NBER WORKING PAPER SERIES

$1000 CASH BACK: ASYMMETRIC INFORMATION IN AUTO MANUFACTURER PROMOTIONS

Meghan Busse
Florian Zettelmeyer
Jorge Silva-Risso

Working Paper 10887
http://www.nber.org/papers/w10887

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2004

We are grateful for helpful comments from Tim Bresnahan, Severin Borenstein, J.P. Dube, Chris Knittel, Fiona Scott Morton, participants at the NBER IO Program Meeting, the HBS Strategy Conference, and seminar participants at UC Berkeley, UC Davis, UC San Diego, 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.silvarisso@ucr.edu The views expressed herein are those of the author(s) and not necessarily those of the National Bureau of Economic Research.

$1000 Cash Back: Asymmetric Information in Auto Manufacturer Promotions
Meghan Busse, Florian Zettelmeyer, Jorge Silva-Risso
NBER Working Paper No. 10887
November 2004
JEL No. D82, L15, L62, L11

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 discounts to dealers, which are not publicized. While the payments nominally go entirely to one party or the other, the real division of the manufacturer-supplied surplus between dealer and customer depends on what price the two parties negotiate. These two types of promotions thus form a natural experiment of the effect of information asymmetry on bargaining outcomes: in the customer rebate case, the parties are symmetrically informed about the availability of the manufacturer-supplied surplus, while in the dealer discount case, the dealer will generally have an informational advantage. The aim of this paper is to compare, in appropriate settings and with appropriate controls, the price outcomes of transactions conducted under these two types of promotions in order to empirically quantify the effect of this information asymmetry. We show that customers obtain approximately 80% of the surplus in cases when they are likely to be well-informed about the promotion (customer rebate), and approximately 35% when they are likely to be uninformed (dealer discount). For a promotion of average size, this difference translates to customers being worse off by $500 when they do not know that the promotion is being offered.

Meghan Busse
UC, Berkeley
meghan@haas.berkeley.edu

Jorge Silva-Risso
UC, Riverside
jorge.silva-risso@ucr.edu

Florian Zettelmeyer
Haas School of Business
UC, Berkeley
Berkeley, CA 94720
and NBER
florian@haas.berkeley.edu

1 Introduction

Although retail demand for an automobile fluctuates due to changing economic conditions, seasonality, and the stage of the model's life cycle, manufacturers rarely vary published retail and invoice prices of a particular model over the course of the model year. The choice to have rigid prices is potentially very costly for auto manufacturers: inventory holding costs for automobiles are high, and so are the costs of changing production schedules to adapt to current demand. As a result, "incentive promotions" play an important role in automobile manufacturers' product market strategies by enabling retail prices to adjust to fluctuating demand conditions. Incentive promotions take a variety of forms. The most common are cash rebates to customers, cash 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 two primary types of cash rebates, which we refer to 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. While the promotion payments are nominally directed to one party or the other, who ultimately receives the benefit of the promotion depends on the outcome of the price negotiation process. For example, if a customer buying during a \$1000 customer cash rebate were to agree to a price that is higher by \$200 than the price he or she would have negotiated without the promotion, then the customer's out-of-pocket expenditure would be lower by only \$800 compared to what he or she would otherwise have paid. The dealer would be reaping \$200 of the benefit of that customer cash promotion. Conversely, if a dealer were induced by a \$1000 dealer cash promotion to agree to a price that is lower by \$500 than he or she otherwise would have, then the customer would obtain \$500 of the benefit of the dealer cash promotion. In short, from an economics perspective, a \$1000 customer cash promotion and a \$1000 dealer cash promotion are both \$1000 of manufacturer-supplied economic surplus that will be divided between the two parties through the bargaining process.

The chief difference between the two promotions is the information environment surrounding them: 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. These two types of promotions thus form a natural experiment of the effect of information asymmetry on bargaining outcomes: in the customer cash case, the parties are symmetrically informed about the availability of the manufacturer-supplied surplus, while in the dealer cash case, the dealer will generally have an informational advantage. The aim of this paper is to compare, in appropriate settings and with appropriate controls, the price outcomes of transactions conducted under these two types of promotions in order to draw inferences about how information asymmetry affects bargaining outcomes. 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]

Our empirical setting overcomes one of the key problems of testing predictions about how asymmetric information between parties affects bargaining outcomes, which is that researchers are often unable to observe the nature of the information asymmetry. In our case, the asymmetry is over a very specific issue—whether there is a promotion or not and how large it is—and we can divide our observations sharply into cases in which there is almost certainly symmetric information and cases in which there is very likely to be asymmetric information.

In examining how information asymmetry between dealers and consumers about manufacturer-supplied surplus affects the division of this surplus in negotiations, we anticipate that, consistent with the results of a broad set of theoretical bargaining models, 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.

We draw on methods from the program evaluation literature to estimate how much of the manufacturer-supplied customer cash and dealer cash is passed on to consumers. Using both a difference-in-differences approach and a regression discontinuity approach, we find that customers obtain 70-90% of the surplus supplied by manufacturers in customer cash promotions, but only 30-40% of the surplus in dealer cash promotions. This is consistent with the theoretical prediction and implies that, for a promotion of average size, consumers receive

[1]While there is experimental work and also some empirical literature on private information and bargaining in labor disputes, to our knowledge there is next to no empirical work in product markets, except for Scott Morton, Zettelmeyer, and Silva-Risso (2004)

approximately \$500 more of the surplus if they know that the promotion is in effect.

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 rates of the different types of promotions. In section 6 we test the validity of assumptions that were maintained when identifying these pass-through rates. In section 7 we consider several extensions to the main result, including how pass-through varies with competition and demographics. We conclude in section 8.

2 Literature Background

There are two strands of literature that inform this paper. One is the literature on game-theoretic models of bargaining under asymmetric information. This literature forms the basis of our prediction that because customers are informed about customer cash and not dealer cash, they will obtain a greater fraction of customer cash than of dealer cash. Second, our paper is also related to a literature on manufacturer promotions. We discuss these two strands in sequence.

2.1 Bargaining under asymmetric information

The segment of the bargaining literature that is relevant to this paper relates information asymmetries between bargaining parties to the division of surplus in the negotiation. These models can be found in the large game-theoretic literature on bargaining with incomplete or private information (see the excellent review papers by Kennan and Wilson (1993) and Ausubel, Cramton, and Deneckere (2002) for an overview). While the primary focus of this literature is on whether economically efficient transactions take place, some of the important papers also make clear predictions with regards to the effect of asymmetric information on the division of surplus.

A natural way to model the car buying process is as a dynamic process of bilateral negotiation. The models that apply to our setting follow the seminal paper by Rubinstein (1982), but assume that one of the bargaining parties has incomplete information about the reservation price of the opponent. To match our case, it is the buyer who should be thought of as the uninformed party. Bargaining is typically considered to follow one of two protocols. In a "buyer-offer game," only the buyer (the uninformed party) is allowed to make offers, which the seller is allowed only to reject or accept. If the seller rejects an offer, the buyer can make another offer. The game ends when the seller accepts an offer. Games of this type can have multiple Bayesian equilibria (see Fudenberg and Tirole (1991), p. 399). However, Fudenberg, Levine, and Tirole (1985) and Gul, Sonnenschein, and Wilson (1986) show that under the "stationary equilibrium" refinement, the buyer-offer game allows the buyer to screen seller types by a series of sequential, increasing price offers. In terms of our empirical prediction, the salient feature of this equilibrium is that, as long as buyers are not infinitely patient, the buyer is not able to perfectly screen among the seller types and is therefore worse off than he or she would be in a situation in which he or she had complete information about the seller's reservation price.

This same prediction also comes out of an "alternating-offer game." In such a game the buyer and seller alternate in making proposals. Ausubel and Deneckere (1998) show in such a model that under the "assuredly perfect equilibrium" refinement, there exists a unique equilibrium in which the buyer is able to screen seller types, albeit imperfectly. As in the buyer-offer game, the buyer's equilibrium payoff is bounded from above by what he or she could extract in the complete-information game.

The prediction common to this class of models is that a negotiating party that has incomplete information about its opponent will obtain a smaller share of the surplus in the negotiation than if that party were better informed. In the context of car promotions, we use the fact that that there is variation between the two types of promotions in how likely consumers are to be informed about the manufacturer-supplied surplus available in the promotion. The prediction from theory is that buyers should be able to obtain a larger share of a customer cash promotion (which they know about) than of a dealer cash promotion (which they are unlikely to know about).

Laboratory experiments support the prediction that an uninformed party obtains less of the surplus (see Kennan and Wilson (1993) and Roth (1995) for comprehensive surveys). For example, in an experiment simulating a real estate market, Valley, Blount White, Neal, and Bazerman (1992) show that transaction prices are lower, conveying greater surplus to the buyer, when the seller's reservation price is common knowledge than when it is the private information of the seller. A similar effect is often found in ultimatum games with one-sided incomplete information. Croson, Boles, and Murnighan (2003) summarize the results from these experiments: "These studies consistently show that proposers make (and responders accept) significantly lower offers when responders do not know the size of the pie and when this lack of information is common knowledge" (p. 145).

Empirical investigations of the effect of asymmetric information on bargaining outcomes in non-laboratory settings have not been concerned primarily with prices in product markets. Instead, the primary area of investigation has been union contract negotiations, and the outcome of interest has been strike activity and strike duration (Kennan and Wilson 1993, Ausubel, Cramton, and Deneckere 2002, Tracy 1986, Tracy 1987). Nevertheless, there are two empirical papers which, like this paper, are concerned with the effects of incomplete information on negotiated prices. Zettelmeyer, Scott Morton, and Silva-Risso (2003) analyze how negotiated prices are affected by whether a buyer used an Internet referral service (Autobytel.com) which makes available to consumers purchase-relevant information, including dealer invoice prices. Scott Morton, Zettelmeyer, and Silva-Risso (2004) relate self-reported measures of how informed a customer was when she purchased a car to the price she paid, finding that consumers who report knowing the invoice price of a dealer pay less.

Other empirical bargaining papers have focused on the demographic factors that affect bargaining outcomes, rather than on the effect of information asymmetries. For example, Harding, Rosenthal, and Sirmans (2003) examine the effect of differences between buyer and seller demographics on the negotiated price of a house. Ayres and Siegelman (1995), Goldberg (1996), and Scott Morton, Zettelmeyer, and Silva-Risso (2003) analyze the effect of buyer's race and gender on negotiated prices.

One final paper that is related to ours examines the effect of asymmetric information on prices, but in the context of auctions instead of bargaining. Hendricks and Porter (1988) use data on bids for drainage leases, which confer the right to extract oil and gas from a particular tract of land. A company that has already drilled on an adjacent tract has an information advantage over other bidders on a particular tract, and Hendricks and Porter find evidence that participants bid strategically in accordance with a model that takes this information asymmetry into account.

2.2 Manufacturer Promotions

Understanding price promotions has been an issue of longstanding interest in the marketing literature. There is a large literature concerned with game theoretic models of promotions. The aspects that have been modelled include the choice by manufacturers to offer wholesale promotions to retailers (Lal 1990, Lal, Little, and Villas-Boas 1996, Gerstner and Hess 1991) and the decisions of retailers to offer promotions to final customers either independently (Narasimhan

1988, Varian 1980, Raju, Srinivasan, and Lal 1990, Rao 1991) or in response to manufacturer promotions (Moorthy 2003, Kumar, Rajiv, and Jeuland 2001, Tyagi 1999).

The segment of the promotions literature that is most relevant to this paper, however, is the segment that concerns empirically measuring how much of a manufacturer promotion gets passed through to final customers in the form of lower retail prices. This is an issue that has been of interest not only to academic researchers, but also to the manufacturers themselves, who are interested in the effectiveness of their promotional activities in lowering prices to final customers. Blattberg and Neslin (1990) give an overview of the the issues surrounding how to measure the effect of promotions on retail prices (see chapter 11). Additional empirical investigations into the rate of pass-through of manufacturer promotions to retail customers include Besanko, Dubé, and Gupta (2004), Chevalier and Curhan (1976), and Walters (1989).

The setting in which the marketing literature has examined rates of pass-through of promotions has been primarily supermarkets. A supermarket setting is much less suited to examining the question at hand in this paper – namely the effect of information asymmetry on bargaining outcomes – for several reasons. First, the prices customers pay in supermarkets are posted rather than negotiated. Second, the promotions used in the supermarket channel do not provide a clean comparison between symmetric and asymmetric information. For example, packaged goods manufacturers offer coupons, about which many customers are informed, and "trade deals" which the final customer can't observe. However, a coupon has a price discriminatory aspect to it because not all customers who have coupons available find it worthwhile to redeem them (Nevo and Wolfram 2002). This does not apply to customer cash for automobiles because the promotional amounts are hundreds or thousands of dollars and redemption happens automatically during closing. Trade deals, on the other hand, have a different disadvantage compared to dealer cash for automobiles for investigating asymmetric information and prices. The issue with trade deals is that they are discounts that apply to all goods *purchased* by a retailer in a given period of time, not to goods *sold* by a retailer in that period of time, which means that retailers can use a trade deal to stock up on inventory without necessarily having any inducement to lower the retail price.[2]

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, that applied to the transaction.

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

[2] A third type of promotion used by manufacturers who sell to supermarkets is the "scan-back," which is analogous to dealer cash in that it pays the retailer the promotional amount for all goods that are sold during the promotion window. Scan-backs have historically been much less common than trade deals, but are being used more frequently.

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.[3] 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 variable

We will use transaction prices as dependent variables in the estimation. The price observed in the dataset is the pre-sales tax price that the customer pays for the vehicle, including factory installed accessories and options, and including any dealer-installed accessories contracted for at the time of sale that contribute to the resale value of the car.[4]

Conceptually, we would like our price variable to measure the customer's total wealth outlay for the car. In order to capture this, we make two modifications to the observed transaction price. First, we subtract off the customer cash rebate amount if the car is purchased under a customer cash rebate since the manufacturer pays that amount on the customer's behalf. Second, we subtract from the purchase price any profit the customer made on his or her trade-in (or add to the purchase price any loss made on the trade-in). The price the dealer pays

[3] Note that these are cash transactions from the perspective of the dealer, but they need not be cash transactions from the perspective of the customer. In particular, if a customer has obtained a loan from a bank, it is a cash transaction from the dealer's perspective.

[4] Dealer-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.

for the trade-in vehicle minus the estimated wholesale value of the vehicle (as booked by the dealer) is called the *TradeInOverAllowance*. Dealers are willing to trade off profits made on the new vehicle transaction and profits made on the trade-in transaction, which is why the *TradeInOverAllowance* can be either positive or negative. When a customer loses money on the trade-in transaction he or she is paying for the new vehicle in part in kind with the trade-in vehicle. By subtracting the *TradeInOverAllowance* we adjust the negotiated (cash) price to include this payment.

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 thus-defined cars after dropping cars with fewer than 200 sales in our sample. 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. Based on 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 also control for the income, education, occupation, and race of buyers by using census data that the data provider matches with the buyer's address from the transaction record. The data 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 (northern or southern California).

3.3 Summary statistics

Table 1 presents summary statistics for the data. 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 $1242 (median $1000). The average amount of dealer cash among transactions that involve dealer cash is $932 (median $700). The average transaction price of a new vehicle in our data is $25,490. The table also presents customer demographics.

4 Estimation approach

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

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

An alternative way to estimate the counterfactual price is to use the a car's own price when it is not on promotion. Conceptually, this is what our second approach, 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, Kane (2003) describes a state-funded college grant program that 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. In a similar example, Van der Klaauw (2002) describes 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.

In this paper, we are interested in estimating the effect of a manufacturer promotion (the treatment) on an outcome measure, price. 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.[5] Thus, in applying the regression discontinuity approach, we restrict our attention to one week on either side of when a promotion begins, ends, or changes in the amount of cash being offered. 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 who purchase just before a promotion starts and just after differed in some way that was related to the outcome 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.

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 promotional payment from the manufacturer only if a car is actually sold. This gives the dealer an incentive to lower

[5] We 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.

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

$$P_{ijt} = \lambda_c CustCash_{jt} + \lambda_d DealCash_{jt} + \boldsymbol{\beta_1 X_i} + \boldsymbol{\beta_2 X_{jt}} + \beta_3 DealerComp_{ij} + \mu_j + \tau_{JT} + \epsilon_{ijt}. \quad (1)$$

$CustCash_{jt}$ and $DealCash_{jt}$ are the amounts of customer cash and dealer cash available for vehicle j at date t. Since manufacturers typically make promotion decisions by nameplate - model - model year (e.g. 1999 Pontiac Grand Am), these variables are unique for a nameplate - model - model year triple. In cases where promotions also vary by region, our promotion variable will record the promotion as being available only in the region in which it is. $\boldsymbol{X_i}$ is a vector of the buyer's individual and neighborhood customer demographic characteristics including sex, race, income, education, employment type, and home ownership. $\boldsymbol{X_{jt}}$ 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). $DealerComp_{ij}$ 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 (e.g. SUVs, compact cars, etc.) that contains car j and T is the week that contains purchase date t. The data cover 121 weeks.

The primary variables of interest are λ_c and λ_d, which measure the extent to which rebates are passed through to customers. If either λ_c or λ_d is equal to 0, that implies that none of the surplus from the respective type of promotion is passed through to customers. In this case the retailer is the sole beneficiary of the promotion. If λ_c or λ_d is equal to -1, then the customer obtains the full amount of the respective rebate in the form of a lower price. One can interpret $100 \cdot |\lambda|$ as the percentage of the rebate the customer obtains.

Table 2 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 difference, statistically significant at the 1% confidence level, 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. In the transactions data, we can identify when a GM Card rebate was applied. This makes an interesting comparison to customer and dealer cash because this is manufacturer-supplied surplus for which the *customer* has an information advantage over the *dealer*. Hence, we expect that consumers can appropriate more of the GM card rebate than of either customer cash or dealer cash. The results follow this prediction: column 1 of Table 2 shows that customers obtain

106% (statistically indistinguishable from 100%) of the surplus from GM Card rebates.[6]

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

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

The regression equation in the regression discontinuity approach is similar to the one used in the difference-in-differences approach. As equation 2 shows, the measures of promotions, demographics ($\boldsymbol{X_i}$), time period controls ($\boldsymbol{X_{jt}}$), and car fixed effects (μ_j) are the same as in equation 1 above. The regression discontinuity specification does not rely on week - vehicle

[6]One alternative explanation for this finding is that GM Card holders are "bargain-hunting, penny-pinching" types and would therefore be likely to negotiate for good prices already. An interpretation of the GM Card that would give the opposite prediction is that GM Card holders hold the GM Card because they have strong preferences for GM vehicles, which would make them less likely to negotiate low prices.

segment fixed effects in order to identify the "treatment effect" of the promotion so τ_{JT} does not appear in equation 2; in its place, we include a week fixed effect, τ_T.

$$P_{ijt} = \lambda_c CustCash_{jt} + \lambda_d DealCash_{jt} + \boldsymbol{\beta_1 X_i} + \boldsymbol{\beta_2 X_{jt}} + \beta_3 DealerComp_{ij} + \mu_j + \tau_T + \epsilon_{ijt} \quad (2)$$

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

The identifying assumption in this approach is that the underlying willingness-to-pay of customers who buy just before and just after a change in a promotion is the same. This would be violated, for example, if there were deal-prone customers with a different willingness-to-pay or bargaining ability, who wait until the begin of a promotion to enter the market. We will test this identifying assumption in section 6.4.

Table 2 reports the estimated coefficients from the two regression discontinuity specifications. Overall, both estimates are slightly smaller than the estimates in the difference-in-differences approach. In column 2a of Table 2, 81% of customer cash is estimated to be passed through to customers, compared to 31% of dealer cash (in column 2b). Notice that the pass-through on GM card rebates cannot be estimated with a regression discontinuity approach since these rebates do not have a start and end date. We include this variable merely to control for the price effects of GM card rebates when estimating the pass-through of customer and dealer cash.

6 Identification issues

Having estimated the pass-through rates of customer and dealer cash using both a difference-in-differences and a regression discontinuity approach, we now test the validity of a series of assumptions that were maintained when identifying these effects. First, we test an assumption maintained in the difference-in-differences approach, namely that cars in the same segment that are not on promotion in a given week are a valid counterfactual for the prices that would have been obtained on the promoted car in the absence of a promotion. Second, we investigate the validity of an assumption maintained in the regression discontinuity approach, namely that the window around a promotion change is sufficiently small that the estimates measure the effect of the promotion but not the effect of changes in demand conditions. Third and fourth, we estimate the validity of another maintained assumption in the regression discontinuity approach, namely that transaction prices during the week just before the promotion starts are a valid counterfactual for transaction prices during the first promotion week. This assumption may be violated if dealers react strategically to an upcoming promotion. This assumption may also be violated if customers who purchase just before a promotion starts and just after it has started differ in some way that is related to pass-through rates. Fifth, we correct for potential bias in our estimations due to a violation of an assumption common to both the difference-in-differences and the regression discontinuity approaches, namely that the distribution of observable characteristics in treated and untreated groups share a common support. Finally, we test an assumption that underlies our interpretation that the differences in pass-through rates between promotion types are due to how well consumers are informed about the existence of each promotion. We will try to rule out that the differences in pass-through rates are due to differences in demand conditions under which manufacturers decide to offer one or the other type of promotion.

6.1 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 two daily time trends of price for each vehicle segment for each month of the sample. One trend is estimated for cars that will go on promotion within 30 days. The other trend is estimated for cars that are not about to go on promotion.

In generating these estimates, we first restrict the sample to transactions that occurred on dates on which the transacted car was on neither a customer cash nor a dealer cash promotion. Next, we calculate for each observation the earliest date after the observation date (t) that the car (j) is on either a customer or dealer cash promotion. We denote this as T_{jt}, the start date of the next promotion for car j at date t. Using T_{jt}, we define an indicator variable $I(T_{jt} \leq t + 30)$, which will equal 1 for all transactions that occur 30 days or less before the start of a promotion. We then define the monthly time trend variable as $\theta_{t,M} = n$ if date t is the n^{th} day of month M and $\theta_{t,M} = 0$ if date t does not fall in month M. Finally we define I_J, an indicator variable that equals one if car j is in segment J.

We then run the following regression:

$$P_{ijt} = \alpha_M + \beta_{M,J}\, I_J \cdot \theta_{t,M} + \gamma_{M,J}\, I_J \cdot I(T_{jt} \leq t + 30) \cdot \theta_{t,M} + \boldsymbol{\delta_1 X_i} + \boldsymbol{\delta_2 X_{jt}} + \mu_j + \eta_{ijt}, \quad (3)$$

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, $\boldsymbol{X_i}$ and $\boldsymbol{X_{jt}}$ are the demographic and time effects from equation 1, and μ_j are fixed car effects. The coefficients $\beta_{M,J}$ will measure the daily trend of prices over the days in month M for segment J. The $\gamma_{M,J}$ coefficients will measure any differences in the daily time trend in month M for prices of cars that are observed within 30 days of the date they will next be on promotion.[7]

Our test of equal trends in the pre-promotion period will be testing whether the $\gamma_{M,J}$ coefficients are equal to zero. In unreported results, 19 of the 111 estimated $\gamma_{M,J}$'s are statistically different from zero at the 5% confidence level, and an additional 5 at the 10% level. While many coefficients are statistically indistinguishable from zero individually, we cannot reject the hypothesis that the $\gamma_{M,J}$'s are jointly statistically different from zero.

In light of this, one might ask whether the estimated differences in pre-promotion trends that *are* statistically different from zero are large enough to explain the price effects that we estimated in Table 2, and which we attributed to pass-through. Of the statistically significant $\gamma_{M,J}$'s, the largest in magnitude that is negative is -11.86 dollars per day; the next largest is -3.47. The coefficients that we estimated in Table 2 imply that the prices of cars on customer

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

cash promotion, controlling for other covariates, are \$1056 less than the prices of cars that are not on promotion (85% estimated pass-through times \$1242 average customer cash promotion amount). In order for this price effect to be explained entirely by the prices for about-to-be-promoted cars drifting downward by \$11.86 more per day than non-promoted cars, it would have to be the case that the promotion lasted 178 days, about 6 months. For a difference of \$3.47 in the daily price trend, the promotion would have to last 609 days, or about 20 months, longer than most cars are even available. Since -\$11.86 and -\$3.47 are the largest of the estimated time trends, most of which are insignificant, and since each is estimated only for one month for one segment, we believe that any bias caused by different price trends over time is not large enough to account for our results.

The estimated price difference between cars promoted on dealer cash and other cars is about \$326 (35% times an average dealer cash promotion amount of \$932). Thus, if price were trending down by \$11.86 per day more for about-to-be-promoted cars than for non-promoted cars, that would account for an estimated price difference of \$326 after 55 days. A differential of \$3.47 could not explain the estimated effect unless the promotion lasted for at least 188 days. While the former is not atypical for the length of a dealer cash promotion, the \$11.86 differential is estimated for only one month-segment combination.

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

6.2 Robustness with regard to window size

One maintained assumption in the regression discontinuity approach is that the estimates of λ_c and λ_d measure the pass-through rate of the promotion but not the effect of changes in demand conditions within the chosen window around a promotion change. This assumption can be violated if the window is chosen too large. To see this, suppose that retail prices for a car are declining at 2% per month due to softening demand and that the car's manufacturer reacts to the softening demand by offering a promotion equal to 4% of the car's last average retail price. Say that half of the promotion amount (2%) is passed on to buyers. A regression discontinuity approach with one month windows before and after the beginning of the promotion would identify that prices have decreased by 4%, leading to the incorrect inference that 100% of the promotion was passed on to consumers. As the window in the regression discontinuity approach narrows, the probability of misattributing the effect of changing demand conditions to promotions decreases.

We find that the results reported so far are robust to changes in the size of the window around the promotion event. We reestimate the regressions in columns 2a and 2b of Table 2 using only two days before and after the change in a customer cash or dealer cash promotion. These results are reported in columns 1a and 1b of Table 3 which show that the estimated pass-through rates change little. We now estimate that 73% and 26% of customer cash and dealer cash, respectively, gets passed through to customers.

6.3 Strategic dealer behavior

Another maintained assumption in the regression discontinuity approach is that transaction prices during the week just before the promotion starts are a valid counterfactual for transaction prices during the first promotion week. This assumption may be violated if dealers react strategically to an upcoming promotion. Suppose that a dealer knows when a dealer cash promotion will begin. Then the dealer should want to sell fewer cars in the days after he or she has learned of the upcoming promotion, but before the promotion actually starts. This would increase prices in the pre-promotion period, leading us to overestimate the effect of the promotion on transaction prices and thereby overestimating the rate of pass-through.

To test whether our maintained assumption is violated, we make use of the fact that dealers typically find out about promotions only 2-3 days before the promotion start date.[8] Hence, we repeat our basic regression discontinuity specification, using the original sample of one week on either side of a promotion change but *excluding* the 3 days directly before and after a promotion change. If dealers react strategically to an upcoming promotion, this should not be reflected in this restricted sample. We find that the results are very similar to the basic regression discontinuity specification (in columns 2a and 2b of Table 2). We estimate that 82% and 36% of customer cash and dealer cash, respectively, gets passed through to customers, compared to 81% and 31% in the basic specification (see columns 2a and 2b in Table 3). This result makes it appear unlikely that pass-through rates are overestimated as a result of strategic dealer behavior ahead of promotions.

6.4 Pass-through by promotion length

As we have discussed in section 4, identification in the regression discontinuity approach would also be upset if the customers who purchase just before a promotion starts differed in some way that was related to negotiated prices from customers who purchase just after the promotion starts. In particular, this would be the case if there are "deal-prone" customers who are particularly effective negotiators, and who wait to purchase a car until a promotion is offered.

[8]From discussions with industry experts.

What this would mean for identification is that the set of customers whom we observe buying before the promotion would pay higher prices on average, with or without a promotion, than the set of customers whom we observe buying during a promotion would pay, with or without a promotion. Thus, our pass-through coefficient would wrongly attribute to promotions what is in fact a price difference due to unobservable buyer characteristics. While our detailed demographic data control for important differences among consumers, we are concerned that such data may not adequately capture "deal-proneness."

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

We build on the difference-in-differences specification from column 1 in Table 2. We split the customer cash and dealer cash variables by whether a promotion is an increase or a decrease from a previous level; we expect deal-prone consumers to be more likely to buy after a promotion change only if the promotion is increased.

In addition, we interact the resulting two customer cash and two dealer cash variables with dummies for whether the transaction occurred when a promotion change had been in place for 0-14 days, 15-30 days, one to two months, three to six months, or six or more months. Consistent with our conjecture that there may be some deal-prone consumers, we find that the pass-through rate for customer cash is greater in weeks 1 and 2 than in all subsequent weeks, but only for new or increased customer cash promotions. When a customer cash promotion is new or an increase, pass-through is estimated to be 96% in the first two weeks in contrast to 75-80% for all subsequent weeks (see Table 4). The difference between the first two weeks and later periods is statistically significant at the 1% level. When the current customer cash promotion is a decrease from a previous level, the rate of pass-through does not show any statistically significant change over time. In contrast to customer cash, the pass-through rate of dealer cash does not show evidence of attracting deal-prone customers. When a dealer cash promotion is an increase, pass-through rates stay statistically unchanged for the first 8 weeks and then rise, while when a dealer cash promotion is a decrease, there is no discernable pattern to pass-through rates over time.

These findings are consistent with some deal-prone consumers waiting to purchase cars until they become aware that a promotion is being offered. This suggests that the regression discontinuity approach somewhat overestimates the pass-through rate of customer cash promotions. However, even if we were to use the customer cash pass-through rate for weeks beyond the second week as estimated in Table 4, which would be 0.80, instead of 0.88 as estimated in column 1 of Table 2, our qualitative finding that customer cash gets passed through to consumers at twice the rate of dealer cash would not change.

6.5 Common support on observables

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 problem that these estimated counterfactuals may not 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 with 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 the distributions of observable characteristics in either car or customer characteristics between promotion and non-promotion groups, even within a region of common support. This could lead us attribute price differences to treatment (promotion) when they are actually due to differences in the average 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 this subsection we use a propensity score approach to minimize the bias arising from non-overlapping support (see, for example, Galiani, Gertler, and Schargrodsky (2004)), the first source of bias described above. To control for bias caused by differences in car and customer observables between promotion and non-promotion subsamples, the second source describe above, we have used in all specifications 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 is the issue we investigated in section 6.4 where we looked for evidence that customers who are inherently good price negotiators are also more likely to buy under promotions.

Difference-in-Differences with common support

We use a propensity score approach to ensure common support in observable characteristics between treatment and control observations. This approach estimates for each observation in the sample the probability (a "propensity score") that, conditional on all observables, the observation is in the treatment as opposed to the control group. The resulting propensity scores are used to ensure that there are no observations in the treatment group whose propensity scores lie outside the range of propensity scores of observations in the control group, and vice versa. The advantage of the propensity score approach is that there need not be a common support between treatment and control group observations on *each* observable characteristic. Instead, to remove the bias from a non-overlapping support in observable characteristics, it suffices to ensure a common support in the distribution of propensity scores between the treatment and control groups.

One complication in our setting is that we have multiple treatments, namely customer cash and dealer cash promotions. Recently, Imbens (2000) and Lechner (2001) have extended the standard model underlying the propensity score approach from one to multiple treatments. The key point is to be able to extend to multiple treatments the property of the propensity score that it can be used as a single dimensional measure by which to find matching observations across treatment and control groups. Both papers show that the propensity score approach can be adapted to preserve this property, what is referred to as the "balancing score property of the propensity score," in a multiple treatment environment. We follow the approach of these papers to restrict the sample to observations with common support in observable characteristics between the multiple promotion states.

We proceed as follows: Using a multinomial logit model, we first estimate the probability that the sold car was in each of the four promotion states $s_i \in S = \{$"no promotion," "only customer cash promotion," "only dealer cash promotion," and "customer and dealer cash promotion"$\}$. We estimate these probabilities as a function of the customer characteristics, the detailed car dummies, the region, and the competition variable used in all previous difference-in-differences specifications. For each pair of states $s_i, s_j \in S, i \neq j$, we then calculate $Prob(s_i|s_i, s_j)$, the probability that a purchase transaction occurs in state s_i conditional on the transaction having occurred either in state s_i or s_j. This way we calculate a propensity score for the promotion state s_i (as the "treatment") relative to each other promotion state s_j (as its "control"). For each such pair of states we identify observations on the common support by (1) excluding all observations in "control" state s_j whose propensity scores are lower than the minimum of the distribution of propensity scores for the observations in "treatment" state s_i and by (2) excluding all "treatment" state observations whose propensity scores are higher than the maximum of the distribution of "control" state propensity scores. Thus, we are left with a set of observations which lie in the common support of *all* promotion states. This procedure ensures that, controlling for observables, every observation in the sample, irrespective of which promotion state it represents, can serve as a control for any other observation. This reduces the sample from 133,424 to 41,533 observations. The decline in the number of observations is mostly due to the fact that our procedure excludes cars that are never promoted, that are always promoted with one but not the other type of cash, or that are always promoted with both types of cash.

In column 3 of Table 3 we re-estimate the difference-in-differences specification reported in column 1 of Table 2, but using as the estimation sample the common support sample described in the previous paragraph. In comparing the columns we see that adjusting the sample so that observations share a common support brings the difference-in-differences estimates of the pass-through rates close to the estimates of the regression discontinuity specification in columns 2a and 2b of Table 2. The estimates indicate that 84% of customer cash is passed through to customers while only 31% of dealer cash reaches buyers, a difference statistically significant at the 1% level.[9]

Regression discontinuity with common support

In the regression discontinuity approach we want to find the region of common support in observables between those transactions that took place in the week on one side of a promotion change and those transactions that took place in the week on the other side of the promotion change. Since we have to estimate the pass-through rates of customer and dealer cash promotions in separate regressions, we calculate the region of common support for each regression separately. For the customer cash regression, we start with the sample of transactions that occurred within a week before or a week after the start, end, or change in size of a customer cash promotion. We then estimate propensity scores from a logit model of the probability that, of the two weeks surrounding a customer cash promotion change, the transaction occurred in the week with the higher promotion amount. If the promotion change is the start of a promotion or a promotion amount increase, the later week will have the higher promotion

[9]This finding is robust to more restrictive ways of defining a common support. In particular, instead of using the minimum and maximum of the distributions of propensity scores to determine the common support, we have also used the 0.5 and 99.5 percentiles, the 1st and 99th percentile, and the 2nd and 98th percentiles as cut-offs. The number of observations in the common support drop to 14,862, 8422, and 2923, respectively. The estimated pass-through rates are 88% (customer cash) and 36% (dealer cash), 86% and 32%, and 67% and 27%, respectively.

amount. If the promotion change is the end of a promotion or a promotion amount decrease, the earlier week will be the week with the higher promotion amount. The explanatory variables are the customer characteristics, the detailed car dummies, the region, and the competition variable. We identify observations on the common support by (1) excluding all observations from the control state (the "zero or lower promotion amount weeks") whose propensity scores are lower than the minimum of the distribution of propensity scores for the observations in treatment state (the "higher promotion amount weeks") and by (2) excluding all treatment state observations whose propensity scores are higher than the maximum of the distribution of control state propensity scores. This reduces the number of observations in the customer cash regression from 6296 to 6185. The procedure for finding the region of common support for the dealer cash regression is analogous and reduces the number of observations in the dealer cash regression from 7046 to 6974.

Notice that the propensity score approach eliminates far fewer observations in the regression discontinuity approach than in the difference-in-differences approach. This is because the regression discontinuity approach already restricts the sample to cars which change their promotion status, leaving out cars that are never promoted or always promoted at the same level. This also indicates that there is close to complete overlap in the observable characteristics of buyers between transactions which occurred in the week before and after a promotion change.

In comparing columns 2a and 2b of Table 2 and columns 4a and 4b of Table 3 we see that adjusting the sample so that observations share a common support changes the estimates very little. The estimate of customer cash pass-through remains 81% while the estimate of dealer cash pass-through increases from 31% to 32%.[10]

[10]As in the difference-in-differences specification, the regression discontinuity results are robust to more restrictive ways of defining a common support. In particular, instead of using the minimum and maximum of the distributions of propensity scores to determine the common support, we have also used the 0.5 and 99.5 percentiles, the 1st and 99th percentile, and the 2nd and 98th percentiles as cut-offs. The number of observations in the common support of the customer cash sample drop to 5959, 5814, and 5553, respectively.

6.6 Selection on market conditions

We have attributed the difference in pass-through rates between the two promotion types to differences in how well consumers are informed about the existence of these promotions. This attribution needs to be treated with caution. Our estimates indicate the average pass-through of a customer cash promotion in periods in which manufacturers chose to have customer cash promotions and the average pass-through of a dealer cash promotion in periods in which manufacturers chose to have dealer cash promotions. So far our results do not necessarily imply that the pass-through of a promotion would increase if a manufacturer were to switch from a dealer cash to a customer cash promotion of equal amount. In particular, if each type of promotion is well suited to a particular configuration of market conditions, and manufacturers want to maximize promotion pass-through, it may be that manufacturers are optimally matching promotions to particular conditions. This is a concern even if we are correctly estimating customer and dealer cash pass-through rates in the states in which they are used. This concern pertains to the interpretation of the difference, not to the estimates of the difference itself.

In this section we will try to rule out that the differences in pass-through rates are due to differences in market conditions under which manufacturers decide to offer one or the other type of promotion. We proceed as follows: We use measures of market conditions to explicitly model the probability that a car on promotion is promoted either with customer cash or with dealer cash. This estimation yields the predicted probability for each transaction that the car is sold on a customer cash as opposed to a dealer cash promotion. We use this probability as a propensity score to exclude observations from the estimation for which there are no comparable market conditions across the two promotion states. This leaves us with a sample

The estimated pass-through rates for customer cash are 81%, 81%, and 83% respectively. The number of observations in the common support of the dealer cash sample drop to 6707, 6594, and 6321, respectively. The estimated pass-through rates for dealer cash are 32%, 33%, and 31%, respectively.

which contains only cars that were sold on promotion, and which share a common support in observable market conditions across customer and dealer cash promotions. Finally, we use this sample to identify pass-through rates while explicitly controlling for the market conditions on which manufacturers base their decision to offer one or the other type of promotion.

Industry sources tell us that the primary variables that manufacturers use when deciding whether to initiate a promotion are vehicle profitability at the dealer level, inventory level, total vehicle sales, and market share within the vehicle's subsegment. Depending on the specific auto manufacturer, these variables are monitored on a weekly or monthly basis.[11] In rare cases, promotions are initiated on the basis of information about market conditions that is as recent as two weeks. More commonly, the lag between the information used and the promotion decision is one to three months.

We want to use these decision variables to model the promotion choice process of manufacturers and as controls in the estimation of promotion pass-through. We do not expect, however, that the above variables will describe perfectly the choice of promotion. In part this is because we know that there are differences among manufacturers in the promotion decision process. In part, this is because manufacturers have additional considerations when choosing promotions besides matching the most effective promotion to the demand condition. Anecdotally, we know that dealers communicate to manufacturers through field reps that they prefer dealer cash to customer cash, arguing that dealer cash gives them more flexibility.[12] While manufacturers and dealers in the automotive sector are separate entities, each party is clearly dependent on the other in many ways, having made multiple relationship specific investments. Hence, while

[11] One of the authors of this paper consulted with several auto manufacturers in improving their promotion decision system.

[12] Evidence for this can also be found in the 1995 complaint of the DOJ against the National Automotive Dealer Association: "On numerous occasions between 1989 and 1992, the NADA urged manufacturers to give franchised dealers, rather than consumers, all of the discounts and incentives offered by manufacturers to induce the purchase of a new car." (Paragraph 14, page 4.)

we expect the promotion decision to be explained in part by observable demand conditions, there will also be unobservable elements driving the decision, including the importance of and current state of dealer-manufacturer relations.

To model the promotion choice process we consider only promotions which followed a period of no promotions. This is so that the lagged measures of market conditions on which manufacturers base their promotion decisions (sales, days of inventory, dealer profitability, etc.) are not themselves affected by the existence of an earlier promotion. To be able to use measures of market conditions which reflect the lag between information and promotion decision commonly found at auto manufacturers, we only consider the first 30 days of any promotion. This way we can use one month lagged measures of market conditions without worrying that these measures are affected by the promotion they triggered. Consequently, our sample to model the promotion choice process consists of transactions for which (1) the car was sold either with a dealer cash promotion or a customer cash promotion, (2) the transaction occurred within the first 30 days of the promotion, and (3) the car was not on any type of promotion preceding the start of the current promotion.

To model the promotion choice process of manufacturers we estimate a logit model where the dependent variable is one if the car is offered with a customer cash promotion and zero if it is offered with a dealer cash promotion. Using our knowledge of the decision making process at manufacturers, our key explanatory variables are monthly changes in dealer profitability, market share (within subsegment), sales, and inventory level; all measures are specific to the transacted model. We include changes in these variables from three months to two months and from two months to one month before a promotions starts, and also control for the levels of these variables one, two, and three months prior to the start of the promotion. We also include customer characteristics, detailed car dummies, region, and competition effects to ensure common support in observable consumer and car characteristics. This ensures that the resulting propensity score shares a common support between customer cash and dealer cash observations with regards to observable car characteristics and consumer characteristics, as well as pre-promotion market conditions.

In the logit model we find that a decrease in dealer profitability is associated with a higher probability that a car will be offered with dealer cash as opposed to customer cash (see Table 5). Since our estimates so far have shown that dealers obtain more of the surplus from dealer cash than they would from an equivalently sized customer cash promotion, it is perhaps not surprising that manufacturers apparently offer dealer cash in response to low dealer margins. This behavior is consistent with an argument by Klein and Murphy (1988) and Klein (1995) in the context of franchising. Klein (1995) argues that an upstream firm must leave rents for the downstream firm if it wants to be able to influence the behavior of the downstream firm by threatening the loss of future rents. Our other three explanatory variables in the promotion prediction logit have to do with sales in some way or other: market share, sales, and inventory levels. The logit coefficients indicate that a decrease in market share is associated with a with a higher probability that a car will be offered with customer cash as opposed to dealer cash. A decrease in sales or an increase in inventory levels, however, is associated with a with a higher probability that a car will be offered with dealer cash. We do not have an explanation based in what we know of the process by which manufacturers plan promotions, why a decrease in market share would lead to customer cash while a decrease in sales or increase in inventory would lead to dealer cash. From a statistical point of view, however, the explanatory variables do a reasonably good job at predicting which type of promotion will be offered: the pseudo R^2 of the logit estimation is 0.58. An estimation of the logit without any measures of market conditions yields a pseudo R^2 of 0.34, indicating that the market condition measures have some power in explaining whether a promotion is offered in form of customer or dealer cash.

The logit yields a propensity score which we use to identify observations on the common support. We identify observations on the common support by (1) excluding all observations in the control state ("car was offered with dealer cash") whose propensity scores are lower than the minimum of the distribution of propensity scores for the observations in the treatment state ("car was offered with customer cash") and by (2) excluding all treatment state observations whose propensity scores are higher than the maximum of the distribution of control state propensity scores. This reduces the number of observations to 3939.

We repeat our basic difference-in-differences pass-through specification with two changes. First, we include week fixed effects instead of week-segment fixed effects since we do not have enough degrees of freedom to identify week-segment fixed effects. Second, we control for the lagged measures of market conditions used to predict whether a car will be offered with dealer cash or customer cash. The results suggest that, controlling for the market conditions under which manufacturers choose one or the other type of promotion, the pass-through rates for customer cash remain at least twice as large as the pass-through rates for dealer cash and statistically significantly different at the 1% level (see Table 6). While the point estimates indicate that 102% of customer cash and 48% of dealer cash is passed through, the confidence intervals are large enough to accommodate, for example, the estimates of 88% and 39% found in the original difference-in-differences specification in column 1 of Table 2.[13]

This finding makes it less likely that differences in market conditions under which manufacturers decide to offer one or the other type of promotion are alone responsible for the differences in pass-through rates between customer and dealer cash. While our somewhat coarse measures of market conditions cannot perfectly capture the promotion generation pro-

[13]As in prior specifications, the qualitative comparison of customer cash and dealer cash pass-through is robust to more restrictive ways of defining a common support. In particular, instead of using the minimum and maximum of the distributions of propensity scores to determine the common support, we have also used the 0.5 and 99.5 percentiles, the 1st and 99th percentile, and the 2nd and 98th percentiles as cut-offs. The number of observations in the common support drop to 3684, 2978, and 1993, respectively. The estimated pass-through rates are 105% (customer cash) and 42% (dealer cash), 94% and 35%, and 107% and 50%, respectively.

cess, we do think they are sufficiently informative measures that if a substantial portion of the difference between dealer cash and customer cash pass-through rates were due to differences in market conditions, adjusting for common support and controlling for market conditions would decrease the difference between the estimated pass-through rates. The fact that this has not happened gives us more confidence that the difference between dealer cash and customer cash pass-through estimates is attributable to differences in how well consumers are informed about the existence of these promotions.

6.7 A non-bargaining explanation

We would like to briefly discuss one simple alternative explanation outside of a bargaining framework for the estimated pass-through rate of dealer cash. Suppose that car dealers have local monopoly power and that transaction prices correspond (to some approximation) to the posted prices set by a monopolistic firm. If this were the case, then the static "marginal revenue equals marginal cost" price optimization would dictate that one half of any cut in wholesale price should be passed on to consumers. If we think of dealer cash as a wholesale price cut, this would predict a 50% pass-through rate for dealer cash in our estimates. A brief review of our estimates so far shows that we cannot reject the hypothesis that our dealer cash estimates are different from 0.5. The problem with this explanation is that it cannot easily explain our customer cash pass-through rate. For example, if we were to interpret also customer cash as a wholesale price cut (it is not obvious one should), we should expect a similar pass-through rate as for dealer cash. This is clearly not what we find; we consistently reject that the customer cash estimate is 0.5. We also reject the hypothesis that, in contrast to wholesale price cut in form of dealer cash, customer cash is simply ignored by dealers in making pricing decisions; in most estimations we reject the hypothesis that consumers receive 100% of the customer cash.

7 Extensions

We now consider a number of extensions to our basic specifications. We first analyze how competition affects pass-through. Next, we investigate how pass-through varies by customer demographics. Finally, we consider possible demand effects from advertising which may accompany customer cash promotions. For brevity, in all cases we use only the difference-in-differences specification.

The effect of competition on pass-through

In column 1 of Table 7 we interact the promotions variables with our measure of dealer competition. We anticipate that a customer who has several dealerships nearby will be able to negotiate lower prices because he or she can easily negotiate – or threaten to negotiate – with multiple dealers. Although our previous results have not shown a statistically significant effect of dealer competition on prices themselves, here we test whether dealer competition leads to higher rates of pass-through of customer cash. We might observe this if the presence of a customer cash rebate increases the amount of surplus a customer believes is up for grabs in the negotiation, and this encourages customers to actually undertake the costly process of negotiating with an additional dealer. Under this intuition, we would not expect competition to increase the rate of pass-through of dealer cash because customers, being uninformed about the availability of dealer cash, would not have any increased incentive to play dealers off one another.

Consistent with our prediction, in column 1 of Table 7, the coefficient estimates imply that while a dealership without competing dealerships of the same nameplate within 10 miles will pass through 84% of customer cash, this rate of pass-through will increase by one percentage point for every additional dealer of the same nameplate in that 10 mile radius. Also, consistent with our expectations, competition has no statistically significant effect on the rate of dealer cash pass-through.

Competing promotions

A nearby dealership of the same nameplate is one kind of outside options that might increase a customer's negotiating leverage; an alternative vehicle in the same subsegment is another kind of outside option that could allow a customer to negotiate more effectively. In this extension we consider the effect on one vehicle's price of promotions currently available on other vehicles in the same segment. We expect that promotions available on alternative vehicles will lower the expected price of those vehicles, lowering the price to which a dealer must agree in order to sell a car for which promoted cars might substitute.

In column 1 of Table 7 we add to the regression variables measuring the number of customer and dealer cash promotions currently available on competing vehicles in the same vehicle segment. In the estimated results, neither the number of competing customer cash offers nor the number of competing dealer cash offers has a statistically significant effect on prices.

In column 2 of Table 7 we use alternative measures to try to capture the same effect, namely the *average level* of customer and dealer cash currently available on competing vehicles in the same vehicle segment. We construct these measures by averaging the customer cash (or dealer cash) available on a given day on competing vehicles in the same segment, weighting the average by each competing car's average market share during the year. In the results reported in column 2, customer cash on competing vehicles has no statistically significant effect on prices, but dealer cash available on competing vehicles lowers transactions prices: the estimated coefficient indicates that 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) would lower the transaction price by about $14. While we expected that both types of competing promotions could 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

Previous empirical studies of automotive retailing have found significant differences in prices paid by customers with different demographic characteristics, especially race and gender (Ayres and Siegelman 1995, Scott Morton, Zettelmeyer, and Silva-Risso 2003). This extension examines whether there are analogous differences in pass-through rates.

In column 3, 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[14]). 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.

Using the fraction of transactions that occur under each type of promotion and the average promotion amount of each type of promotion, we can calculate how much of the *direct* effect of these demographic factors on transaction prices is due to female, black, and Hispanic buyers

[14]See Scott Morton, Zettelmeyer, and Silva-Risso (2003) for a detailed discussion of this interpretation.

getting lower rates of pass-through than white male or Asian buyers. The average price effect of pass-through rate differences is roughly equal to differences between columns 1 and 3 in the estimated coefficients on *Female*, *%Black*, and *%Hispanic*. Approximately 17% of the higher price paid by women, and approximately 20% of the higher price paid by blacks appears to be attributable to lower rates of pass-through on promotions. Although Hispanics are predicted to receive lower rates of pass-through than whites, the direct effect of being Hispanic on prices estimated in column 3 is negative, making the combined average effect of being Hispanic on prices statistically zero.

Further investigation into why women and racial minorities obtain lower rates of pass-through of promotional surplus is beyond the scope of this paper. However, these results are consistent with the findings of Scott Morton, Zettelmeyer, and Silva-Risso (2003) on the effect of such demographic characteristics on overall prices.

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 the demand for a car. In column 4 of Table 7, we investigate the effect of advertising on our estimated customer cash pass-through rates.

In order to identify the effect of advertising on pass-through rates, 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 4 we introduce a variable called *CustomerCash*CustomerCash 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.[15] The idea behind the interaction term is that the variable *CustomerCash* will estimate the pass-through rate of customer cash when it is accompanied by advertising, while *CustomerCash*CustomerCash decrease* will measure how much the pass-through rate differs when the promotion is unadvertised. An analogous variable is defined for dealer cash.

The estimates in column 4 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 raises customers' reservation prices in their negotiations with dealers. Since there is less advertising when promotion amounts decrease, customers who buy during these periods will have lower reservation prices and will obtain more of the promotion amount. 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 prices, 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 4, 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

[15]For example, if there were no promotion in January, then a $750 promotion for the month of February, followed by a $500 promotion for the month of March, then *CustomerCash * CustomerCash decrease* would be zero for the months of January and February and $500 for the month of March.

smaller than either estimate of the customer cash pass-through.

8 Conclusion

This paper has analyzed how the information asymmetry between dealers and consumers about manufacturer-supplied surplus affects the division of this surplus in customer-dealer negotiations. Overall, our results tell a remarkably consistent story across approaches and across specifications. Customers obtain 70-90% of the surplus supplied by manufacturers in customer cash promotions, but only 30-40% of the surplus in dealer cash promotions. Customers also obtain all the surplus available through the GM Card.

We have tested the validity of a series of assumptions that were maintained when identifying pass-through rates using both a difference-in-differences and a regression discontinuity approach. First, we have analyzed whether non-promoted cars in the same segment are a valid counterfactual for promoted cars — an assumption maintained in the difference-in-differences approach. We have found that non-promoted cars are not a perfect control for promoted cars, however, that the potential bias is not large enough to change our conclusions. Second, we have analyzed whether the window around a promotion change is sufficiently small that the estimates measure the effect of the promotion but not the effect of changes in demand conditions — an assumption maintained in the regression discontinuity approach. We confirm that this is the case: the results obtained with a one-week window are substantially the same as those obtained with a very small (2 day) window. Third, we have investigated whether there is any evidence that dealers behave strategically by encouraging customers who might buy just before a promotion starts to come back and buy during the promotion. If this were happening, it could increase the observed pre-promotion prices, increasing our estimated pass-through rates in the regression discontinuity approach. We do not find evidence that strategic dealer behavior influences our estimates. Fourth, we have investigated whether deal-prone consumers wait to purchase a car until a customer cash promotion is on — the regression discontinuity approach assumes that they do not. We have found evidence consistent with promotions attracting deal-prone consumers, however, the potential bias in the pass-through rate is small and does not change our substantive findings. Fifth, we have analyzed whether promotion and non-promotion observations are drawn from a common support. We have found that not all observations are, but that our results do not change once the sample is restricted to observation on a common support. Finally, we have analyzed whether the differences in promotion pass-through rates are attributable to promotions being matched to market condition. Using industry information on how promotions are chosen, we find no such evidence.

We conclude that the difference between dealer cash and customer cash pass-through estimates are most likely attributable to differences in how well consumers are informed about the existence of these promotion. This is consistent with the theoretical prediction that when customers are at an information disadvantage, they are disadvantaged in negotiations. In the setting of car manufacturer promotions this information disadvantage is substantial: for a promotion of average size, consumers receive $500 less of the surplus if they do not know that the promotion is on the table. To our knowledge, this is one of very few measurements of how important information asymmetries in product markets are in practice.

References

Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002): "Bargaining with Incomplete Information," in *Handbook of Game Theory*, ed. by R. J. Aumann, and S. Hart, vol. 3. Elsevier, Amsterdam.

Ausubel, L. M., and R. J. Deneckere (1998): "Bargaining and Forward Induction," mimeo, University of Maryland University of Wisconsin-Madison.

Ayres, I., and P. Siegelman (1995): "Race and Gender Discrimination in Bargaining for a New Car," *American Economic Review*, 85(3), 304–321.

Besanko, D., J.-P. Dubé, and S. Gupta (2004): "Own-brand and Cross-brand Retail Pass-through," *Marketing Science*.

Blattberg, R. C., and S. A. Neslin (1990): *Sales Promotion: Concepts, Methods, and Strategies*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey.

Chevalier, M., and R. C. Curhan (1976): "Retail Promotions as a Function of Trade Promotions: A Descriptive Analysis," *Sloan Management Review*, 18(3), 19–32.

Croson, R., T. Boles, and J. K. Murnighan (2003): "Cheap talk in bargaining experiments: lying and threats in ultimatum games," *Journal of Economic Behavior and Organization*, 51, 143–159.

Fudenberg, D., D. K. Levine, and J. Tirole (1985): "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in *Game Theoretic Models of Bargaining*, ed. by A. Roth. Cambridge University Press, Cambridge, England.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. The MIT Press, Cambridge, Massachusetts.

Galiani, S., P. Gertler, and E. Schargrodsky (2004): "Water for Life: The Impact of the Privatization of Water Services on Child Mortality," *Journal of Political Economy*.

Gerstner, E., and J. D. Hess (1991): "A Theory of Channel Price Promotions," *American Economic Review*, 81(4), 872–886.

Goldberg, P. K. (1996): "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey," *Journal-of-Political-Economy*, 104(3), 622–654.

Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39, 155–190.

Hahn, J., P. Todd, and W. Van der Klaauw (2001): "Identification and Estimation of Treatment Effects with a Regression Discontinuity Design," *Econometrica*, 69(1), 201–209.

Harding, J., S. Rosenthal, and C. Sirmans (2003): "Estimating Bargaining Power In The Market For Existing Homes," *Review of Economics and Statistics*, 85(1), 178–188.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997): "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *Review of Economic Studies*, 64(4), 605–654.

Hendricks, K., and R. Porter (1988): "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review*, 78(5), 865–883.

Imbens, G. (2000): "The Role of Propensity Score Estimating Dose-Response Functions," *Biometrica*, 87, 706–710.

Imbens, G. W., and J. D. Angrist (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62(2), 467–475.

Kane, T. J. (2003): "A Quasi-Experimental Estimate of the Impact of Financial Aid on College-Going," Working Paper 9703, NBER.

Kennan, J., and R. Wilson (1993): "Bargaining with Private Information," *Journal of Economic Literature*, 31(1), 45–104.

Klein, B. (1995): "The Economics of Franchise Contracts," *Journal of Corporate Finance: Contracting, Governance and Organization*, 2, 9–37.

Klein, B., and K. Murphy (1988): "Vertical Restraints as Contractual Enforcement Mechanisms," *Journal of Law and Economics*, 31, 265–297.

Kumar, N., S. Rajiv, and A. Jeuland (2001): "Effectiveness of Trade Promotions: Analyzing the Determinants of Retail Pass Through," *Marketing Science*, 20(4), 382–404.

Lal, R. (1990): "Manufacturer Trade Deals and Retail Price Promotions," *Journal of Marketing Research*, 27, 428–444.

Lal, R., J. D. C. Little, and J. M. Villas-Boas (1996): "A Theory of Forward Buying, Merchandising, and Trade Deals," *Marketing Science*, 15(1), 21–37.

Lechner, M. (2001): "Identification and Estimation of Causal Effects of Multiple Treatments under the Conditional Independence Assumption," in *Econometric Evaluation of Active Labor Market Policies in Europe*, ed. by M. Lechner, and F. Pfeiffer, pp. 43–58. Physica/Springer, Heidelberg.

Moorthy, S. (2003): "A General Theory of Pass-through in Channels with Category Management," mimeo, University of Toronto.

Narasimhan, C. (1988): "Competitive Promotional Strategies," *Journal of Business*, 61(4), 427–449.

Nevo, A., and C. Wolfram (2002): "Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals," *RAND Journal of Economics*, 33(2), 319–339.

Raju, J. R., V. Srinivasan, and R. Lal (1990): "The Effects of Brand Loyalty on Competitive Price Promotional Strategies," *Management Science*, 36(3), 276–304.

Rao, R. C. (1991): "Pricing and Promotion in Asymmetric Duopolies," *Marketing Science*, 10(2), 131–144.

Roth, A. E. (1995): "Bargaining Experiments," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth. Princeton University Press, Princeton, N.J.

Rubinstein, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97–109.

Scott Morton, F., F. Zettelmeyer, and J. 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), 65–92.

——— (2004): "A Test of Bargaining Theory in the Auto Retailing Industry," Discussion paper, UC Berkeley.

Tracy, J. S. (1986): "An Investigation into the Determinants of U.S. Strike Activity," *American Economic Review*, 76, 423–436.

——— (1987): "An Empirical Test of an Asymmetric Information Model of Strikes," *Journal of Labor Economics*, 5, 149–173.

Tyagi, R. (1999): "A Characterization of Retailer Response to Manufacturer Trade Deals," *Journal of Marketing Research*, 65(2), 253–272.

Valley, K., S. Blount White, M. Neal, and M. Bazerman (1992): "Agents as Information Brokers: The Effects of Information Disclosure on Negotiated Outcomes," *Organizational Behavior and Human Decision Processes*, 51, 220–236.

Van der Klaauw, W. (2002): "Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression Discontinuity Approach," *International Economic Review*, 43(4), 1249–1287.

Varian, H. R. (1980): "A Model of Sales," *American Economic Review*, 70(4), 651–659.

Walters, R. G. (1989): "An Empirical Investigation into Retailer Response to Manufacturer Trade Promotions," *Journal of Retailing*, 48, 253–272.

Zettelmeyer, F., F. Scott Morton, and J. Silva-Risso (2003): "Cowboys or Cowards: Why are Internet car prices lower?," Discussion paper, Berkeley, CA.

Table 1: Summary statistics

Variable	N	Mean	Median	St. Dev.	Min	Max
Price	133424	25490	23487	10382	5988	109755
Customer Cash	34296	1242	1000	669	10	7805
GM Card	1204	1934	1785	1215	2	7043
Dealer Cash	24620	932	700	819	200	5000
Sales Manager Incen.	2319	141	50	181	20	500
Sales Rep Incen.	3601	147	75	154	25	500
# CustCash Prom. in Seg.	133424	6	6	4.7	0	17
# DealCash Prom. in Seg.	133424	2.9	2	2.5	0	13
Avg. CustCash Prom. in Seg.	133424	231	191	209	0	1254
Avg. DealCash Prom. in Seg.	133424	112	53	155	0	2305
# Competing Dealers	133424	3.7	3	2.8	0	24
Female	133424	0.28	0	0.45	0	1
%Asian	133424	0.11	0.064	0.12	0	0.97
%Black	133424	0.033	0.011	0.083	0	1
%BlueCollar	133424	0.23	0.2	0.15	0	1
%CollegeGrad	133424	0.36	0.35	0.18	0	1
%Hispanic	133424	0.13	0.096	0.1	0	0.55
%LessHighSchool	133424	0.1	0.074	0.1	0	1
%HouseOwnership	133424	0.69	0.76	0.24	0.0043	1
%Executives	133424	0.19	0.19	0.083	0	1
%Professional	133424	0.19	0.18	0.091	0	1
%Technicians	133424	0.031	0.028	0.022	0	1
Income	133424	6.5	6.2	2.8	1.1	15
Income^2	133424	49	38	43	1.1	225
MediaHHSize	133424	2.9	2.8	0.57	1.5	6
MedianHouseVal.	133424	2.6	2.4	1.2	0.075	5
Weekend	133424	0.27	0	0.44	0	1
EndOfMonth	133424	0.22	0	0.41	0	1
EndOfYear	133424	0.026	0	0.16	0	1
SouthernCal	133424	0.48	0	0.5	0	1

† For Customer Cash, GM Card, Dealer Cash, Sales Manager Incentives, and Sales Rep Incentives, "N" reports the number of non-zero observations. Hence, the summary statistics reflect observations with non-zero values.

Table 2: Price effects, basic results†

	(1)	(2a)	(2b)
	Diff-in-Diff	Reg. Disc.	
Customer Cash	-0.88 (0.03)**	-0.81 (0.07)**	-0.78 (0.12)**
Dealer Cash	-0.39 (0.07)**	-0.38 (0.14)**	-0.31 (0.07)**
GM Card	-1.06 (0.03)**	-1.13 (0.10)**	-1.13 (0.09)**
Competition	-7.66 (5.60)	-15.03 (9.59)	-18.63 (9.63)+
Female	144.17 (12.01)**	139.74 (45.20)**	203.25 (57.59)**
%Asian	-261.75 (61.95)**	-188.40 (239.89)	-115.40 (173.48)
%Black	474.48 (92.72)**	605.65 (381.97)	470.96 (251.97)+
%BlueCollar	208.11 (89.62)*	766.91 (366.29)*	398.17 (355.60)
%College	-283.87 (86.40)**	-103.70 (300.31)	-132.31 (271.87)
%Hispanic	-36.43 (86.37)	-936.23 (352.61)**	-68.13 (343.59)
%LessHighSchool	-166.54 (111.15)	-276.23 (422.61)	-363.69 (451.70)
%HouseOwnership	14.57 (35.88)	-59.68 (161.61)	52.05 (142.31)
%Executive	253.73 (113.07)*	542.61 (501.99)	-404.04 (517.10)
%Professional	217.33 (119.99)+	-268.76 (406.52)	231.95 (426.55)
%Technicians	147.06 (236.15)	-101.90 (1166.86)	-500.29 (1190.77)
MedianHHIncome	-21.17 (11.07)+	-30.22 (55.34)	-49.37 (46.16)
(MedianHHInc.)^2	3.01 (0.60)**	2.99 (3.19)	3.79 (2.32)
MedianHHSize	-24.95 (13.59)+	22.45 (59.12)	54.37 (47.24)
MedianHouseVal.	-8.10 (10.35)	-0.80 (36.52)	7.91 (36.05)
Weekend	-27.51 (16.79)	-124.92 (76.11)	-107.97 (49.35)*
EndOfMonth	-55.57 (16.61)**	-90.90 (71.22)	-93.05 (58.23)
EndOfYear	14.61 (69.49)	133.47 (321.75)	-1265.44 (192.47)**
ModelMonth5-13	31.46 (36.42)	231.02 (181.81)	-156.72 (183.77)
ModelMonth14+	-99.24 (55.96)+	156.11 (268.65)	-172.80 (322.56)
SouthernCal	-243.03 (48.64)**	-227.94 (107.87)*	-274.99 (107.55)*
Constant	26610.52 (163.18)**	23482.33 (555.17)**	24598.77 (798.96)**
Car fixed effects	Yes	Yes	Yes
Other fixed effects	Week*Segment	Week	Week
Observations	133424	6296	7046
Adj. R-squared	0.97	0.95	0.95

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses.

† MedianHouseValue in $100,000. Income in $1000.

Table 3: Price effects, identification issues (I)†

	(1a)	(1b)	(2a)	(2b)	(3)	(4a)	(4b)
	Reg. Disc. 2 day window		Reg. Disc. 3 days excluded		Diff-in-Diff	Reg. Disc.	
					Common Support		
Customer Cash	-0.73 (0.10)**	-0.91 (0.23)**	-0.82 (0.10)**	-0.73 (0.13)**	-0.84 (0.05)**	-0.81 (0.07)**	-0.79 (0.12)**
Dealer Cash	-0.38 (0.19)*	-0.26 (0.11)*	-0.34 (0.17)*	-0.36 (0.09)*	-0.31 (0.06)**	-0.40 (0.13)**	-0.32 (0.08)**
GM Card	-1.14 (0.12)**	-1.10 (0.12)**	-1.19 (0.13)**	-1.09 (0.10)**	-1.10 (0.07)**	-1.13 (0.10)**	-1.13 (0.09)**
Competition	-10.17 (21.66)	-12.48 (19.15)	-15.59 (14.04)	-16.62 (11.42)	-6.00 (7.40)	-13.68 (9.41)	-19.22 (9.71)*
Weekend	-270.01 (134.93)*	-165.10 (161.72)	-133.32 (98.29)*	-111.61 (77.17)	-12.38 (37.13)	-117.75 (77.04)	-101.04 (49.79)*
EndOfMonth	-2.79 (170.11)	-259.34 (110.69)*	-29.92 (131.55)	151.06 (123.16)*	-66.89 (26.58)*	-95.35 (70.86)	-99.63 (57.86)+
EndOfYear	-2161.76 (1826.86)	-1962.62 (881.00)*	136.75 (382.81)	-1346.84 (251.90)*	156.98 (118.98)	85.53 (320.19)	-1,271.96 (191.88)**
ModelMonth5-13	553.27 (430.82)	681.95 (677.67)	-18.36 (198.90)	-165.51 (232.87)	112.86 (66.01)+	262.79 (186.66)	-150.75 (174.87)
ModelMonth14+	680.05 (577.90)	1021.50 (987.31)	-56.88 (327.29)	-183.71 (407.46)	95.06 (95.99)	161.85 (276.07)	-167.23 (316.93)
SouthernCal	-206.83 (178.41)	-300.83 (205.10)	-271.58 (129.36)	-207.98 (128.42)	-344.32 (94.09)**	-216.15 (106.41)*	-269.45 (109.50)*
Constant	22610.47 (802.37)**	25148.66 (761.60)**	23057.55 (1215.90)**	28811.47 (588.46)**	25638.43 (322.51)**	23,392.12 (609.74)**	25,832.64 (555.13)**
Car fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Other fixed effects	Week	Week	Week	Week	Week*Seg.	Week	Week
Observations	2017	2099	3476	3999	41533	6181	6914
Adj. R-squared	0.95	0.95	0.95	0.95	0.95	0.96	0.96

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses.

† 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, black, or Hispanic; who are female; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value.

Table 4: Price effects, identification issues (II)†

	(1)		continued...
Customer Cash (increase) (weeks 1 and 2)	-.96 (.03)**	GM Card	-1.07 (.033)**
Customer Cash (increase) (weeks 3 and 4)	-.80 (.04)**	Competition	-7.35 (5.59)
Customer Cash (increase) (weeks 5 to 8)	-.80 (.03)**	Weekend	-27.94 (16.77)+
Customer Cash (increase) (weeks 9 to 26)	-.77 (.04)**	EndOfMonth	-56.09 (16.61)**
Customer Cash (increase) (weeks 26 +)	-.75 (.06)**	EndOfYear	15.41 (69.17)
Customer Cash (decrease) (weeks 1 and 2)	-.84 (.08)**	ModelMonth5-13	31.94 (36.03)
Customer Cash (decrease) (weeks 3 and 4)	-.88 (.08)**	ModelMonth14+	-103.04 (57.39)+
Customer Cash (decrease) (weeks 5 to 8)	-.92 (.06)**	SouthernCal	-243.49 (48.56)**
Customer Cash (decrease) (weeks 9 to 26)	-.91 (.07)**	Constant	26644.62 (164.92)**
Customer Cash (decrease) (weeks 26 +)	-1.27 (.11)**		
Dealer Cash (increase) (weeks 1 and 2)	-.28 (.06)**		
Dealer Cash (increase) (weeks 3 and 4)	-.35 (.06)**		
Dealer Cash (increase) (weeks 5 to 8)	-.38 (.08)**		
Dealer Cash (increase) (weeks 9 to 26)	-.48 (.08)**		
Dealer Cash (increase) (weeks 26 +)	-.42 (.14)**		
Dealer Cash (decrease) (weeks 1 and 2)	-.33 (.19)+		
Dealer Cash (decrease) (weeks 3 and 4)	-.49 (.20)*		
Dealer Cash (decrease) (weeks 5 to 8)	-.22 (.23)		
Dealer Cash (decrease) (weeks 9 to 26)	-1.05 (.13)**		
Dealer Cash (decrease) (weeks 26 +)	-.42 (.28)		
		Car fixed effects	Yes
		Other fixed effects	Week*Segment
		Observations	133642
		R-squared	0.97

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses.

† 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, black, or Hispanic; who are female; and who own their homes; as well as census block-level median income and median income squared, median household size, and median house value.

Table 5: Choice of promotion†

Logit of customer cash (1) vs. dealer cash (0)			
	(1)		continued...
Female	0.06 (0.12)	Δ Veh. Profit Margin (T-2 to T-1)	66.87 (4.74)**
%Asian	0.18 (0.58)	Δ Veh. Profit Margin (T-3 to T-2)	4.45 (3.38)
%Black	-0.56 (0.70)	Veh. Profit Margin (T-1)	-0.01 (0.00)**
%BlueCollar	1.72 (0.88)+	Veh. Profit Margin (T-2)	0.01 (0.00)**
%CollegeGrad	1.33 (0.81)	Veh. Profit Margin (T-3)	0.00 (0.00)
%Hispanic	-1.11 (0.87)	Δ Segm. Market Share (T-2 to T-1)	-20.29 (6.12)**
%LessHighSchool	-0.75 (1.13)	Δ Segm. Market Share (T-3 to T-2)	-71.79 (6.03)**
%HouseOwnership	0.43 (0.37)	Segm. Market Share (T-1)	98.45 (17.94)**
%Executives	-0.64 (1.19)	Segm. Market Share (T-2)	105.98 (19.03)**
%Professional	-0.79 (1.17)	Segm. Market Share (T-3)	-57.22 (13.82)**
%Technicians	0.14 (2.85)	Δ Sales (T-2 to T-1)	12.39 (2.97)**
Income	-0.08 (0.12)	Δ Sales (T-3 to T-2)	35.82 (3.15)**
Income^2	0.00 (0.01)	Sales (T-1)	-0.01 (0.00)*
MediaHHSize	0.14 (0.15)	Sales (T-2)	-0.03 (0.00)**
MedianHouseVal.	0.08 (0.09)	Sales (T-3)	0.00 (0.00)
# Competing Dealers	-0.09 (0.02)**	Δ Inventory (T-3 to T-2)	-9.29 (2.03)**
SouthernCal	-0.03 (0.16)	Inventory (T-1)	0.01 (0.00)**
Model Age	-0.32 (0.05)**	Inventory (T-2)	-0.04 (0.00)**
ModelMonth1-4	1.85 (0.61)**	Inventory (T-3)	0.04 (0.00)**
ModelMonth5-13	1.49 (0.44)**	Constant	3.98 (1.41)**
		Observations	4114
		Pseudo. R-squared	0.58

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses. MedianHouseValue in $100,000. Income in $1000.

† Sample includes only transactions for which (1) the car was sold either with a dealer cash promotion or a customer cash promotion, (2) the transaction occurred within the first 30 days of the promotion, and (3) the car was not on any type of promotion preceding the start of the current promotion.

Table 6: Pass-through controlling for market conditions and common support†

	(1)		continued...
Customer Cash	-1.02 (0.08)**	Δ Veh. Profit Margin (T-2 to T-1)	-3,296.14 (2,006.64)
Dealer Cash	-0.48 (0.12)**	Δ Veh. Profit Margin (T-3 to T-2)	-1,131.45 (1,677.42)
# Competing Dealers	-9.16 (16.96)	Veh. Profit Margin (T-1)	0.74 (0.65)
Female	104.78 (68.51)	Veh. Profit Margin (T-2)	-0.85 (0.70)
%Asian	-560.90 (273.24)*	Veh. Profit Margin (T-3)	-0.72 (0.45)
%Black	-27.04 (498.02)	Δ Segm. Market Share (T-2 to T-1)	-7,255.56 (4,028.62)+
%BlueCollar	-127.28 (458.51)	Δ Segm. Market Share (T-3 to T-2)	6,188.64 (4,472.44)
%CollegeGrad	-218.91 (506.91)	Segm. Market Share (T-1)	14,131.38 (13,497.33)
%Hispanic	17.62 (368.77)	Segm. Market Share (T-2)	-22,247.33 (10,538.93)*
%LessHighSchool	-769.80 (596.04)	Segm. Market Share (T-3)	-983.61 (7,067.90)
%HouseOwnership	83.25 (187.48)	Δ Sales (T-2 to T-1)	3,178.21 (2,040.37)
%Executives	-271.80 (587.30)	Δ Sales (T-3 to T-2)	-2,715.93 (2,303.78)
%Professional	-178.75 (749.31)	Sales (T-1)	-3.07 (2.39)
%Technicians	344.52 (1,646.84)	Sales (T-2)	4.48 (2.40)+
Income	-126.74 (84.68)	Sales (T-3)	-0.10 (1.74)
Income^2	9.79 (4.90)+	Δ Inventory (T-2 to T-1)	0.00 (0.00)
MediaHHSize	10.31 (72.24)	Δ Inventory (T-3 to T-2)	-419.76 (1,393.21)
MedianHouseVal.	-18.46 (44.35)	Inventory (T-1)	0.04 (1.62)
Weekend	30.11 (65.45)	Inventory (T-2)	-1.28 (3.20)
EndOfMonth	99.21 (122.39)	Inventory (T-3)	3.61 (3.18)
EndOfYear	82.50 (446.66)	Constant	30,821.22 (2,320.52)**
ModelMonth5-13	-162.02 (209.63)		
ModelMonth14+	-228.91 (360.43)		
SouthernCal	-180.85 (132.32)		
		Car fixed effects	Yes
		Week fixed effects	Yes
		Observations	3939
		R-squared	0.96

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses. MedianHouseValue in $100,000. Income in $1000.

† Sample includes only transactions for which (1) the car was sold either with a dealer cash promotion or a customer cash promotion, (2) the transaction occurred within the first 30 days of the promotion, (3) the car was not on any type of promotion preceding the start of the current promotion, and (4) the transaction is in the common support on observable market conditions, consumer characteristics, and car characteristics.

Table 7: Extensions[†]

	(1)	(2)	(3)	(4)
	Diff-in-Diff	Diff-in-Diff	Diff-in-Diff	Diff-in-Diff
Customer Cash	-0.84 (0.04)**	-0.82 (0.04)**	-0.90 (0.04)**	-0.74 (0.04)**
Dealer Cash	-0.42 (0.07)**	-0.52 (0.09)**	-0.47 (0.08)**	-0.41 (0.07)**
GM Card	-1.03 (0.04)**	-1.02 (0.04)**	-1.02 (0.04)**	-1.04 (0.04)**
CustCash*Competition	-0.01 (0.004)*	-0.01 (0.005)*	-0.01 (0.005)*	-0.01 (0.004)*
DealCash*Competition	0.01 (0.01)	0.02 (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)
Num. DealCash Prom. in Seg.	-15.26 (16.94)		-14.73 (16.80)	-11.32 (16.77)
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)
# Competing Dealers	-4.87 (6.75)	-5.79 (6.80)	-4.15 (6.82)	-4.48 (6.75)
Female	144.40 (12.00)**	146.34 (12.05)**	120.72 (13.50)**	143.78 (11.97)**
%Black	473.90 (92.79)**	476.36 (94.39)**	365.69 (93.47)**	473.38 (93.14)**
%Hispanic	-32.62 (86.62)	-32.24 (86.96)	-229.67 (94.66)*	-27.99 (86.40)
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	26565 (165)**	26607 (180)**	26609 (165)**	26627 (166)**
Other fixed effects	Week*Segment	Week	Week*Segment	Week*Segment
Car fixed effects	Yes	Yes	Yes	Yes
Observations	133424	133424	133424	133424
Adj. R-squared	0.97	0.97	0.97	0.97

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses.

† 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.