

THE PSYCHOLOGICAL EFFECT OF WEATHER ON CAR PURCHASES*

MEGHAN R. BUSSE
DEVIN G. POPE
JAREN C. POPE
JORGE SILVA-RISSO

When buying durable goods, consumers must forecast how much utility they will derive from future consumption, including consumption in different states of the world. This can be complicated for consumers because making intertemporal evaluations may expose them to a variety of psychological biases such as present bias, projection bias, and salience effects. We investigate whether consumers are affected by such intertemporal biases when they purchase automobiles. Using data for more than 40 million vehicle transactions, we explore the impact of weather on purchasing decisions. We find that the choice to purchase a convertible or a four-wheel-drive is highly dependent on the weather at the time of purchase in a way that is inconsistent with classical utility theory. We consider a range of rational explanations for the empirical effects we find, but none can explain fully the effects we estimate. We then discuss and explore projection bias and salience as two primary psychological mechanisms that are consistent with our results. *JEL* Codes: D03; D12.

I. INTRODUCTION

People make many decisions that require them to evaluate not only current benefits and costs but also future utility. For example, choosing a job, deciding where to live, planning a vacation, deciding whether to have a baby, and purchasing a durable good are all important life decisions that require an individual to think about utility that will accrue in the future. The standard economic model assumes that individuals are able to accurately

*We are grateful to Chris Bruegge and Ezra Karger for valuable research assistance. We also thank Stefano DellaVigna, Glenn Ellison, Emir Kamenica, Ulrike Malmendier, Ted O'Donoghue, Loren Pope, Joe Price, Mathew Rabin, Chad Syverson, Dick Thaler, Florian Zettelmeyer, and seminar participants at the Behavioral Economics Annual Meeting, the ISMS Marketing Science Conference, the NBER Summer Institute, the Psychonomic Society Annual Meeting, Brigham Young University, Columbia University, Cornell University, DePaul University, Harvard Business School, Indiana University, MIT, Simon Fraser University, Stanford University, UC Berkeley, UC Los Angeles, UC San Diego, UC Santa Barbara, University of British Columbia, University of Chicago, University of Illinois, University of Kentucky, University of Zurich, the Wharton School, and Yale University for helpful suggestions. We are also very grateful for helpful input from the editors and three anonymous referees.

© The Author(s) 2014. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2015), 371–414. doi:10.1093/qje/qju033.
Advance Access publication on November 20, 2014.

estimate future benefits and costs and thereby make decisions that maximize intertemporal utility. Evidence from psychology, however, suggests that individuals may make systematic errors when making intertemporal decisions. Cautions against such systematic errors are contained in the familiar advice to never shop on an empty stomach, to sleep on it before making an important decision, and to decide what you are going to buy before walking into the store. Recent psychological models such as present-biased preferences (Laibson 1997; O'Donoghue and Rabin 1999), projection bias (Loewenstein, O'Donoghue, and Rabin 2003), salience (Bordalo, Gennaioli, and Shleifer 2013), and others provide underlying mechanisms for why consumers make decisions that are too heavily influenced by their mental and/or emotional state at the time of the decision.

In this article, we test whether consumers are overly influenced by conditions at the time of purchase in one particular high-stakes environment: the car market. Since vehicles are durable goods, consumers must predict at the time of purchase which vehicle will generate the highest intertemporal utility across the future states of the world. We posit that consumers may mistakenly purchase a vehicle that has a high perceived utility at the time of purchase, but whose realized utility is systematically lower. Specifically, we test the extent to which weather variation at the time of purchase can cause consumers to overweigh the value they place on certain vehicle characteristics. We predict that consumers will overvalue warm-weather vehicle types (e.g., convertibles) when the weather is warm and sunny at the time of purchase and overvalue cold-weather vehicle types (e.g., four-wheel-drive vehicles) when the weather is cold and snowy at the time of purchase. We choose to focus our attention on a large and long-lived durable good for two reasons. First, the weather on the day of purchase will have very little effect on the total intertemporal utility consumers obtain from owning a vehicle. This means that a fully rational, utility-maximizing consumer should respond very little to the current weather when buying a car (although he or she might very reasonably put great weight on the current weather when buying an ice cream cone or a cup of hot cocoa). Second, consumers—knowing that vehicles are very expensive and that they will likely keep the vehicle for several years—attempt to make the correct long-term decision. Thus, finding bias in this setting is particularly

compelling evidence of the importance of psychological biases on individual decision making.

We explore this hypothesis using transaction-level data for more than 40 million new and used vehicles from dealerships around the United States. We find that the sales of convertibles and four-wheel-drives 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 fraction of cars being sold that are convertibles. Controlling for seasonal sales patterns, our estimates suggest that a location that experiences a temperature that is 10°F degrees higher than normal will experience a 2.7% increase in the fraction of cars sold that are convertibles. We find large and significant effects both in the spring and in the fall (e.g., an atypically warm day in November increases the fraction of vehicles sold that are convertibles). Importantly, we also show that atypically warm weather does not impact the fraction of cars sold that are convertibles when the temperature is already high (above 75°F or 80°F). Purchases of four-wheel-drive vehicles are also very responsive to idiosyncratic weather variation—particularly snowfall. Our results suggest that a snow storm of approximately 10 inches will increase the fraction of vehicles sold that have four-wheel-drive by about 6% over the next two to three weeks.

In the article we consider ways these effects could arise from consumers behaving as standard, rational economic agents. The data allow us to rule out that these standard rational-agent models can fully explain our findings. For example, a distributed lag model indicates that the increase in convertible and four-wheel-drive sales due to idiosyncratic weather variation cannot be explained by short-run substitutions in vehicle purchases across days (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.

We next consider psychological mechanisms that might explain our empirical findings. We show that our results are

consistent with both projection bias and salience, psychological effects that are closely related in this context.

Our findings are significant for several reasons. First, vehicles are one of the highest value purchases that most households make. Identifying and potentially correcting systematic errors in this market can have important welfare implications. Perhaps more important, our results suggest that focusing too much on conditions at the time of decision may be a mistake that is prevalent in other contexts (getting married, buying a house, choosing a job, etc.) that are similarly distinguished by having large stakes, state-dependent utility, and low-frequency decision making.

Our article is related to a growing literature that uses field data to test models from behavioral economics (see DellaVigna 2009 for a review). Our study is most similar to the work of Conlin, O'Donoghue, and Vogelsang (2007) who test for intertemporal bias in catalog order purchases. 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°F lower, consumers are 0.57 percentage point more likely to return a purchased item (a 3.95% increase relative to the average return rate). They argue that these empirical findings fit the predictions made by a model of projection bias. Our article complements this earlier work.

The article proceeds as follows. In Section II, we describe the vehicle market and weather data. In Section III, we estimate the effect of short-term weather fluctuations on vehicle purchasing. In Section IV, we consider a variety of rational explanations for the weather effects we estimate in Section III, including the empirical evidence in support of each. In Section V, we discuss several psychological explanations for our estimated weather effects, and evaluate the empirical evidence for each. Section VI concludes.

II. 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 United States 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 United States 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 those in more rural regions. Examples of DMAs in our data include Phoenix, Arizona; Tulsa, Oklahoma; Lansing, Michigan; and Billings, Montana.¹

We 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.² Data were collected on temperature, precipitation, precipitation type, and cloud cover. Temperature is measured as the daily high temperature, measured in degrees Fahrenheit. Precipitation is measured as the cumulative liquidized inches in a day. If the only precipitation type reported for the day is rain, we classify the precipitation as rainfall (measured in inches). If the only precipitation type reported during the day is snow, we classify the precipitation as snowfall (measured in liquidized inches). If both rain and snow are reported on a day, we classify the precipitation as slushfall (measured in liquidized inches). Cloud cover is a daily measure of the fraction of the sky covered by clouds.

The data show that vehicle transactions occur year round but are most common during the summer months. Of primary interest in this article is the seasonal trend in convertible and four-wheel-drive purchases. In Figure A1, Panel A in the online appendix, we illustrate the percentage of total vehicle transactions that are 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

1. A list of all the DMAs covered by our data is available from the authors.

2. If the weather station did not have weather data available for at least 90% of the 4,745 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.)

of vehicles sold that are convertibles is highest in the early spring, peaking in April in seven out of the eight years in the data. Although 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 this percentage are well over half the magnitude of the corresponding spring peaks. We note that these seasonal differences in convertible purchases are consistent with a standard model of state-dependent preferences: consumers who buy a convertible in the spring will be able to immediately consume several months of warm-weather driving, whereas fall buyers will have to wait a few months to consume their convertible in its ideal weather. This makes the total discounted utility for spring convertible buyers higher than that of fall convertible buyers.

Similarly, Figure A1, Panel B in the online appendix illustrates the percentage of total vehicle transactions that are four-wheel-drive vehicles by month and year. Four-wheel-drive transactions range between 20% and 35% of total vehicle transactions. Panel B shows a seasonal pattern in which four-wheel-drive vehicles are particularly popular in the early winter months (purchases usually peak in December).³ As was the case for convertibles, this is not yet strong evidence for intertemporal bias since a standard model of state-dependent preferences would predict that the discounted utility of a four-wheel-drive is highest at the beginning of the winter.

One might expect that the seasonal sales patterns of the two different types of vehicles would differ with geography because of differences in climate. If we divide our DMAs into two groups, those with above-median monthly temperature variation⁴ (such as Chicago) and those with below-median monthly temperature variation (such as Miami), we see differences between the groups, but the overall patterns are similar. Perhaps surprisingly, 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 in the percentage of vehicles sold that are convertibles is higher in the variable-temperature DMAs.

3. There is a mid-summer peak in 2005 that 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.)

4. 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.

For four-wheel-drive vehicles, there is a large level difference in the percentage of such vehicles purchased in the two types of DMAs, and once again the variable temperature areas appear to have a more pronounced seasonal pattern. (Figure A2 in the online appendix shows these differences.)

Our identification strategy involves testing whether idiosyncratic weather conditions (controlling for time of year to eliminate seasonal purchasing patterns) are correlated with variation in the sales share of convertible and four-wheel-drive vehicles. To do this, we create indicator variables for each transaction denoting whether the vehicle purchased was a convertible or whether it was a four-wheel-drive. Regressing these indicators on measures of the weather in the DMA where the transaction occurred will enable us to test whether atypical weather leads to variation in convertible and four-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 consumers have systematic behavioral biases, and respond accordingly. Our estimates identify the combined effect of changes in consumers' behavior and dealers' responses to those changes.⁵

III. ESTIMATION OF WEATHER EFFECTS ON VEHICLE PURCHASING

In this section, we estimate the effect of local, daily weather on the types of vehicles purchased. We begin by estimating the effect of weather on convertible purchases, and then on four-wheel-drive purchases.

5. In the extreme, the effect could be driven entirely by the supply side if, for example, salespeople enjoy test-driving convertibles on sunny days, and buyers are influenced by the salesperson's extra effort. For the only effect to be a supply-side effect, buyers would have to be immune to the effect of good weather that makes salespeople want to test-drive convertibles on sunny days.

III.A. Effect of Weather on Convertible Purchases

We estimate the effect of weather on convertible sales using the following specification:

$$(1) \quad I(\text{Convertible})_{irt} = \alpha_0 + \alpha_1 \mathbf{Weather}_{rt} + \mu_{rT} + \tau_Y + \epsilon_{irt}.$$

$I(\text{Convertible})$ is an indicator variable equal to 1 if the vehicle sold in transaction i in DMA r on day t was a convertible. $\mathbf{Weather}$ is a vector of weather variables for DMA r on day t —temperature, rainfall, snowfall, slushfall, and cloud cover—defined in the previous section. (Summary statistics can be found in Table I.) μ_{rT} are DMA*week-of-the-year fixed effects and τ_Y are year fixed effects. μ_{rT} will absorb the average seasonal variation in convertible sales at the week level, separately for each DMA. τ_Y will absorb year-to-year changes in consumer tastes for convertibles.⁶

Our main coefficients of interest will be α_1 , the vector of weather coefficients. Each element of α_1 can be interpreted as the effect of a one-unit change in the corresponding weather variable on the probability that a particular transaction is a convertible, or—more suitably for our application—on the fraction of vehicles sold on a given day that are convertibles.

Table II reports the results of estimating equation (1). Column (1) indicates that when the temperature is 1°F higher than expected in a given DMA, the DMA experiences on average a 0.007 percentage point increase in the fraction of total vehicles sold that are convertibles. Thus a day with a temperature that is 10°F higher than average for that DMA in that week of the year would be predicted to see 0.07 percentage point more convertibles sold as a percentage of the total number of vehicles sold.⁷ This would be a 2.7% increase relative to the weighted base rate of 2.6% of vehicles sold being convertibles. Liquid inches of rain, snow, and slushfall all have negative effects on the fraction of vehicles sold that are convertibles, 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

6. To be able to use very granular fixed effects, we estimate linear probability models rather than probit or logit models.

7. We use a 10°F change in temperature as a benchmark to help understand the size of the effects that we find. How does 10°F compare to typical temperature fluctuations? The standard deviation for temperature within a DMA and week of the year is 7.87°F. A 10°F higher or lower temperature than average for a DMA and week of the year occurs on approximately 20% of days.

TABLE I
SUMMARY STATISTICS FOR RETAIL VEHICLE SALES AND WEATHER

	Mean	Std. dev.	Min.	Max.
Vehicle characteristics				
Convertible indicator	.022	0.145	0	1
Four-wheel-drive indicator	.267	0.442	0	1
Weather variables				
Max temperature	71.4	17.8	-38.9	120.9
Rainfall	0.202	0.758	0	33.9
Snowfall	0.004	0.058	0	8.53
Slushfall	0.011	0.148	0	22.0
Cloud cover	0.526	0.322	0	1

Notes. Observations: 40,164,136. Summary statistics are reported for all vehicle sales in our data set between 2001 and 2008.

TABLE II
EFFECT OF WEATHER ON CONVERTIBLE PURCHASES

	(1)	(2)	(3)	(4)	(5)
	Dependent variable: indicator equal to 1 if purchase was a convertible				
	Full year	Quarter 1	Quarter 2	Quarter 3	Quarter 4
Temperature	.007** (.0004)	.010** (.001)	.008** (.001)	.002* (.001)	.005** (.001)
Rainfall	-.016** (.004)	-.035** (.010)	-.010 (.007)	-.014** (.006)	-.016** (.007)
Snowfall	-.053 (.036)	-.021 (.053)	-.193 (.126)	-8.178 (8.517)	-.047 (.048)
Slushfall	-.043** (.015)	-.036 (.021)	-.067* (.033)	.019 (.033)	-.058* (.027)
Cloud cover	-.126** (.010)	-.114** (.018)	-.206** (.022)	-.120** (.020)	-.087** (.017)
Year fixed effects	X	X	X	X	X
DMA*week-of-the-year fixed effects	X	X	X	X	X
R-squared	0.003	0.003	0.002	0.002	0.003
Observations	39,984,509	9,150,047	10,205,673	10,608,527	10,020,262

Notes. Coefficient values and clustered standard errors are presented from OLS regressions of an indicator for whether a car sold was a convertible on weather variables: temperature (°F), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Fixed effects are included for each year and for DMA*week-of-the-year (week 1–week 52). The first column uses all the data and the next four columns present results separately for the four quarters of the year. All coefficients and standard errors have been multiplied by 100 for ease of presentation. Thus each coefficient represents the percentage point change in probability of purchasing a convertible. Standard errors are clustered at the DMA*day. * significant at 5%; ** significant at 1%.

important for convertible sales. As the sky goes from completely clear to completely cloudy, convertible sales as a fraction of total vehicles sold decreases by 0.126 percentage point. Thus, a clear sky (relative to completely overcast) increases convertible sales by the same amount as an increase in temperature of 18°F.

The next four columns in Table II estimate the effect of temperature and other weather variables on convertible sales by quarter of the year. The effect of temperature is large and statistically significant in quarters 1, 2, and 4, but smaller in quarter 3. During quarter 3 (July, August, September), the baseline temperature is already quite warm in most areas, which may explain why an increase in temperature has less effect on convertible sales than in other quarters. We find it particularly noteworthy that high temperatures continue to have a fairly large effect on convertible sales in quarter 4 (October, November, December) because this is the time of year when the rational discounted utility of buying a convertible should be the lowest. Cloud cover—which is arguably important to the perceived utility of driving a convertible no matter what time of year—has a large and significant effect in all quarters (including quarter 3).

The differences across quarters in the estimated effects of temperature on convertible purchases suggest that the effect might be heterogeneous across the range of temperature. To assess the extent of this heterogeneity, we reestimate equation (1), replacing the linear measure of temperature with indicator variables for the 5° bin into which the daily high temperature falls. We use 15 indicator variables for temperatures in bins from (25–29.9°F) to (95–99.9°F). We group all days with daily highs of 100°F or more into a single indicator, and leave days with high temperatures below 25°F as our left-out group.

We report the estimated coefficients of these indicator variables graphically in Figure I. For each bin, we plot the estimate coefficient as a dot, with “whiskers” representing the 95% confidence interval around the estimated coefficient. Higher temperatures increase the percentage of transactions that are convertibles fairly steadily from temperatures below 25°F up until some point between 70°F and 80°F. At that point, the effect flattens out, suggesting that beyond 80°F, higher temperatures are no longer associated with increases in the share of vehicles sold that are convertibles.




FIGURE I

Flexible Relationship between Temperature and Convertible Purchases

This figure provides the coefficient values and 95% confidence intervals for the effect of daily high temperature on the probability of purchasing a convertible. Each dot plots the coefficient from the regression reported in column (1) of Table II, but with dummy variables for 5°F bins of temperature used in place of the linear temperature term. The omitted temperature category is less than 25°F.

TABLE III
EFFECT OF WEATHER ON FOUR-WHEEL-DRIVE PURCHASES

	Dependent variable: indicator equal to 1 if purchase was a four-wheel-drive				
	Full year	Quarter 1	Quarter 2	Quarter 3	Quarter 4
Temperature	-.032** (.001)	-.038** (.002)	-.018** (.003)	-.029** (.004)	-.038** (.003)
Rainfall	.084** (.014)	.119** (.036)	.081** (.026)	.054** (.023)	.132** (.032)
Snowfall	1.81** (.26)	1.67** (.33)	.72 (.82)	125* (53)	2.11** (.48)
Slushfall	.504** (.077)	.540** (.110)	.27 (.22)	-.029 (.167)	.769** (.166)
Cloud cover	.461** (.039)	.337** (.072)	.512** (.077)	.383** (.087)	.598** (.076)
Year fixed effects	X	X	X	X	X
DMA*week-of-the-year fixed effects	X	X	X	X	X
R-squared	0.086	0.086	0.074	0.084	0.097
Observations	39,984,509	9,150,047	10,205,673	10,608,527	10,020,262

Notes. Coefficient values and clustered standard errors are presented from OLS regressions of an indicator for whether a car sold was a four-wheel-drive on weather variables: temperature (°F), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Fixed effects are included for each year and for DMA*week-of-the-year (week 1–week 52). The first column uses all the data and the next four columns present results separately for the four quarters of the year. All coefficients and standard errors have been multiplied by 100 for ease of presentation. Thus each coefficient represents the percentage point change in probability of purchasing a convertible. Standard errors are clustered at the DMA*day. * significant at 5%; ** significant at 1%.

III.B. Effect of Weather on Four-Wheel-Drive Purchases

Although buying a convertible may seem especially attractive on a warm day, it is cold and snowy days that make four-wheel-drive vehicles seem like an especially good idea. Table III presents our estimates of the impact of weather variation on the percentage of vehicles sold that are four-wheel-drives. We obtain these estimates by substituting $I(4WheelDrive)$, an indicator for whether a given transaction was for a four-wheel-drive vehicle, on the left hand side of equation (1) to obtain the following estimation equation:

$$(2) \quad I(4WheelDrive)_{irt} = \beta_0 + \beta_1 \mathbf{Weather}_{rt} + \mu_{rT} + \tau_Y + \epsilon_{irt}.$$

$I(4WheelDrive)$ equals 1 if transaction i that occurred in DMA r on day t is for a four-wheel-drive vehicle. All other variables are as

defined for equation (1). Our main coefficients of interest will be β_1 , the vector of coefficients that represent the effect of weather on the fraction of vehicles sold in a given DMA on a given day that are four-wheel-drive. The results are reported in Table III.

As we expected, the results we find are roughly the opposite of what we found for convertibles. We find that colder temperature values lead to more four-wheel-drive purchases. For example, on a day with temperatures that are 10°F below the average for that DMA and week of year, the estimated coefficient predicts that the fraction of vehicles sold that are four-wheel-drives would be 0.32 percentage point higher than otherwise. This represents a 0.85% increase relative to the weighted baseline of 33.5% of vehicles sold with four-wheel-drive. We also find a large, positive effect of snow and slush on four-wheel-drive transactions. One inch of liquidized snow (about 10 inches of snow) leads to a 1.81 percentage point increase in the percentage of total vehicles sold with four-wheel-drive. The effects for snowfall are statistically significant in quarters 1 and 4. The effect size in quarter 2 is still reasonably large, but the standard error is much higher than in quarters 1 and 4. The impossibly large effect estimated for quarter 3 and its accompanying large standard error is clearly driven by the lack of snowfall variation that exists in the data during quarter 3. The effect of snowfall is slightly larger in quarter 4 than in quarter 1. However, the significant effect of snowfall in quarter 1 suggests that even a snow storm that occurs toward the end of the winter season can have a powerful impact on four-wheel-drive purchase behavior.

As we did when analyzing convertible purchasing, we want to allow the effect of temperature to be nonlinear in its effect on four-wheel-drive purchasing. We do this by reestimating equation (2), replacing the linear measure of temperature with indicators for the same 5° bins reported in Figure I. We report the estimated coefficients on these indicator bins graphically in Figure II. As in Figure I, the dot represents the estimated coefficient and the whiskers represent the 95% confidence interval. The results in Figure II indicate that changes in temperature have little effect on four-wheel-drive purchasing for temperatures above 50°F. However, for temperatures below that, each 5°F drop in temperature significantly increases the percentage of vehicles purchased that are four-wheel-drives.

Although our estimates show that the fraction of vehicles sold that have four-wheel-drive rises in cold and snowy weather,




FIGURE II

Flexible Relationship between Temperature and Four-Wheel-Drive Purchases

This figure provides the coefficient values and 95% confidence intervals for the effect of daily high temperature on the probability of purchasing a four-wheel-drive vehicle. Each dot plots the coefficient from the regression reported in column (1) of Table III, but with dummy variables for 5°F bins of temperature used in place of the linear temperature term. The omitted temperature category is less than 25°F.

this does not necessarily mean that the total *number* of four-wheel-drive sales rises. One way we could obtain our results is if total vehicle sales fall in cold, snowy weather, but sales of non-four-wheel-drive vehicles fall by more than the sales of four-wheel-drive vehicles. We investigate whether this is the case by aggregating the data to the day level and regressing the log of the number of convertibles sold and the log of the number of four-wheel-drives sold during that day on our weather measures. The result show that on warm, sunny days the total number of vehicles sold rises, but the number of convertibles sold rises by even more, which is why the percentage increases. Unit sales of four-wheel-drives, however, fall on cold snowy, days, but by proportionally less than purchases of vehicles without four-wheel-drive. Thus, it is worth noting that the four-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 four-wheel-drive vehicle appears to be more motivated go to the dealership than are buyers of non-four-wheel-drive vehicles.

III.C. Effect of Weather on Vehicle Prices

We have shown in the previous two subsections that weather affects the equilibrium sales of vehicles of different types. Specifically, we have shown that the percentage of vehicles sold that are convertibles is higher on days with warm and sunny weather, whereas the percentage of vehicles sold that are four-wheel-drives are higher on days with cold, snowy weather. In this section, we investigate the effect of weather on the equilibrium prices of these types of vehicles.

The most intuitive way to explain in supply-and-demand terms our estimated effects of weather on purchasing is that weather causes a change in daily demand for different vehicle types. If this is so, what effect would we expect weather to have on equilibrium prices? The answer depends, of course, on the shape of the supply curve.⁸ In a simple supply-and-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 this simple model is not an ideal fit for the car industry. First, from a dealer's perspective, the supply

8. See Busse (2012) for a discussion of how different supply relationships affect the equilibrium price and quantity that arise from changes in demand.

of vehicles is not upward-sloping. Dealers can order vehicles from manufacturers at 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 a competitive market, the effect of an increase in demand should be increased sales, with essentially no increase in price.⁹

Second, a competitive price-taking market is not a very good description of the retail car industry. Individual consumers negotiate a price for a specific vehicle with the dealer. Whether the incremental buyers who enter (or leave) the market in response to a change in weather will obtain higher prices or lower prices than the inframarginal buyers who are in the market at all times depends on the reservation prices and bargaining characteristics of the incremental buyers relative to the inframarginal buyers. One might argue that the incremental buyers must have higher reservation prices than buyers on average, because they are being strongly swayed by temporary weather conditions. Similarly, one might argue that consumers who can buy “on impulse” must have high liquidity, and therefore likely higher incomes and higher reservation prices, than inframarginal buyers.

On the other hand, one might argue that if incremental buyers are buying this vehicle because of the weather (but would not be buying it on a day with different weather), the weather must have nudged them just above their point of indifference about buying. In this case, they might well have lower reservation prices than inframarginal buyers. Similarly, if dealers recognize which buyers are the incremental buyers who have come into the dealership because of the temporary weather condition, 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.¹⁰

Overall, we conclude that it is an empirical question whether prices for convertibles and four-wheel-drives will be

9. 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 that daily vehicle supply is not fixed. For more on how dealer supply and inventory affects prices, see Zettelmeyer, Scott Morton, and Silva-Risso (2007).

10. We thank Glenn Ellison for suggesting this point.

higher on the same days that warm and sunny weather or cold and snowy weather leads to an increased sales share for these types of vehicles. We estimate the effect of weather on the prices of convertibles and four-wheel-drives using the following specification.

$$(3) \quad \begin{aligned} Price_{ijrt} = & \gamma_0 + \gamma_1 \mathbf{Weather}_{rt} + \gamma_2 \mathbf{PurchaseTiming}_{it} \\ & + f(\text{Odometer}_i, \gamma_3) + \mu_{rT} + \tau_Y + \phi_j + \epsilon_{ijrt} \end{aligned}$$

Price measures the price paid in transaction *i* for vehicle *j* that occurred on day *t* in DMA *r*. (To make our measure of price represent a customer's total wealth outlay for the vehicle, we define price as the contract price for the vehicle agreed on by the consumer and the dealer, minus any manufacturer rebate the buyer received, plus any loss [minus any gain] the consumer received in negotiating a price for his or her trade-in.) **Weather** is a vector containing the temperature, rainfall, snowfall, slush-fall, and cloud cover on day *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 to hit sales volume targets. The specification also includes DMA*week-of-year (μ_{rT}), year (τ_Y), and vehicle type (ϕ_j) 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 (3) separately for new convertibles, used convertibles, new four-wheel-drives, and used four-wheel-drives. The specifications that estimate the effect of weather on used vehicle prices also include a linear spline in the vehicle's odometer with knots at 10,000-mile increments, which allows vehicle prices to depreciate over time in a reasonably flexible way. (See Busse, Knittel, and Zettelmeyer 2013 for use of a similar specification to estimate price effects in similar data.)

Table IV reports the results of estimating equation (3). 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. We do not find any evidence of weather affecting the transaction prices of convertibles. The weather coefficients, reported in column (1) for new convertibles and in column (2) for

TABLE IV
EFFECT OF WEATHER ON CONVERTIBLE AND FOUR-WHEEL-DRIVE PURCHASE PRICE

	(1) (2) (3) (4)			
	Dependent variable: vehicle sales price (less rebate)			
	Convertibles		Four-wheel-drives	
	New	Used	New	Used
Mean of dependent variable	\$40,001	\$22,222	\$31,845	\$19,132
Temperature	0.91 (0.89)	0.74 (0.92)	0.23 (0.25)	0.27 (.18)
Rainfall	-3.15 (8.77)	-11.25 (9.61)	1.55 (2.93)	6.95** (2.17)
Snowfall	156.74 (128.05)	86.41 (138.33)	11.69 (28.43)	-13.82 (19.56)
Slushfall	-41.66 (48.32)	1.06 (47.42)	-1.39 (12.21)	-4.33 (8.22)
Cloud cover	4.95 (24.04)	-31.24 (25.39)	21.69** (7.34)	-23.79** (5.25)
Year fixed effects	X	X	X	X
DMA*week-of-the-year fixed effects	X	X	X	X
Purchase timing fixed effects	X	X	X	X
Vehicle-type fixed effects	X	X	X	X
Odometer value spline		X		X
Observations	385,771	371,790	5,405,663	4,076,024

Notes. Coefficient values and standard errors are presented from OLS regressions of vehicle transaction prices on weather variables: temperature ($^{\circ}$ F), 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 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 used vehicle specifications (columns (2) and (4)) include a linear spline in odometer values with knots at 10,000-mile increments. * significant at 5%; ** significant at 1%.

used convertibles are both small and statistically insignificant. For four-wheel-drives, the results are similarly small. In column (3), we estimate that cloud cover leads to a statistically significant increase in the price of new four-wheel-drives, but by only \$21.69 for a change from completely clear to completely overcast skies. For used four-wheel-drives, in column (4), we estimate that one inch of rainfall increases the average transaction price by \$6.95, and that cloud cover reduces the price by \$23.79, the latter result being in the opposite direction of what we would expect. In addition to being statistically insignificant, these effects are very small compared to average transaction prices of \$31,845 for new four-wheel-drives and \$19,132 for used four-wheel-drives.

IV. RATIONAL EXPLANATIONS FOR WEATHER EFFECTS

In this section, we consider several potential explanations for the weather effects we estimate that do not rely on any psychological mechanisms. First, we consider whether our estimated weather effects are the result of short-term intertemporal shifts in demand. For example, it could be that consumers who have decided to buy a convertible may wait for a nice day to actually make the purchase. Second, we consider whether consumers are more likely to purchase a vehicle with a weather-related characteristic such as a convertible roof or four-wheel-drive on a day whose weather enables them to test-drive the vehicle in that weather. Third, we consider whether consumers who test-drive a convertible or a four-wheel-drive on a day with the complementary weather learn thereby about their utility for the vehicle type and that this is what drives the increased fraction of convertible and four-wheel-drive sales on the relevant weather days.

IV.A. Shifts in Purchase Timing

Our results in Section III indicate that the fraction of vehicles sold on a given day that are either convertibles or four-wheel-drives is influenced by daily weather variation. One possible explanation for this empirical finding is that weather fluctuations may appear to be incrementally increasing purchases of certain types of vehicles, but instead just cause short-run intertemporal substitutions in vehicle purchasing behavior. In other words, weather shifts when, but ultimately not what, people buy. An example of this harvesting story is that a consumer may be interested in purchasing a convertible some time in the next month and then actually makes her purchase whenever it happens to be a nice day outside.¹¹ Our previously noted finding that atypically warm weather in November can affect convertible purchases and a snow storm in February can affect four-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.

To assess the extent to which there is short-run intertemporal substitution of purchases with respect to daily weather

11. The fact that more convertibles are bought in spring than winter, and the reverse for four-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.

fluctuations, we estimate a distributed lag model. We do so by reestimating equations (1) and (2) with 60 daily lags of each weather variable added to the estimating equation.

$$(4) \quad \begin{aligned} I(\text{Convertible})_{irt} &= \alpha_0 + \alpha_1 \mathbf{Weather}_{rt} \\ &+ \sum_{j=1}^{60} \alpha_{1,-j} \mathbf{Weather}_{r,t-j} + \mu_{rT} + \tau_Y + \epsilon_{irt} \end{aligned}$$

$$(5) \quad \begin{aligned} I(4\text{WheelDrive})_{irt} &= \beta_0 + \beta_1 \mathbf{Weather}_{rt} \\ &+ \sum_{j=1}^{60} \beta_{1,-j} \mathbf{Weather}_{r,t-j} + \mu_{rT} + \tau_Y + \epsilon_{irt} \end{aligned}$$

$\mathbf{Weather}_{r,t-j}$ is the vector of weather variables (temperature, rainfall, snowfall, slushfall, and cloud cover) j days before the transaction date. $\alpha_{1,-j}$ and $\beta_{1,-j}$ are vectors of coefficients that estimate the effect of weather j days ago on day t purchases of convertibles and four-wheel-drives, respectively.¹² By including lagged variables, we are able to test whether having cold or hot days leading up to the day of purchase influences how the current weather affects behavior. For example, in the convertible scenario, negative coefficients on the lagged variables would be interpreted as evidence of harvesting via the following argument. A negative coefficient on, say, the three-day lag of temperature would indicate that if the weather three days ago was hot, sales today are lower by some amount than they otherwise would have been. Additionally, this implies that if the weather today is hot, sales three days 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 today, how much lower will sales be in subsequent days?” The one-day lag gives us an estimate for the effect of hot weather today on sales one day from now, the two-day lag estimates the effect of hot weather today on sales two days 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 today are made up entirely of sales displaced from the following

12. This specification is identical to instead regressing each of the weather variables on the fixed effects, obtaining residual values for each weather variable, and then including these weather residuals and 60 days’ worth of lags of weather residuals in the regression.

60 days. More generally, the sum of the lag coefficients tells us how much of our estimated current period effect is due to intertemporal substitution.¹³

Figures III and IV present the results of this dynamic analysis for convertible purchases estimated by equation (4). Figure III plots the estimated coefficient and confidence intervals for the current day's temperature (daily lag = 0) and for 60 daily lags of temperature. Figure IV plots the coefficient and standard errors for the current day's cloud cover and 60 daily lags of cloud cover. The results once again show a large and significant effect of current weather on convertible purchases. The estimated coefficient on the current day's high temperature is 0.005, with a standard error of 0.0006, similar to the coefficient of 0.007 reported in Table II. The coefficients on the lag variables are all small relative to the current temperature coefficient, and all but four of the 60 lag coefficients are statistically insignificant. Most important for the question at hand, more of the coefficients are positive than negative, especially among the most recent two or three weeks of lags. The sum of the 60 lag coefficients is 0.012; if our estimated warm weather effect were intertemporal shifts in sales, the coefficients would be expected instead to sum to the negative of the current day coefficient of 0.005. We can test the null hypothesis that the sum of the first X lags are equal to the negative of the current day coefficient by imposing this as a linear restriction on the estimation. If we do so, we reject this null hypothesis with a p -value $< .001$ for X equal to 7, 14, 21, 28, 45, or 60 days of lags. If anything, it appears that warm weather over the past several weeks leads to an even higher fraction of vehicles sold today being convertibles.

Figure IV shows the estimated coefficients for the current day's cloud cover and 60 daily lags of cloud cover. In this specification, the estimated coefficient on the current day's cloud cover is -0.125 with a standard error of 0.013, almost identical to the -0.126 reported in Table II. The sum of the 60 lag coefficients is 0.199, which means that we cannot conclude that the cloud cover effects on convertible sales are not due to short-term shifts in

13. See Jacob, Lefgren, and Moretti (2007) for a similar analysis that tests for intertemporal substitution of crime using weather shocks and Deschenes and Moretti (2009) who test for intertemporal substitution of mortality using weather shocks. See these papers also for a discussion of summing up coefficient values to test for evidence of intertemporal substitution.




FIGURE III

Distributed Lag Analysis of Temperature and Convertible Purchases

This figure provides the coefficient values and 95% confidence intervals for the effect of daily high temperature on the probability of purchasing a convertible. Each dot plots the coefficient of a lag temperature variable from the same regression as reported in column (1) of Table II, but with the addition of 60 daily lags of each of the weather variables.




FIGURE IV

Distributed Lag Analysis of Cloud Cover and Convertible Purchases

This figure provides the coefficient values and 95% confidence intervals for the effect of cloud cover (fraction of the sky) on the probability of purchasing a convertible. Each dot plots the coefficient of a lag cloud cover variable from the same regression as reported in column (1) of Table II, but with the addition of 60 daily lags of each of the weather variables.

purchase timing. However, most of the positive lag coefficients that contribute to this result are in the longer lags. We can reject that the sum of the first X lag coefficients equal the negative of the current day coefficient for X equal to 7, 14, and 21 days with a p -value $< .001$ and for 28 days with a p -value of $.026$. For longer sums of lags, we cannot reject the null hypothesis that our results are due to intertemporal shifts in purchase timing. However, this means that to the extent that our estimated result is the consequence of intertemporal shifts, it is shifts coming mostly from more than a month ago, not shifts coming from the previous several weeks.

Figures V and VI provide a similar analysis for four-wheel-drive purchases. Figure V shows the current day and 60 daily lag coefficients for temperature, and Figure VI shows them for snowfall. The estimated coefficient of current day temperature on the fraction of vehicles sold that are four-wheel-drive is -0.015 , with a standard error of 0.0027 , which is somewhat smaller than the -0.032 reported in Table III. We can reject at a p -value $< .001$ that the sum of the lagged coefficients equal the negative of the current day coefficient over the 7-, 14-, 21-, 28-, 45-, or 60-day horizon. The lag coefficients, while not statistically significant, are more suggestive that cold weather in the two or three previous weeks increases four-wheel-drive sales today.

The last set of distributed lag results is shown in Figure VI. The coefficients in this figure indicate that the current day's snowfall has a positive and significant effect on the fraction of vehicles sold today that are four-wheel-drive vehicles, but so does snowfall on almost any of the days of the previous two weeks. These coefficients are not consistent with shifts in purchase timing that would explain our empirical result from Section III. (We can again reject at a p -value $< .001$ that the sum of the lagged coefficients equal the negative of the current day coefficient over the 7-, 14-, 21-, 28-, 45-, or 60-day horizon.) Instead, the results here suggest that although a snowstorm today has the biggest effect on sales today, it will contribute to an increased fraction of sales of four-wheel-drive vehicles for almost two weeks.

IV.B. *Test-drive Timing*

One aspect of vehicle purchasing that may lead to a correlation between weather and vehicle purchase timing, particularly




FIGURE V
Distributed Lag Analysis of Temperature and Four-Wheel-Drive Purchases

This figure provides the coefficient values and 95% confidence intervals for the effect of daily high temperature on the probability of purchasing a four-wheel-drive vehicle. Each dot plots the coefficient of a lag temperature variable from the same regression as reported in column (1) of Table III, but with the addition of 60 daily lags of each of the weather variables.




FIGURE VI

Distributed Lag Analysis of Snowfall and Four-Wheel-Drive Purchases

This figure provides the coefficient values and 95% confidence intervals for the effect of daily snowfall on the probability of purchasing a four-wheel-drive vehicle. Each dot plots the coefficient of a lag snowfall variable from the same regression as reported in column (1) of Table III, but with the addition of 60 daily lags of each of the weather variables.

for convertibles, is the desire of most customers to test-drive a vehicle before buying. Suppose a customer is considering buying a convertible, and she is able to forecast accurately her long-term utility from owning a convertible. Now suppose that before she buys the convertible, she would like to be able to test out various features of it: how convenient it is to put the top up and down, how much wind or road noise she experiences with the top down, and so on. 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 nonconvertible instead.

The behavior of both types of customers would lead to a higher percentage of vehicles sold on warm days being convertibles relative to the percentage on cold days for reasons other than errors in forecasting intertemporal utility. The first type of customer would lead to harvesting (customers wait until a warm week to buy a convertible so that they can test-drive the vehicle). In the previous section, we already discussed and ruled out harvesting effects for convertibles with regard to temperature and also with regard to cloud cover within a four-week window. However, the second customer type is not ruled out by our distributed lag model. Several pieces of evidence, however, argue against a test-drive learning story. For example, Figure I indicates that an extra degree of warm weather results in more convertible purchases even when the baseline temperature is in the 60°F–80°F range. This is a range of temperature for which it is clearly possible for someone to test out the convertible 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 fraction 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, psychological mechanisms that cause consumers to over-respond to current mental and emotional states should lead to warm, sunny days being those on which people are particularly likely to buy convertibles, rather than warm, cloudy days. Indeed, if we examine the results in Table II, we find that unusually cloudy days have a significant negative effect on the percentage of vehicles sold that are convertibles, consistent with one of several psychological mechanisms. It is particularly noteworthy that cloudy days have a statistically significant 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, clear days not because it is more possible to test-drive them, but because it seems more attractive to own a convertible on such days.

IV.C. Consumer Learning

Another alternative hypothesis that would explain our findings is that customers need to test-drive a vehicle with weather-related features on a day that has the relevant weather (warm and sunny or cold and snowy) to actually learn what their utility will be from owning either a convertible or a four-wheel-drive in such weather conditions. Under this hypothesis, a warm, sunny day does not lead a customer to overestimate the utility she will get from owning a convertible; instead, it enables her to learn for the first time how high her true utility will be from owning a convertible in such weather states. Before considering this as an alternative hypothesis, we note that this type of extreme learning story—in which vehicle buyers cannot 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—is in itself an impediment to correctly forecasting intertemporal utility. As such, learning is a rational mechanism that in its extreme form is very closely related to the psychological mechanisms we discuss in the next section.

Despite the similarity between learning as described here and psychological biases, 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 owned already either a convertible or a four-wheel-drive vehicle. Previous convertible owners are less likely to need to learn about what it is like to drive a convertible during a warm weather state, and similarly for previous four-wheel-drive owners and cold or snowy states, so evidence that idiosyncratic weather affects these consumers is particularly strong evidence against learning as an explanation.

If we look within the subset of transactions that use a convertible as a trade-in, we find that approximately 25% of these consumers purchase another convertible whereas 75% purchase a nonconvertible vehicle. Column (1) of Table V reports the results of our baseline specification if we restrict the sample to consumers who are trading in a convertible. Although 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 sales. 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 comparable in percentage terms because the convertible purchase rate in this sample (25%) is so much higher.¹⁴

In column (2) of Table V, we estimate the effect of weather on consumers who are trading in a four-wheel-drive vehicle. Overall, 78% of people who trade in a four-wheel-drive vehicle purchase another four-wheel-drive vehicle. In column (2) we continue to find strong and statistically significant effects of all five weather measures on four-wheel-drive purchases for buyers who traded in a four-wheel-drive vehicle. The estimated effects are substantially smaller in percentage terms than in the full sample, in large part because the unconditional probability of buying a four-wheel-drive vehicle is so high in this sample.¹⁵

14. The full sample results indicate that a 10°F increase in temperature increases the percentage of vehicles sold that are convertibles by 0.07 percentage point in the full sample, a 2.7% increase relative to a base percentage of 2.6%. In the “convertible trade-in” subsample, the effect is a 0.65 percentage point increase, a 2.6% increase relative to a base percentage of 25%.

15. The full sample results indicate that a 10°F decrease in temperature increases the percentage of vehicles sold that are four-wheel-drive by 0.32 percentage point in the full sample, a 0.85% increase relative to a base percentage of 33.5%. In the “four-wheel-drive trade-in” subsample, the effect is a 0.43 percentage point increase, a 0.55% increase relative to a base percentage of 78%. The full sample results indicate that one liquidized inch of snowfall increases the percentage of vehicles sold that are four-wheel-drives by 0.81 percentage points in the full sample, a 2.4% increase relative to a base percentage of 33.5%. In the “four-

TABLE V
EFFECT OF WEATHER ON CONVERTIBLE AND FOUR-WHEEL-DRIVE PURCHASES FOR
CONSUMERS TRADING IN A CONVERTIBLE OR FOUR-WHEEL-DRIVE VEHICLE

	(1)	(2)
	Dep. var.: indicator equal to 1 if purchase was a convertible/four-wheel-drive	
	Convertibles	Four-wheel-drives
Temperature	.065** (.011)	-.043** (.003)
Rainfall	-.196 (.102)	.129** (.039)
Snowfall	-2.38 (1.48)	1.07** (.40)
Slushfall	-.760 (.544)	.501** (.143)
Cloud cover	-.284 (.294)	.701** (.099)
Year fixed effects	X	X
DMA*week-of-the-year fixed effects	X	X
R-squared	0.049	0.076
Observations	336,091	4,099,091

Notes. Coefficient values and standard errors are presented from OLS regressions of an indicator for whether a car was a convertible (column (1)) or a four-wheel-drive (column (2)) on weather variables: temperature (°F), rain (inches), snow (liquidized inches), slush (liquidized inches), and cloud cover (fraction of sky covered). Fixed effects are included for 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 four-wheel-drive (column (2)). Standard errors are clustered at the DMA*day. * significant at 5%; ** significant at 1%.

The fact that we find effects of idiosyncratic weather variation in precisely the subsample of buyers who would seem to have the *least* to learn about their utility from owning either a convertible or a four-wheel-drive vehicle casts doubt on the learning story being a key explanation of the effects that we find.¹⁶

wheel-drive trade-in” subsample, the effect is a 1.07 percentage point increase, a 1.4% increase relative to a base percentage of 78%.

16. These results also speak to a separate selection concern that one might have. Perhaps the primary reason the fraction of vehicles sold that are convertibles is higher on sunny days is that consumers who like convertibles (and already own one) are more likely to be out shopping on a warm, sunny day. The “trade-in sample” analysis restricts the sample to a set of people who all own the same type of vehicle

V. PSYCHOLOGICAL MECHANISMS

In the previous two sections, we provided evidence that vehicle consumers are highly sensitive to the weather on the day of purchase, effects that cannot be justified by a standard model of consumer choice and cannot be explained by a variety of other rational mechanisms, such as purchase timing or learning. In this section we explore psychological theories that predict being overly sensitive to the state of the world at the time of purchase. We focus on projection bias and salience as the two most likely candidates and separately discuss the evidence against and in favor of each of these mechanisms.¹⁷

V.A. *Projection Bias*

Projection bias has received significant attention in the economics and psychology literature. This bias, which is based on earlier psychological work, was formalized by Loewenstein, O'Donoghue, and Rabin (2003). In their model, Loewenstein, O'Donoghue, and Rabin assume that a person has state-dependent utility such that her instantaneous utility of consumption, c , in state s can be represented as $u(c, s)$. They then consider an individual who is currently in state s' and attempting to predict her future instantaneous utility of consumption, c , in state s ; the utility prediction is denoted $\tilde{u}(c, s | s')$. An accurate prediction would be represented by $\tilde{u}(c, s | s') = u(c, s)$.

Loewenstein, O'Donoghue, and Rabin 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

$$(6) \quad \tilde{u}(c, s | s') = (1 - \alpha)u(c, s) + \alpha(u(c, s')),$$

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

and thus helps ease this particular concern (in conjunction with some of the other evidence that we present, such as the results from the distributed lag model).

17. Two additional psychological theories that are worth mentioning are present-biased preferences and mood effects. In an earlier version of this article, we provided a calibration that indicated that the level of impatience required to explain our effects are outside the range of parameter values typically found using a model of present-biased preferences. We also find no evidence that individuals simply purchase more expensive vehicles or sporty vehicles (e.g., coupe body style) on warm days, which are plausible predictions of a mood effect story.

individual perceives her future utility to reflect a convex combination of her true future utility and the utility that consumption c would provide in her current state s' .

This simple model of projection bias can be extended easily to an intertemporal choice framework. Consider, for example, the utility that a person receives from purchasing a convertible at time t ($conv_t$) and owning it until period T . Her true utility can be represented by

$$(7) \quad U^t(conv_t, \dots, conv_T) = \sum_{\tau=t}^T \delta^{(\tau-t)} u(conv_\tau, s_\tau),$$

where $0 \leq \delta \leq 1$ is her standard discount factor. Once again, following Loewenstein, O'Donoghue, and Rabin (2003), a person with projection bias perceives her intertemporal utility to be

$$(8) \quad \tilde{U}^t(conv_t, \dots, conv_T | s_t) = \sum_{\tau=t}^T \delta^{(\tau-t)} \tilde{u}(conv_\tau, s_\tau | s_t),$$

where \tilde{u} represents the perceived instantaneous utility described by Equation (6).

This framework illustrates how an individual's perceived intertemporal utility of purchasing a convertible at time t , \tilde{U}^t , can be overly influenced by s_t . Specifically, this framework 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 owning the convertible than when s_t is a bad state of the world for consuming a convertible (cold, cloudy weather).

One question that arises from this model is whether individuals correctly anticipate the path of states (s_t, \dots, 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.¹⁸ Loewenstein, O'Donoghue, and 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) lead to similar

18. 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 and Critcher 2011; Li, Johnson, and Zaval 2011).

incorrect predictions of future utility. Thus, it is difficult to separate these different types of projection bias and our analysis does not attempt to do so. However, the prevalence of weather information that is available to people during the time of our study, including their own experience of local weather patterns, argues against projection bias of states as the underlying mechanism. It is much harder to find information about future utility than it is to find information about the likelihood of future weather states.

Unlike a model of present-biased preferences, which would require that the entire weather-related effect that we estimated be driven by consumers overweighing the value they place on consumption in the first day or two of owning the vehicle, projection bias suggests that consumers mispredict the value they will receive from owning the convertible in every future state of the world when the weather is different from the weather at the time of purchase. Because the bias affects almost all periods rather than just one period, this model can more easily predict effects of the size that we estimate here than can a model of present-biased preferences. While we cannot conclusively determine that projection bias is the sole contributing factor in the results that we find, our results are consistent with this psychological model of behavior.

Conlin, O'Donoghue, and Vogelsang (2007) propose that one can test directly for projection bias by finding evidence of consumers deciding *ex post* that a decision was a mistake. Specifically, projection bias suggests that people who make a choice in one state of the world may realize the mistake when the state of the world changes. Conlin, O'Donoghue, and Vogelsang (2007) test for such mistakes by analyzing whether cold-weather clothing (boots, gloves, etc.) purchased by mail order was more likely to be returned if the purchase was made during very cold weather. In the vehicle market, projection bias mistakes might be identified by seeing vehicles that were purchased on atypical weather days reappear in the market (either as trade-ins or as subsequent used car sales) more quickly than vehicles that were purchased on days when the weather was seasonally typical. 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 that testing for early returns in the vehicle market is much harder than doing so for catalog orders. The first reason is simply a data limitation.

Although our data are extensive and represent a 20% sample of all new car dealerships in the United States, 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 dealership, and no chance of seeing it if that transaction happens person to person. Second, and perhaps more important, 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 winter 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 a convertible for a hardtop will require the individual to 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 (and again undergo the initial depreciation if the replacement vehicle is new). 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 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 data set, 2.37% of them reappear within one year as a trade-in or subsequent sale, 5.03% within two years, and 7.16% within three years.¹⁹ On average in the United States, owners keep their vehicles for just over five years (Polk 2010).

Our empirical strategy is to estimate whether convertibles that were purchased when the weather was atypically warm and four-wheel-drive vehicles that were purchased when the weather was atypically 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 VI report results for regressions in which the outcome variable is an indicator that equals 1 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, one, two, or three years. We control for DMA*week fixed effects to eliminate seasonal and geographic differences in how quickly vehicles are returned. Table VI shows

19. Unique identification numbers corresponding to individual vehicle identification numbers are used to track vehicles over time.

TABLE VI
EFFECT OF WEATHER ON QUICKLY TRADING IN A VEHICLE

	(1)	(2)	(3)
	Dep. var.: dummy variable if returned within 1–3 years		
	1 year	2 years	3 years
Percent of cars returned in window	2.37	5.03	7.16
Convertible	1.272** (.019)	2.302** (.030)	2.905** (.042)
Convertible interacted with:			
Temperature	.006 (.004)	.017** (.007)	.006 (.009)
Rainfall	.008 (.009)	.002 (.015)	–.018 (.021)
Snowfall	.181 (.131)	–.041 (.222)	–.142 (.289)
Slushfall	.063 (.053)	.028 (.094)	–.116 (.131)
Cloud cover	–.197 (.138)	–.036 (.228)	.332 (.312)
Four-wheel-drive	.285** (.006)	.929** (.006)	1.634** (.014)
Four-wheel-drive interacted with:			
Temperature	–.003* (.001)	–.005* (.002)	–.013** (.003)
Rainfall	–.005 (.003)	–.005 (.006)	.001 (.008)
Snowfall	.000 (.035)	.063 (.058)	.004 (.076)
Slushfall	.002 (.016)	–.019 (.028)	–.048 (.038)
Cloud cover	.006 (.047)	–.109 (.078)	–.124 (.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 transacted vehicle reappears in our data set (as a trade-in car or as a used-car sale) within one, two, or three years from the date of purchase on a dummy variable for whether the transacted vehicle was a convertible or a four-wheel-drive and an interaction between these dummies and weather variables at the time of purchase: temperature (°F), 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 data set is also restricted so as to eliminate all truncation (the data eliminate the last 1–3 years of car sales in the sample, respectively). All coefficients and standard errors have been multiplied by 100 for ease of presentation. * significant at 5%; ** significant at 1%.

that convertibles are overall 1.272 percentage points more likely to be returned within a year than other types of vehicles; four-wheel-drive vehicles are also more likely to be returned (by 0.285 percentage point) 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 on days with atypically warm weather. However, this result is statistically significant only in column (2). The point estimates suggest that when the weather is 10°F warmer than average for that DMA and week of the year, convertibles are 0.17 percentage point more likely to be returned within two years than hardtops (a 2.3% change relative to the baseline convertible return rate of 7.332%). The temperature interaction with four-wheel-drive vehicles is more consistently statistically significant, and indicates that a four-wheel-drive vehicle is more likely to be returned within one, two, or three years if it is purchased on an atypically cold day. Overall, our results for the effect of weather on returning vehicles, though clearly suggestive, is less strong than our evidence for the effect on purchasing vehicles. An important constraint we face is that the number of vehicles we see sold and then see reappear within our data is simply not that high. As a consequence, we have limited ability to identify differences in the rates at which vehicles are returned under different circumstances.

V.B. Salience

The second psychological mechanism that we consider in detail is salience. In a consumer context, salience refers to the idea that a customer's attention may be systematically directed toward certain features of a product and that those features will receive disproportionate weight in purchase decisions. The idea of salience is not new in the psychology literature, but it has been formalized and received renewed attention in the past few years (see, for example, Bordalo, Gennaioli, and Shleifer 2012, 2013; Koszegi and Szeidl 2013).²⁰

20. There have also been many recent empirical papers that find evidence of salience/attention. See, for example, Gabaix and Laibson (2006), Chetty, Looney, and Kroft (2009), Finkelstein (2009), Brown, Hossain, and Morgan (2010), Malmendier and Lee (2011), Lacetera, Pope, and Sydnor (2012), and Hastings and Shapiro (2013).

To understand how salience could predict the effects that we find, we consider the formalization of salience provided by Bordalo, Gennaioli, and Shleifer (2013). Their model predicts that consumers will place greater weight on product attributes that are salient at the time of purchase. The salience of an attribute is determined by how that attribute compares to the attributes of the products in a consumer's "choice context." The choice context includes the products among which a consumer can currently choose, the "choice set."²¹ In a dynamic setting, the choice context may also include products that the consumer expects to be available, even if those choices aren't currently available. For example, if prices have recently increased or decreased, products at their past prices could be part of the choice context. A consumer is assumed to have a reference good whose attributes equal the average of the attributes in the choice context. A particular attribute of a good is salient (and will therefore be weighed more heavily) when it differs by more from the reference good than other attributes do. For example, the price of a good will be more salient than the quality of a good if the price is more different from the price of the reference good than the quality is different from the quality of the reference good.

Hastings and Shapiro (2013) contain an application of this model. They consider whether salience can explain consumers' choices between gasoline grades that differ both in quality (octane rating) and price. They suppose that past prices inform consumers' price expectations and therefore are part of consumers' choice contexts when considering gasoline choices.²² Thus, if gasoline prices increase suddenly, the current prices of gasoline will differ substantially from the reference price for gasoline defined by the choice context, which includes the former, lower gasoline prices. This will increase the salience of prices, increasing the utility weight that consumers place on the price attribute relative to the quality attribute of gasoline, and thereby decreasing the premium they are willing to pay for higher octane gas relative to what they were willing to pay when the overall price levels were lower.

21. The choice context could include only a subset of choices actually available if a consumer restricts attention to a "consideration set" instead of the full choice set.

22. Hastings and Shapiro use the phrase "evoked set" to refer to a choice context, referring to the set of products that are evoked in a consumer's mind while she is considering a particular purchase decision.

In Hastings and Shapiro (2013), as in many of the examples described in Bordalo, Gennaioli, and Shleifer (2013), the value of the attribute itself varies across contexts and causes the attribute's salience to vary (for example, gasoline prices at different points in time, wine prices at stores versus restaurants, etc.). In our context, a convertible roof does not become "more convertible" in some contexts than others. Instead, there is another variable, weather, which varies over time and which can make the roof style of the car become more or less salient at different points in time.

For example, consider a situation in which a consumer is choosing between one of two cars: a high-priced convertible or a low-priced sedan. In a given purchase situation, a consumer may find the price to be the salient feature of the vehicle or may find the roof style to be salient. On a day with especially beautiful weather, having a convertible will seem particularly attractive compared to how it would seem on days with less nice weather. In this sense, the value of the convertible attribute today will be unusual compared to the value of the convertible attribute on "average" days, making it the salient feature, and making consumers value it more relative to price than on days with less good weather.²³

V.C. Distinguishing between salience and projection bias

As a coarse generalization, both salience and projection bias would predict that there should be more convertibles sold on

23. How product attributes are defined is an important part of specifying this model. If one defines the product attribute as "roof style" then both a warm, sunny day (which will make a convertible seem especially delightful to drive) and a cold, wet, blustery day (which will make a convertible seem drafty and unpleasant to drive) should make the roof style more salient, because both weather types will lead to a value of the roof type that is very different from the reference value. If the convertible style of the roof is the generally preferred style, then when weather makes it more salient, consumers should put more weight on roof style and be more willing to buy a convertible whether the increased salience arose from unusually good weather or unusually bad weather. This is inconsistent with our results, and probably inconsistent with most people's intuition. To have salience explain both increased sales of convertibles in good weather and decreased sales of convertibles in bad weather, we could define two product attributes that are associated with roof style: "enjoyableness" (an attribute experienced in good weather) and "draftiness" (an attribute experienced in bad weather). Good weather would make the enjoyableness attribute more salient (without affecting the salience of draftiness), leading to increased sales of convertibles on good weather days. Bad weather, however, would make the draftiness attribute more salient (without affecting the attribute of enjoyableness), and would lead to decreased sales on bad weather days.

sunny days and more four-wheel-drives sold on snowy days. Therefore, one might well ask if there is any way to distinguish whether salience or projection bias is a more likely explanation for the effects we have found. One difference between the mechanisms is that projection bias predicts that people will buy more convertibles when the weather is nice in an absolute sense, and salience predicts that people will buy more convertibles when the weather is unusually nice relative to some benchmark (since it is in comparison to the average that makes an attribute salient). Said another way, projection bias predicts something about the effect of weather “levels” and salience predicts something about the effect of weather “differences” or “surprises.”

Whether this distinction enables us to tease apart these two potential mechanisms depends on what kinds of weather “surprises” one thinks are relevant for salience. One definition of a weather surprise would be weather that differs from the average weather at that time of year. In other words, consumers are surprised when the weather is different from the expectation they have formed based on their experience of weather in past years. Another definition of a weather surprise would be weather that differs from what the weather has been recently, in the last few days or weeks, for example. In this case, consumers would be surprised when the weather changes from what it has been recently. In terms of the formal model of Bordalo, Gennaioli, and Shleifer, the distinction is whether the reference good is “driving a convertible in weather that is typical for this time of year” or whether the reference good is “driving a convertible in the weather experienced over the last few days or weeks.”

We cannot distinguish between projection bias and salience under the “surprise relative to expectations” interpretation. The reason for this has been described already. Summarizing briefly, weather follows a seasonal trend, which means that there is a very plausible discounted utility explanation for a correlation between weather and car sales. Therefore, we would not want to attribute a correlation in the seasonal pattern of weather and car sales to an intertemporal bias. The only empirical effect that we could attribute to a psychological bias (either projection bias or salience) is an effect of *atypical* weather on car sales (since a temporary weather fluctuation shouldn’t affect the discounted utility of owning a long-lived durable good). Unfortunately, the correlation of atypical weather with car sales would be predicted

by both mechanisms, and therefore doesn't enable us to distinguish between the two.

In contrast, under the "weather surprise relative to recent" interpretation of when particular attributes of a car should be salient, we may be able to distinguish between the two mechanisms. In the language of Bordalo, Gennaioli, and Shleifer, suppose that the "choice set" part of a consumer's choice context includes both convertibles and nonconvertibles on the day of purchase (with its associated weather) and the "expectations" part of the choice context includes the same vehicles in the days leading up to the actual purchase (which might have had different weather states). If this is the case, then the roof style of a car should be particularly salient when the weather on the date of purchase is not only nice but nice *relative to* the weather in the days leading up to the purchase. For example, if the weather was very cool and cloudy in the week leading up to the purchase, then warm and sunny weather on the day of purchase makes the convertible feature particularly unusual relative to the reference good in the choice context. Projection bias, on the other hand, does not predict that convertible sales would be higher on a good weather day that follows a string of bad weather days than on a good weather day that follows a string of other good weather days.

This prediction was tested in Section IV.A when we provided results from a distributed lag model.²⁴ For convertibles, we found that conditional on the weather at the time of purchase, colder weather in the days leading up to the purchase did not lead to increased sales relative to warmer weather in the days leading up to the purchase. Although this finding goes against the prediction the salience model would make if the previous days' weather is what the choice context expectation is based on, it is also possible that buying a car a few days earlier is not in a consumers' choice context (or maybe it is only a small part of the overall choice context that consumers use, making this effect hard to identify).

A prediction that could differ between a model of projection bias and salience is as follows. When the weather is extremely cold (e.g., 25°F), it could be the case that increasing the temperature by a small amount (e.g., to 35°F) would have no impact on the current utility associated with owning and driving a convertible. Thus, a model of projection bias would not predict that a

24. This prediction of a model of salience provides the same prediction as inter-temporal substitution of car purchases.

movement of temperature from 25°F to 35°F would have any impact on convertible sales. However, moving from 25°F to 35°F could change the salience placed on a convertible feature when making a car purchase even if current utility doesn't change. A model of salience could therefore result in a higher fraction of convertibles purchased as the temperature increased from 25°F to 35°F. Based on Figure I, we find that car sales are positively correlated with temperature even at very low temperature levels. This finding could be seen as evidence for salience rather than projection bias if one believes that differences in temperatures within the low range of temperatures will not affect the perceived utility of driving a convertible.

Overall, it is very difficult to distinguish between projection bias and salience in this domain; our primary results are consistent with either model.

VI. CONCLUSION

Many of the most important decisions that we make in life are those with long-term implications. Our ability to make these decisions well depends in part on our ability to forecast accurately our intertemporal utility. In this article, we have examined purchasing behavior in the car market and found evidence that consumers are affected by projection bias or salience, limiting their ability to make the decision that will maximize their long-term utility. We argue that our results imply that behavioral biases can have important implications for large-stakes markets and that these biases merit 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 (Lowenstein, O'Donoghue, and Rabin, 2003). Such laws could also provide incentives for sellers to help buyers be in a "cool" state before an important transaction or contract is made.²⁵ The Federal Trade Commission has an explicit "Cooling-Off Rule" that applies to situations when "[you] buy an item in your home

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

or at a location that is not the seller's permanent place of business."²⁶ 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 automobile sales even though there clearly can be high-pressure sale situations for this important durable good. Although 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 intertemporal behavioral biases can impact an important consumer durable good market, there are many questions about the behavioral bias we have documented 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 providing consumers with information about how intertemporal choices can be influenced by behavioral biases, 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 behavioral biases in intertemporal choice in various other important empirical contexts such as the decision of a couple to have a baby, whether to get married, which house to buy, or whether to accept a given job offer. Studying intertemporal choice bias in other important settings such as these may offer additional empirical tests that could shed light on the relative strength and importance of the underlying behavioral mechanisms, such as projection bias and salience, that have been hypothesized in the literature in these different choice contexts.

NORTHWESTERN UNIVERSITY AND NBER
UNIVERSITY OF CHICAGO AND NBER
BRIGHAM YOUNG UNIVERSITY
UNIVERSITY OF CALIFORNIA, RIVERSIDE

26. More information on the Federal Trade Commission's Cooling-Off Rule can be found at: <http://www.ftc.gov/bcp/edu/pubs/consumer/products/pro03.shtm>.

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournals.org).

REFERENCES

- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience Theory of Choice under Risk," *Quarterly Journal of Economics*, 127 (2012), 1243–1285.
- , "Salience and Consumer Choice," *Journal of Political Economy*, 121 (2013), 803–843.
- Brown, Jennifer, Tanjim Hossain, and John Morgan, "Shrouded Attributes and Information Suppression: Evidence from the Field," *Quarterly Journal of Economics*, 125 (2010), 859–876.
- Busse, M., "When Supply and Demand Just Won't Do: Using the Equilibrium Locus to Think about Comparative Statics," (Mimeo, Northwestern University, 2012).
- Busse, M., C. Knittel, and F. Zettelmeyer, "Are Consumers Myopic?: Evidence from New and Used Car Purchases," *American Economic Review*, 103 (2013), 220–256.
- Busse, M., D. Simester, and F. Zettelmeyer, "'The Best Price You'll Ever Get': The 2005 Employee Discount Pricing Promotions in the U.S. Automobile Industry," *Marketing Science*, 29 (2010), 268–290.
- Camerer, C., S. Issacharoff, G. Loewenstein, T. O'Donoghue, and M. Rabin, "Regulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism'," *University of Pennsylvania Law Review*, 151 (2003), 1211–1254.
- Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99 (2009), 1145–1177.
- Conlin, M., T. O'Donoghue, and T. Vogelsang, "Projection Bias in Catalog Orders," *American Economic Review*, 97 (2007), 1217–1249.
- Deschenes, O., and E. Moretti, "Extreme Weather Events, Mortality, and Migration," *Review of Economics and Statistics*, 91 (2009), 659–681.
- DellaVigna, S., "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47 (2009), 315–372.
- Finkelstein, Amy, "E-ZTax: Tax Salience and Tax Rates," *Quarterly Journal of Economics*, 124 (2009), 969–1010.
- Gabaix, Xavier, and David Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Market," *Quarterly Journal of Economics*, 121 (2006), 505–540.
- Hastings, J., and J. Shapiro, "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks," *Quarterly Journal of Economics*, 128 (2013), 1449–1498.
- Jacob, B., L. Lefgren, and E. Moretti, "The Dynamics of Criminal Behavior: Evidence from Weather Shocks," *Journal of Human Resources*, 42 (2007), 489–527.
- Koszegi, Botond, and Adam Szeidl, "A Model of Focusing in Economic Choice," *Quarterly Journal of Economics*, 128 (2013), 53–107.
- Lacetera, Nicola, Devin Pope, and Justin Sydnor, "Heuristic Thinking and Limited Attention in the Car Market," *American Economic Review*, 102 (2012), 2206–2236.
- Laibson, D., "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112 (1997), 443–477.
- Li, Y., E. Johnson, and L. Zaval, "Local Warming: Daily Temperature Change Influences Belief in Global Warming," *Psychological Science*, 22 (2011), 454–459.
- Loewenstein, G., T. O'Donoghue, and M. Rabin, "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics*, 118 (2003), 1209–1248.
- Malmendier, Ulrike, and Young Han Lee, "The Bidder's Curse," *American Economic Review*, 101 (2011), 749–787.
- O'Donoghue, T., and M. Rabin, "Doing It Now or Later," *American Economic Review*, 89 (1999), 103–124.

- Polk, "Consumers Continuing to Hold onto Vehicles Longer," November 3, 2010, available at https://www.polk.com/company/news/consumers_continuing_to_hold_onto_vehicles_longer_according_to_polk.
- Risen, J., and C. Critcher, "Visceral Fit: While in a Visceral State, Associated States of the World Seem More Likely," *Journal of Personality and Social Psychology*, 100 (2011), 777–793.
- Zettelmeyer, F., F. Scott Morton, and J. Silva-Risso, "Scarcity Rents in Car Retailing: Evidence from Inventory Fluctuations at Dealerships," Mimeo, Northwestern University, 2007.