

\$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions

By MEGHAN BUSSE, JORGE SILVA-RISSO, AND FLORIAN ZETTELMEYER*

Automobile manufacturers frequently use promotions involving cash incentives. While payments are nominally directed to either customers or dealers, the ultimate beneficiary of the promotion depends on the outcome of price negotiation. We use program evaluation methods to compare the incidence of these two types of promotions. Customers obtain 70 to 90 percent of a customer rebate, but only 30 to 40 percent of a dealer discount promotion, a \$500 difference for a typical promotion. Our leading hypothesis is that pass-through rates differ because of information asymmetries: customer rebates are well-publicized to customers, while dealer discount promotions are not. (JEL D82, L11, L15, L62, L81, M31)

Although retail demand for an automobile fluctuates due to changing economic conditions, seasonality, and the stage of the model's life cycle, manufacturers rarely vary published retail and invoice prices of a particular model over the course of the model year. The choice to have rigid prices is potentially very costly for auto manufacturers: inventory holding costs for automobiles are high, and so are the costs of changing production schedules to adapt to current demand. As a result, "incentive promotions" play an important role in automobile

manufacturers' product market strategies by enabling retail prices to adjust to fluctuating demand conditions. Incentive promotions take a variety of forms. The most common are cash rebates to customers, cash payments to dealers, subsidized interest rates for customers who finance through the manufacturer's captive lending arm, and lease incentives. In this paper, we focus on the two primary types of cash incentives, which we refer to by their industry terminology, "customer cash" and "dealer cash."

"Customer cash" refers to cash rebates that are directed to customers: once the customer and the dealer have negotiated the purchase price, the dealer hands the customer a check from the manufacturer for the promotion amount. The customer endorses the check over to the dealer, and the amount is immediately applied to the negotiated price. "Dealer cash," on the other hand, is a payment from the manufacturer directly to the dealer, and is paid for every vehicle the dealer sells during the specified promotion period. Both promotions typically range between \$300 and \$3,000.

From an economist's perspective, these two types of promotions provide an interesting comparison. While the promotion payments are nominally directed to one party or the other, who ultimately receives the benefit of the promotion depends on the outcome of the price negotiation process. For example, if a customer buying during a \$1,000 customer cash rebate were to agree to a price that is higher by \$200 than the price he or she would have negotiated without the promotion, then the customer's

* Busse: Haas School of Business, University of California, Berkeley, Berkeley, CA 94720 (e-mail: meghan@haas.berkeley.edu); Silva-Risso: The A. G. Anderson Graduate School of Management, University of California, Riverside, Riverside, CA 92521 (e-mail: jorge.silva-risso@ucr.edu); Zettelmeyer: Haas School of Business, University of California, Berkeley, Berkeley, CA 94720, and National Bureau of Economic Research (e-mail: florian@haas.berkeley.edu). Busse and Zettelmeyer are grateful for financial support from the National Science Foundation under awards SES-0550508 and SES-0550911. Any opinions, findings, and conclusions expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation. We are grateful for helpful comments from Tim Bresnahan, Severin Borenstein, J. P. Dubé, Wes Hartmann, Guido Imbens, Chris Knittel, Fiona Scott Morton, Kathleen Johnson, participants at the NBER IO Program Meeting and the HBS Strategy Conference, and seminar participants at UC Berkeley, UC Davis, UC San Diego, Stanford University, the Federal Trade Commission, and the University of East Anglia. We also thank two anonymous referees for their suggestions. We are especially grateful to Paul Gertler for extensive and valuable suggestions.

out-of-pocket expenditure would be lower by only \$800 compared to what he or she would otherwise have paid. The dealer would be reaping \$200 of the benefit of that customer cash promotion. Conversely, if a dealer were induced by a \$1,000 dealer cash promotion to agree to a price that is lower by \$500 than he or she otherwise would have, then the customer would obtain \$500 of the benefit of the dealer cash promotion. In short, from an economics perspective, a \$1,000 customer cash promotion and a \$1,000 dealer cash promotion are both \$1,000 of manufacturer-supplied economic surplus that will be divided between the two parties through the bargaining process.

The question of how promotional cash is divided between buyer and seller is analogous to the incidence of a per-unit tax or subsidy. A familiar result of elementary economics is that the incidence of a subsidy is invariant to whether it is directed to the buyer or the seller of a good. In other words, a subsidy of a given size will lead to the same increase in the total amount the seller receives and the same decrease in the buyer's out-of-pocket expenditure regardless of whether it is the buyer or seller who nominally receives the subsidy. Translating to the automobile promotions case, the seller is an automobile dealer, and the subsidizing third party is an auto manufacturer (required by dealer franchise laws to be a separate entity from the dealer). Customer cash is then a subsidy directed to the buyer, and dealer cash a subsidy directed to the seller. If the invariance result holds, then the rate of pass-through—the proportion of the promotional amount by which the customer's out-of-pocket expense falls—should be the same under customer cash and dealer cash promotions.

This paper investigates whether the invariance result holds in automobile promotions. We use a program evaluation approach to estimate the effect of a promotion on the price buyers pay, paying careful attention to the counterfactual against which the effect is measured in order to isolate the "treatment effect" of the promotion. Using both a difference-in-differences approach and a regression discontinuity approach, we find that customers obtain 70 to 90 percent of the surplus supplied by manufacturers in customer cash promotions, but only 30 to 40 percent of the surplus in dealer cash promotions. This is counter to the simple invariance of incidence analogy.

We propose that pass-through rates differ by promotion type because of differences in the information environments surrounding customer and dealer cash. Customer cash promotions are always publicized to potential customers, often in prime-time television advertisements. In contrast, dealer cash promotions are not advertised by manufacturers. While it is possible for a customer to find out about dealer cash promotions, they need to search specialized publications or Web sites. Overall, consumers are much less likely to be informed about dealer cash promotions than about customer cash promotions.

The game-theoretic bargaining literature suggests that a negotiating party who has incomplete information about his opponent will obtain a smaller share of the surplus in the negotiation than if that party were better informed. We suggest that these promotions can be viewed as a natural experiment on the effect of information asymmetry on bargaining outcomes: in the customer cash case, the parties are symmetrically informed about the availability of the manufacturer-supplied surplus, while in the dealer cash case, the customer will generally be uninformed about the available surplus. Our evidence is consistent with an information asymmetry explanation for the bargaining outcomes.¹ We also consider alternative explanations of our findings, some (although not all) of which we are able to rule out.

We think that understanding promotions in the automotive industry is also of interest because of their economic importance, and because of the central role promotions play in marketing. As such, this paper is related to an extensive literature which empirically measures how much of a manufacturer promotion gets passed through to final customers in the form of lower retail prices. Robert C. Blattberg and Scott A. Neslin (1990) give an overview of the issues surrounding how to measure the effect of promotions on retail prices (see chap. 11). Additional empirical investigations into the rate of pass-through of manufacturer promotions to retail customers include David Besanko et al. (2005), Michel Chevalier and Ronald C. Curhan (1976), and Rockney G. Walters (1989). These

¹ While there is experimental work and also some empirical literature on private information and bargaining in labor disputes, to our knowledge there is next to no empirical work in product markets, except for Fiona Scott Morton et al. (2006).

studies have been primarily of packaged goods sold in supermarkets, a setting that is important to study in its own right. The setting, however, is not suited to analyzing the effect of asymmetric information on bargaining. This is because, first, supermarket prices are posted rather than negotiated and, second, coupons (the equivalent of customer cash) have a price discriminatory function which customer cash, where redemption happens automatically during closing, does not have (Aviv Nevo and Catherine Wolfram, 2002). Finally, a few papers analyze promotions specifically in the automobile industry. Norris Bruce et al. (2005) show conditions under which manufacturers would want to offer trade promotions to their dealers, and Bruce et al. (forthcoming) introduce one reason why manufacturers would want to offer cash rebates, namely that the rebate helps consumers compensate for negative equity in the vehicle they are trading in.

The paper proceeds as follows. In Section I we present the data. In Section II we discuss relevant estimation issues and our empirical approach. In Section III we estimate the pass-through rates of the different types of promotions. In Section IV we test the validity of the key assumptions that were maintained when identifying these effects. In Section V we discuss the interpretation of the empirical results. We conclude in Section VI.

I. Data

We have combined two types of data for this analysis. The first is data on automobile transactions from a sample of 15 to 20 percent of the dealerships in California from September 1, 1998, to December 31, 2000. The data were collected by a major market research firm and include every transaction within the time period for the dealers in the sample. For each transaction, we observe the exact vehicle purchased, the price paid for the car, the dealer's cost of obtaining the car from the manufacturer, information on a potential trade-in vehicle, and demographic information on the customer. We have supplemented these transaction data with promotion listings that detail all types of promotions available during the sample period, including customer and dealer cash. For each promotion, we observe the promotion amount, the starting and ending dates of the promotion, and any restrictions on the promotion's appli-

cation (typically by region and/or specific equipment package, or "trim level").

In this paper, we restrict our attention to cash transactions, namely transactions that are not leases and are not financed through manufacturer-backed financing.² The reason is that for these transactions, customer cash and dealer cash are the promotions that are relevant for pricing. This leaves us with 133,424 transactions.

A. Dependent Variable

We will use transaction prices as dependent variables in the estimation. The price observed in the dataset is the pre-sales tax price that the customer pays for the vehicle, including factory-installed accessories and options, and including any dealer-installed accessories contracted for at the time of sale that contribute to the resale value of the car.³

Conceptually, we would like our price variable to measure the customer's total wealth outlay for the car. In order to capture this, we make two modifications to the observed transaction price. First, we subtract the customer cash amount if the car is purchased under a customer cash rebate since the manufacturer pays that amount on the customer's behalf. Second, we subtract from the purchase price any profit the customer made on his or her trade-in (or add to the purchase price any loss made on the trade-in). The price the dealer pays for the trade-in vehicle minus the estimated wholesale value of the vehicle (as booked by the dealer) is called the *TradeInOverAllowance*. Dealers are willing to trade off profits made on the new vehicle transaction and profits made on the trade-in transaction, which is why the *TradeInOverAllowance* can be either positive or negative. A customer who loses money on the trade-in transaction pays for the new vehicle partially in-kind with the trade-in vehicle. By subtracting the *TradeInOverAllowance*, we adjust the negotiated (cash) price to include this payment.

² Note that a cash transaction from the perspective of the dealer need not be a cash transaction from the perspective of the customer. If a customer has obtained a loan from a bank, it is a cash transaction from the dealer's perspective.

³ Dealer-installed accessories that contribute to the resale value include such items as upgraded tires or a sound system, but would exclude such options as undercoating or waxing.

B. Controls

We control for car fixed effects. A “car” in our sample is the interaction of make, model, model year, body type, transmission, displacement, doors, cylinders, and trim level. We drop from the sample any thus-defined cars that have fewer than 200 sales in our sample, since the smaller number of observations limits what we learn from these cars, and because we want to be able to estimate car fixed effects to control for many of the factors that contribute to the price of a car. The remaining sample contains 942 cars.

To control for time variation in prices, we define a dummy *EndOfMonth* that equals one if the car was sold within the last five days of the month. A dummy variable *WeekEnd* specifies whether the car was purchased on a Saturday or Sunday. In addition, we use dummies for each week in our 121-week sample period (September 1998 to December 2000) to control for seasonal effects and for inflation. If there are volume targets or sales on weekends, near the end of the month, or seasonally, we will pick them up with these variables.

We control for the number of months between a car’s introduction and when it was sold. This proxies for how new a car design is and also for the dealer’s opportunity cost of not selling the car. We distinguish between sales in the first four months, months 5 to 13, and month 14 and later with a dummy variable for each category.

We control for the competitiveness of each dealer’s market. For each dealership we count the number of dealerships with the same nameplate that fall in a zip code that is within a ten-mile radius of the zip code of the focal dealership. We control for cases where one owner owns several franchises by counting only the number of separately controlled entities.

We also control for the income, education, occupation, and race of buyers by using census data that the data provider matches with the buyer’s address from the transaction record. The data are at the level of “block groups,” which, on average, contain about 1,100 people. Finally, we control the geographic region in which the car was sold (northern or southern California).

Table 1 presents summary statistics for the data. Twenty-six percent of transactions in our

sample involve customer cash, and 18 percent involve dealer cash. The average amount of customer cash observed for transactions in our sample that involve customer cash is \$1,242 (median \$1,000). The average amount of dealer cash among transactions that involve dealer cash is \$932 (median \$700). The average transaction price of a new vehicle in our data is \$25,490. The table also presents customer demographics.

II. Estimation Approach

The aim of our paper is to estimate the “treatment effect” of promotions on prices. Our primary empirical problem is to find the correct counterfactual against which to measure this effect. At an intuitive level, estimating the treatment effect of a promotion means comparing pricing with a promotion to dealer pricing without a promotion. The chief complication in doing this, however, is that manufacturers might be more likely to instigate incentive promotions when prices are either low or declining due to a slump in demand. This means that the price observed in the periods in which a manufacturer chose not to have a promotion is not necessarily what the price *would have been* in the periods in which a manufacturer chose to have a promotion had it chosen instead not to have a promotion. If we do not correctly account for this in choosing the counterfactual, we could overestimate the rate at which surplus is passed through to customers because we would attribute a low customer price to the promotion when part of the low price might have been attributable to demand conditions.

We use two different empirical approaches in this paper, a difference-in-differences approach and a regression discontinuity approach. Conceptually, the difference-in-differences approach uses the prices of similar cars that are not on promotion to estimate the counterfactual price of a car that is on promotion at a given time. This is implemented by incorporating week–vehicle segment fixed effects to control for underlying changes in price. Thus, the estimated change in price that is attributed to the promotion is the change that is net of the contemporaneous change in prices of other cars within the same vehicle segment.

An alternative way to estimate the counterfactual price is to use a car’s own price when it is not on promotion. Conceptually, this is what

TABLE 1—SUMMARY STATISTICS[†]

Variable	<i>N</i>	Mean	Median	St. dev.	Min	Max
Price	133,424	25,490	23,487	10,382	5,988	109,755
Customer cash (CC)	34,296	1,242	1,000	669	10	7,805
GM Card	1,204	1,934	1,785	1,215	2	7,043
Dealer cash (DC)	24,620	932	700	819	200	5,000
Sales manager incen.	2,319	141	50	181	20	500
Sales rep incen.	3,601	147	75	154	25	500
# CC prom. in seg.	133,424	6	6	4.7	0	17
# DC prom. in seg.	133,424	2.9	2	2.5	0	13
Avg. CC prom. in seg.	133,424	231	191	209	0	1,254
Avg. DC prom. in seg.	133,424	112	53	155	0	2,305
# Competing dealers	133,424	3.7	3	2.8	0	24
Female	133,424	0.28	0	0.45	0	1
%Asian	133,424	0.11	0.064	0.12	0	0.97
%Black	133,424	0.033	0.011	0.083	0	1
%BlueCollar	133,424	0.23	0.2	0.15	0	1
%CollegeGrad	133,424	0.36	0.35	0.18	0	1
%Hispanic	133,424	0.13	0.096	0.1	0	0.55
%LessHighSchool	133,424	0.1	0.074	0.1	0	1
%HouseOwnership	133,424	0.69	0.76	0.24	0.0043	1
%Executives	133,424	0.19	0.19	0.083	0	1
%Professional	133,424	0.19	0.18	0.091	0	1
%Technicians	133,424	0.031	0.028	0.022	0	1
Income	133,424	6.5	6.2	2.8	1.1	15
Income ²	133,424	49	38	43	1.1	225
MediaHHSIZE	133,424	2.9	2.8	0.57	1.5	6
MedianHouseVal.	133,424	2.6	2.4	1.2	0.075	5
Weekend	133,424	0.27	0	0.44	0	1
EndOfMonth	133,424	0.22	0	0.41	0	1
EndOfYear	133,424	0.026	0	0.16	0	1
SouthernCal	133,424	0.48	0	0.5	0	1

[†] For customer cash, GM Card, dealer cash, sales manager incentives, and sales rep incentives, “*N*” reports the number of nonzero observations. Hence, the summary statistics reflect observations with nonzero values.

our second approach, regression discontinuity, does. This approach has been used primarily in applications in which researchers need to evaluate programs in which participants are not randomly assigned to treatment and control groups. Such situations include when subjects self-select into treatment, or when treatment is assigned on the basis of need or some other characteristic related to the outcome that is the aim of the program. This means that estimating the treatment effect by regressing outcomes on indicators of treatment is likely to produce biased estimates.

The regression discontinuity approach utilizes the fact that even in programs without random assignment, there are often discontinuities in treatment among subjects who are otherwise similar in the characteristics that influence the outcome of interest, and that these discontinuities are likely to lead to discontinuities in outcomes. The average differences in outcomes

between groups just to one side and just to the other side of a treatment discontinuity can give a consistent estimate of the average treatment effect (Jinyong Hahn et al., 2001; Guido W. Imbens and Joshua D. Angrist, 1994).⁴

⁴ In contrast to the standard application of regression discontinuity, our treatment variables (manufacturer promotions) are not binary, since promotions can be of different amounts. Notice, however, that identification in the regression discontinuity approach does not require the treatment to be *binary*. Identification requires only that there be a *discontinuity* in the treatment, and that there not be coincident discontinuities in other covariates at the same point. Other papers have also used a regression discontinuity approach with a continuous treatment and covariates. For example, Sandra E. Black (1999) measures the effect of test scores on housing prices. The paper uses a continuous treatment variable (test scores) and covariates (house characteristics). The discontinuity in the paper lies in houses that are neighboring each other but fall across school district boundaries. While test scores change discretely at the boundaries, the value of test scores will vary from boundary

In this paper, we are interested in estimating the effect of a manufacturer promotion (the treatment) on an outcome measure (price). We do not believe that manufacturers apply the treatment randomly: promotions are likely to be applied at times when sales are slow, and when customers are willing to purchase only if offered relatively low prices. We believe, however, that the underlying demand conditions that determine the rate of sales and the customers' price elasticity are likely to change fairly little over the course of a few weeks.⁵ Thus, in applying the regression discontinuity approach, we restrict our attention to one week on either side of when a promotion begins, ends, or changes in the amount of cash being offered.⁶ The treatment discontinuity we make use of is a temporal one: the amount of customer cash available changes abruptly between one day and the next. Even if underlying demand conditions are trending one way or another within this short window, as long as there is no discontinuous change, except for what is the result of the change in promotion, the regression discontinuity approach will consistently estimate at least the local average treatment effect.⁷

III. Promotion Pass-Through

In this section we estimate how much of the manufacturer-supplied surplus is passed through to customers, depending on the type of promotion used to supply the surplus. We begin with the difference-in-differences approach to estimate the rate of pass-through, and then turn to the regression discontinuity approach.

to boundary. Thus, the paper regresses housing prices on the actual level of test scores, not just a binary indicator for whether they are low or high.

⁵ We know from industry sources that there is typically a gap of several weeks between the most recent demand and inventory information that was used in making the decision to initiate a promotion and the start of the promotion.

⁶ Note that one can think of the regression discontinuity approach as akin to a robustness check on a subsample.

⁷ The average treatment effect estimated by regression discontinuity is local in the sense that it is measured in the neighborhood of the treatment discontinuity. See Hahn et al. (2001).

A. Difference-in-Differences

We wish to estimate the effect of customer cash promotions and dealer cash promotions on the out-of-pocket price that customers pay. Our dependent variable is therefore the vehicle price net of rebate, which we denote P_{ijt} , the price that customer i pays for vehicle j at date t . The specification we estimate is

$$(1) \quad P_{ijt} = \lambda_c \text{CustCash}_{jt} + \lambda_d \text{DealCash}_{jt} \\ + \beta_1 \mathbf{X}_i + \beta_2 \mathbf{X}_j + \beta_3 \text{DealerComp}_{ij} \\ + \mu_j + \tau_{JT} + \varepsilon_{ijt}.$$

CustCash_{jt} and DealCash_{jt} are the amounts of customer cash and dealer cash available for vehicle j at date t . Since manufacturers typically make promotion decisions by nameplate, model, and model year (e.g., 1999 Pontiac Grand Am), these variables are unique for a nameplate-model-model year triple. In cases where promotions also vary by region, our promotion variable will record the promotion as being available only for transactions in that region. \mathbf{X}_i is a vector of the buyer's individual and neighborhood customer demographic characteristics including sex, race, income, education, employment type, and home ownership. \mathbf{X}_j is a vector of control variables, some of which are defined only by t (weekend, end of month, and end of year) and some of which depend on j and t (time since model introduction). DealerComp_{ij} is a measure of how competitive the dealer at which customer i purchased vehicle j is; specifically, the measure used is the number of competing dealers of the same nameplate as vehicle j within a ten-mile radius of this dealer. μ_j are car fixed effects. For these fixed effects, we use a very fine definition of a car, namely the cross product of make, model, model year, body type, transmission, displacement, doors, cylinders, and trim level. τ_{JT} is a week-vehicle segment fixed effect where J is the segment (e.g., SUVs, compact cars, etc.) that contains car j , and T is the week that contains purchase date t . The data cover 121 weeks.

The primary variables of interest are λ_c and λ_d , which measure the extent to which rebates are passed through to customers. If either λ_c or λ_d is equal to 0, that implies that none of the surplus from the respective type of promotion is passed through to customers. In this case, the retailer is the sole beneficiary of the promotion.

If λ_c or λ_d is equal to -1 , then the customer obtains the full amount of the respective rebate in the form of a lower price. One can interpret $100 \cdot |\lambda|$ as the percentage of the rebate the customer obtains.

The identifying assumption in this approach is that other cars in the same segment that are not under promotion in a given week are a valid counterfactual for the prices that would have been obtained on the promoted car in the absence of a promotion. We will investigate this identifying assumption in Section IVA.

Table 2 reports the results of estimating this specification. In column 1, the customer cash coefficient implies that 88 percent of customer cash is passed through to customers, while 39 percent of dealer cash is passed through. This difference, statistically significant at the 1-percent confidence level, indicates that customers receive a substantially greater share of the surplus from a customer cash promotion than from a dealer cash promotion.⁸

The pass-through rates estimated by difference-in-differences are very similar to what is reported in this subsection if, instead of week-segment fixed effects, we use month-subsegment fixed effects (fewer degrees of freedom longitudinally but more cross-sectionally) or even week-subsegment fixed effects.

B. Regression Discontinuity

In the regression discontinuity approach, we analyze only transactions that occur within a window of one week on either side of a promotion change from zero to some positive amount, from one amount to another amount, or from

⁸ A possible concern in this specification is that the distribution of observable characteristics in “customer cash,” “dealer cash,” and “no cash” observations may not share a common support, leading to bias in the pass-through estimates. This concern appears to be unfounded. In a prior version of this paper, we use a propensity score approach to adjust the sample so that observations share a common support. In this specification, the estimates are similar, indicating that 84 percent of customer cash is passed through to customers, while only 31 percent of dealer cash reaches buyers. In this propensity score estimation, we account for the fact that there are multiple treatments (customers can buy with or without customer cash, and with or without dealer cash). We do not, however, explicitly account for promotion amounts of different sizes. See Busse et al. (2004) for a detailed description of the procedure and results.

some positive amount to zero.⁹ In contrast to the difference-in-differences approach, we have to determine the pass-through of customer cash and dealer cash promotions in separate estimations: to identify the effect of, for example, customer cash promotions, the regression discontinuity approach dictates that we use data only immediately before and after a change in *customer cash* promotions but *not* data surrounding changes in dealer cash promotions. The analogous procedure applies to estimating dealer cash promotions.

The regression equation in the regression discontinuity approach is identical to the one used in the difference-in-differences approach, except that the specification does not rely on week-vehicle segment fixed effects in order to identify the treatment effect of the promotion, so τ_{JT} does not appear; in its place, we include a week fixed effect, τ_T , for each of the 121 weeks of the sample.

When we restrict the sample to observations in the windows surrounding changes in customer cash promotions, the coefficient of interest is the coefficient on *CustomerCash* (λ_c), which has the same interpretation as in the difference-in-differences approach. When using this sample, the coefficient on *DealerCash* (λ_d) cannot be interpreted as the pass-through rate identified by regression discontinuity, since λ_d is not identified by observations immediately before and after a dealer cash change. Instead, we include *DealerCash* when using the customer cash window sample merely to control for the price effects of dealer cash. The converse applies when we restrict the sample to observations in the windows surrounding changes in dealer cash promotions. Then the coefficient of interest is the coefficient on *DealerCash* (λ_d), while *CustomerCash* merely controls for the price effects of customer cash.

The identifying assumption in this approach is that the underlying willingness to pay of customers who buy just before and just after a change in a promotion is the same. This would be violated if, for example, there were deal-prone customers

⁹ The results are robust to changes in the size of the window around the promotion event. When we estimate pass-through rates using only two days before and after the change in a customer cash or dealer cash promotion, the estimated pass-through rates change little. See Busse et al. (2004) for details.

TABLE 2—PRICE EFFECTS, BASIC RESULTS[†]

	Difference-in-differences		Regression discontinuity	
	(1)	(2A)	(2B)	
Customer cash	-0.88 (0.03)**	-0.81 (0.07)**	-0.78 (0.12)**	
Dealer cash	-0.39 (0.07)**	-0.38 (0.14)**	-0.31 (0.07)**	
GM Card	-1.06 (0.03)**	-1.13 (0.10)**	-1.13 (0.09)**	
Competition	-7.66 (5.60)	-15.03 (9.59)	-18.63 (9.63) ⁺	
Female	144.17 (12.01)**	139.74 (45.20)**	203.25 (57.59)**	
%Asian	-261.75 (61.95)**	-188.40 (239.89)	-115.40 (173.48)	
%Black	474.48 (92.72)**	605.65 (381.97)	470.96 (251.97) ⁺	
%BlueCollar	208.11 (89.62)*	766.91 (366.29)*	398.17 (355.60)	
%College	-283.87 (86.40)**	-103.70 (300.31)	-132.31 (271.87)	
%Hispanic	-36.43 (86.37)	-936.23 (352.61)**	-68.13 (343.59)	
%LessHighSchool	-166.54 (111.15)	-276.23 (422.61)	-363.69 (451.70)	
%HouseOwnership	14.57 (35.88)	-59.68 (161.61)	52.05 (142.31)	
%Executive	253.73 (113.07)*	542.61 (501.99)	-404.04 (517.10)	
%Professional	217.33 (119.99) ⁺	-268.76 (406.52)	231.95 (426.55)	
%Technicians	147.06 (236.15)	-101.90 (1166.86)	-500.29 (1190.77)	
MedianHHIncome	-21.17 (11.07) ⁺	-30.22 (55.34)	-49.37 (46.16)	
(MedianHHInc.) ²	3.01 (0.60)**	2.99 (3.19)	3.79 (2.32)	
MedianHHSIZE	-24.95 (13.59) ⁺	22.45 (59.12)	54.37 (47.24)	
MedianHouseVal.	-8.10 (10.35)	-0.80 (36.52)	7.91 (36.05)	
Weekend	-27.51 (16.79)	-124.92 (76.11)	-107.97 (49.35)*	
EndOfMonth	-55.57 (16.61)**	-90.90 (71.22)	-93.05 (58.23)	
EndOfYear	14.61 (69.49)	133.47 (321.75)	-1,265.44 (192.47)**	
ModelMonth5-13	31.46 (36.42)	231.02 (181.81)	-156.72 (183.77)	
ModelMonth14+	-99.24 (55.96) ⁺	156.11 (268.65)	-172.80 (322.56)	
SouthernCal	-243.03 (48.64)**	-227.94 (107.87)*	-274.99 (107.55)*	
Constant	26,610.52 (163.18)**	23,482.33 (555.17)**	24,598.77 (798.96)**	
Car fixed effects	Yes	Yes	Yes	
Other fixed effects	Week × Segment	Week	Week	
Observations	133,424	6,296	7,046	
Adj. R-squared	0.97	0.95	0.95	

Note: Robust SEs in parentheses.

[†] MedianHouseValue in \$100,000. Income in \$1,000.

⁺ Significant at 10-percent level.

* Significant at 5-percent level.

** Significant at 1-percent level.

with a different willingness-to-pay or bargaining ability, who wait until the beginning of a promotion to enter the market. We will test this identifying assumption in Section IVB.

Table 2 reports the estimated coefficients from the two regression discontinuity specifications. Overall, both estimates are slightly smaller than the estimates in the difference-in-differences approach. In column 2A of Table 2, 81 percent of customer cash is estimated to be passed through to customers, compared to 31 percent of dealer cash (in column 2B).¹⁰ Again we conclude that customers obtain a greater share of surplus from customer cash promotions than from dealer cash promotions.

IV. Identification Issues

Having estimated the pass-through rates of customer and dealer cash using both a difference-in-differences and a regression discontinuity approach, we now test the validity of the two key assumptions that were maintained when identifying these effects.

A. Difference-in-Differences: Test of Prepromotion Trends

The identifying assumption of the difference-in-differences approach is that the prices of cars in the same segment that are not on promotion in a given week are a valid counterfactual for the prices that would have been obtained on the promoted car in the absence of a promotion. Although we cannot observe this directly, we can examine the trends of promoted and non-promoted cars in the period just prior to the promotion. If the trends are similar between cars that are soon to be promoted and cars that are not, that gives some assurance that the non-promoted cars are a valid counterfactual in the promotion period.

¹⁰ As in the difference-in-differences specification, the distribution of observable characteristics in “customer cash,” “dealer cash,” and “no cash” observations may not share a common support, leading to bias in the pass-through estimates. This concern appears to be unfounded. Adjusting the sample so that observations share a common support changes the estimates very little. The estimate of customer cash pass-through remains 81 percent, while the estimate of dealer cash pass-through increases from 31 percent to 32 percent. See Busse et al. (2004) for a detailed description of the procedure and results.

To test this, we estimate for each vehicle segment in each month of the sample two separate daily time trends of price. One trend is estimated for cars that will go on promotion within 30 days. The other trend is estimated for cars that are not about to go on promotion.

We first restrict the sample to transactions that occurred with neither customer cash nor dealer cash. Next, we calculate for each observation the earliest date after the observation date (t) that the car (j) is on either a customer or dealer cash promotion. We denote this as T_{jt} , the start date of the next promotion for car j at date t . Using T_{jt} , we define an indicator variable $I(T_{jt} \leq t + 30)$, which will equal one if a transaction occurs 30 days or less before the start of a promotion for that car. We then define monthly time trend variables for each month as $\theta_{t,M} = n$ if date t is the n^{th} day of month M and $\theta_{t,M} = 0$ if date t does not fall in month M . Finally we define I_J , an indicator variable that equals one if car j is in segment J .

We then run the following regression:

$$(2) \quad P_{ijt} = \alpha_M + \beta_{M,J} I_J \cdot \theta_{t,M} + \gamma_{M,J} I_J \cdot I(T_{jt} \leq t + 30) \cdot \theta_{t,M} + \delta_1 \mathbf{X}_i + \delta_2 \mathbf{X}_{jt} + \mu_j + \eta_{ijt},$$

where P_{ijt} is the price paid by customer i for car j at date t in month M , α_M is a fixed month effect, \mathbf{X}_i and \mathbf{X}_{jt} are the demographic and time effects from equation (1), and μ_j are car fixed effects. The coefficient $\beta_{M,J}$ will measure the daily trend of prices over the days in month M for segment J . The $\gamma_{M,J}$ coefficient will measure any differences in the daily time trend in month M for prices of cars that are observed within 30 days of the date they will next be on promotion.¹¹

Our test of equal trends in the prepromotion period is to test whether the $\gamma_{M,J}$ coefficients are equal to zero. In unreported results, 19 of the

¹¹ We also restrict the sample to transactions from month-segment combinations where we observe at least 10 transactions of cars in that month from that segment which will be on promotion within 30 days and where we also observe at least 10 transactions of cars in that month from that segment that will not be on promotion within 30 days.

111 estimated $\gamma_{M,j}$'s are statistically different from zero at the 5-percent confidence level, and an additional 5 at the 10-percent level. While many coefficients are statistically indistinguishable from zero individually, we can reject the hypothesis that the $\gamma_{M,j}$'s are jointly statistically equal to zero.

We next ask whether the estimated differences in prepromotion trends that *are* statistically different from zero are large enough to explain the price effects we estimated in Table 2, and which we attributed to pass-through. The largest negative, statistically significant estimated $\gamma_{M,j}$ is $-\$11.86$ per day; the next largest is $-\$3.47$. The estimated coefficients in Table 2 imply that cars on customer cash promotion, controlling for other covariates, sell for $\$1,056$ less than the prices of cars that are not on promotion (85 percent estimated pass-through times $\$1,242$ average customer cash promotion amount). In order for this price effect to be explained entirely by the prices for about-to-be-promoted cars drifting downward by $\$11.86$ more per day than nonpromoted cars, it would have to be the case that the promotion lasted 178 days, about 6 months. For a difference of $\$3.47$ in the daily price trend, the promotion would have to last 609 days, or about 20 months, longer than most cars are even available. Since $-\$11.86$ and $-\$3.47$ are the largest of the estimated time trends, most of which are insignificant, and since each is estimated only for one month for one segment, we believe that any bias caused by different price trends over time is not large enough to account for our results.

The analogous calculation for dealer cash indicates that the extra $-\$11.86$ in downward trend for about-to-be-promoted cars could account for the price effect estimated in Table 2 after 55 days, and a differential of $\$3.47$ after 188 days. While the former is not atypical for the length of a dealer cash promotion, the $\$11.86$ differential is estimated for only one month-segment combination.

These estimates provide little support for the argument that what we estimate as promotional pass-through in Table 2 can be explained as differences in price trends. The finding that about-to-be-promoted cars in most cases do not have different price trends from other cars in the segment just before a promotion starts is also supportive of our use of the prices of

nonpromoted cars to control for underlying price trends of promoted cars in the difference-in-differences specification.

B. Regression Discontinuity: Pass-Through by Promotion Length

The key maintained assumption in the regression discontinuity approach is that transaction prices during the week just before the promotion starts are a valid counterfactual for transaction prices during the first week of a promotion. This would be violated if the customers who purchase just before a promotion starts differed in some way that was related to negotiated prices from customers who purchase just after the promotion starts.¹² In particular, this would be the case if there are "deal-prone" customers, who are particularly effective negotiators, and who wait to purchase a car until a promotion is offered. This would mean that the set of customers whom we observe buying before the promotion would pay higher prices on average, *with or without a promotion*, than the set of customers whom we observe buying during a promotion would pay, *with or without a promotion*. Thus, our pass-through coefficient would wrongly attribute to promotions what is in fact a price difference due to unobservable buyer characteristics. While our detailed demographic data control for important differences among consumers, we are concerned that such data may not adequately capture "deal-proneness."

To estimate whether the pass-through estimate in the regression discontinuity approach is biased by consumers who "wait for a deal," we use the difference-in-differences approach to es-

¹² The maintained assumption in the regression discontinuity approach would also be violated if dealers change their pricing behavior strategically in response to an upcoming promotion. This seems not to be the case. Dealers typically find out about promotions only two to three days before the promotion start date. In Busse et al. (2004), we reestimate our basic regression discontinuity specification *excluding* the three days directly before and after a promotion change. If dealers react strategically to an upcoming promotion, this should not be reflected in this restricted sample. We find that 82 percent and 36 percent of customer cash and dealer cash, respectively, gets passed through to customers, a small change compared to 81 percent and 31 percent in the basic specification.

timate how the pass-through rate changes over the life of the promotion. Assuming that the pent-up demand from customers waiting for a deal comes into the market early and that these customers are effective negotiators, we should observe a higher pass-through rate at the beginning of a new or increased promotion than when the promotion has been offered for some time or has been decreased from the previous level. However, a higher initial pass-through rate should be expected only for customer cash promotions. This is because consumers are usually not aware when dealer cash promotions are offered. Hence, deal-prone consumers are less likely to time their purchase to coincide with the start of a dealer cash promotion.

We build on the difference-in-differences specification from column 1 in Table 2. We split the customer cash and dealer cash variables by whether a promotion is an increase or a decrease from a previous level; we expect deal-prone consumers to be more likely to buy after a promotion change only if the promotion is increased. In addition, we interact the resulting two customer cash and two dealer cash variables with dummies for whether the transaction occurred when a promotion change had been in place for 0 to 14 days, 15 to 30 days, 1 to 2 months, 3 to 6 months, or 6 months or more. Consistent with our conjecture that there may be some deal-prone consumers, we find that the pass-through rate for customer cash is greater in weeks 1 and 2 than in all subsequent weeks, but only for new or increased customer cash promotions. When a customer cash promotion is new or an increase, pass-through is estimated to be 96 percent in the first two weeks in contrast to 75 to 80 percent for all subsequent weeks (see Table 3). The difference between the first two weeks and later periods is statistically significant at the 1-percent level. When the current customer cash promotion is a decrease from a previous level, the rate of pass-through does not show any statistically significant change over time. In contrast to customer cash, the pass-through rate of dealer cash does not show evidence of attracting deal-prone customers. When a dealer cash promotion is an increase, pass-through rates stay statistically unchanged for the first eight weeks and then rise, while when a dealer cash promotion is a decrease, there is no discernible pattern to pass-through rates over time.

These findings are consistent with some deal-prone consumers waiting to purchase cars until they become aware that a promotion is being offered. This suggests that the regression discontinuity approach somewhat overestimates the pass-through rate of customer cash promotions. However, even if we were to use the customer cash pass-through rate for weeks beyond the second week as estimated in Table 3, which would be 0.80, instead of 0.88 as estimated in column 1 of Table 2, our qualitative finding that customer cash gets passed through to consumers at twice the rate of dealer cash would not change.

V. Interpretation

A \$1,000 customer cash promotion and a \$1,000 dealer cash promotion both consist of \$1,000 of manufacturer-supplied economic surplus that will be divided between the two parties through a bargaining process. The party to whom the manufacturer transfers the money will not necessarily obtain all the benefit from it; the division of the benefits depends on the outcome of the bargaining process. Indeed, the basic “tax-incidence” motivation described in the introduction would suggest that the division of the surplus would be invariant to how the surplus is nominally directed. In the preceding empirical results, however, we find dramatically different outcomes for the amount of surplus customers obtain between the two cases. What accounts for this difference? In this section we propose that this finding results from different information environments surrounding customer cash and dealer cash. We also discuss several alternative explanations.

A. Asymmetric Information

One striking difference between customer cash and dealer cash is that the two bargaining parties are symmetrically informed about customer cash, while the customer is typically uninformed about dealer cash. We propose that this difference could explain the different rates of customer pass-through that we estimate for the two different types of promotions.

A rich bargaining theoretic literature relates information asymmetries between bargaining

TABLE 3—PRICE EFFECTS, IDENTIFICATION ISSUES[†]

Customer cash (increase) (weeks 1 and 2)	-0.96 (0.03)**	Dealer cash (decrease) (weeks 3 and 4)	-0.49 (0.20)*
Customer cash (increase) (weeks 3 and 4)	-0.80 (0.04)**	Dealer cash (decrease) (weeks 5 to 8)	-0.22 (0.23)
Customer cash (increase) (weeks 5 to 8)	-0.80 (0.03)**	Dealer cash (decrease) (weeks 9 to 26)	-1.05 (0.13)**
Customer cash (increase) (weeks 9 to 26)	-0.77 (0.04)**	Dealer cash (decrease) (weeks 26+)	-0.42 (0.28)
Customer cash (increase) (weeks 26+)	-0.75 (0.06)**	GM Card	-1.07 (0.033)**
Customer cash (decrease) (weeks 1 and 2)	-0.84 (0.08)**	Competition	-7.35 (5.59)
Customer cash (decrease) (weeks 3 and 4)	-0.88 (0.08)**	Weekend	-27.94 (16.77) ⁺
Customer cash (decrease) (weeks 5 to 8)	-0.92 (0.06)**	EndOfMonth	-56.09 (16.61)**
Customer cash (decrease) (weeks 9 to 26)	-0.91 (0.07)**	EndOfYear	15.41 (69.17)
Customer cash (decrease) (weeks 26+)	-1.27 (0.11)**	ModelMonth5-13	31.94 (36.03)
Dealer cash (increase) (weeks 1 and 2)	-0.28 (0.06)**	ModelMonth14+	-103.04 (57.39) ⁺
Dealer cash (increase) (weeks 3 and 4)	-0.35 (0.06)**	SouthernCal	-243.49 (48.56)**
Dealer cash (increase) (weeks 5 to 8)	-0.38 (0.08)**	Constant	26,644.62 (164.92)**
Dealer cash (increase) (weeks 9 to 26)	-0.48 (0.08)**	Car fixed effects	Yes
Dealer cash (increase) (weeks 26+)	-0.42 (0.14)**	Other fixed effects	Week × Segment
Dealer cash (decrease) (weeks 1 and 2)	-0.33 (0.19) ⁺	Observations	133,424
		R-squared	0.97

Note: Robust SEs in parentheses.

[†] Unreported demographic characteristics include: census block percentages of residents who are college graduates; who have less than high-school education; who are blue-collar workers, executives, professionals, or technicians; who are Asian, black, or Hispanic; who are female; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value.

⁺ Significant at 10-percent level.

* Significant at 5-percent level.

** Significant at 1-percent level.

parties to the division of surplus in the negotiation.¹³ To match our case, it is the buyer who should be thought of as the possibly uninformed party. We consider two different protocols for dynamic bilateral negotiation. In a “buyer-offer game,” only the buyer (the uninformed party) is allowed to make offers, which the seller is allowed only to reject or accept. If the seller rejects an offer, the buyer can make another offer. The game ends when the seller accepts an

offer. Drew Fudenberg et al. (1985) and Faruk Gul et al. (1986) show that under the “stationary equilibrium” refinement, the buyer-offer game allows the buyer to screen seller types by a series of sequential, increasing price offers. In terms of our empirical prediction, the salient feature of this equilibrium is that, as long as buyers are not infinitely patient, the buyer is not able to perfectly screen among the seller types and is therefore worse off than he or she would be in a situation in which he or she had complete information about the seller’s reservation price. This same prediction also comes out of an “alternating-offer game.” In such a game the buyer and seller alternate in making proposals.

¹³ For a detailed review of the game-theoretic literature on bargaining with incomplete or private information, see the excellent review papers by John Kennan and Robert Wilson (1993) and Lawrence M. Ausubel et al. (2002).

Ausubel and Raymond J. Deneckere (1998) show in such a model that under the “assuredly perfect equilibrium” refinement, there exists a unique equilibrium in which the buyer is able to screen seller types, albeit imperfectly. As in the buyer-offer game, the buyer’s equilibrium payoff is bounded from above by what he or she could extract in the complete-information game.

The prediction common to this class of models is that a negotiating party that has incomplete information about its opponent will obtain a smaller share of the surplus in the negotiation than if that party were better informed. The prediction from theory in the context of car promotions is that buyers should be able to obtain a larger share of a customer cash promotion (which they know about) than of a dealer cash promotion (which they are unlikely to know about). Our results are clearly consistent with this prediction: customers obtain 70 to 90 percent of the surplus supplied by manufacturers in customer cash promotions, but only 30 to 40 percent of the surplus in dealer cash promotions.¹⁴

We can test an additional connection between information and pass-through rates by exploiting a particular market institution. General Motors offers the “GM Card,” a credit card that accumulates a rebate toward the purchase of a new GM car. In the transactions data, we can identify when a GM Card rebate was applied. This makes an interesting comparison to customer and dealer cash because this is manufacturer-supplied surplus for which the *customer* has an information advantage over the *dealer*. Hence, we expect that consumers can appropriate more of the GM card rebate than of either customer cash or dealer cash. The results follow this prediction: column 1 of Table 2 shows that customers obtain 106 percent (statistically indistinguishable from 100 percent) of the surplus from GM Card rebates.¹⁵

¹⁴ Notice that the fact that customers obtain some of the dealer cash surplus is not inconsistent with customers not knowing that dealer cash is on the table. This is because the dealer will get the promotional payment from the manufacturer only if a car is actually sold. This gives the dealer an incentive to lower prices to customers more than he or she would have without the promotion, in order to entice customers to complete the sale.

¹⁵ One alternative explanation for this finding is that GM Card holders are “bargain-hunting, penny-pinching” types and would therefore be likely to negotiate for good prices. An interpretation of the GM Card that would give the opposite prediction is that GM Card holders hold the GM

B. Differences in Market Conditions

We have argued so far that the difference in the information environments surrounding each type of promotion can explain our results. We now consider the first in a series of alternative explanations for the finding that customer cash gets passed through to consumers at a higher rate than dealer cash.

One alternative explanation is that the differences in pass-through rates are due to differences in market conditions under which manufacturers decide to offer one or the other type of promotion. Our estimates indicate the average pass-through of a customer cash promotion in periods in which manufacturers chose to have customer cash promotions and the average pass-through of a dealer cash promotion in periods in which manufacturers chose to have dealer cash promotions. Thus far, our results do not necessarily imply that the pass-through of a promotion would increase if a manufacturer were to switch from a dealer cash to a customer cash promotion of equal amount. In particular, if each type of promotion is well suited to a particular configuration of market conditions, and manufacturers want to maximize promotion pass-through, it may be that manufacturers are optimally matching promotions to particular conditions.

To investigate this alternative explanation, we proceed as follows: we use measures of market conditions to explicitly model the probability that a car on promotion is promoted either with customer cash or with dealer cash.¹⁶ This estimation yields the predicted probability for each transaction that the car is sold on a customer cash as opposed to a dealer cash promotion. We use this probability as a propensity score to exclude observations from the estimation for which there are no comparable market

Card because they have strong preferences for GM vehicles, which would make them less likely to negotiate low prices.

¹⁶ Industry sources tell us that the primary variables that manufacturers use when deciding whether to initiate a promotion are vehicle profitability at the dealer level, inventories, total vehicle sales, and market share within the vehicle’s segment. These variables are monitored on a weekly to monthly basis. The lag between the information used and the promotion decision is typically one to three months, and in rare cases as short as two weeks.

conditions across the two promotion states. This leaves us with a sample which contains only cars that were sold on promotion, and which shares a common support in observable market conditions across customer and dealer cash promotions. We use this sample to identify pass-through rates while explicitly controlling for the market conditions on which manufacturers base their decision to offer one or the other type of promotion (see Busse et al., 2004, for details of the analysis).

The results suggest that, controlling for the market conditions under which manufacturers choose one or the other type of promotion, the pass-through rates for customer cash remain at least twice as large as the pass-through rates for dealer cash and statistically significantly different at the 1-percent level. While the point estimates indicate that 102 percent of customer cash and 48 percent of dealer cash is passed through, the confidence intervals are large enough to accommodate the estimates of 88 percent and 39 percent found in the original difference-in-differences specification in column 1 of Table 2.

While our somewhat coarse measures of market conditions cannot perfectly capture the promotion generation process, we do think they are sufficiently informative measures that if a substantial portion of the difference between dealer cash and customer cash pass-through rates were due to differences in market conditions, adjusting for common support and controlling for market conditions would decrease the difference between the estimated pass-through rates. The fact that this has not happened makes it less likely that differences in market conditions are responsible for the estimated differences in pass-through rates between customer and dealer cash.

C. Prospect Theory

Prospect theory proposes that the effect on utility of an outcome will depend on an individual's reference point, and whether the outcome is a loss or a gain relative to that reference point. One of the elements of this theory is "loss aversion," meaning that a loss of a given size will have a negative effect on utility that is greater in magnitude than the increase in utility that would result from a gain of the same amount.

This suggests an alternative explanation for our pass-through result based on the fact that customer cash is nominally directed to the customer, whereas dealer cash is nominally directed to the dealer. Suppose that a customer's reference point when customer cash is offered is to obtain the entire amount and when dealer cash is offered to obtain none of it; in other words, suppose that customers believe that customer cash "belongs" to them and dealer cash "belongs" to dealers. If that is the case, then prospect theory would suggest that customers should negotiate harder to obtain the surplus from customer cash than from dealer cash because the former is a loss avoided and the latter is a gain achieved.

With the data at hand, we are logically unable to distinguish between prospect theory and asymmetric information as an explanation. This is because the condition "parties are symmetrically informed" coincides perfectly with "promotion is labeled as belonging to customer," and "customer is at an information disadvantage" coincides perfectly with "promotion is labeled as belonging to dealer." In order to distinguish between the two explanations, we would need a condition where the two models make different predictions; for example, a promotion that is directed to the dealer, but only the customer knows about it. We could construct such a scenario in a laboratory experiment, but the actual configuration of promotions that we observe does not enable us to distinguish between the explanations.¹⁷

D. Differences in Customer Types

The different customer and dealer cash pass-through estimates could also be due to differences in the elasticities of demand of customers who buy under each promotion type. If particularly elastic buyers are attracted to the market by customer cash, then one might expect prices to be lower under customer cash than under dealer cash, in part because of the difference in elasticity.

Since we cannot estimate these elasticities, we cannot test this hypothesis directly. We can,

¹⁷ Of course, the two explanations are not mutually exclusive; if parties were symmetrically informed, the differences might persist with smaller magnitude.

however, present four simple tests to assess whether buyers who purchase with customer cash differ from those who purchase with dealer cash.

Differences in Observable Demographic Characteristics.—In the first test we determine differences in observable demographic characteristics between the two types of buyers. If customer cash attracts more elastic buyers, we should observe that demographic characteristics differ in systematic ways between buyers who purchase with customer cash and with dealer cash. For example, we would expect customer cash buyers to have on average lower incomes than dealer cash buyers, holding the purchased car constant.

To determine whether this is the case, we restrict the sample to transactions that were made with either customer cash or with dealer cash, approximately 39,000 observations. We then compare the mean of each demographic variable for customer cash transactions with the mean for the variable for dealer cash transaction, within the detailed definition of a “car.”

The means of the two customer groups do not differ for income, percent of college graduates, percent blue collar, percent executives, percent professionals, percent technicians, median house value, percent black, or gender. We do find statistically significant differences for percent Hispanic (1.5 percentage points lower for customer cash), percent Asian (0.5 percentage points lower for customer cash), percent with less than high-school education (0.9 percentage points lower for customer cash), household size (0.03 persons lower for customer cash), and also percent who own their own home (1.3 percentage points higher for customer cash). Notice, however, that the sign of the differences does not support the conventionally assumed relation between demographics and elasticities. We would normally expect that Hispanic consumers, those with less education, and those with larger families to be more price elastic; however, they are represented less among those who buy with customer cash. We have no prior regarding the price elasticity of Asians or homeowners. While we cannot conclude that there are no observable demographic differences between customer cash and dealer cash buyers, the observed differences are small and the sign of the differences does not readily suggest that

buyers who buy under customer cash are more elastic than those who buy under dealer cash.

Differences in Observed Behavior.—In the second test we determine differences in observable behavior between the two types of buyers. We begin by recategorizing “cars” by name plate, model, model year, body type, doors, transmission type, drive-train type, trim level, but *not the number of cylinders or engine displacement*. If customer cash attracts more elastic buyers, we should observe that, within this new, broader “car” definition, customer cash buyers would opt for less expensive engines (namely, fewer cylinders and smaller displacement) than dealer cash buyers.

To test whether this is the case, we again restrict the sample to transactions that were made with either customer cash or with dealer cash, but not with both. Among these transactions, 17.4 percent are for cars which offer an option for the number of cylinders, almost all with two choices. For these cars, we compare the percent of transactions with the larger number of cylinders for buyers who purchased with customer cash versus dealer cash, holding the “car” constant. We do not find a significant difference. We get the same result if we repeat the procedure with engine displacement.

Elasticity Differences from Intertemporal Price Discrimination.—In the third test, we allow elasticity differences between customers to be unobservable, and test implications of one model of how such differences might arise. In the model of intertemporal price discrimination by Joel Sobel (1991), a seller faces a steady inflow of both high valuation (i.e., inelastic) and low valuation (i.e., elastic) customers. The seller sells at high prices to the inelastic customers for some period of time, and then has a sale to sell to the cumulated low valuation customers, clearing them out of the market. This might be the motivation for auto manufacturers to use customer cash: a publicized sale to draw elastic customers into the market for a period of time. This would lead to the customers who buy during customer cash promotions being more elastic than customers who buy during other times (including dealer cash). In this case, early on during the customer cash promotion there should be a relatively high proportion of elastic buyers (who have been waiting for a sale). As

the promotion wears on, more and more of the pent-up high elasticity buyers will have purchased; the proportion of high elasticity buyers will fall to the proportion represented in the steady inflow.

Our earlier results on pass-through by promotion length show that pass-through rates are higher (meaning prices are lower for a given promotion amount) for the first two weeks of a customer cash promotion, but that there is no decrease in pass-through (meaning no rise in price for a given promotion) over the course of months after that. Even if we disregard the pass-through rates estimated for the first two weeks as being the result of elastic buyers coming into the market, the pass-through rate for the rest of the time period is 80 percent, about twice the pass-through rate estimated for dealer cash.

Differences in the Substitution between Price and Nonprice Selling Inducements.—Our final test is based on the idea that different customers may trade off price and nonprice selling inducements differently. Suppose that dealer cash allows dealers to adjust the mix they offer to customers of price discounts and other inducements such as advertising, additional sales effort, discounted service contracts, or after-market options which are “thrown in” with the sale. Price-inelastic customers—those who are willing to accept nonprice inducements in place of price discounts—would then be relatively more easily persuaded to buy with dealer cash; these customers would receive part of the promotional surplus in cash and part “in kind” of some sort. Meanwhile, customer cash would attract price-sensitive customers who prefer to receive promotional surplus as cash, i.e., a discount on the price of the car in question.

We observe two items that may serve as inducements and are not related to the price of the car as we measure it. First, we observe whether the transaction includes after-market options. These are items, such as “underbody wax coating,” that do not add to the resale value of the car and are not included in our price variable. If dealer cash buyers are more likely to value nonprice inducements than customer cash buyers, we would expect that dealers “throw in” after market options to dealer cash buyers instead of reducing the price (excluding after-market options) of the vehicle. This should result in dealer cash transactions being more

likely to include after-market options than customer cash transaction. Instead, we find that 53.38 percent of dealer cash transactions and 53.39 percent of customer cash transactions include after-market options, a difference that is not statistically significant.¹⁸ This is again based on the restricted sample of transactions that were made with either customer cash or with dealer cash, but not with both.

Second, we observe the profits that dealers realize on extended service contracts that they sell in connection with new vehicles. Service agreements are bought by 17 percent of customers, and conditional on purchase, give the dealer an average of \$581 in profits (median is \$500, standard deviation is \$357). If dealer cash buyers are more likely to value nonprice inducements than customer cash buyers, we would expect dealers to “throw in” discounted service contracts to dealer cash buyers instead of reducing the price of the vehicle. We find that dealer profits on those contracts are no lower, however, for transactions that occur with dealer cash than with customer cash.

In summary, we have presented evidence (albeit necessarily incomplete) indicating that the price elasticity of customer cash buyers seems not to be higher than that of dealer cash buyers. While we cannot reject definitively that this explains part of our observed result, we think that the magnitudes of the effects suggest that it cannot explain the entire empirical result.

VI. Conclusion

This paper has analyzed how the outcomes of customer-dealer price negotiations differ between promotions that direct manufacturer rebates to customers and to dealers. Overall, our results tell a remarkably consistent story across approaches and across specifications. Customers obtain 70 to 90 percent of the surplus supplied by manufacturers in customer cash promotions, but only 30 to 40 percent of the surplus in dealer cash promotions.

We have tested the validity of the key assumptions that were maintained when identifying pass-through rates using both a difference-in-differences and a regression discontinuity approach. First, we

¹⁸ Unfortunately, we cannot tell how many such options were included with the car and what profit the dealer made.

have analyzed whether nonpromoted cars in the same segment are a valid counterfactual for promoted cars—an assumption maintained in the difference-in-differences approach. We have found that nonpromoted cars are not a perfect control for promoted cars; however, the potential bias is not large enough to change our conclusions. Second, we have investigated whether deal-prone consumers wait to purchase a car until a customer cash promotion is on—the regression discontinuity approach assumes that they do not. We have found evidence consistent with promotions attracting deal-prone consumers; however, the potential bias in the pass-through rate is small and does not change our findings.

Our leading hypothesis to explain why customers obtain more of the surplus available under customer cash than under dealer cash is that consumers are informed about the availability of customer cash promotions, but generally uninformed about whether dealer cash is available. We consider several alternatives to the asymmetric information hypothesis. We conclude that the empirical result is not due to differences in the demand conditions under which the two types of promotions are used. We also find little evidence that the result is due to differences in underlying demand elasticities between consumer types. With our data, however, we are unable to rule out prospect theory as an alternative explanation.

Our leading hypothesis is consistent with the theoretical prediction that when customers are at an information disadvantage, they are disadvantaged in negotiations. Under this hypothesis, the information disadvantage of consumers regarding car manufacturer promotions is substantial: for a promotion of average size, consumers receive \$500 less of the surplus if they do not know that the promotion is on the table. This is, to our knowledge, one of the first field measurements of the importance of information asymmetries in product markets.

REFERENCES

- Ausubel, Lawrence M.; Cramton, Peter and Deneckere, Raymond J.** “Bargaining with Incomplete Information,” in Robert J. Aumann and Sergiu Hart, eds., *Handbook of game theory with economic applications*. Vol. 3. Amsterdam: Elsevier Science, North-Holland, 2002, pp. 1897–1945.
- Ausubel, Lawrence M. and Deneckere, Raymond J.** “Bargaining and Forward Induction.” Unpublished Paper, 1998.
- Besanko, David; Dubé, Jean-Pierre and Gupta, Sachin.** “Own-Brand and Cross-Brand Retail Pass-Through.” *Marketing Science*, 2005, 24(1), pp. 123–37.
- Black, Sandra E.** “Do Better Schools Matter? Parental Valuation of Elementary Education.” *Quarterly Journal of Economics*, 1999, 114(2), pp. 577–99.
- Blattberg, Robert C. and Neslin, Scott A.** *Sales promotion: Concepts, methods, and strategies*. Englewood Cliffs, NJ: Prentice-Hall, Inc., 1990.
- Bruce, Norris; Desai, Preyas and Staelin, Richard.** “The Better They Are, the More They Give: Trade Promotions of Consumer Durables.” *Journal of Marketing Research*, 2005, 42(1), pp. 54–66.
- Bruce, Norris; Desai, Preyas and Staelin, Richard.** “Enabling the Willing: Consumer Rebates for Durable Goods.” *Marketing Science* (forthcoming).
- Busse, Meghan; Zettelmeyer, Florian and Silva-Risso, Jorge.** “\$1000 Cash Back: Asymmetric Information in Auto Manufacturer Promotions.” National Bureau of Economic Research, Inc., NBER Working Papers: No. 10887, 2004.
- Chevalier, Michel and Curhan, Ronald C.** “Retail Promotions as a Function of Trade Promotions: A Descriptive Analysis.” *Sloan Management Review*, 1976, 18(1), pp. 19–32.
- Fudenberg, Drew; Levine, David K. and Tirole, Jean.** “Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information,” in Alvin E. Roth, ed., *Game theoretic models of bargaining*. Cambridge: Cambridge University Press, 1985, pp. 73–98.
- Gul, Faruk; Sonnenschein, Hugo and Wilson, Robert.** “Foundations of Dynamic Monopoly and the Coase Conjecture.” *Journal of Economic Theory*, 1986, 39(1), pp. 155–90.
- Hahn, Jinyong; Todd, Petra and Van der Klaauw, Wilbert.** “Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design.” *Econometrica*, 2001, 69(1), pp. 201–09.
- Imbens, Guido W. and Angrist, Joshua D.** “Identification and Estimation of Local Average Treatment Effects.” *Econometrica*, 1994, 62(2), pp. 467–75.

- Kennan, John and Wilson, Robert.** "Bargaining with Private Information." *Journal of Economic Literature*, 1993, 31(1), pp. 45–104.
- Nevo, Aviv and Wolfram, Catherine.** "Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals." *RAND Journal of Economics*, 2002, 33(2), pp. 319–39.
- Scott Morton, Fiona; Zettelmeyer, Florian and Silva-Risso, Jorge.** "What Matters for Price Negotiations: Evidence from the US Auto Retailing Industry." Unpublished Paper, 2006.
- Sobel, Joel.** "Durable Goods Monopoly with Entry of New Consumers." *Econometrica*, 1991, 59(5), pp. 1455–85.
- Walters, Rockney G.** "An Empirical Investigation into Retailer Response to Manufacturer Trade Promotions." *Journal of Retailing*, 1989, 65(3), pp. 46–62.