

\$1000 Cash Back:
Asymmetric Information in Auto Manufacturer Promotions*

Meghan Busse

University of California, Berkeley

Jorge Silva-Risso

University of California, Riverside

Florian Zettelmeyer

University of California, Berkeley and NBER

February 2004

*We are grateful for helpful comments from Severin Borenstein and seminar participants at UC Berkeley and Stanford. We are especially grateful to Paul Gertler for extensive and valuable suggestions. Addresses for correspondence: Haas School of Business, UC Berkeley, Berkeley CA 94720-1900. E-mail: meghan@haas.berkeley.edu, florian@haas.berkeley.edu, jorge.silva-risso@ucr.edu

\$1000 Cash Back: Asymmetric Information in Auto Manufacturer Promotions

Abstract

Automobile manufacturers make frequent use of promotions that give cash-back payments. Two common types of cash-back promotions are rebates to customers, which are widely publicized to potential customers, and rebates to dealers, which are not publicized. While the payments nominally go entirely to one party or another, the real division of the manufacturer-supplied surplus between dealer and customer depends on what price the two parties negotiate. We show that customers obtain most of the surplus in cases when they are likely to be well-informed about the promotion (customer rebate), and about half the surplus when they are likely to be uninformed (dealer rebate). However, manufacturers ultimately do not care about the price effect of a promotion; a thousand dollar promotion costs a manufacturer the same whether it is directed to customers or dealers. Instead, manufacturers are interested in whether a given promotion leads to the sale of an additional car. We find that customer cash promotions, consistent with the pass-through results, have a larger dollar-for-dollar effect on sales than do dealer cash promotions.

1 Introduction

Promotions play a very large role in automobile manufacturers' product market strategies. Several distinctive features of the automobile industry lead to promotions being an important tool in reconciling inflexible prices, fluctuating demand, and costly inventory storage. The first of these features is that the ability of automobile manufacturers to influence retail prices for automobiles is restricted by state-level dealer franchise laws. These laws require that manufacturers sell cars to dealers at a constant per unit price, preventing them from offering volume discounts to high-volume dealers. Manufacturers are also largely forbidden from price discriminating; they are required to offer any given car model at equal prices to all dealers in a market. Furthermore, manufacturers by custom do not vary published retail and invoice prices of a particular model over the course of the model year. While published manufacturer prices are fixed, retail demand for automobiles is very volatile due to changing economic conditions and seasonality. This makes price rigidity potentially very costly for auto manufacturers: inventory holding costs for automobiles are high, and so are changes to production schedules to adapt to current demand. As a result, auto manufacturers use "incentive promotions" to adjust to fluctuating demand conditions. Such promotions take a variety of forms. The most common are cash rebates to customers, cash rebates 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 effect of cash rebates. We refer to these rebates by their industry terminology, namely "customer cash" for cash rebates that are directed to customers and "dealer cash" for cash rebates that are directed to dealers.

Customer cash promotions are always publicized to potential customers, often in prime-time television advertisements by regional dealer associations. The size of the rebate ranges typically from \$500 to \$2000. In practice, customer cash promotions are administered as follows: If a customer buys the specified vehicle during the time window of the promotion, then 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 then endorses the check over to the dealer, and the amount is immediately applied to the agreed-upon purchase price of the vehicle.

In contrast, dealer cash promotions are not advertised by manufacturers. While it is possible for a customer to find out if a dealer promotion is currently available, customers will not be informed about their existence unless they specifically search in specialized publications, or more recently through websites such as Edmunds.com. Overall, consumers are much less likely to be informed about dealer cash promotions than customer cash promotions.

From an economist's perspective, these two types of promotions provide an interesting

comparison. A \$1000 customer cash promotion and a \$1000 dealer cash promotion are both \$1000 of manufacturer-supplied surplus that will be divided between the two parties through the bargaining process. However, in the customer cash case, both parties know that the surplus is on the table, while in the dealer cash case, generally only the dealer knows that the surplus is on the table. This forms the basis for our first research question: How does the information asymmetry between dealers and consumers about manufacturer-supplied surplus affect the division of this surplus in customer-dealer negotiations?

While this question is of interest to economists, it is ultimately not the question that manufacturers care about. From the perspective of a manufacturer, \$1000 of customer cash and \$1000 of dealer cash are equally costly, irrespective of how the promotion amount gets divided between customer and dealer. What manufacturers care about is selling an incremental car. This motivates our second research question: Do different types of promotions have different incremental sales effects?

In answer to the first question, we expect that customers will obtain a greater share of the surplus, the better informed they are about the surplus that is to be divided. This implies that a given promotion amount in the form of customer cash should lower the final transaction price by more than the same promotion amount in the form of dealer cash. Consistent with this prediction we find that customers obtain 80-90% of the surplus supplied by manufacturers in customer cash promotions, but only 40-50% of the surplus in dealer cash promotions.

If promotions increase sales purely due to a price reduction to customers, our first result suggests that customer cash promotions should be more effective in increasing sales than dealer cash promotions. However, automobiles are a product for which there is substantial sales effort, and for which we believe that the salesperson can potentially have a large effect on whether a car is sold or not. If dealers do indeed retain more of the surplus when manufacturers offer dealer cash, and this better motivates sales effort, then it is possible that dealer cash will be more effective at increasing car sales. We find that this is not the case: while a \$500 increase in customer cash increases vehicle sales by 8%-13%, we find that \$500 of dealer cash increases sales by only 3-7%. Anecdotally, we know that dealers communicate to manufacturers through field reps that they prefer dealer cash to customer cash. The results of this paper suggest that dealer cash may indeed be more profitable for dealers. But the results also suggest that dealer cash may not be the instrument that best motivates sales. Manufacturers have traditionally strongly relied on their field reps and their dealers to assess the success of promotions. Our findings suggest that they might be better served by directly evaluating promotions with transactions data that has recently become available in the automotive industry.

We think that understanding incentive promotions is of independent interest since they play a key role in the automotive industry. More importantly, however, we think that this analysis

provides a rare opportunity to examine the effect of asymmetric information on bargaining outcomes in a product market.¹ Finally, we also think that this paper can shed additional light on the larger question of the effectiveness of “push” vs. “pull” promotions, which so far have been studied only in the supermarket industry.

The paper proceeds as follows. In section 2 we discuss the relevant literature. In section 3 we present the data. In section 4 we discuss relevant estimation issues and our empirical approach. In section 5 we estimate the pass-through rate of the different types of promotions. In section 6 we analyze the sales effect of these promotions. In section 7 we present a preliminary analysis of the promotion generation process. We conclude in section 8.

2 Promotions

Promotions are typically classified as “push” vs. “pull” promotions. Push promotions are intended to motivate the retailer to promote the product to customers, for example, through end-aisle displays or sales. Pull promotions are intended to motivate customers to seek out the product, for example, through coupons. The question of the extent to which manufacturer promotions are passed through to retail customers, and what the sales effects of such promotions are is not a novel question, although the question has gotten more attention in the marketing literature than in the economics literature. Blattberg and Neslin (1990) give an overview of the issues surrounding how to measure the effect and—from a manufacturer’s perspective—effectiveness of promotions (see chapter 11).

The literature on promotions has fallen into three primary strains. The first strain are models, primarily game theoretic models, of the decision by manufacturers to offer promotions to retailers who then sell the good to final customers (Lal (1990), Lal, Little, and Villas-Boas (1996), Gerstner and Hess (1991)).

These models are linked to a second strain of literature, also game theoretic, which models just the retailer’s decisions about setting promotional, discounted, or sale prices to retail customers. This strain contains both papers that model retail price promotions without any manufacturer promotion (Narasimhan (1988), Varian (1980), Raju, Srinivasan, and Lal (1990), and Rao (1991)) and with such manufacturer promotions (Moorthy (2003), Kumar, Rajiv, and Jeuland (2001), and Tyagi (1999)).

A third strain of literature attempts to measure empirically the rate of pass-through of manufacturer promotion to customers in the form of lower retail prices, and the effect of promotions

¹While 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 IO, except for Scott Morton, Zettelmeyer, and Silva-Risso (2003).

on sales (Besanko, Dubé, and Gupta (forthcoming), Chevalier and Curhan (1976), and Walters (1989)). Recently, promotions have also been of interest to industrial organization economists. Nevo and Wolfram (2002) analyze whether couponing by manufacturers is consistent with price discriminatory behavior.

All three strains of literature are set in the context of supermarkets. This is explicitly true for the empirical literature, which generally uses supermarket based data, and is often the motivation or application given by theoretical papers.

In analyzing the differential effect of push vs. pull promotions, auto retailing has several advantages over the supermarket context. First, a pull promotion in supermarkets, such as couponing, often has a price discriminatory aspect to it (not all customers will clip and use a coupon). This is not the case for customer cash promotions in automotive retailing since all consumers receive the promotion. Second, contracts between supermarkets and packaged goods suppliers can have many kinds of transfer payments which are not observable to researchers. As Berto Villas-Boas (2003) shows, using only per unit prices and ignoring other transfer payments can lead to erroneous conclusions. In the auto retailing context, dealer franchise laws restrict contracts to linear per unit prices, with the exception of the promotions which we observe. Hence, since we observe all monetary flows and since pull and push promotions in auto retailing are identical expect for the party to whom the promotion is targeted, we can better isolate what is purely due to “pushing” vs. “pulling” than is possible in the context of supermarkets.

3 Data

We have combined two types of data for this analysis. The first is data on automobile transactions from a sample of 15-20% of the dealerships in California from September 1998 to December 2000. The data are 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 (nameplate, model, model year, trim level, body type, number of doors, engine, etc.). We also observe the price paid for the car, the dealer’s cost of obtaining the car from the manufacturer, demographic information on the customer, detailed information on the trade-in vehicle if the customer used a trade-in, and the profitability of the car to the dealership. We also observe in these data the amount of customer cash rebate, if any, the customer received.

In these data, however, we do not observe the dealer cash rebates which were available at the time of sales. We have thus supplemented these transaction data with promotion listings. In this second set of data, we observe all types of promotions available during the sample

period, including customer cash, dealer cash, subsidized interest rates (APR incentives), lease incentives, and incentives given directly to sales managers and sales reps. For each promotion, we observe the promotion amount, the starting and ending dates of the promotion, and any restrictions on the promotion’s application. The most common kinds of restrictions are that the promotion is available only in certain regions of the country (or has varying promotion amounts in different areas of the country), or that the promotion is available only for certain trim levels (or have varying promotion amounts based on 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. 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. In future work, we hope to explore the effect of APR incentives and of lease incentives.

3.1 Dependent variables

We will use prices, weekly sales by model, and model market shares within subsegments as dependent variables in the estimation.

The price observed in the dataset is the price that the customer pays for the vehicle including factory installed accessories and options and the dealer-installed accessories contracted for at the time of sale that contribute to the resale value of the car.² The *Price* variable we use as the dependent variable is this price net of rebates, minus what is known as the *TradeInOverAllowance*. *TradeInOverAllowance* is the difference between the trade-in price paid by the dealer to the consumer and the estimated wholesale value of the trade-in vehicle (as booked by the dealer). We adjust for this amount to account for the possibility, for example, that dealers may offer consumers a low price for the new car because they are profiting from the trade-in.

We calculate weekly sales by nameplate, model and model year (e.g. 1998 Honda Accord, 1999 Toyota Camry) from the number of cash transactions in our data. To obtain market shares we divide these sales by the weekly sales of all cars in the “subsegment” classification employed by our data provider (e.g. Compact SUV, Luxury SUV, Fullsize SUV, Entry Compact, Premium Compact, etc.).

3.2 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. This leaves 942

²Dealer-installed accessories that contribute to the resale value include items such as upgraded tires or a sound system, but would exclude options such as undercoating or waxing.

thus-defined cars after dropping cars with fewer than 200 sales in California. We exclude these data because 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.

To control for time variation in prices, we define a dummy *EndOfMonth* that equals 1 if the car was sold within the last 5 days of the month. A dummy variable *WeekEnd* specifies whether the car was purchased on a Saturday or Sunday to control for a similar, weekly effect. In addition, we use weekly dummies for each week in our 121 week sample period (September 1998-December 2000) to control for other 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. Judging by the distribution of sales after car introductions, we distinguish between sales in the first four months, months 5-13, and month 14 and later and assign a dummy variable to 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 10 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 control for competitive promotions by including in our regression the average customer cash and the average dealer cash offered by competitors on a particular day in the same subsegment, weighted by the average yearly market share of each car in the subsegment. We also construct variables that simply count the number of competing customer cash and number of competing dealer cash promotions in the same subsegment.

We also control for the income, education, occupation, and race of buyers by using census data that MRI matches with the buyer’s address from the transaction record. The data is on the level of a “block group,” which makes up about one fourth of the area and population of a census tract. On average, block groups have about 1100 people in them. Finally, we control the geographic region in which the car was sold.

3.3 Summary statistics

Twenty-six percent of transactions in our sample involve customer cash, and 18% involve dealer cash. The average amount of customer cash observed for transactions in our sample that involve customer cash is \$1188 (median \$1000). The average amount of dealer cash among transactions

that involve dealer cash is \$992 (median \$700). The average transaction price of a new vehicle in our data is \$25,490. In our sample of cash transactions the average car sells 7.85 per week. In interpreting this number note that our data covers only a 15-20% sample of dealerships and is restricted to cash transactions.

4 Estimation approach

The aim of our paper is to estimate the “treatment effect” of promotions on prices and on sales. 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 or sales with a promotion to dealer price or sales without a promotion. However, manufacturers might be more likely to instigate incentive promotions when prices and/or sales are either low or declining due to variation in demand. This means that the price or sales observed in the periods in which a manufacturer chose not to have a promotion are not necessarily what price or sales *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. Analogously, we could underestimate the treatment effect of promotions on sales because the condition of reduced demand would work against the sales-increasing effect of the promotion.

4.1 Difference-in-differences and regression discontinuity

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 and quantity. Thus the estimated change in price or in sales that is attributed to the promotion is the change that is net of the contemporaneous change in prices or sales of other cars within the same vehicle segment.

An alternative way to estimate the counterfactual price is to use the a car’s own price when it is not on promotion. Conceptually, this is what regression discontinuity does. This approach has been used primarily in the program evaluation literature, particularly in education and job training applications. In those applications, researchers must often 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 which is likely to be 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 takes advantage of 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 (Hahn, Todd, and Van der Klaauw, 2001, Imbens and Angrist, 1994). For example, a state-funded college grant program has a specific GPA cutoff for eligibility which is unknown to applicants at the time they apply. Thus, applicants who are close to the cutoff on either side can be used to estimate the effect of receiving a grant on the probability of enrolling in college (Kane, 2003). In a similar example, a particular college's financial aid program uses a scoring method to sort students into financial aid categories using a formula based on SAT score and GPA. Students near the cutoff for each category can be used to estimate the effect of the size of a financial aid offer on the probability of enrolling in a that college (Van der Klaauw, 2002).

In this paper, we are interested in estimating the effect of a manufacturer promotion (the treatment) on two different outcome measures, price and sales. As is often the case in the program evaluation literature, we do not believe that manufacturers apply the treatment randomly. In particular, we expect that 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.

However, we believe that the underlying demand conditions which determine the rate of sales and the customers' price elasticity is likely to change fairly little over the course of several weeks.³ Thus, in applying the regression discontinuity approach, we restrict our attention to two weeks around when a promotion begins, ends, or changes in the amount of cash being offered. 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 promotion, the regression discontinuity approach will consistently estimate at least the local average treatment effect.

Within the regression discontinuity approach, identification would be upset if the customers

³We know from industry sources, that there is typically a gap of several weeks between the most recent information that was used in making the decision to initiate a promotion and the start of the promotion.

who purchase just before a promotion starts and just after differed in some way that was related to the outcomes of interest. This would be the case, for example, if there are deal-prone customers, who are particularly effective negotiators, and who wait to purchase a car until a promotion is offered.

4.2 Propensity score

As in the previous subsection, the difference-in-differences approach and the regression discontinuity approach give us two ways to estimate the counterfactual prices that would have been obtained for cars that were purchased during a promotion had the promotion not been offered. We now address the question of how confident we can be that these estimated counterfactuals will lead us to an unbiased estimate of the true effect of “treatment on the treated.”

Heckman, Ichimura, and Todd (1997) decompose the bias that can arise from using the outcome in non-treated groups to stand in for the (unobservable) outcome in treated groups had the treatment not occurred. One component of bias that they identify is non-overlapping support in observable characteristics. There may be cars that are promoted for which there are no comparable cars which are not promoted, and vice versa, cars that are unpromoted for which there are no comparable cars that are promoted. There may also be customers who buy under a promotion for whom there are no comparable customers that do not buy without a promotion, and vice versa, customers that buy without a promotion for whom there are no customers who buy under a promotion. This might lead us to conclude that part of the estimated price difference we observe is attributable to treatment when it is in fact attributable to differences between the support of cars that are and are not promoted, or customers who buy with and without a promotion.

A second source of bias is differences in distributions of observable characteristics in either car or customer characteristics between promotion and non-promotion groups, even within a region of common support. Again, this could lead us attribute to treatment price differences that are actually due to differences in the characteristics of “treated” and “non-treated” cars and customers.

The third and final component of bias that Heckman, Ichimura, and Todd identify is differences in outcomes that arise *within* a region of common support *and* conditioning on observable characteristics. This component is what is usually referred to as selection bias.

In section 5.2, we use a propensity score approach to minimize the bias arising from non-overlapping support. To control for bias caused by differences in car and customer observables between promotion and non-promotion subsamples we use a large set of consumer characteristics and detailed car fixed effects as covariates.

Heckman, Ichimura, and Todd argue that frequently these two components are larger than the bias arising from selection on unobservables. Selection bias would arise in our case if customers “select into treatment” on the basis of the gains they would obtain. This would be the case if, for example, customers who are inherently good price negotiators are also more likely to buy under promotions. We do not currently have good instruments that would allow us to account for this, so for now it is a maintained assumption that such differences do not exist.

5 Rebate 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 that was used to supply the surplus. We hypothesize that in the price negotiation process, customers will obtain a greater share of the surplus when they are better informed about the existence and size of the promotion; in other words, we anticipate greater pass-through of customer cash than of dealer cash. However, even if customers do not know that dealer cash is on the table, we expect that they may still obtain some of the dealer cash surplus. This is because the dealer will get the rebate 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.

We begin with the difference-in-differences approach to estimate the rate of pass-through, and then turn to the regression discontinuity approach.

5.1 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

$$P_{ijt} = \lambda_c \text{CustCash}_{jt} + \lambda_d \text{DealCash}_{jt} + \beta_1 X_i + \beta_2 X_{jt} + \beta_3 \text{DealerComp}_{ij} + \mu_j + \tau_{JT} + \epsilon_{ijt}. \quad (1)$$

CustCash_{jt} and DealCash_{jt} are the amount of customer cash and dealer cash available for vehicle j at date t . Since manufacturers typically make promotion decisions by model - 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 reflect such variation. X_i is a vector of individual and neighborhood customer demographic

characteristics including sex, race, income, education, employment type, and home ownership. X_{jt} 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). $DealerComp_{it}$ is a measure of how competitive is the dealer at which customer i purchased his or her vehicle; specifically, the measure used is the number of competing dealers of the same nameplate within a 10 mile radius of this dealer. μ_j are vehicle fixed effects. For these fixed effects, we use a very fine definition of a vehicle, 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 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 towards a lower price. One can interpret $100 \cdot |\lambda|$ as the percentage of the rebate the customer obtains.

Pass-through of manufacturer-supplied surplus

Table 1 reports the results of estimating this specification. In column 1, the customer cash coefficient implies that 88% of customer cash is passed through to customers, while 39% of dealer cash is passed through. This is consistent with our prediction that customers will obtain more of the manufacturer-supplied surplus the better informed they are of its existence.

In this column, we also test one 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. The amount of rebate a user accumulates is proportional to the amount charged to the card (otherwise the GM Card is a normal credit card). 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 1 shows that customers obtain 106% (statistically indistinguishable from 100%) of the surplus from GM Card rebates.

The effect of competition on pass-through

In column 2 of Table 1 we interact the promotions variables with our measure of dealer competition. We expect to see that rates of pass-through increase for customers who have the option

of easily visiting multiple dealers of the same nameplate. Consistent with this prediction, the coefficient estimates imply that a dealership without competing dealerships of the same nameplate within 10 miles will pass through 84% of customer cash, and that this rate of pass-through will increase by one percentage point for every additional dealer of the same nameplate in that 10 mile radius. Competition has no statistically significant effect on the rate of dealer cash pass-through, which is consistent with customers having no additional incentive to play dealers off each other to obtain a greater share of the promotional cash if they do not know that the promotion is being offered, as will typically be the case with dealer cash.

Competing promotions

In column 2 we also introduce variables measuring the number of customer and dealer cash promotions currently available on competing vehicles in the same vehicle segment (e.g. SUVs, compact cars, etc.). Neither the number of competing customer cash offers nor the number of competing dealer cash offers has a statistically significant effect on prices.

Perhaps a better measure is a variable measuring the *average level* of customer and dealer cash currently available on competing vehicles in the same vehicle segment. The results are in column 3. Again, customer cash on competing vehicles has no statistically significant effect on prices, but dealer cash available on competing vehicles lowers transactions prices: increasing the average dealer cash available on competing vehicles by \$50 (for example, if one of 10 competing vehicles introduced a \$500 dealer cash promotion) lowers the transaction price by about \$14. While we expected that both types of competing promotions would lower prices, we find such an effect only for competing dealer cash promotions. Notice that in this specification we include only week fixed effects since these ‘average competing promotion in segment’ variables together with the vehicle’s own promotion are colinear with the week - segment fixed effects. For this reason we continue with the ‘number of customer and dealer cash promotions’ variable in subsequent specifications.

Pass-through by demographics

In column 4, we interact the customer cash and dealer cash measures with three demographic measures, namely whether the customer is female, and the percentage of the customer’s census block group that is black or Hispanic (which can be thought of as the probability that the customer is black or Hispanic). The positive coefficients estimated for these variables indicate that women, blacks, and Hispanics obtain less promotional surplus in the negotiation process. The estimated effect for women is 4 percentage points less pass-through of customer cash and 7 percentage points less pass-through of dealer cash than for men. The effects for blacks and Hispanics are much larger: 32 and 44 percentage points respectively less pass-through of customer cash, and for Hispanics 28 percentage points less pass-through of dealer cash. These

effects are large, in particular if one keeps in mind that in this specification the direct effect of demographics on price is already controlled for.

Demand effects from advertising of customer cash

In comparing the rates of pass-through of customer cash and of dealer cash, we have interpreted the estimated differences to be differences in the effect of information on the price negotiation process. Part of what has enabled us to do so is that customer cash deals are widely advertised on television, radio, and in newspapers. However, we recognize that advertising itself is likely to have an effect on demand for a car. In column 5, we try to get at the effect of advertising on our estimated customer cash pass-through rates. In identifying the pass-through rate, we make use of intertemporal variation in customer cash; manufacturers not only begin and end promotions, they also adjust promotion amounts up and down. We believe that the beginnings of promotions or increases in promotion amounts are likely to be advertised, but that the end of a promotion or a decrease in the amount of a promotion is much less likely to be advertised. Therefore, in column 5 we introduce a variable called *CustCash*CustCash decrease* which is an interaction of the currently available customer cash amount, and an indicator variable for whether this amount is a decrease from the customer cash amount available immediately preceding this promotion.⁴ An analogous variable is defined for dealer cash.

The estimates in column 5 indicate that when the current customer cash promotion is an increase from what was previously offered, 74% is passed through to customers. If the current promotion is a decrease from what was previously offered, 96% of the surplus is passed through to customers. There are two ways to interpret this finding. One is that the advertising that accompanies customer cash promotions makes customers less price sensitive in their negotiations with dealers. Since there is less advertising when promotion amounts decrease, such customers are more price sensitive and will appropriate more of the surplus. Alternatively, our finding could be the result of customers coming into the dealership expecting to receive a previously advertised larger promotion that is no longer available. This expectation could reduce their reservation price, enabling them to extract more of the promotional surplus from the dealer.

Regardless of the interpretation of these findings, including the promotion decrease indicators does not alter our conclusion that customer cash pass-through is higher than dealer cash pass-through. In column 5, 41% of dealer cash is passed through, and there is no statistically significant difference when the current promotion level is a decrease. This is significantly smaller than either estimate of the customer cash pass-through.

Sales incentives

⁴For example, if there is no promotion in January, then a \$750 promotion in February, followed by a \$500 promotion in March, the indicator variable would be 1 for the month of March.

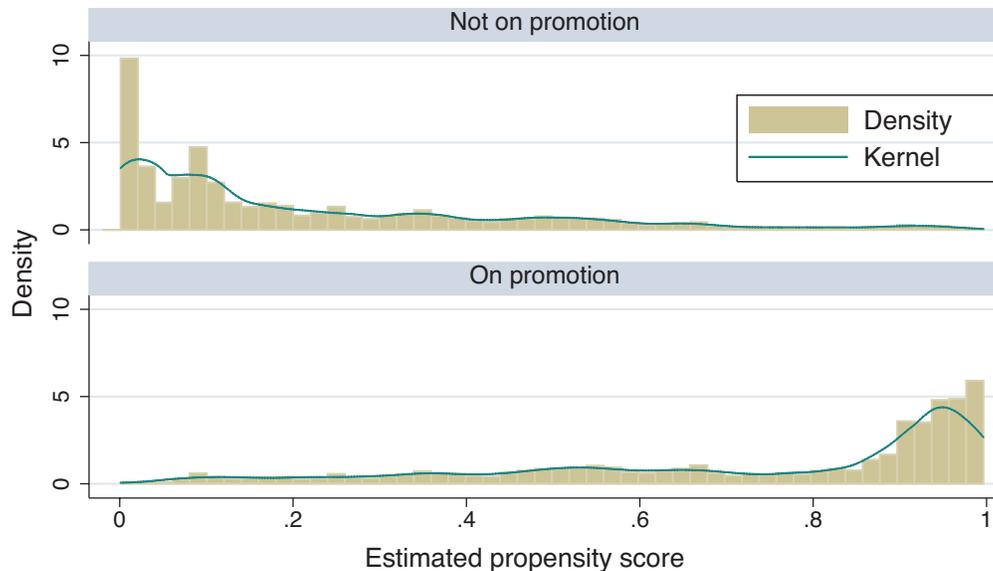
In column 6, we include measures of two additional manufacturer promotions, sales manager incentives and sales representative incentives. Like customer cash and dealer cash, these are cash payments to the respective parties. These promotions do not have a statistically significant effect on transaction prices (but they do have, as we will later show, an effect on sales).

Propensity score

As we described in section 4.2, the estimated treatment effect can be biased if the distribution of observable characteristics in treated and untreated groups do not share a common support. We use a propensity score approach to mitigate this source of bias (Galiani, Gertler, and Schargrofsky, 2002).

We estimate propensity scores from a probit model of the probability that a transaction occurred in the week immediately following a promotion change as a function of the customer characteristics, the detailed car dummies, the region, and the competition variable used in all previous regression difference-in-differences specifications. This yields the probability that a purchase transaction will occur while a customer or dealer cash promotion is offered. Figure 1 graphs the distribution of propensity scores for the control group (transactions for a car that is not on promotion) and the distribution of propensity score for the treatment group (transactions for a car that is on promotion).

Figure 1: Distribution of propensity scores for the control group (transactions for a car that is not on promotion) and the treatment group (transactions for a car that is on promotion)



The fat tails of the propensity score distribution are mostly due to the fact that we estimate

the propensity score using detailed car dummies. Since there are numerous cars that are never promoted and also cars that are almost always promoted, there are many observations with a very low propensity score in the control group and many other observations with a very high propensity score in the treatment group.

We identify observations on the common support by (1) excluding all control (non-promotion) observations whose propensity scores are lower than the 1st percentile of the distribution of propensity scores for the treatment (promotion) observations and by (2) excluding all treatment observations whose propensity scores are higher than the 99th percentile of the distribution of control propensity scores. This reduces the sample from 133,424 to 68,278 observations. The decline in the number of observations is mostly due to the fact that our procedure excludes cars that are never promoted and cars that are always promoted.

In comparing Tables 1 and 2 we see that adjusting the sample so that observations share a common support changes some of the estimates. While the major result, that customer cash is passed through to customers at a higher rate than dealer cash, is unchanged, the estimated rates of customer cash pass-through are higher by a few percentage points in most specifications in Table 2. The other major difference is in the competition interaction results. In the common support specification in Table 2, dealer competition no longer increases the rate of customer cash pass-through, which it did in the results reported in Table 1. There is also now evidence that a competing dealer cash promotions in the segment decrease prices, by about \$50 for each competing promotion.

Test of pre-promotion trends

The maintained assumption underlying the validity of the difference-in-differences 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. While we have no way of observing directly whether this assumption is valid, we can examine the trends of promoted and non-promoted cars in a period prior to the promotion. If the trends are similar between cars that are soon to be promoted and other cars, that gives some assurance that the non-promoted cars may be a valid counterfactual in the promotion period.

To test this, we estimate a daily time trend for each month of the sample, separately for cars that will and will not be on promotion during the next month. First, we restrict the sample to observations of cars during months in which the car is not on promotion. We then generate an indicator variable which is equal to one if the car is on promotion any time in the succeeding month. Formally, $I(M+1)_{jt} = 1$ for car j at date t in month M if car j goes on promotion at

some t' in month $M + 1$. We then run the following regression

$$P_{ijt} = \alpha_M + \beta_M \theta_{t,M} + \gamma_M I(M + 1)_{jt} \cdot \theta_{t,M} + \mu_j + \eta_{ijt}, \quad (2)$$

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, $\theta_{t,M}$ is a daily time trend for the days in month M , and μ_j are fixed car effects. The coefficients β_M will measure the daily trend of prices over the days in month M , and the γ_M coefficients will measure any differences in the daily time trend in month M for cars that will be on promotion in month $M + 1$. Our test of equal trends in the pre-promotion period will be testing whether the γ_M coefficients are equal to zero.

In unreported results, 6 of the 25 γ_M 's are statistically different from zero at the 5% confidence level, and an additional 2 at the 10% level. Of the statistically significant γ_M 's, the largest in magnitude that is negative is -1.6 dollars per day. The implication of this coefficient is that the difference in prices between promoted and non-promoted cars that could arise due to differences in trends is at most \$50 over a 30-day time period. Recall that the estimated amount of customer cash passed through to customers would be around \$1000 for an average customer cash promotion (85% times an average level of \$1188), and around \$450 for an average dealer cash promotion (45% times an average level of \$992). This means that the average customer cash promotion would have to last for 20 months, longer than the lifespan of most car models, and the average dealer cash promotion would have to last 9 months for differences in time trends to explain the estimated pass-through effect.

Nevertheless, a test of joint significance rejects the null hypothesis that the γ_M coefficients are jointly zero. In other words, we reject that the time trends of prices for cars that will go on promotion in the next month is the same as that of cars that will not go on promotion in the next month. While this is not the result we had hoped for in support of the difference-in-differences specification, the magnitudes of the estimated time trend differences lead us to believe that any bias is not substantial in size. Nevertheless, our rejection of identical trends between promoted and non-promoted cars gives us additional impetus for employing the regression discontinuity approach.

5.2 Regression discontinuity

The regression discontinuity approach consists of analyzing only transactions that occur immediately before and immediately after a change in a promotion. The idea is that demand conditions do not change within a short window that include an event of interest. We choose 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 some positive amount to zero.

The regression equation in the regression discontinuity approach is similar to the one used in the difference-in-differences approach. As equation 3 shows, the measures of promotions, demographics (X_i), time period controls (X_{jt}), and car fixed effects (μ_j) are the same as in equation 1 above. The regression discontinuity specification does not rely on week - vehicle segment fixed effects in order to identify the “treatment effect” of the promotion; in its place, we include a week fixed effect, τ_T .

$$P_{ijt} = \lambda_c \text{CustCash}_{jt} + \lambda_d \text{DealCash}_{jt} + \beta_1 X_i + \beta_2 X_{jt} + \beta_3 \text{DealerComp}_{ij} + \mu_j + \tau_T + \epsilon_{ijt} \quad (3)$$

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, for example, if there are deal-prone customers with a different willingness-to-pay or bargaining ability, who wait until the begin of a promotion to enter the market.

As in the difference-in-differences approach, the coefficients of primary interest are λ_c and λ_d , which have the same interpretation as before. Table 3 reports the estimated coefficients from the regression discontinuity approach. Overall, they are very similar to our estimates in the difference-in-differences approach.

Pass-through of manufacturer-supplied surplus

In column 1 of Table 3, 81% of customer cash is estimated to be passed through to customers, compared to 38% of dealer cash. Customers are estimated to obtain 112% of GM Card cash, the manufacturer-supplied surplus for which customers have the information advantage over dealers. This estimate is not statistically different from 100% at the 5% significance level.

The effect of competition on pass-through

In column 2, we find that the effect of dealership competition on the rates of customer and dealer cash pass-through is less easily interpretable than in the difference-in-differences approach. The $\text{CustCash} \cdot \text{Competition}$ interaction term is now statistically insignificant, while the $\text{DealerCash} \cdot \text{Competition}$ interaction term is now statistically significant. The estimated coefficient on the latter term is 0.03, which implies, counterintuitively, that more competition among dealers means less pass-through of dealer cash by 3 percentage points. We currently have no good explanation for this result.

In this specification we again find that available dealer cash on other vehicles in the same segment is associated with lower prices. The estimated coefficient of -0.52 implies that an increase by \$50 of the average dealer cash available on other cars in the segment lowers the price for this vehicle by \$26. This is almost twice the size of the effect estimated in the difference-in-differences specification.

Pass-through by demographics

In column 3 we introduce sex- and race-based interaction terms with the customer cash and dealer cash variables. In the difference-in-differences specification, we found that, almost across the board, women and racial minorities were able to appropriate less of both customer and dealer cash than other consumers. In the regression discontinuity specification, we find statistically weaker evidence of such effects. This specification indicates that 92% of customer cash gets passed through to non-black, non-Hispanic men. Only Hispanics are found to get statistically significantly lower rates of pass-through of customer cash, by 43 percentage points. In this same specification, non-black, non-Hispanic men are estimated to obtain 54% of dealer cash. Only women are found to get lower rates of pass-through of dealer cash, by 18 percentage points. In interpreting these results one should keep in mind that the sample size is reduced by a factor of 10 compared to the difference-in-differences approach.

Demand effects from advertising of customer cash

Column 4 allows for asymmetric pass-through depending on whether the current promotion is an increase or decrease compared to the promotion available immediately prior to the current promotion. The estimates indicate that when the current customer cash promotion is an increase (either from zero or from a lower promotional level), customers obtain 79% of the surplus. When the current promotion is a decrease from a prior level, customers obtain 17 percentage points more, although this effect is statistically significant only at the 90% confidence level. This is the same direction as the effect found in the difference-in-differences specification. Unlike the earlier approach, however, in the regression discontinuity approach, dealer cash is found to have asymmetric pass-through rates. The estimates imply that 47% of dealer cash is passed through when the promotion is an increase, and an additional 30% of surplus is passed through to customers if the current dealer cash level is a decrease. The original motivation for including these asymmetric variables was to estimate the effect of advertising on customers' negotiating behavior. However, advertising effects would predict that customer cash would have asymmetric effects, while dealer cash would not. This finding is anomalous to that explanation.

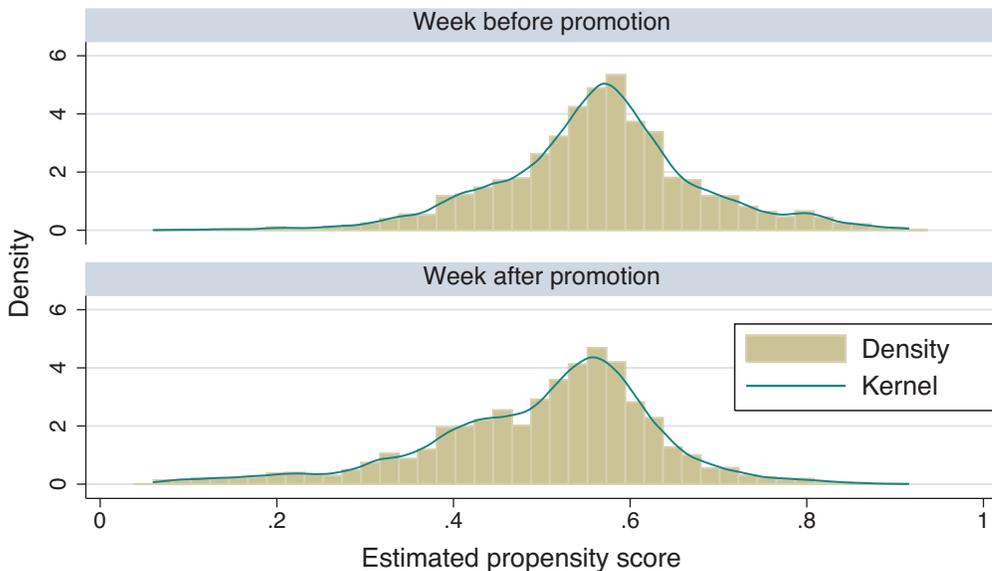
Sales incentives

Finally, in column 5, we include the effect of sales manager and sales representative promotions on customer prices. There is no statistically significant effect, which could be an artifact of the sample chosen for the regression discontinuity approach. In particular, we chose the windows around changes in customer and dealer cash promotions, however, not around changes in sales manager or sales representative incentives. Therefore, it is not surprising that the regression could not detect an effect of these incentives. We plan to estimate the effect of these incentives with a more appropriately chosen sample in a later version of the paper.

Propensity score

We repeat the propensity score approach used in section 5.1 for our regression discontinuity specifications. We estimate propensity scores from a probit model of the probability that a transaction occurred in the week immediately following a promotion change as a function of the customer characteristics, the detailed car dummies, the region, and the competition variable used in all previous regression discontinuity specifications. This yields the probability that a purchase transaction will occur in the week just before or just after a promotion change. Figure 2 graphs the distribution of propensity scores for the control group (transactions in the pre-promotion week) and the distribution of propensity scores for the treatment group (transactions in the during-promotion week).

Figure 2: Distribution of propensity scores for the control group (transactions within one week before a promotion change) and the treatment group (transactions within one week after a promotion change)



We identify observations in the common support following the procedure described in section 5.1. Notice that the distribution of propensity scores for the control and treatment groups is much more similar than in the difference-in-differences specification. This is because the regression discontinuity approach restricts the sample to cars which change their promotion status, leaving out cars that are never or always promoted at the same level.

In comparing tables 3 and 4 we see that adjusting the sample so that observations share a common support changes the estimates very little.

Robustness with regard to window size

The results reported so far appear to be robust to changes in the size of the window around the promotion event. We reestimated the regressions in Table 3 using only two days prior and after the change in a customer cash or dealer cash promotion. We restrict ourselves to observations on a common support using the same propensity score approach as we used in the previous subsection. Figure 3 graphs the distribution of propensity scores for the control group and the distribution of propensity score for the treatment group. Table 5 shows that most coefficient estimates of interest are very similar. The only substantive change is that the demographic interactions are insignificant, perhaps not surprisingly since the sample in this specification has decreased to only 3218 observations.

The robustness of the results to the window size is important because it makes less likely a possible source of mismeasurement. Suppose that dealers knew when a dealer cash promotion will begin. Then the dealer should want to sell fewer cars in the weeks leading up to the promotion. This would increase prices in the pre-promotion weeks, hence leading us to overestimate the effect of the promotion on transaction prices and thereby on pass-through. We can test for this by comparing the pass-through estimates in Tables 3 and 5. This is because industry participants say that dealers typically only find out about promotions 2-3 days before they occur. Since the pass-through estimates in the 3 days on either side window are no lower than those in the 1 week on either side window, it seems unlikely that pass-through rates are overestimated due to strategic dealer behavior ahead of promotions.

5.3 Summary

Overall, the results of this section tell a remarkably consistent story across approaches and across specifications. Customers obtain 80-90% of the surplus supplied by manufacturers in customer cash promotions, but half or less of the surplus in dealer cash promotions. Customers also obtain all the surplus available through the GM Card. This is consistent with the hypothesis that when customers are at an information disadvantage, they are disadvantaged in negotiations.

The results of this section also indicate, although less uniformly, that greater dealer competition may help customers get higher rates of pass-through and that women and racial minorities may get less pass-through. These results on women and racial minorities are consistent with the findings of Scott Morton, Zettelmeyer, and Silva-Risso (2003) on the effect of such demographic characteristics on overall prices.

6 Sales effect of promotions

Estimating the price effect of different promotions yields results about the effect of asymmetric information on bargaining in a product market. While these results are of interest to economists, manufacturers are primarily concerned with the effect of promotions on sales. This is because equal-sized customer cash and dealer cash promotions cost the same amount to the manufacturer. Manufacturers care about which type of promotion leads to more incremental sales.

We use two different dependent variables to measure the sales effect. One measure is the weekly unit sales of a given vehicle (defined by make, model, and model year). The other is the weekly subsegment market share. This measure is defined as weekly unit sales divided by the total number of vehicles sold in that week in the vehicle's subsegment.⁵ We use subsegment market share in addition to sales for two reasons. First, market share points are more directly comparable across vehicles than sales. Otherwise we are comparing apples to oranges, or, literally, Escorts to Suburbans. Second, we know that many manufacturers have explicit goals for market share and market share will thus be used as a benchmark in determining a promotion's success.

We use both a difference-in-differences approach and a regression discontinuity approach to estimate specifications with each dependent variable. One important change we must make as we move from price regressions to quantity regressions is that we can no longer use individual transactions as the unit of observation. Instead, our unit of observation is a vehicle - week, or, more precisely a make - model - model year - week. We indicate other differences in the specifications below.

6.1 Weekly sales

We start by estimating the effect of promotions on weekly sales, using both the difference-in-differences and regression discontinuity approaches. Our dependent variable in both cases is S_{jT} , the number of units of car j (defined by make - model - model year) sold in week T . For the difference-in-differences approach, the basic specification is

$$S_{jT} = \delta_c \text{CustCash}_{jT} + \delta_d \text{DealCash}_{jT} + \gamma_1 X_{jT} + \gamma_2 \text{DealerComp}_{ij} + \mu_j + \tau_T + \epsilon_{jT} \quad (4)$$

As with the pass-through regressions, CustCash_{jT} and DealCash_{jT} are the amounts of customer cash and dealer cash available in week T for car j . Unlike the pass-through regressions,

⁵A subsegment is a finer breakdown than segment. For example, compact cars are broken down into economy compact, such as Hyundai Accent, Ford Fiesta, and Toyota Echo, and premium compact, such as Honda Civic.

there is no longer a vector of customer demographics. However, X_{jT} —the controls such as end-of-month and time since model introduction—are still included. The vehicle fixed effects, μ_j are now defined by make - model - model year, instead of the finer vehicle fixed effects used in the pass-through specifications. τ_T is a weekly fixed effect.

The regression discontinuity specification is identical, but is estimated using only a subset of the data. The sample contains data only from the week just before a promotion starts or changes and the week just after. In the pass-through regressions, where we used individual transactions as the unit of observation, the regression discontinuity specifications used data from the seven days before and the first seven days of a promotion. In the sales effect regressions, where we must use weekly sales, we have selected into the sample the last full calendar week before a promotion starts or changes, and the first full calendar week after a promotion starts or changes, leaving out the week during which the change actually happens.

In both specifications, the coefficients of primary interest are δ_c and δ_d . These coefficients estimate the increase in sales, measured in number of vehicles, of a one dollar increase in a customer or dealer cash promotion.

Difference-in-differences The results for the difference-in-differences approach is reported in Table 6. In the first column, the estimated coefficient of 0.0017 for customer cash implies that a \$500 increase in customer cash would increase sales by 0.85 cars per month, or a 10.8% increase compared to the mean vehicle sales per week of 7.85.⁶ Meanwhile, an additional \$500 in dealer cash increases sales by 0.25 cars per week, or by 3.2% of the average level of weekly vehicle sales. Thus, the same amount of promotional expenditure on the part of manufacturers has less than a third the effect when it is spent as dealer cash as opposed to customer cash.

In column 2, we add measures of the average customer cash and dealer cash promotions available on other cars in the segment, but these have statistically insignificant effects. In this column, the dealer cash measure is no longer statistically significant, although the point estimate is unchanged from column 1. The customer cash effect on sales is unchanged. In column 3, we allow the effects of customer cash and dealer cash to be asymmetric depending on whether they are increases or decreases relative to previous cash levels, but these also have statistically zero effects.

Finally, in column 4, we add variables measuring the size of sales manager and sales rep incentive promotions. Sales manager promotions are estimated to have a very large effect on sales: a \$50 incentive would increase weekly sales by 1.9 cars per week, or by 24%. This seems

⁶This may seem like a low number for the average sales of a given make, model, and model year defined car in California in a given week. Recall that the data sample includes only cash sales at a 15-20% subset of California dealers.

an implausibly large effect. Fewer than 2% of transactions are made with a sales manager incentives; by comparison, 36% of transactions are made with customer cash. In our data, sales manager incentives range between \$20 and \$500, with a median incentive of \$50.

Regression discontinuity The regression discontinuity approach, reported in Table 7, yields comparable results to the difference-in-differences approach. In column 1, the customer cash coefficient of 0.0021 implies that a \$500 increase in customer cash would increase sales by 1.05 vehicles per week, or 13.4% compared to the mean. The estimated dealer cash coefficient is only statistically significant at the 90% confidence level; taking the point estimate at face value nevertheless, the predicted effect of an additional \$500 in dealer cash is a 0.55 vehicle per week increase in sales, or 7%.

In column 2, the average level of competing customer cash and dealer cash promotions in the segment has no statistically significant effect on price. However, in this column, the estimated sales effect of customer cash increases slightly to 14% and the effect of dealer cash is now statistically significant at the 95% confidence level.

As in the difference-in-differences estimates of the weekly sales effects, allowing promotions to have asymmetric effects has no effect, as reported in column 3. Finally, unlike the results in Table 6, sales manager and sales rep incentives have no statistically significant effect on sales. Notice that the sample in this specification is 1628 as opposed to 21614 in the difference-in-differences approach.

As with the pass-through results, the two approaches give similar results. We think that the estimated sales effects are large, but plausible.

6.2 Weekly market share

In estimating the sales effect using market share as a dependent variable, we use exactly the same specification as for sales (equation 4), except that we substitute for the dependent variable M_{jT} , defined as

$$M_{jT} = \frac{S_{jT}}{\sum_{h \in H} S_{hT}}, j \in H, \quad (5)$$

where S_{jT} is weekly sales of vehicle j in week T and j is a vehicle in subsegment H .

Difference-in-differences Table 8 reports the results of the difference-in-differences approach. The estimated coefficients are small because the dependent variable is market share. In column 1, the estimated coefficient for customer cash implies that a \$500 increase in customer cash would increase market share by 1 percentage point, or a 9% increase on the average within-segment market share of 11.2%. This estimate is very close to the estimate of a 10.8% increase

when sales are measured as weekly unit sales. Dealer cash promotions are not estimated to have a statistically significant effect on market share.

Column 2 examines the effect of competing promotions offered on other vehicles in the same segment. The point estimate indicates that a \$50 increase in the average customer cash available on competing vehicles (the equivalent, for example, of one of ten competing vehicles introducing a \$500 promotion) will reduce market share by two-tenths of a percentage point of market share, or by 1.8% relative to the average within-segment vehicle market share. However, this coefficient must be interpreted with some caution. It is not, as it might appear at first glance, evidence that a competing promotion steals business from a given manufacturer. If a firm's competitor increases its unit sales, while the firm's unit sales remain unchanged, the competitor's market share will go up, while the firm's market share (unit sales as a fraction of the now higher total market sales) will go down. At the mean of the data, about half of the estimated effect of competing customer cash deals is this purely mechanical effect of market expansion.⁷

As with all the previous sales effect regressions, allowing promotions to have asymmetric effects has no statistically significant effect, as shown in column 3. In column 4, sales manager incentives are estimated to have a large effect on sales, as in the difference-in-differences specification using unit sales as the dependent variable. Here, a \$50 sales manager incentive is predicted to increase market share by 1 percentage point, or 8.8% compared to the average market share of 11.2%. In considering these results, it seems unlikely that the following three facts can be simultaneously true: (1) \$500 customer cash and a \$50 of sales manager incentive have about the same effect on sales, (2) only 2% of transactions are made with sales manager incentives, and (3) manufacturers are well-informed and profit maximizing. We are investigating this result further.

Regression discontinuity We report the estimates of the regression discontinuity approach with weekly market share as the dependent variable in Table 9. In column 1, the estimated

⁷At the mean of the data, each nameplate sells 8 cars per week in a segment in which it operates and has an 11.2% market share, equivalent to 9 equal sized firms. Suppose that one firm in the segment, Firm 1, offers a \$500 customer cash promotion. The estimated coefficients from Tables 6 and 7 imply that by doing so, Firm 1 should it sells approximately one additional car. But the \$500 is estimated to have no effect on the unit sales of the other 8 firms. This means the segment goes from selling 72 cars per week to 73 cars per week when Firm 1 introduces the \$500 promotion. This would mean that Firm 1 goes from a market share of 11.1% (8 divided by 72) to a market share of 12.3% (9 divided by 73). The other firms go from market share of 11.1% to 11% (8 divided by 73).

The market share regressions of 8 predict that a \$500 customer cash deal should increase Firm 1's market share by 1.1 percentage points, similar to what is predicted in the sales regression. This promotion by Firm 1 will increase the average customer cash available on competing vehicles by \$55 for each of the other eight firms in the segment. According to the estimated coefficients in Table 8, this should lower each of their market shares by 0.2 percentage points.

effect of a \$500 increase in customer cash is to increase market share by 0.85 of a percentage point, which is a 7.8% increase compared to the average segment market share. The dealer cash effect is not statistically significant.

In column 2 we add measures of the average level of competing (within segment) customer cash and dealer cash promotions. The estimated effect is statistically significant only for competing customer cash. The point estimate is that an increase of \$50 in the average level of customer cash available on competing vehicles would reduce a manufacturer's own segment market share by four-tenths of a percentage point, or by 3.7% compared to the average segment market share. This a larger effect than was estimated using the difference-in-differences specification. By itself, this estimate suggests that promotions do steal sales from other vehicles in the segment; however, the unit sales specifications, which should be a more direct test of this point, failed to find this effect.

Column 3 allows promotions to have asymmetric effects, and column 4 adds sales manager and sales rep incentives. As with other sales effect specifications, there is no estimated asymmetry. Column 4, however, introduces a result not seen in previous specifications. Sales representative incentives are predicted to significantly increase sales. According to the point estimate, a \$50 sales rep incentive would increase market share by 0.9 percentage points, or 7.7% compared to the average vehicle market share. Sales rep incentives are employed more frequently than sales manager incentives (4.4% of transactions involve a sales rep incentive), but nowhere nearly as frequently as customer cash or dealer cash. Observed sales rep incentives range between \$25 and \$500, with a median size of \$75 when offered.

6.3 Summary

The strongest result of this section is that customer cash promotions significantly increase a manufacturer's vehicle sales. This result holds across both dependent variables and both estimation approaches used in this section. The estimate across specifications is that a \$500 increase in customer cash increases vehicle sales by 7.8-13.4%. The effect of dealer cash is less robustly estimated. In the specifications in which it is statistically significant, the estimates indicate that \$500 of dealer cash increases sales by 3.2-7%.

The effect of such promotions on competitors' sales is not conclusive. When sales are measured as unit sales, neither a competitor's customer cash nor its dealer cash promotions appear to affect a manufacturer's sales. A customer cash promotion by a competitor does decrease a manufacturer's market share, however, and the magnitude of the coefficients suggests that the effect is larger than what would result if promotions were purely market-expanding, not business-stealing.

Finally, both sales manager and sales rep incentives are estimated to have an effect on sales, although the effects are not consistent across specifications, and are larger than seem plausible.

7 Promotion generation process

In interpreting the results of the previous section, we have drawn conclusions about the effectiveness of customer cash vs. dealer cash in promoting sales, making comparisons of the dollar-per-dollar effectiveness of the two promotions. These conclusions must be interpreted with caution. They indicate the average effect of a customer cash promotion on sales in periods in which manufacturers chose to have customer cash promotions and the average effect of a dealer cash promotion in periods in which manufacturers chose to have dealer cash promotions. Our results do not necessarily imply that a manufacturer would increase sales by switching from dealer cash to customer cash. For example, if each type of promotion is well suited to a particular configuration of market conditions, it may be that manufacturers are optimally matching promotions to particular conditions. For example, given the results reported so far, indicating that customer cash has a greater sales promotion effect, it may be that manufacturers would prefer to use customer cash in periods in which they would like to liquidate excess inventory. Conversely, given that dealer cash leaves more surplus with dealers, manufacturers may choose to use dealer cash in periods in which they are under pressure to support dealer profits.

In order to explore this question, we hope to model the promotion generation process. For now, we report a simple comparison of means which is informative, but falls short of what we hope to do in future.

Industry sources tell us that the primary variables that manufacturers take into account when deciding whether to initiate a promotion are vehicle profitability at the dealer level, days of inventory (current inventory holdings divided by sales rate), and market share within the vehicle's subsegment. In Table 10 we compare the average of these measures over eight week periods leading up to manufacturers' decisions to either offer or not offer a customer cash promotion. To be specific, we generate a weekly indicator for every car indicating whether the car changed that week from having zero customer cash to having a positive amount of customer cash. Such weeks form our "week with customer cash" sample in column 2 of Table 10. We also generate a weekly indicator for weeks in which there was no customer cash promotion. Such weeks form our "week without customer cash" sample in column 3. We then report the vehicle profitability, days of inventory, and subsegment market share for the eight weeks preceding the weeks included in each sample. The first panel of Table 10 indicates that cars are less profitable by \$250 to \$300 in the eight weeks leading up to the initiation of a customer cash promotion.

Such cars also have 3 to 4 more days of inventory in the system in the five weeks preceding a promotion than cars in weeks not leading up to promotion. Subsegment market share is less consistently different between the two samples.

Table 11 reports the same analysis for dealer cash. In the first panel, vehicles are also less profitable in weeks leading up to dealer cash promotions than in weeks not leading up to a dealer cash promotion, although the difference is smaller than for customer cash. Differences in days of inventory and subsegment market share are generally not statistically significant.

While these results are preliminary, they suggest that manufacturers may apply different kinds of promotions to different situations. In particular, we find, consistent with our intuition, that the onset of a customer cash promotion is preceded by relatively high days of inventory, while the onset of a dealer cash promotion is not. However, while we do find that the onset of a dealer cash promotion is preceded by low vehicle margins, the same is true for the onset of a customer cash promotion. Further analysis is needed before drawing any but preliminary conclusions.

8 Conclusion

In the introduction, this paper proposed to answer two questions. First, how does the information asymmetry between dealers and consumers about manufacturer-supplied surplus affect the division of this surplus in customer-dealer negotiations? Second, do different types of promotions have different incremental sales effects?

In answer to the first question, we find that information asymmetry has a large effect on the division of surplus. Customers were able to obtain almost twice the surplus in circumstances in which they were very likely to know what was on the table compared to circumstances in which they were unlikely to know. In answer to the second question, we find that customer cash promotions have a larger dollar-for-dollar effect on sales than do dealer cash promotions.

In the automotive sector, manufacturers and dealers are separate entities, as dealer franchise laws require. Each party, however, is clearly dependent on the other in many ways, having made multiple relationship specific investments. Anecdotally, we know that dealers communicate to manufacturers through field reps that they prefer dealer cash to customer cash. The results of this paper suggest that dealer cash may indeed be more profitable for dealers. But the results also suggest that dealer cash may not be the most effective instrument for manufacturers in increasing. Manufacturers have traditionally strongly relied on their field reps and their dealers to assess the success of promotions. Our findings suggest that they might be better served by directly evaluating promotions with transactions data that has recently become available in the automotive industry.

What can we learn from these results about promotions in general? We note that customer cash is accompanied by advertising, and by the lure of low prices, both of which are likely to encourage customers to visit dealers and to consider purchasing a new car. On the other hand, since a greater proportion of dealer cash than of customer remains with the dealer, dealer cash has the potential to, dollar-for-dollar, motivate more dealer sales effort. The sales effect results suggest that, at least at levels currently prevailing in the auto industry, the marginal sales effect of getting more customers through the door is greater than that of motivating dealers. In marketing parlance, a “pull” promotion is more effective than a “push” promotion.

References

- [1] Besanko, David, Jean-Pierre Dubé, and Sachin Gupta (forthcoming) “Own-brand and Cross-brand Retail Pass-through,” *Marketing Science*.
- [2] Berto Villas-Boas, Sofia (2003) “Vertical Contracts between Manufacturers and Retailers: An Empirical Analysis,” mimeo, University of California, Berkeley.
- [3] Blattberg, Robert C. and Scott A. Neslin (1990) *Sales Promotion: Concepts, Methods, and Strategies*, Englewood Cliffs, New Jersey:Prentice-Hall, Inc.
- [4] Chevalier, Michel and Ronald C. Curhan (1976) “Retail Promotions as a Function of Trade Promotions: A Descriptive Analysis,” *Sloan Management Review*, 18 (3) pp. 19-32.
- [5] Gerstner, Eitan and James D. Hess (1991) “A Theory of Channel Price Promotions,” *American Economic Review*, 81 (4) pp. 872-886.
- [6] Galiani, Sebastian, Paul Gertler, and Ernesto Schargrotsky (2002) “Water for Life: The Impact of the Privatization of Water Services on Child Mortality,” mimeo, University of California, Berkeley.
- [7] Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw (2001) “Identification and Estimation of Treatment Effects with a Regression Discontinuity Design,” *Econometrica*, 69 (1) pp. 201-209.
- [8] Heckman, James J., Hidehiko Ichimura, and Petra E. Todd (1997) “Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme,” *Review of Economic Studies*, 64 (4) pp. 605-654.
- [9] Imbens, Guido W. and Joshua D. Angrist (1994) “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 62 (2) pp. 467-475.
- [10] Kane, Thomas J. (2003) “A Quasi-Experimental Estimate of the Impact of Financial Aid on College-Going,” *NBER Working Paper #9703*.
- [11] Kumar, Nanda, Surendra Rajiv, and Abel Jeuland (2001) “Effectiveness of Trade Promotions: Analyzing the Determinants of Retail Pass Through,” *Marketing Science*, 20 (4) pp. 382-404.
- [12] Lal, Rajiv (1990) “Manufacturer Trade Deals and Retail Price Promotions,” *Journal of Marketing Research*, 27 () pp. 428-444.

- [13] Lal, Rajiv, John D. C. Little, and J. Miguel Villas-Boas (1996) “A Theory of Forward Buying, Merchandising, and Trade Deals,” *Marketing Science*, 15 (1) pp. 21-37.
- [14] Moorthy, Sridhar (2003) “A General Theory of Pass-through in Channels with Category Management,” mimeo, University of Toronto.
- [15] Narasimhan, Chakravarthi (1988) “Competitive Promotional Strategies,” *Journal of Business*, 61 (4) pp. 427-449.
- [16] Nevo, Aviv and Catherine Wolfram (2002) “Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals,” *RAND Journal of Economics*, 33 (2) pp. 319-339.
- [17] Raju, Jagmohan R., V. Srinivasan, and Rajiv Lal (1990) “The Effects of Brand Loyalty on Competitive Price Promotional Strategies,” *Management Science*, 36 (3) pp. 276-304.
- [18] Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Risso (2003) “Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?,” *Quantitative Marketing and Economics*, 1 (1) pp. 65-92.
- [19] Rao, Ram C. (1991) “Pricing and Promotion in Asymmetric Duopolies,” *Marketing Science*, 10 (2) pp. 131-144.
- [20] Tyagi, Rajiv (1999) “A Characterization of Retailer Response to Manufacturer Trade Deals,” *Journal of Marketing Research*, 65 (2) pp. 253-272.
- [21] Van der Klaauw (2002) “Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression Discontinuity Approach,” *International Economic Review*, 43 (4) pp. 1249-1287.
- [22] Varian, Hal R. (1980) “A Model of Sales,” *American Economic Review*, 70 (4) pp. 651-659.
- [23] Walters, Rockney G. (1989) “An Empirical Investigation into Retailer Response to Manufacturer Trade Promotions,” *Journal of Retailing*, 48 () pp. 253-272.

Table 1: Price effect, difference-in-differences[†]

	(1)	(2)	(3)	(4)	(5)	(6)
Customer Cash	-0.88 (0.03)**	-0.84 (0.04)**	-0.82 (0.04)**	-0.90 (0.04)**	-0.74 (0.04)**	-0.84 (0.04)**
GM Card	-1.06 (0.03)**	-1.03 (0.04)**	-1.02 (0.04)**	-1.02 (0.04)**	-1.04 (0.04)**	-1.03 (0.04)**
Dealer Cash	-0.39 (0.07)**	-0.42 (0.07)**	-0.52 (0.09)**	-0.47 (0.08)**	-0.41 (0.07)**	-0.42 (0.07)**
CustCash*Competition		-0.01 (0.004)*	-0.01 (0.005)*	-0.01 (0.005)*	-0.01 (0.004)*	-0.01 (0.004)*
DealCash*Competition		0.01 (0.01)	0.02 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)
Num. CustCash Prom. in Seg.		21.53 (14.21)		21.65 (14.19)	19.40 (14.09)	21.61 (14.22)
Num. DealCash Prom. in Seg.		-15.26 (16.94)		-14.73 (16.80)	-11.32 (16.77)	-15.32 (16.93)
Avg. CustCash Prom. in Seg.			0.18 (0.13)			
Avg. DealCash Prom. in Seg.			-0.28 (0.14)*			
CustCash*CustCash decrease					-0.22 (0.04)**	
DealCash*DealCash decrease					-0.16 (0.12)	
Sales Manager Incen.						0.10 (0.58)
Sales Rep Incen.						0.04 (0.61)
# Competing Dealers	-7.66 (5.60)	-4.87 (6.75)	-5.79 (6.80)	-4.15 (6.82)	-4.48 (6.75)	-4.86 (6.75)
Female	144.17 (12.01)**	144.40 (12.00)**	146.34 (12.05)**	120.72 (13.50)**	143.78 (11.97)**	144.37 (12.00)**
%Black	474.48 (92.72)**	473.90 (92.79)**	476.36 (94.39)**	365.69 (93.47)**	473.38 (93.14)**	474.11 (92.77)**
%Hispanic	-36.43 (86.37)	-32.62 (86.62)	-32.24 (86.96)	-229.67 (94.66)*	-27.99 (86.40)	-32.69 (86.61)
CustCash*Female				0.04 (0.02)*		
CustCash*Pct.Black				0.32 (0.10)**		
CustCash*Pct.Hispanic				0.44 (0.10)**		
DealCash*Female				0.07 (0.03)**		
DealCash*Pct.Black				0.003 (0.12)		
DealCash*Pct.Hispanic				0.28 (0.14)*		
Constant	26611 (163)**	26565 (165)**	26607 (180)**	26609 (165)**	26627 (166)**	26565 (165)**
Week*Segment fixed effects	Yes	Yes	No (week)	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	133424	133424	133424	133424	133424	133424
Adj. R-squared	0.97	0.97	0.97	0.97	0.97	0.97

[†] Unreported demographic characteristics include: census block percentages of residents with college graduates; with less than high school education; who are blue collar workers, executives, professionals, or technicians; who are Asian; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value. Unreported controls include whether the transaction occurred on the weekend, at the end of the month, or at the end of the year; the time since the model introduction; and whether the transaction took place in Southern California.

Table 2: Price effect, difference-in-differences after sample restriction to propensity score common support[†]

	(1)	(2)	(3)	(4)	(5)	(6)
Customer Cash	-0.88 (0.04)**	-0.88 (0.05)**	-0.87 (0.05)**	-0.94 (0.05)**	-0.86 (0.05)**	-0.88 (0.05)**
GM Card	-1.04 (0.03)**	-1.04 (0.04)**	-1.03 (0.04)**	-1.03 (0.04)**	-1.04 (0.04)**	-1.04 (0.04)**
Dealer Cash	-0.34 (0.07)**	-0.42 (0.08)**	-0.59 (0.10)**	-0.47 (0.08)**	-0.42 (0.08)**	-0.42 (0.08)**
CustCash*Competition		-0.002 (0.006)	-0.002 (0.006)	-0.004 (0.006)	-0.003 (0.006)	-0.002 (0.01)
DealCash*Competition		0.022 (0.012)+	0.035 (0.011)**	0.022 (0.012)+	0.022 (0.012)+	0.022 (0.01)+
Num. CustCash Prom. in Seg.		33.70 (19.30)+		33.62 (19.21)+	32.21 (19.24)+	33.91 (19.31)+
Num. DealCash Prom. in Seg.		-54.70 (21.13)**		-53.63 (20.96)**	-54.04 (21.33)*	-54.88 (21.12)**
Avg. CustCash Prom. in Seg.			0.24 (0.19)			
Avg. DealCash Prom. in Seg.			-0.33 (0.22)			
CustCash*CustCash decrease					-0.16 (0.06)	
DealCash*DealCash decrease					-0.34 (0.21)	
Sales Manager Incen.						-0.15 (0.93)
Sales Rep Incen.						0.29 (1.00)
# Competing Dealers	-3.76 (7.46)	-6.67 (8.39)	-9.83 (8.41)	-5.99 (8.51)	-6.46 (8.37)	-6.66 (8.40)
Female	172.66 (19.26)**	172.76 (19.33)**	177.06 (19.43)**	148.38 (22.62)**	172.41 (19.27)**	172.71 (19.31)**
%Black	542.30 (127.48)**	537.82 (127.72)**	549.01 (129.71)**	473.10 (144.76)**	535.64 (127.92)**	538.11 (127.67)**
%Hispanic	-15.17 (119.66)	-15.91 (119.37)	0.51 (120.32)	-229.03 (131.33)+	-13.51 (119.27)	-15.98 (119.36)
CustCash*Female				0.02 (0.03)		
CustCash*Pct.Black				0.35 (0.15)*		
CustCash*Pct.Hispanic				0.43 (0.13)**		
DealCash*Female				0.08 (0.04)*		
DealCash*Pct.Black				-0.22 (0.19)		
DealCash*Pct.Hispanic				0.32 (0.19)+		
Constant	26045 (239.45)**	26022 (238.79)**	25896 (205.31)**	26064 (238.71)**	26012 (238.64)**	26021 (238.86)**
Week*Segment fixed effects	Yes	Yes	No (week)	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	68278	68278	68278	68278	68278	68278
Adj. R-squared	0.96	0.96	0.96	0.96	0.96	0.96

[†] Unreported demographic characteristics include: census block percentages of residents with college graduates; with less than high school education; who are blue collar workers, executives, professionals, or technicians; who are Asian; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value. Unreported controls include whether the transaction occurred on the weekend, at the end of the month, or at the end of the year; the time since the model introduction; and whether the transaction took place in Southern California.

Table 3: Price effect, regression discontinuity[†]

	(1)	(2)	(3)	(4)	(5)
Customer Cash	-0.81 (0.07)**	-0.84 (0.08)**	-0.92 (0.09)**	-0.79 (0.08)**	-0.84 (0.08)**
GM Card	-1.12 (0.06)**	-1.14 (0.07)**	-1.13 (0.07)**	-1.15 (0.07)**	-1.14 (0.07)**
Dealer Cash	-0.38 (0.08)**	-0.48 (0.10)**	-0.54 (0.12)**	-0.47 (0.09)**	-0.48 (0.10)**
CustCash*Competition		0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)
DealCash*Competition		0.03 (0.01)**	0.03 (0.01)*	0.03 (0.01)**	0.03 (0.01)*
Avg. CustCash Prom. in Seg.		0.28 (0.23)	0.29 (0.23)	0.29 (0.23)	0.28 (0.23)
Avg. DealCash Prom. in Seg.		-0.52 (0.18)**	-0.50 (0.18)**	-0.49 (0.18)**	-0.52 (0.18)**
CustCash*CustCash decrease				-0.17 (0.09)+	
DealCash*DealCash decrease				-0.30 (0.13)*	
Sales Manager Incen.					-1.74 (2.43)
Sales Rep Incen.					-0.94 (0.96)
# Competing Dealers	-18.54 (6.98)**	-35.47 (9.64)**	-33.09 (9.53)**	-34.10 (9.38)**	-35.35 (9.68)**
Female	157.10 (36.14)**	158.20 (36.07)**	55.18 (49.33)	158.25 (36.10)**	157.65 (36.05)**
%Black	481.46 (235.64)*	473.58 (232.32)*	17.65 (348.37)	465.87 (232.53)*	478.67 (232.47)*
%Hispanic	-480.91 (234.25)*	-456.38 (231.82)*	-722.65 (314.65)*	-452.69 (231.16)+	-453.71 (232.08)+
CustCash*Female			0.06 (0.05)		
CustCash*Pct.Black			0.61 (0.37)+		
CustCash*Pct.Hispanic			0.43 (0.21)*		
DealCash*Female			0.18 (0.07)**		
DealCash*Pct.Black			0.17 (0.45)		
DealCash*Pct.Hispanic			0.08 (0.39)		
Constant	23,883 (434)**	23,966 (434)**	24,043 (434)**	24,098 (452)**	23,965 (434)**
Week fixed effects	Yes	Yes	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	12572	12572	12572	12572	12572
R-squared	0.96	0.96	0.96	0.96	0.96

[†] Unreported demographic characteristics include: census block percentages of residents who are college graduates; with less than high school education; who are blue collar workers, executives, professionals, or technicians; who are Asian; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value. Unreported controls include whether the transaction occurred on the weekend, at the end of the month, or at the end of the year; the time since the model introduction; and whether the transaction took place in Southern California.

Table 4: Price effect, regression discontinuity after sample restriction to propensity score common support[†]

	(1)	(2)	(3)	(4)	(5)
Customer Cash	-0.81 (0.06)**	-0.85 (0.08)**	-0.94 (0.09)**	-0.82 (0.07)**	-0.85 (0.08)**
GM Card	-1.13 (0.08)**	-1.17 (0.09)**	-1.15 (0.09)**	-1.17 (0.08)**	-1.17 (0.09)**
Dealer Cash	-0.41 (0.09)**	-0.51 (0.10)**	-0.55 (0.12)**	-0.50 (0.10)**	-0.51 (0.10)**
CustCash*Competition		0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)
DealCash*Competition		0.03 (0.02)*	0.03 (0.02)*	0.03 (0.02)*	0.03 (0.02)*
Avg. CustCash Prom. in Seg.		0.38 (0.27)	0.40 (0.27)	0.38 (0.27)	0.39 (0.27)
Avg. DealCash Prom. in Seg.		-0.46 (0.21)*	-0.44 (0.21)*	-0.42 (0.21)*	-0.46 (0.21)*
CustCash*CustCash decrease				-0.23 (0.11)*	
DealCash*DealCash decrease				-0.43 (0.16)**	
Sales Manager Incen.					-1.91 (3.03)
Sales Rep Incen.					-1.11 (1.01)
# Competing Dealers	-19.72 (7.67)*	-38.38 (10.59)**	-36.43 (10.36)**	-37.27 (10.05)**	-38.27 (10.61)**
Female	166.91 (39.63)**	167.30 (39.66)**	44.69 (57.93)	167.62 (39.36)**	166.63 (39.59)**
%Black	463.26 (243.49)+	451.72 (238.72)+	15.52 (368.78)	444.70 (238.32)+	455.13 (238.75)+
%Hispanic	-463.84 (253.51)+	-438.65 (251.09)+	-641.15 (381.17)+	-429.62 (250.76)+	-435.01 (252.10)+
CustCash*Female			0.07 (0.06)		
CustCash*Pct.Black			0.63 (0.38)+		
CustCash*Pct.Hispanic			0.45 (0.25)+		
DealCash*Female			0.22 (0.08)**		
DealCash*Pct.Black			0.13 (0.50)		
DealCash*Pct.Hispanic			-0.10 (0.52)		
Constant	23,821.34 (693.41)**	23,912.97 (698.01)**	23,992.98 (703.45)**	23,821.80 (707.04)**	23,905.91 (697.34)**
Week fixed effects	Yes	Yes	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	11157	11157	11157	11157	11157
R-squared	0.96	0.96	0.96	0.96	0.96

[†] Unreported demographic characteristics include: census block percentages of residents who are college graduates; with less than high school education; who are blue collar workers, executives, professionals, or technicians; who are Asian; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value. Unreported controls include whether the transaction occurred on the weekend, at the end of the month, or at the end of the year; the time since the model introduction; and whether the transaction took place in Southern California.

Table 5: Price effect, regression discontinuity (2 days on either side of promotion change) after sample restriction to propensity score common support[†]

	(1)	(2)	(3)	(4)	(5)
Customer Cash	-0.78 (0.12)**	-0.91 (0.12)**	-0.87 (0.15)**	-0.83 (0.12)**	-0.91 (0.12)**
GM Card	-1.09 (0.08)**	-1.17 (0.09)**	-1.18 (0.09)**	-1.20 (0.09)**	-1.17 (0.09)**
Dealer Cash	-0.38 (0.08)**	-0.50 (0.13)**	-0.53 (0.18)**	-0.50 (0.13)**	-0.50 (0.13)**
CustCash*Competition		0.03 (0.02)*	0.03 (0.02)*	0.03 (0.02)*	0.03 (0.02)*
DealCash*Competition		0.04 (0.03)	0.04 (0.03)	0.04 (0.03)	0.04 (0.03)
Avg. CustCash Prom. in Seg.		0.05 (0.36)	0.05 (0.36)	0.10 (0.36)	0.07 (0.35)
Avg. DealCash Prom. in Seg.		-0.49 (0.26)+	-0.46 (0.27)+	-0.46 (0.27)+	-0.49 (0.26)+
CustCash*CustCash decrease				-0.34 (0.18)+	
DealCash*DealCash decrease				0.23 (0.32)	
Sales Manager Incen.					-5.19 (9.21)
Sales Rep Incen.					-0.35 (1.85)
# Competing Dealers	-14.26 (14.80)	-49.70 (18.34)**	-49.93 (17.38)**	-48.96 (17.73)**	-48.88 (18.21)**
Female	120.61 (59.44)*	123.42 (59.06)*	148.68 (118.53)	120.64 (58.84)*	124.35 (59.35)*
%Black	1,408.25 (524.53)**	1,370.43 (514.98)**	1,110.67 (892.66)	1,351.55 (517.73)**	1,377.43 (515.44)**
%Hispanic	-311.87 (512.35)	-286.67 (517.42)	-166.54 (764.06)	-250.83 (517.02)	-295.08 (519.07)
CustCash*Female			-0.11 (0.15)		
CustCash*Pct.Black			0.17 (0.78)		
CustCash*Pct.Hispanic			-0.18 (0.71)		
DealCash*Female			0.11 (0.14)		
DealCash*Pct.Black			0.28 (0.83)		
DealCash*Pct.Hispanic			-0.01 (1.03)		
Constant	24,376.03 (707.02)**	24,499.17 (703.06)**	24,519.24 (707.75)**	24,450.87 (722.13)**	24,501.61 (702.24)**
Week fixed effects	Yes	Yes	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	3218	3218	3218	3218	3218
R-squared	0.96	0.96	0.96	0.96	0.96

[†] Unreported demographic characteristics include: census block percentages of residents who are college graduates; with less than high school education; who are blue collar workers, executives, professionals, or technicians; who are Asian; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value. Unreported controls include whether the transaction occurred on the weekend, at the end of the month, or at the end of the year; the time since the model introduction; and whether the transaction took place in Southern California.

Table 6: Weekly sales effect, difference-in-differences

	(1)	(2)	(3)	(4)
Customer Cash	0.0017 (0.0001)**	0.0017 (0.0003)**	0.0017 (0.0003)**	0.0017 (0.0003)**
Dealer Cash	0.0005 (0.0001)**	0.0005 (0.0003)	0.0005 (0.0003)	0.0005 (0.0003)
Avg. CustCash Prom. in Seg.		-0.0005 (0.0012)	-0.0005 (0.0012)	-0.0005 (0.0012)
Avg. DealCash Prom. in Seg.		-0.0005 (0.0010)	-0.0005 (0.0010)	-0.0004 (0.0010)
CustCash Decrease			0.0000 (0.0003)	
DealCash Decrease			-0.0002 (0.0005)	
Sales Manager Incen.				0.0377 (0.0111)**
Sales Rep Incen.				-0.0023 (0.0043)
Model Age	-0.7877 (0.0962)**	-0.7853 (0.4418)+	-0.7855 (0.4418)+	-0.8036 (0.4461)+
Model Age 5-13 Months	1.9853 (0.1773)**	1.9771 (0.2609)**	1.9775 (0.2615)**	1.9888 (0.2634)**
Model Age over 14 Months	-2.8933 (0.3193)**	-2.8959 (0.7056)**	-2.8963 (0.7051)**	-2.8730 (0.7067)**
SouthernCal	-0.2336 (0.0983)*	-0.2334 (0.2220)	-0.2329 (0.2219)	-0.2182 (0.2148)
Constant	-0.1762 (0.8917)	-0.1470 (2.6399)	-0.1472 (2.6398)	-0.2739 (2.6587)
Week fixed effects	Yes	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes
Observations	21614	21614	21614	21614
R-squared	0.71	0.71	0.71	0.71

Table 7: Weekly sales effect, regression discontinuity

	(1)	(2)	(3)	(4)
Customer Cash	0.0021 (0.0006)**	0.0022 (0.0006)**	0.0022 (0.0006)**	0.0022 (0.0006)**
Dealer Cash	0.0011 (0.0006)+	0.0011 (0.0006)*	0.0011 (0.0006)*	0.0011 (0.0006)*
Avg. CustCash Prom. in Seg.		-0.0026 (0.0021)	-0.0026 (0.0021)	-0.0026 (0.0021)
Avg. DealCash Prom. in Seg.		-0.0024 (0.0016)	-0.0025 (0.0016)	-0.0025 (0.0016)
CustCash Decrease			0.0001 (0.0004)	
DealCash Decrease			-0.0001 (0.0011)	
Sales Manager Incen.				-0.0078 (0.0177)
Sales Rep Incen.				0.0022 (0.0040)
Model Age	0.0393 (0.4133)	0.0601 (0.4077)	0.0587 (0.4076)	0.0586 (0.4095)
Model Age 5-13 Months	2.3621 (1.0178)*	2.2592 (0.9756)*	2.2671 (0.9764)*	2.2522 (0.9837)*
Model Age over 14 Months	-2.9867 (1.7403)+	-3.1565 (1.7178)+	-3.1468 (1.7187)+	-3.1748 (1.7295)+
SouthernCal	1.0570 (0.4154)*	1.1036 (0.4121)**	1.1000 (0.4124)**	1.1086 (0.4105)**
Constant	8.8024 (2.9777)**	9.2616 (3.0224)**	9.1508 (3.0103)**	9.2593 (3.0339)**
Week fixed effects	Yes	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes
Observations	1628	1628	1628	1628
R-squared	0.85	0.85	0.85	0.85

Table 8: Market share effect, difference-in-differences

	(1)	(2)	(3)	(4)
Customer Cash	0.000020 (0.000004)**	0.000022 (0.000004)**	0.000022 (0.000004)**	0.000022 (0.000004)**
Dealer Cash	0.000006 (0.000005)	0.000007 (0.000005)	0.000007 (0.000005)	0.000007 (0.000005)
Avg. CustCash Prom. in Seg.		-0.000044 (0.000017)*	-0.000044 (0.000017)*	-0.000044 (0.000017)*
Avg. DealCash Prom. in Seg.		-0.000017 (0.000012)	-0.000016 (0.000012)	-0.000017 (0.000012)
CustCash Decrease			-0.000003 (0.000003)	
DealCash Decrease			-0.000009 (0.000012)	
Sales Manager Incen.				0.000197 (0.000086)*
Sales Rep Incen.				0.000053 (0.000044)
Model Age	-0.005602 (0.002665)*	-0.005425 (0.002671)*	-0.005448 (0.002674)*	-0.005538 (0.002700)*
Model Age 5-13 Months	0.022824 (0.003731)**	0.022032 (0.003692)**	0.022149 (0.003695)**	0.022194 (0.003701)**
Model Age over 14 Months	-0.026312 (0.007420)**	-0.026730 (0.007395)**	-0.026557 (0.007389)**	-0.026644 (0.007397)**
SouthernCal	0.001770 (0.001710)	0.001801 (0.001709)	0.001830 (0.001706)	0.001916 (0.001665)
Constant	0.167276 (0.019739)**	0.169348 (0.019769)**	0.170451 (0.019915)**	0.168369 (0.019860)**
Week fixed effects	Yes	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes
Observations	21614	21614	21614	21614
R-squared	0.73	0.73	0.73	0.73

Table 9: Market share effect, regression discontinuity

	(1)	(2)	(3)	(4)
Customer Cash	0.000017 (0.000006)**	0.000019 (0.000006)**	0.000019 (0.000006)**	0.000019 (0.000006)**
Dealer Cash	0.000004 (0.000007)	0.000005 (0.000007)	0.000005 (0.000007)	0.000005 (0.000007)
Avg. CustCash Prom. in Seg.		-0.000083 (0.000035)*	-0.000083 (0.000035)*	-0.000081 (0.000033)*
Avg. DealCash Prom. in Seg.		-0.000038 (0.000025)	-0.000038 (0.000025)	-0.000038 (0.000025)
CustCash Decrease			-0.000003 (0.000006)	
DealCash Decrease			-0.000004 (0.000015)	
Sales Manager Incen.				-0.000128 (0.000207)
Sales Rep Incen.				0.000172 (0.000072)*
Model Age	0.002302 (0.004393)	0.002417 (0.004167)	0.002461 (0.004161)	0.002122 (0.004092)
Model Age 5-13 Months	0.030065 (0.015329)+	0.027151 (0.014973)+	0.026940 (0.014987)+	0.027816 (0.014978)+
Model Age over 14 Months	-0.012930 (0.024843)	-0.017708 (0.024398)	-0.017948 (0.024404)	-0.018107 (0.024397)
SouthernCal	0.010559 (0.006275)+	0.011032 (0.006309)+	0.011147 (0.006286)+	0.011645 (0.006429)+
Constant	0.222006 (0.032316)**	0.230333 (0.032786)**	0.233436 (0.032404)**	0.228833 (0.032603)**
Week fixed effects	Yes	Yes	Yes	Yes
Car fixed effects	Yes	Yes	Yes	Yes
Observations	1628	1628	1628	1628
R-squared	0.85	0.85	0.85	0.86

Table 10: T-tests of of difference in customer cash prediction variables

Vehicle profitability

Number of weeks priorto week with customer cash	...to week without customer cash	Difference	T-test p-value
1	1440	1711	-271	0.000
2	1473	1725	-252	0.000
3	1490	1738	-248	0.000
4	1469	1750	-281	0.000
5	1484	1763	-279	0.000
6	1498	1777	-279	0.000
7	1546	1792	-246	0.000
8	1514	1813	-299	0.000

Days of inventory

Number of weeks prior	...to week with customer cash	...to week without customer cash	Difference	T-test p-value
1	17.5	14.2	3.3	0.004
2	17.7	14.7	3.0	0.009
3	18.5	15.4	3.1	0.017
4	19.6	16.6	3.0	0.060
5	22.4	18.3	4.1	0.075
6	22.4	21.2	1.2	0.672
7	23.5	28.6	-5.1	0.138
8	29.6	32.3	-2.7	0.632

Subsegment market share (percentage points)

Number of weeks prior	...to week with customer cash	...to week without customer cash	Difference	T-test p-value
1	10.1	9.5	0.6	0.334
2	10.3	9.3	1.0	0.166
3	10.3	9.2	1.1	0.137
4	10.4	9.1	1.3	0.064
5	10.2	9.0	1.2	0.089
6	10.2	8.8	1.4	0.053
7	9.7	8.6	1.1	0.106
8	10.4	8.5	1.9	0.005

Table 11: T-tests of of difference in dealer cash prediction variables

Vehicle profitability

Number of weeks priorto week with dealer cash	...to week without dealer cash	Difference	T-test p-value
1	1452	1624	-172	0.000
2	1476	1639	-163	0.000
3	1557	1647	-90	0.045
4	1538	1657	-119	0.015
5	1511	1671	-160	0.000
6	1574	1681	-107	0.026
7	1597	1695	-98	0.052
8	1562	1709	-147	0.003

Days of inventory

Number of weeks prior	...to week with dealer cash	...to week without dealer cash	Difference	T-test p-value
1	16.1	14.4	1.7	0.043
2	16.3	14.7	1.6	0.089
3	16.8	15.2	1.6	0.121
4	17.4	16.0	1.4	0.206
5	20.9	17.6	3.3	0.121
6	22.1	19.8	2.3	0.446
7	29.1	26.5	2.6	0.660
8	23.9	29.6	-5.7	0.153

Subsegment market share (percentage points)

Number of weeks prior	...to week with dealer cash	...to week without dealer cash	Difference	T-test p-value
1	9.7	10.0	-0.3	0.575
2	10.1	10.0	0.1	0.849
3	10.2	9.9	0.3	0.652
4	10.2	9.8	0.4	0.474
5	10.4	9.7	0.7	0.225
6	10.4	9.6	0.8	0.154
7	10.5	9.4	1.1	0.064
8	10.4	9.2	1.2	0.028

Figure 3: Distribution of propensity scores for the control group (transactions within 2 days before a promotion change) and the treatment group (transactions within 2 days after a promotion change)

