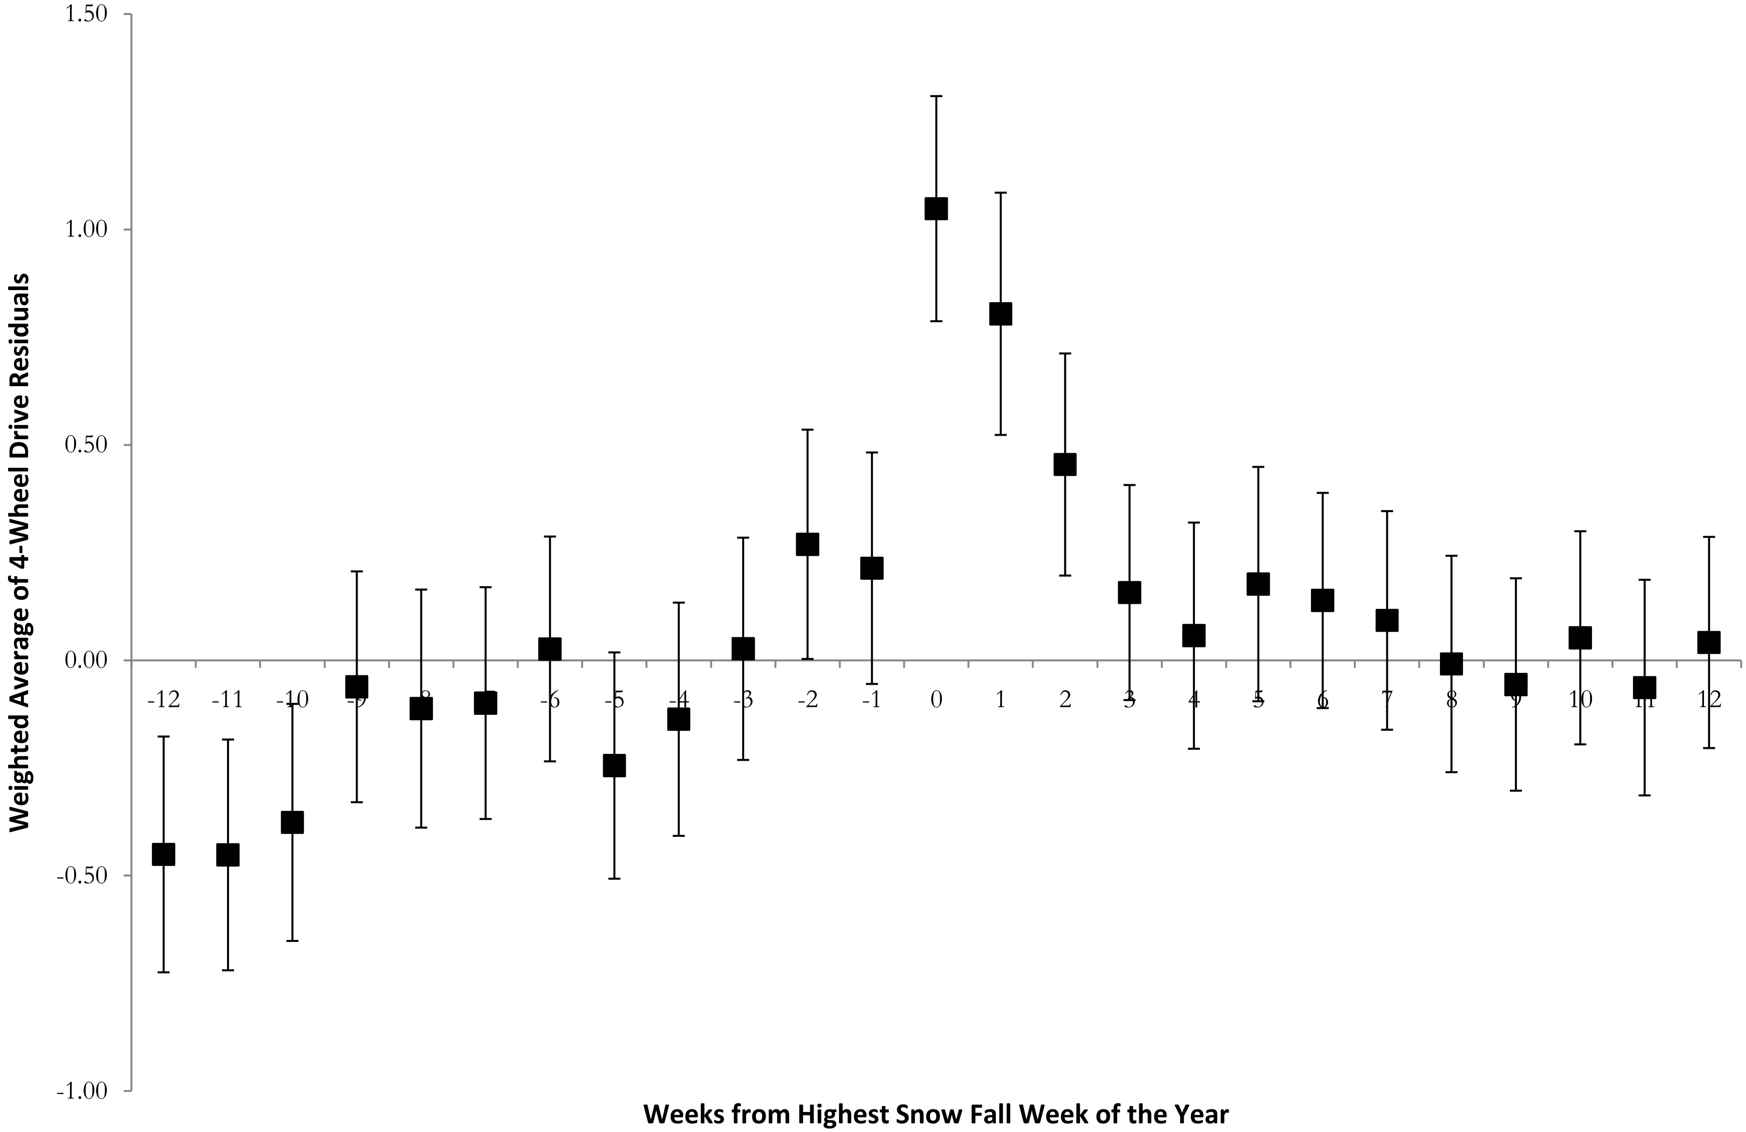

Figure 11 - Seasonal Value of a Swimming Pool. Panel A shows the average residual values for homes with swimming pools that go under contract during each month of the year. Panel B shows the estimated effect of a swimming pool on a house's residual sales price, conditional on other house characteristics, as estimated by Equation (7). 95% confidence intervals are also presented.

Panel A. Residuals by Month

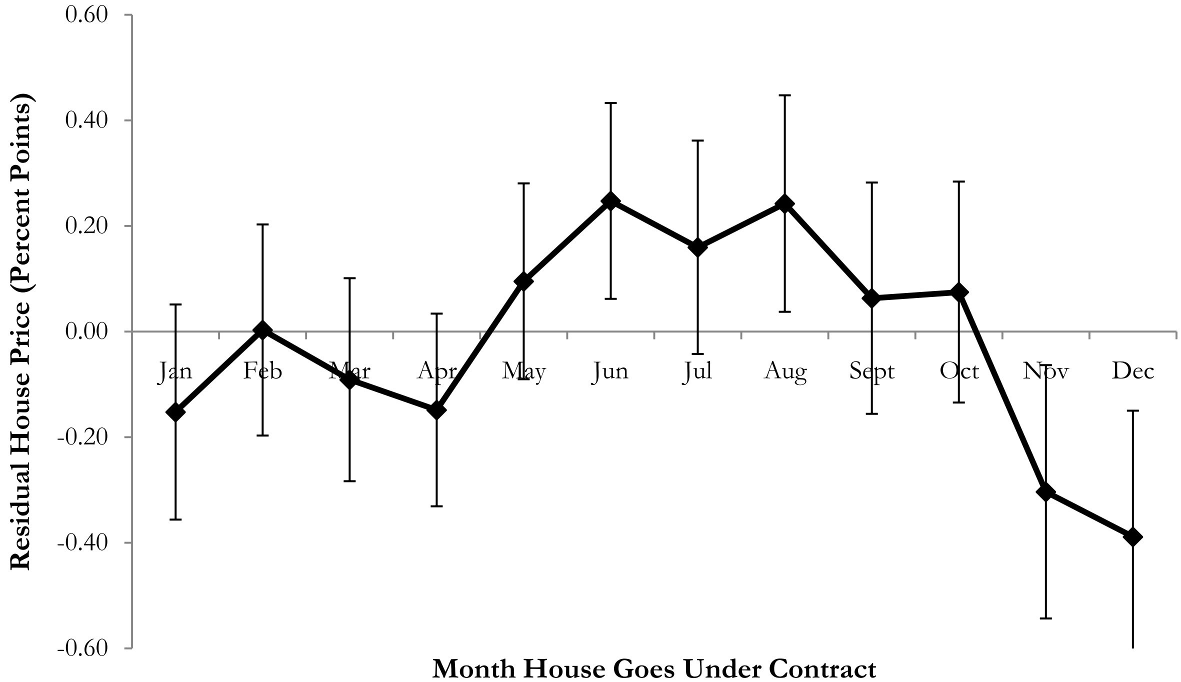

Panel B. Conditional Effect of a Swimming Pool by Month

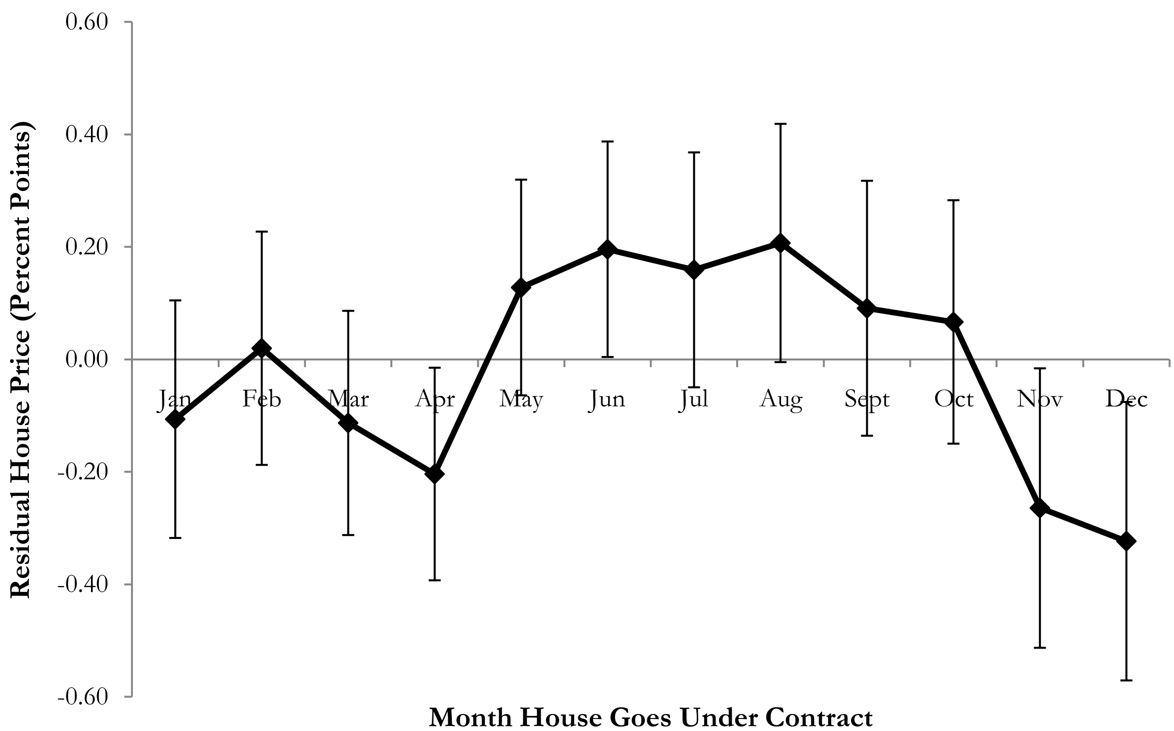

Figure 12 - Seasonal Value of a Swimming Pool - Trimming. Panel A provides the estimated effect of a swimming pool on a house's residual sales price, conditional on other house characteristics, as estimated by Equation (7), after eliminating residuals in the top and bottom 1%. Panel B shows the same estimated effects after eliminating residuals in the top and bottom 5%. 95% confidence intervals are also presented.

Panel A. Conditional Effect of a Swimming Pool by Month - 1% Trim

Panel B. Conditional Effect of a Swimming Pool by Month - 5% Trim

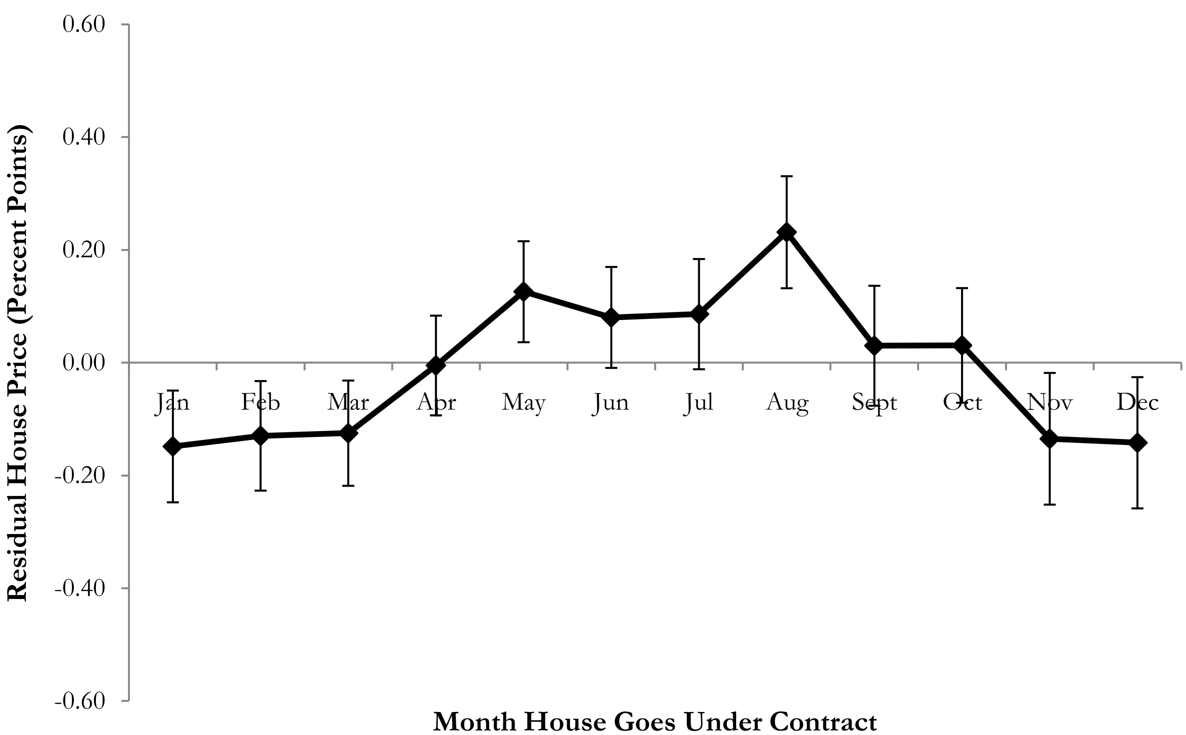

Panel A. Conditional Effect of Central Air by Month

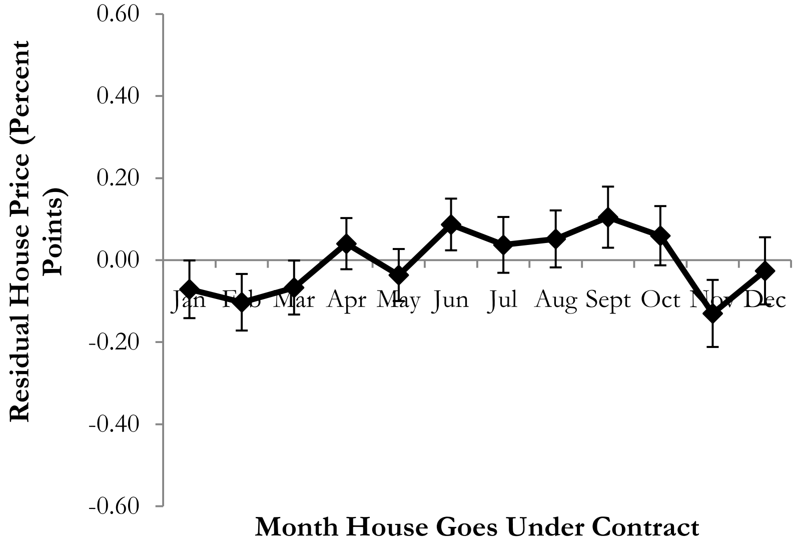

Panel C. Conditional Effect of Lot Size by Month

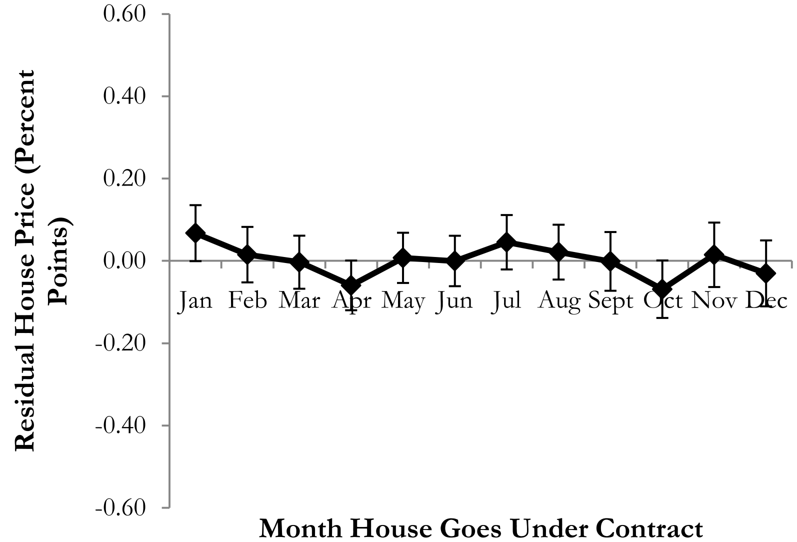

Panel B. conditional Effect of a Fireplace by Month

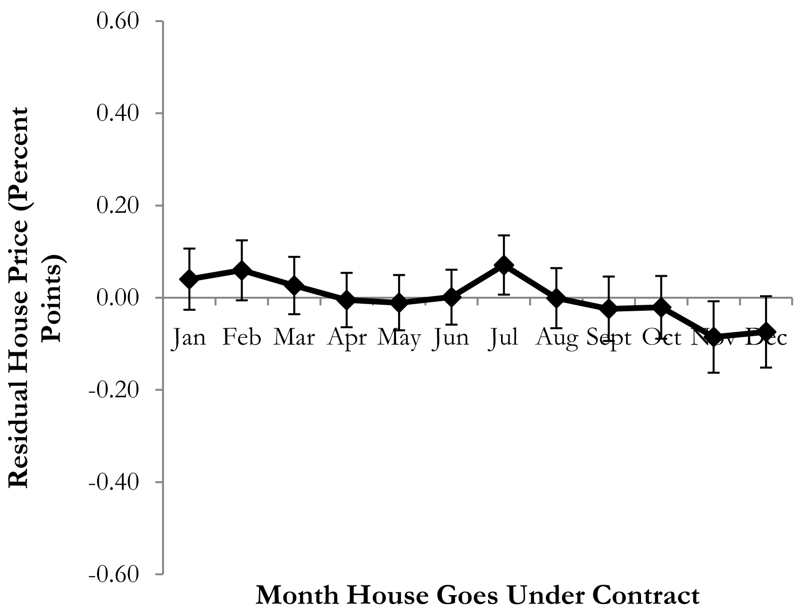

Figure 13 - Seasonal Value of Other Seasonal Housing Characteristics. This figure provides the estimated effect of a various characteristics on a house's residual sales price, conditional on other house characteristics, as estimated by Equation (7). Panel A shows the effect of central air, Panel B the effect of a fireplace, and Panel C the effect of lot size. The top and bottom 5% of residuals are removed. 95% confidence intervals are also presented.

Panel A. Conditional Effect of Number of Bedrooms by Month

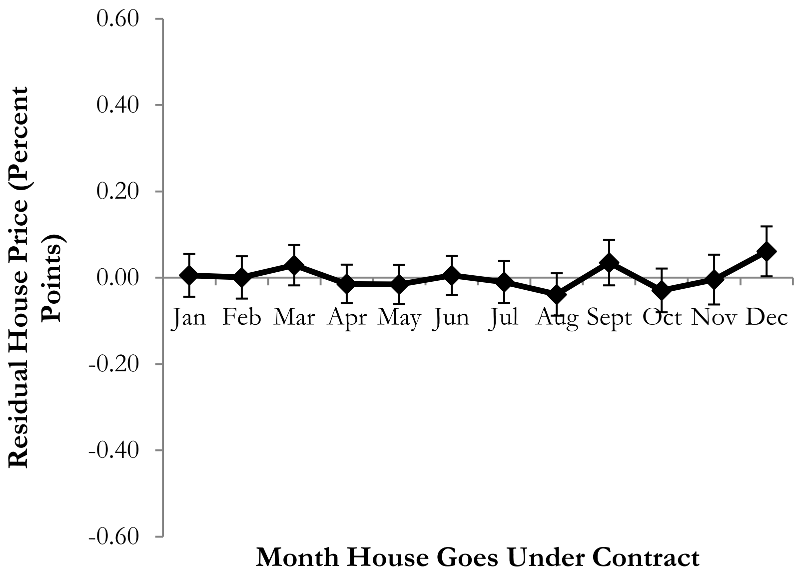

Panel C. Conditional Effect of Square Footage by Month

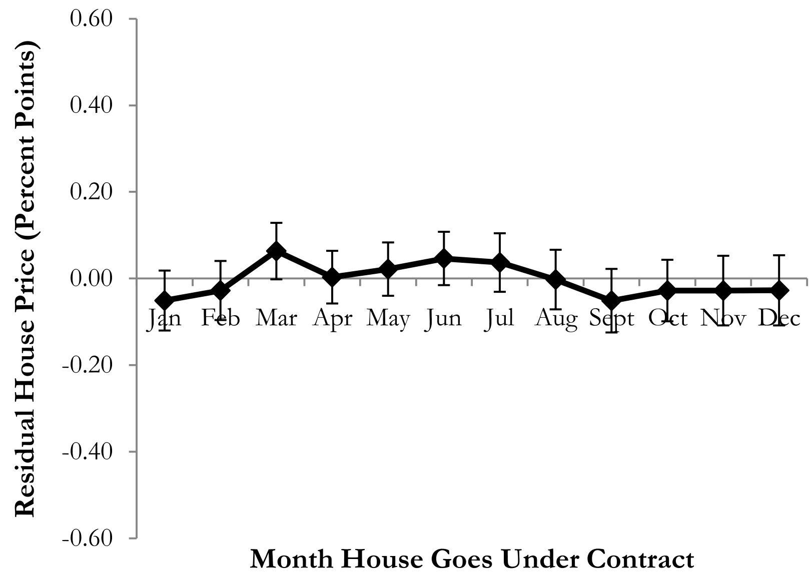

Panel B. Conditional Effect of Number of Bathrooms by Month

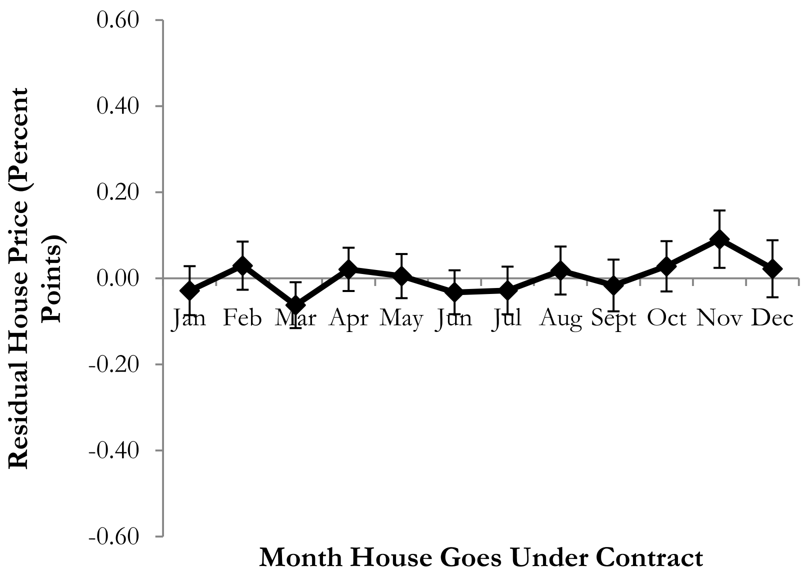

Figure 14 - Seasonal Value of Non-Seasonal Housing Characteristics. This figure provides the estimated effect of a various characteristics on a house's residual sales price, conditional on other house characteristics, as estimated by Equation (7). Panel A shows the effect of number of bedrooms, Panel B the effect of a number of bathrooms, and Panel C the effect of square footage. The top and bottom 5% of residuals are removed. 95% confidence intervals are also presented.

Table 1. Impact of Weather on Convertible Purchases

	Dep. Var.: Convertible Percentage of Total Vehicles Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Temperature	.011**	.014**	.010**	.002	.011**
	(.000)	(.001)	(.002)	(.002)	(.001)
Rain Fall	-.005**	-.017**	-.006	-.000	-.003
	(.002)	(.004)	(.003)	(.003)	(.003)
Snow Fall	-.022	-.006	-.082	-	-.034
	(.024)	(.032)	(.106)	-	(.034)
Slush Fall	-.028**	-.020	-.028	-.026	-.033
	(.009)	(.014)	(.020)	(.026)	(.018)
Cloud Cover	-.172**	-.125*	-.342**	-.171**	-.108*
	(.027)	(.053)	(.057)	(.052)	(.044)
DMA*Year F.E.s	X	X	X	X	X
DMA*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 vehicles sold on weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is a DMA-Week and is weighted by the total number of vehicles sold. Fixed effects are included for DMA*Year and for DMA*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 Weather on 4-Wheel Drive Purchases

	Dep. Var.: 4-Wheel Drive Percentage of Total Vehicles Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Temperature	-.050**	-.069**	-.024**	-.039**	-.063**
	(.002)	(.003)	(.004)	(.007)	(.004)
Rain Fall	.014*	.021	.029**	.031**	-.003
	(.006)	(.014)	(.010)	(.010)	(.012)
Snow Fall	1.02**	.73**	-.18	-8.11	1.18**
	(.05)	(.07)	(.20)	(25.3)	(.08)
Slush Fall	.24**	.24**	.12*	-.14	.45**
	(.02)	(.04)	(.04)	(.09)	(.05)
Cloud Cover	.378**	.351*	1.030**	.265	.405*
	(.082)	(.155)	(.158)	(.186)	(.150)
DMA*Year F.E.s	X	X	X	X	X
DMA*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 vehicles sold on weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is a DMA-Week and is weighted by the total number of vehicles sold. Fixed effects are included for DMA*Year and for DMA*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 Weather on Black Vehicle Purchases

	Dep. Var.: Black Percentage of Total Vehicles Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Temperature	-.013**	-.012**	-.018**	-.006	-.010**
	(.001)	(.002)	(.003)	(.004)	(.002)
Rain Fall	.002	-.003	-.004	-.002	.011
	(.003)	(.007)	(.004)	(.005)	(.006)
Snow Fall	.097**	.102*	.273	-.145	.087
	(.032)	(.045)	(.151)	(.348)	(.050)
Slush Fall	.013	.021	.058*	.049	.013
	(.013)	(.020)	(.028)	(.044)	(.027)
Cloud Cover	.71**	.65**	.67**	.93**	.65**
	(.04)	(.09)	(.09)	(.094)	(.08)
DMA*Year F.E.s	X	X	X	X	X
DMA*Week-of-the-Year F.E.s	X	X	X	X	X
R-Squared	0.812	0.822	0.834	0.842	0.835
Observations	66,219	15,940	16,601	16,803	16,875

Notes: Coefficient values and standard errors are presented from OLS regressions of the percentage of total vehicles sold that are black in color on weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is a DMA*Week and is weighted by the total number of vehicles sold. Fixed effects are included for DMA*Year and for DMA*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 Weather on Convertible Purchases - Dynamic Analysis

	Dep. Var.: Convertible Percentage of Total Vehicles Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Temperature Lead 1	.001	.003	.001	.005*	-.000
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature	.011**	.015**	.006**	.004	.010**
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 1	.001	.005**	-.005*	.001	.003
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 2	.003	.007**	-.001	.003	.000
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 3	.001	.002	-.001	.007**	-.001
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 4	.001	.002	-.003	.005	.001
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 5	-.001	.000	-.003	.003	-.001
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 6	.002*	-.001	.001	.010**	.004
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 7	.002*	.004**	.001	.002	-.003
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 8	.004**	.000	.006**	.008**	.003
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 9	.003**	.002	.000	.004	.004
	(.001)	(.002)	(.002)	(.002)	(.002)
Temperature Lag 10	-.000	.001	-.002	.004*	-.003
	(.001)	(.002)	(.001)	(.002)	(.003)
Temperature Lag 11	.000	-.002	-.000	.004	.007*
	(.001)	(.002)	(.001)	(.002)	(.003)
Temperature Lag 12	.000	.001	-.004**	.005*	.000
	(.001)	(.002)	(.001)	(.002)	(.003)
DMA*Year F.E.s	X	X	X	X	X
DMA*Week-of-the-Year F.E.s	X	X	X	X	X
Rain Fall (with Lead and Lags)	X	X	X	X	X
Snow Fall (with Lead and Lags)	X	X	X	X	X
Slush Fall (with Lead and Lags)	X	X	X	X	X
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 vehicles sold on weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Both the current weather as well as the lead and 12 weekly lag 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 a DMA-Week and is weighted by the total number of vehicles sold. Fixed effects are included for DMA*Year and for DMA*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 Weather on 4-Wheel Drive Purchases - Dynamic Analysis

	Dep. Var.: 4-Wheel Drive Percentage of Total Vehicles Sold				
	Full Year	**Quarter 1**	**Quarter 2**	**Quarter 3**	**Quarter 4**
Snow Fall Lead 1	.18**	-.01	-.90	3.22**	.27**
	(.06)	(.06)	(.29)	(1.2)	(.09)
Snow Fall	1.01**	.71**	-.19	22.9	1.22**
	(.06)	(.09)	(.24)	(81.3)	(.10)
Snow Fall Lag 1	.85**	.60**	.04	21.6	1.29**
	(.06)	(.08)	(.23)	(121.5)	(.11)
Snow Fall Lag 2	.26**	.19*	.37	-247.8	.74**
	(.06)	(.07)	(.21)	(193.4)	(.13)
Snow Fall Lag 3	.12	.26**	-.04	147.4	.40**
	(.06)	(.08)	(.15)	(432.5)	(.14)
Snow Fall Lag 4	-.09	.01	-.39**	-250.9*	.26
	(.07)	(.08)	(.14)	(117.2)	(.17)
Snow Fall Lag 5	-.14*	-.07	-.35**	-156.3	.05
	(.07)	(.08)	(.12)	(116.2)	(.22)
Snow Fall Lag 6	-.26**	-.17	-.31**	-320.3**	-.07
	(.06)	(.08)	(.10)	(105.6)	(.39)
Snow Fall Lag 7	-.15*	-.15	-.05	-71.1*	-.59
	(.06)	(.08)	(.09)	(34.8)	(.52)
Snow Fall Lag 8	-.09	-.02	-.12	-39.4	1.58*
	(.06)	(.09)	(.09)	(32.6)	(.65)
Snow Fall Lag 9	-.09	-.05	-.11	-1.4	1.65*
	(.06)	(.09)	(.08)	(1.9)	(.74)
Snow Fall Lag 10	-.09	-.18	-.12	1.4	1.23
	(.06)	(.09)	(.08)	(.87)	(.77)
Snow Fall Lag 11	-.09	.05	-.25**	.01	-.33
	(.06)	(.10)	(.07)	(.75)	(.76)
Snow Fall Lag 12	-.13*	-.01	-.30**	.41	-1.68
	(.06)	(.09)	(.08)	(.29)	(1.2)
DMA*Year F.E.s	X	X	X	X	X
DMA*Week-of-the-Year F.E.s	X	X	X	X	X
Temperature (with Lead and Lags)	X	X	X	X	X
Rain Fall (with Lead and Lags)	X	X	X	X	X
Slush Fall (with Lead and Lags)	X	X	X	X	X
Cloud Cover (with Lead and Lags)	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 vehicles sold on weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Both the current weather as well as the lead and 12 weekly lag 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 a DMA-Week and is weighted by the total number of vehicles sold. Fixed effects are included for DMA*Year and for DMA*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 6. Impact of 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 Vehicles Sold	
	Convertibles	**4-Wheel Drives**
Temperature	.060**	-.044**
	(.020)	(.004)
Rain Fall	.007	.003
	(.042)	(.011)
Snow Fall	-.23	.61**
	(.60)	(.09)
Slush Fall	-.45	.19**
	(.24)	(.04)
Cloud Cover	-1.26	.89**
	(.679)	(.15)
DMA*Year F.E.s	X	X
DMA*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 vehicles sold (Column 2) on weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is a DMA-Week and is weighted by the total number of vehicles sold. Fixed effects are included for DMA*Year and for DMA*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 7. Impact of 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:			
Temperature	.006%	.017%**	.006%
	(.004%)	(.007%)	(.009%)
Rain Fall	.008%	.002%	-.018%
	(.009%)	(.015%)	(.021%)
Snow Fall	.181%	-.041%	-.142%
	(.131%)	(.222%)	(.289%)
Slush Fall	.063%	.028%	-.116%
	(.053%)	(.094%)	(.131%)
Cloud Cover	-.197%	-.036%	.332%
	(.138%)	(.228%)	(.312%)
4-Wheel Drive	.285%**	.929%**	1.634%**
	(.006%)	(.006%)	(.014%)
4-Wheel Drive Interacted With:			
Temperature	-.003%*	-.005%*	-.013%**
	(.001%)	(.002%)	(.003%)
Rain Fall	-.005%	-.005%	.001%
	(.003%)	(.006%)	(.008%)
Snow Fall	.000%	.063%	.004%
	(.035%)	(.058%)	(.076%)
Slush Fall	.002%	-.019%	-.048%
	(.016%)	(.028%)	(.038%)
Cloud Cover	.006%	-.109%	-.124%
	(.047%)	(.078%)	(.106%)
DMA*Week 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 vehicle types and 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 individual vehicle level and DMA*Week 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 8. Impact of Weather on Convertible and 4-Wheel Drive Purchase Price

	Dependent Variable: Vehicle Sales Price (Less Rebate)			
	Convertibles		4-Wheel Drives	
	New	Used	New	Used
Mean of Dependent Variable	$40,001	$22,222	$31,845	$19,132
Temperature	1.22	3.98**	.83*	2.03**
	(1.46)	(1.50)	(.38)	(.29)
Rain Fall	1.56	2.67	-.05	.74
	(3.16)	(3.40)	(.98)	(.80)
Snow Fall	50.68	-114.48*	3.33	-23.80**
	(46.02)	(46.45)	(8.77)	(6.91)
Slush Fall	-6.92	-30.39	-5.33	5.03
	(17.83)	(17.90)	(4.16)	(3.11)
Cloud Cover	25.24	-64.52	36.44**	-54.06**
	(49.69)	(51.68)	(13.96)	(11.04)
DMA*Year F.E.s	X	X	X	X
DMA*Week-of-the-Year F.E.s	X	X	X	X
Purchase Timing F.E.s	X	X	X	X
Vehicle-type F.E.s	X	X	X	X
Odometer Value Spline		X		X
Observations	391,438	377,321	5,495,657	4,152,489

Notes: Coefficient values and standard errors are presented from OLS regressions of vehicle transaction prices on weather variables - temperature (degrees Fahrenheit), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Each observation is an individual transaction. Fixed effects are included for DMA*Year and for DMA*Week-of-the-Year (Week 1 - Week 52), and for detailed vehicle types. Purchase Timing indicates whether a vehicle was purchased on a weekend or at the end of the month. The first two columns present results for new and used convertibles, respectively, while the second two columns present results for new and used 4-wheel drives. The used vehicle specifications (columns 2 and 4) include an linear spline in odometer values with knots at 10,000 mile increments.

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

Table 9. Summary Statistics for Retail Vehicle Sales, by DMA*Week

	Mean	St. Dev.	Min	Max
Vehicle Characteristics				
Number of Convertibles Sold	12.4	25.4	0	287
4-Wheel Drives Sold	153.0	292.3	0	6220
Total Vehicles Sold	574.9	1029.9	1	11633
Percentage Convertibles	1.8%	2.6%	0%	100%
Percentage 4-Wheel Drives	30.2%	18.5%	0%	100%
Percentage Black Vehicles	11.2%	6.2%	0%	100%
Weather Variables				
Temperature	70.1	18.3	-26.1	115.8
Rain Fall	1.4	2.4	0	52.9
Snow Fall	.04	.23	0	10.5
Slush Fall	.10	.48	0	22.0
Cloud Cover	.47	.23	0	1
Observations	70,790	70,790	70,790	70,790

Notes: Summary statistics reported for DMA*Week observations.

Table 10. 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 11. 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	.013**	.010**	.23**	.18**	.27**	.13**	.41	.37**
	(.004)	(.002)	(.06)	(.03)	(.09)	(.04)	(.30)	(.14)
Fire Place	.0006	.0013	-.01	.02	.06	.02	.18	-.12
	(.0024)	(.0011)	(.05)	(.02)	(.07)	(.04)	(.24)	(.11)
Lot Acre	.0006	.0006	-.06	-.05*	.00	-.02	-.68*	-.28
	(.0024)	(.0012)	(.05)	(.02)	(.09)	(.04)	(.29)	(.15)
Central Air	.0002	.0060**	.02	.10**	.02	.03	-.13	-.23
	(.0025)	(.0012)	(.05)	(.02)	(.07)	(.04)	(.25)	(.12)
Square Footage (1,000s)	.0043	.0004	.10*	.04	.09	-.03	.42	.22
	(.0025)	(.0012)	(.05)	(.02)	(.08)	(.04)	(.27)	(.13)
Number of Baths	-.0034	-.0008	-.03	-.03	-.05	.03	.07	.09
	(.0020)	(.0010)	(.04)	(.02)	(.07)	(.03)	(.25)	(.12)
Number of Bedrooms	.0004	-.0009	.01	-.02	.02	.00	-.34	-.09
	(.0018)	(.0009)	(.03)	(.02)	(.06)	(.03)	(.17)	(.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 (6) 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. All coefficients are multiplied by 100 to make them easier to read (see text). Thus, the coefficients can be interpretted as approximate percentage-point changes.

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