

Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence[†]

By SUSAN E. WOODWARD AND ROBERT E. HALL*

Mortgage loans are leading examples of transactions where experts on one side of the market take advantage of consumers' lack of knowledge and experience. We study the compensation that borrowers pay to mortgage brokers for assistance from application to closing. Two findings support the conclusion that confused borrowers overpay for brokers' services: (i) A model of effective shopping shows that borrowers sacrifice at least \$1,000 by shopping from too few brokers. (ii) Borrowers who compensate their brokers with both cash and a commission from the lender pay twice as much as similar borrowers who pay no cash. (JEL D12, D14, G21)

We study a large and financially complex expenditure that the majority of consumers make only a few times in their lives, the payment for mortgage origination services. In these transactions, information is not only imperfect, but traditional market practices have likely contributed to consumer confusion. These payments range from 0 to \$30,000 for mortgages of normal size. Mortgage brokers receive payments from both borrowers and from wholesale lenders on most loans. Broker fees paid in cash by the borrower are usually broken down into more than a dozen individual items such as “origination fee,” “discount points,” “funding fee,” and “doc prep.” In addition to these cash payments, about half of the value that the borrower transfers to the broker takes the form of a payment from the wholesale lender to the mortgage broker. The borrower bears the burden of this part of the payment in the form of a higher interest rate on the loan.

These payments compensate mortgage brokers for arranging mortgage loans. We stress that this paper is about how and how much borrowers pay their brokers for origination services. The broker is an administrator of the process of loan origination. The broker bears none of the risk of default on the mortgage, so that risk is not a determinant of the broker's compensation under standard theories of pricing.

Because consumers enter the mortgage market infrequently and because features of the market make it difficult to learn the best price, mortgage origination pricing

*Woodward: Sand Hill Econometrics, 1682 Oak Avenue, Menlo Park, CA 94025 (e-mail: swoodward@sandhillecon.com); Hall: Hoover Institution, Stanford University, Stanford, CA 94305 (e-mail: rehall@stanford.edu). Disclosure: Woodward served as a consultant to the US Department of Housing and Urban Development on issues relating to mortgage origination charges—see Woodward (2008). She also served as a litigation expert for the Massachusetts Attorney General on mortgage brokers' charges.

[†] To view additional materials, visit the article page at <http://dx.doi.org/10.1257/aer.102.7.3249>.

is a market where one suspects that many consumers pay well above the best price. Our results confirm this suspicion.

We reach our conclusion by studying the distribution of origination charges for a large sample of mortgages involving brokers, where federally mandated disclosures report the entire amount of the broker's revenue, showing both the cash charges to the borrower plus the additional amount the lender pays the broker. We consider a minimal shopping strategy that borrowers might pursue in trying to find the best price—getting quotes from two brokers, asking the one with the higher proposed price to beat the lower price of the opponent, and continuing this process until one broker is unwilling to improve on the other's best proposal. This process is an English auction, in which the lower-cost broker gets the business and the charge equals the cost of the losing broker, according to standard auction principles.

It's a standard statistical exercise to find the distribution of a variable from the distribution of the larger of a pair of draws from the variable. We perform that exercise to calculate the distribution of broker cost. We find that the implied cost is generally quite high, but more important, the upper tail of the cost distribution is thick—a significant fraction of mortgages appear to cost the broker more than \$5,000 to originate. When we repeat the exercise for shopping among three and four brokers, we find that the implied distribution of cost is most implausible, with an even larger fraction implied to cost more than \$5,000. The distribution has an implausible shape as well. We conclude that among our shopping models, only the one where borrowers shop from just a pair of brokers is close to reasonable. Our conclusion that borrowers consider no more than two mortgages draws support from surveys of borrower behavior as well.

Given this conclusion, we ask what benefit a borrower who shopped from only two brokers passed up by not shopping from three or four. The answers are so large that we believe that most borrowers must have been unaware of the likely benefits of more shopping. For example, for a mortgage with \$100,000 principal, a borrower would save a median of \$981 by adding one more broker to the mix and \$1,393 by adding two. And with \$200,000 principal, the savings are \$1,866 and \$2,664. Because we do not believe that borrowers would intentionally pass up such large benefits just to avoid talking to another broker, we conclude that confusion about how this market works caused borrowers to shop too little. We doubt that borrowers understand either the nature of the payments made by wholesale lenders to brokers and their implications for the rate-cash trade-off, or the incentives these payments present to the mortgage broker.

Our second approach to studying confusion among mortgage borrowers is to compare (i) the total origination charges for loans where the borrower pays a higher interest rate to fund the origination charge to (ii) the total charges for loans where the borrower pays all of those charges in cash. The first group pays somewhat lower total charges than the second, but the important finding is that both those groups pay far less than borrowers who use both types of funding in roughly equal proportion. This evidence is consistent with the hypothesis that borrowers treat the two charges independently, failing to recognize that a borrower who pays more cash is entitled to a lower interest rate and vice versa.

Earlier research has shown that mortgage charges are higher for less-educated borrowers, members of minorities, borrowers who pay high interest rates, and

those who borrow larger principal amounts—we review that research in the online Appendix. We confirm these findings. The research has not shown whether the borrowers paying higher charges did so because arranging the mortgage broker's cost was higher or because those borrowers suffered exploitation due to their lack of knowledge of the best available charge, which should be little higher than cost. Our results suggest that large fractions of the higher charges are the result of limited shopping rather than higher cost.

Our data come from a sample of mortgages insured by the Federal Housing Administration (FHA) during a six-week period in 2001. FHA insures mortgages of fairly creditworthy borrowers for modestly priced houses, up to loan limits that vary by geography. The borrower pays FHA an insurance premium. All of these loans are 30-year fixed-rate mortgages for the purchase of a house. No loans have prepayment penalties. FHA bears essentially all of the default risk, so pricing of default risk is not a concern for wholesale lenders and mortgage brokers.

Our econometric approach is mainly non-parametric. We represent the full distribution of origination charges in terms of 299 quantiles, conditional on observed borrower and loan characteristics. From the quantiles, we make calculations of the distribution of broker cost and of the charges that borrowers would have paid under more effective shopping strategies. We provide bootstrap standard errors for all of our calculations.

I. Economics of Mortgage Origination

A mortgage is a loan secured by a house. The typical mortgage provides for monthly payments over a term of 30 years. The amount of the loan, called the loan *principal*, passes to the seller of the house at the moment when the borrower takes ownership of the house, a moment called the loan closing. The borrower's *coupon rate* is applied to the principal amount of the loan to calculate the borrower's monthly payment. The loans are fully amortizing, so there is no final repayment of principal, in contrast to the typical bond. The borrower has the option to pay off a mortgage before 30 years, subject to a pre-payment penalty, which is limited by law in all states. Because borrowers sometimes move, change houses, extract appreciated equity, and take advantage of lower interest rates to refinance, most mortgages pre-pay prior to their 30-year maturity. Nonetheless, roughly 30 percent of owner-occupied homes have no mortgage, and nearly all got to this state by paying off a 30-year loan—see US Census Department (2001). Our FHA sample contains only 30-year fixed-rate mortgages; there are no investor loans and no refinancings among them and none has a prepayment penalty.

A. Brokers and the Yield-Spread Premium

Mortgage brokers perform the service of *originating* a mortgage. A broker matches a borrower with a wholesale lender. The broker is not a party to the resulting financial contract between borrower and lender and thus bears none of the default risk of the mortgage. The broker helps the borrower prepare an application and arranges for the services of an appraiser and an agent to close the transaction. The broker serves lenders by finding potential borrowers and helping them complete the

necessary paperwork. Most large lenders have retail origination operations and also use brokers as originating agents. A broker usually has relationships with a number of wholesale lenders.

The borrower deals with a broker under conditions comparable to a purchaser dealing with a retailer of an expensive item. Like the retailer and purchaser, the broker and borrower negotiate the terms of the transaction without participation from the upstream wholesaler. The wholesaler provides funds for the loan under quoted terms. The borrower receives a specified amount of cash, the principal amount of the mortgage, delivered by the lender at the time of the closing, in exchange for a mortgage at an interest rate resulting from the negotiation. The borrower also usually pays the broker a negotiated amount of cash at the closing; we call this the *cash origination charge*. Berndt, Hollifield, and Sandås (2010) discuss the bargaining problem between borrower and broker in the framework of the Nash bargain in its alternating offer form.

The broker may also receive a payment from the lender called the *yield-spread premium* or YSP, which is typically about half of the broker's compensation. The terms offered by the wholesale lender appear in a *rate sheet*, a document the broker receives from the lender at least daily. The rate sheet shows the YSP the lender offers to pay the broker for originating a mortgage. The YSP is an increasing function of the coupon rate and principal amount of the loan, and decreasing in the number of days (15, 30, 45, or 60) for which the loan is locked (the time the broker and borrower have to complete the loan). Because lenders can sell loans with higher rates for higher prices, the YSP rises as the rate on the loan rises, but at a decreasing rate because higher rate loans are likely to prepay sooner than lower rate loans, so higher rates are generally earned for shorter periods. This function is determined by expectations about movements in interest rates in the competitive wholesale mortgage market. We take it as given.

The following example, representative of the transactions in our data, illustrates the operation of the YSP: A borrower pays her broker \$1,800 as a cash closing payment. In addition, the lender pays the broker a YSP of \$2,300. The broker's all-in cost, mainly the value of his time, is \$2,400. The borrower has paid $\$1,800 + \$2,300 = \$4,100$ in origination charges for a loan when she could have pressed the broker to do the loan for close to \$2,400. A savvy borrower could have insisted that the broker charge her only \$100 in cash, which, together with the yield-spread premium of \$2,300, would have just covered his cost of \$2,400.

Figure 1 shows an example of an actual rate sheet. The figures show the amount that the lender will deliver at closing on behalf of the borrower (always taken to be 100) plus the premium to be paid to the broker, such as 2.25 percent of the principal for a loan at a coupon rate of 8.875 percent with a lock period of 30 days. Notice that below the solid line, the YSP becomes negative. For loans at these low interest rates, the borrower (not the broker) pays the specified amount of cash into the closing to make up the difference.

The coupon rate that corresponds to a zero YSP is called the *par* rate and is a useful benchmark rate. It corresponds to the row in the figure just below the solid line.

A recent change in mortgage law bars lenders from making payments to origination agents (loan officers and mortgage brokers) that are tied to the interest rate on the loan. The new rule attempts to ban both payments of YSPs to mortgage brokers,

Rate	Lock period			
	15 days	30 days	45 days	60 days
8.875%	102.625	102.500	102.375	102.250
8.750%	102.375	102.250	102.125	102.000
8.625%	102.000	101.875	101.750	101.625
8.500%	101.500	101.375	101.250	101.125
8.375%	101.000	100.875	100.750	100.625
8.250%	100.625	100.500	100.375	100.250
8.125%	100.250	100.125	100.000	99.875
8.000%	99.750	99.625	99.500	99.375

FIGURE 1. AN ACTUAL RATE SHEET FOR APRIL 2000

who were required to disclose YSPs, and to agents of correspondent lenders, who, under the law prevailing in 2001, were not required to report YSPs. The ban has no effect on the underlying logic that a borrower faces two ways to compensate a lender for origination services, by paying cash up front or accepting a higher interest rate. It imposes no limit on the interest rate that a lender can charge. Vertically integrated lenders have never faced any requirement to disclose the internal equivalent of a YSP or any limit on its amount. The evidence in this paper based on observed YSPs is directly relevant to the structure the new law created. The new law changes the organization of mortgage origination in a way that obscures the YSP without changing the economic forces that cause lenders to charge borrowers for origination services in part through higher interest rates.

B. *Bargaining over Mortgage Terms*

The YSP lets the home-buyer borrow funds to compensate the broker, provided the borrower remains within the lender's payment-to-income limits at the higher interest rate and corresponding higher payment. The value flowing from the borrower to the broker is the cash origination charges plus the YSP. For the broker, the mix of the two components is immaterial, as the broker receives both the cash origination charge and the YSP as cash at the same time. For a borrower who is cash-constrained and knows she is likely to pay off the mortgage fairly soon, the cost of the higher interest is less than what the broker receives as YSP, so their efficient bargain should pay the broker entirely with the YSP and the origination charge should be zero (or even negative, which occurs in a small fraction of our sample). Similarly, for a borrower expecting to have the loan for a long time, the cost of the higher rate is greater than the YSP and then the efficient deal with the broker will not involve borrowing any of the broker's compensation; rather, the borrower will pay it in cash, and even pay a negative YSP to lower the interest rate further. The online Appendix contains a formal model of the bargain between broker and borrower.

C. *Potential Consumer Confusion*

Our predictions about the efficient deal between borrower and broker rest on the assumption that the borrower is able to recognize when she is paying a higher interest rate and thus creating value that could be used to pay the broker in place of

a cash origination charge. Another possibility is that some borrowers are unable to spot when an interest rate is high enough to deserve a reduction in the cash charges and thus wind up paying both normal cash charges and giving the brokers handsome YSPs as well. A profit-maximizing broker will try to keep borrowers uninformed about the availability of mortgages at lower interest rates, in the hope of earning high compensation from both components.

The borrower's understanding of the mortgage market is often incomplete. For mortgages on newly purchased homes, the buyer's real-estate broker often refers the purchaser to a mortgage broker, who explains that he has access to mortgages from dozens of competing lenders and that he can help the borrower choose the one with the best terms. The mortgage broker does not mention that each of these lenders links the interest rate to the YSP that the broker will receive. He proposes one or more mortgages, each with a rate that includes as big a YSP as he thinks the borrower will accept. His proposal includes a rate and detailed cash origination charges. His offer will reflect how careful and informed a shopper he believes the borrower is.

Mortgage law in effect at the time of our sample required the broker to provide the borrower at this stage in the process a good-faith estimate of the origination charge, but not of the YSP. Thus the borrower did not know prior to making a deal with a mortgage broker how much the broker will earn in total fees. The borrower might have eventually learned the amount of the YSP, in principle, because the law required its disclosure at the time of the closing, on the HUD-1 form that is part of the closing document package. The terminology of the disclosure, the location of the YSP entry on the form, and the bewildering nature of the form itself likely inhibited many borrowers from understanding the YSP. In any case, the information arrived long after the mortgage deal was made.

D. Points

One important source of confusion among borrowers is the labeling of parts of cash origination charges as "points." A great deal of commentary on mortgage economics shares this confusion. A point is a component of the broker's origination charge calculated as a percent of the principal amount of the loan. From the perspective of the economics of a mortgage origination, a charge for points is just part of the total amount that the borrower pays to the broker. Borrowers may believe that paying points "buys down" the interest rate. In principle, this should be true. If a broker receives part of his cash origination charge as points, just as in any other form, the borrower should be able to bargain for a lower interest rate and thus a lower YSP. But this only happens if the borrower uses the added bargaining power that paying more cash to the broker ought to give the borrower. Nothing stops the broker from offering an interest rate that earns him a handsome YSP and also including a charge for points as part of the pricing of the origination. In our sample, charges for points on loans are present on roughly 30 percent of the loans with positive YSPs.

When the coupon rate is sufficiently low that the YSP is negative, the borrower will face a charge for points that are paid to the lender. If the charge for points is the amount on the rate sheet for the negative YSP (below the horizontal line in Figure 1), the charge genuinely brings a lower interest rate. But here too, the broker may charge

more for points than the lender receives—nothing stops the broker from charging more for the negative YSP than appears on the rate sheet and keeping the excess as part of his own compensation. In any case, negative YSPs are rare in our sample, despite the high frequency of charges labeled as points.

Even the borrower's understanding of the broker's cash fees may be limited by the practice of dividing the charge into many elements. Figure 2 is an example containing terms found frequently in our data. The good-faith disclosure form prevailing when these loans were written did not require brokers to state the sum of the charges (new rules do require disclosure of the sum), so comparison of loans was challenging. Some of the challenge remains, because the disclosure law perversely allows for points to be separated from other origination charges, suggesting that even the regulators fail to understand that points are just another name for broker or lender charges. Borrowers may not recognize that only the sum—including points—is meaningful. Lacko and Pappalardo (2007) found that few borrowers had any idea what the term "discount points" means, and worse, that one-third believed it was a discount they were receiving instead of an amount they were paying.

E. Industrial Organization Analysis of Mortgage Brokerage

Our findings later in the paper suggest that an important fraction of borrowers leave a lot of money on the table for their brokers. As we show later in the paper, the upper tail of the distribution of origination charges brokers receive contains some astonishingly large values. We use the term *margin* to mean the difference between the revenue from originating a mortgage and the cost of executing the transaction. Here we include only costs incurred after the broker meets the borrower and strikes a deal. The margin is the amount at issue when the broker bargains with the borrower over the charge. Thus we believe that mortgage brokers earn substantial margins from many but not all of their origination efforts.

Mortgage brokerage involves only small out-of-pocket costs for the broker. A mortgage broker does not usually outsource any of the origination process to other professionals. Further, though charges labeled "origination fee," "funding fee," and "commitment fee" sound as if they passed through to the lender, they remain firmly in the broker's pocket. Thus the costs we have in mind are largely the value of the broker's time.

The equilibrium we describe is inefficient. We do not believe that mortgage brokers earn any important rents from their origination activities. Entry to mortgage brokerage is close to free. Most states license mortgage brokers or require them to have licenses as real-estate agents, but these licenses are not difficult to obtain. Mortgage brokers dissipate the anticipated margin from confused borrowers by spending effort and other resources trying to find these customers. Because borrowers seldom seek out several brokers to compete for the borrower's business, brokers have to work hard to find brokerage customers. In Lacko and Pappalardo (2007), the modal number of loans considered by borrowers was two. In Federal Reserve Board (2009), more than half of all borrowers looked seriously at only one loan.

We noted earlier that mortgage brokers often receive referrals from real-estate brokers and that it appears that many borrowers accept a proposal from such a

Borrower signs up for 8.875% with a 60-day close, \$100,000 loan	
Lender delivers \$100,000 on behalf of borrower	
Broker gets	
From lender:	
\$2,300	yield-spread premium (paid outside of closing)
From borrower:	
\$ 800	origination fee
75	warehouse wire fee
50	fee for preparing amortization table for loan
200	document processing charge
150	funding fee
25	fax fee
200	document review fee
300	commitment fee
\$4,100	TOTAL

FIGURE 2. EXAMPLE OF A LOAN TRANSACTION

mortgage broker without further shopping. Hence the referrals are valuable and mortgage brokers cultivate real-estate brokers hoping to receive their referrals. Federal law—the Real Estate Settlement Practices Act—prohibits mortgage brokers from paying referral fees, but little is known about the effectiveness or enforcement of that prohibition. In any case, we believe that a zero-profit equilibrium prevails in the mortgage-brokerage business, where new entrants expect a zero net payoff given the costs of finding customers. But once found, some of the customers are very profitable. From the social point of view, the resources dissipated by brokers in their search for overpaying customers are wasted.

II. Data and Data Description Model

A. Description of the FHA data

Table 1 describes the variables in our analysis. We excluded loans with rates below 7 percent and those with interest rates not at 0.125 point ticks as presumptively subsidized. Interest rates are fairly tightly clustered around 7.5 percent, as shown in the first column of Table 2. All the loans were made at essentially the same time, so the variation arises in the cross section of borrowers and not from changes in credit markets. The total origination charges—cash origination charge plus yield-spread premium—average \$4,111, but have substantial dispersion. The cash component is typically a little under half of the total charge and the YSP a little more than half. The principal is generally around \$100,000 and rarely exceeds \$200,000. The average credit score of the borrower is 659, which is just below the median of the national distribution of scores. The fractions of the borrowers who are members of minorities are close to the US average in the population, at 11 percent African American and 14 percent Latino. The last statistic is the fraction of the adult population in the borrower's census tract who hold a BA degree—its average level is 21 percent. We have no measure of the education of individual borrowers, but are able to measure the education of their neighbors because we know the addresses of the borrowers.

TABLE 1—DESCRIPTIVE STATISTICS FOR BROKERED LOANS

	Mean	SD
Total closing charge, dollars	4,111	2,291
Percent of closing cost paid in cash	40	42
Principal, dollars	112,907	39,891
Credit score	659	63
Percent African American	11	
Percent Latino	14	
Percent of neighbors with BA degrees	21	12
Number of loans	1,525	

TABLE 2—AVERAGE YIELD-SPREAD PREMIUMS IN THE FHA DATA, BY INTEREST RATE

Interest rate in category, percent	Percent of sample	Average rate in category, percent	Average yield-spread premium per \$100 principal	SD of YSP ratio
7 to 7.375	37	7.17	1.33	0.60
7.5	39	7.50	2.32	0.59
7.625 to 7.875	13	7.78	2.68	0.83
8 and higher	11	8.11	3.18	1.13

B. *The Yield-Spread Premium*

The YSP is a payment from the wholesale lender to the broker based on the coupon rate for the mortgage and on the amount of the principal. The lender pays the premium to the broker outside the closing. Earlier we reviewed what we believe is highly reliable evidence about the formula for the YSP from rate sheets. Here we examine the information in the FHA data on the YSP, which provides a different view of the formula. The source of the FHA broker compensation data is the HUD-1 settlement statement, required by the Real Estate Settlement Practices Act (RESPA). This source is potentially imperfect, because there seems to be relatively little monitoring of the accuracy of a broker's disclosure and the broker may try to conceal a large premium from a borrower. Table 2 shows the average yield-spread premium for the brokered loans in our sample, by interest rate.

Figure 3 compares the relation between the coupon rate and the YSP in the FHA loans and in the rate sheets we discussed earlier. We adjusted the coupon rate from the rate sheets downward by 99 basis points to account for the general decline in rates between 2000 and 2001, as reported by the Federal Reserve Board. The error bars are the standard errors of the average in each category (hardly visible for the FHA data because of the large number of observations). The slope of the premium is generally lower in the FHA data. The difference in the slope could reflect changes in expectations about pre-payments and changes in the slope of the yield curve between 2000 and 2001. It could also reflect some tendency for brokers to understate their actual premiums when they are high. Because the yield-spread premiums reported in the FHA data are not obviously at odds with those in the rate sheets, and because we are unable to adjust the curve from 2000 apart from shifting it to the left, we accept the reported premiums for the rest of our analysis.

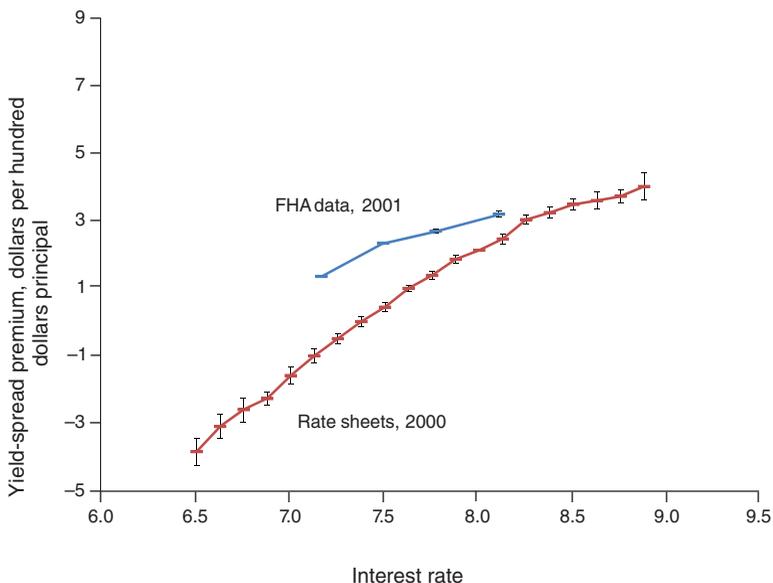


FIGURE 3. COMPARISON OF YIELD-SPREAD PREMIUMS IN FHA LOANS AND LENDERS' RATE SHEETS

C. Descriptive Model

A broker receives revenue

$$(1) \quad \tau = L + Y(r)$$

from originating a mortgage at coupon rate r . In Section IB, we noted that brokers are surely unconcerned about the mix between the cash from the borrower, L , and the cash from the lender, $Y(r)$, and that the typical borrower should be close to indifferent about the mix as well. The near-indifference suggests we should start by studying the sum.

All earlier research on mortgage terms has examined the expectation of τ or its components conditional on a vector of observed characteristics, via regression. Our interest extends to the entire distribution of τ conditional on characteristics, a vastly more complicated object than the expectation. Our approach is to estimate the quantiles $\tau_i(x)$ of the distribution as functions of the observed characteristics x . The quantile gives the value of τ such that the probability that τ is no greater than $\tau_i(x)$ has a designated value, q_i . It is the inverse of the cumulative distribution function $F(\tau)$:

$$(2) \quad F(\tau_i(x)) = q_i,$$

where q_i is the designated probability for quantile i . We take the probability values to be: $q_i = i/300$ for $i = 1, \dots, 299$. We fit the equation

$$(3) \quad \tau_i(x) = x\beta_i,$$

TABLE 3—ESTIMATES OF THE PARAMETERS OF THE FUNCTION DESCRIBING THE MEDIAN OF THE TOTAL ORIGATION CHARGE

Characteristic	Coefficient (bootstrap standard error)
Intercept	0.96 (0.21)
Principal, hundreds of thousands of dollars	2.84 (0.14)
African American	0.94 (0.21)
Latino	0.91 (0.15)
Fraction of neighbors with BAs	-2.51 (0.48)
Credit score, hundreds	-0.05 (0.10)

where β_i is the vector of parameters describing the i th percentile of the distribution of τ , conditional on x . See Koenker and Hallock (2001) for a discussion of the estimation of quantiles. Our use of quantile estimation is only to transform the data into a form that is useful for further analysis. The transformation has the same kind of role that a Fourier transform of the data has in time-series analysis.

Table 3 shows the equation for the median ($i = 150$) of the origination charges for the mortgages in the FHA data, using the variables from Table 1. The median rises sharply with the amount of the principal. The median charge is \$2,842 higher for a mortgage with a principal of \$200,000 compared to one with a principal of \$100,000. One might expect that broker fees are strongly related to loan amount simply because of the mechanical relationship between the loan amount and YSP; however, the cash fees are nearly as strongly related to loan amount as the YSPs are.

Table 3 also shows that African American borrowers pay median origination charges that are \$939 above those of otherwise similar white borrowers, and Latino borrowers pay \$912 more than white borrowers. Borrowers from census tracts with higher educational attainment pay smaller total origination charges—an increase of one standard deviation (12 percentage points) lowers the median cost by \$300. Finally, the borrower's credit score has a small negative relation to the median—100 extra points lowers the median charge by \$55. The last result is consistent with our observation earlier that the broker has no stake in the mortgage itself. The results in this table are similar to those found in earlier work in the regression framework. We estimated the standard errors in this and subsequent tables by bootstrap.

Table 4 defines five illustrative cases we carry through the rest of the paper, together with the estimated median for each. Each case perturbs the base specification along one dimension of the explanatory variables. We do not include a case for a Latino borrower because our results show little difference between African Americans and Latinos. Our base case, in the first column, is a white borrower with \$100,000 principal, living in a census tract where 21 percent of adults have BAs. The median total origination charge for the base case is \$2,185. Note that all of

TABLE 4—FIVE ILLUSTRATIVE CASES

	Base (1)	African American borrower (2)	Highly educated neighbors (3)	High principal (4)	High credit score (5)
Principal, dollars	100,000	100,000	100,000	200,000	100,000
Race	White	African American	White	White	White
Percent of neighbors with BA degree	21	21	45	21	21
Credit score	650	650	650	650	800
Estimated median total charge, dollars	2,185 (44)	2,790 (157)	1,720 (90)	4,127 (161)	2,191 (78)
Difference from base case, dollars		605 (157)	-465 (90)	1,942 (161)	7 (78)

Note: Bootstrap standard errors in parentheses.

the differences between the cases and the base case are statistically unambiguous, except for the high credit-score case.

Figure 4 and Figure 5 illustrate how we handle the entire distribution implied by our descriptive model, for the base case defined in Table 4. We compare the raw distribution implied by the model to a smoothed distribution that we use for subsequent analysis. The raw cumulative distribution is the graph of the percentiles on the vertical axis and the fitted values for the base case from our quantile estimates on the horizontal axis. This graph demonstrates how a set of quantile estimates describes a distribution conditional on a given point in the space of characteristics, x . Notice that it has small wiggles arising from sampling variation. The wiggles in the cumulative distribution turn into huge sawtooth fluctuations in the raw density. It is a well-known statistical problem that the estimate of a density at a point is not consistent. But a kernel smoother of a density is consistent. Figure 5 shows the standard kernel smoother of the raw density. Then looking back at Figure 4, one can see that applying the smoother does little to alter the shape of the cumulative distribution.

Figure 6 shows the smoothed densities that the descriptive model implies for the base case in Table 4, along with the smoothed density for the actual distribution of total origination charges in the FHA data. The model pictures the actual distribution as the mixture of the distributions conditional on the x -characteristics. These distributions gain their shapes almost entirely from the underlying data. Our fitted model, with $299 \times 6 = 1,794$ estimated coefficients, is hardly restrictive. The actual distribution has a bulge around \$6,000 relative to the distribution for the base case, reflecting the presence of loans with higher principal and other factors associated with higher total origination charges than in the base case.

The dotted lines in Figure 6 show bootstrap one-standard-error bands for the smoothed actual density. Bands for the other densities shown in other figures are similarly tight and are omitted.

Figure 7 compares the distribution for African American borrowers to the distribution for the base case. The mode for the African Americans is only about \$700 higher than for the white borrowers with otherwise similar characteristics in the base case, but the upper tail is much thicker for the African Americans. Their density is

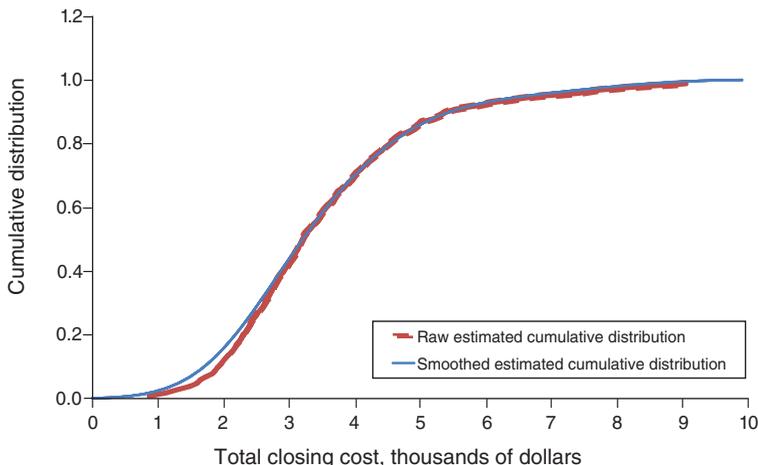


FIGURE 4. RAW AND SMOOTHED CUMULATIVE DISTRIBUTION FUNCTIONS FOR THE BASE CASE

more than double that of the base-case borrowers at a total origination charge of \$6,000, a great deal of money for originating a loan for \$100,000.

Figure 8 compares the distribution for borrowers having more-educated neighbors to the distribution for the base case. These borrowers have a substantial advantage over those in the base case. We believe that the advantage comes mainly from the likelihood that the borrower or a family member is better educated, but living among better-educated people may also confer an advantage because the environment contains more people who grasp some principles of mortgage pricing. The group with more educated neighbors has a very much lower density than does the base case for the range from \$5,000 to \$7,000. Even among those with educated neighbors, a significant minority pay what appear to be extreme total origination charges.

Figure 9 compares the distribution for loans with high principal (\$200,000) to the distribution for the base case (\$100,000). The distribution for the larger mortgages lies far to the right of the distribution for the base case. The powerful relationship with the amount of the transaction is a puzzling feature of many financial services. Little of the work of a mortgage broker scales with the amount of the principal, but the modal and median charge for a \$200,000 mortgage is about double that of a \$100,000 mortgage. The upper tail for the bigger principal is quite thick—a few people pay astronomical origination charges for the larger mortgages.

The difference between the base case and the case with a higher credit score, shown in Figure 10, is instructive. For all but the high values of the origination charge, the two densities are similar, reaching peaks at about the same value, but with lower probability for all of the lower values in the case of the higher credit score. All of the difference is at the upper end, where the base case has much more probability above $\tau = \$7,000$. Apparently borrowers with better credit scores are able to avoid the mistake of agreeing to such high charges. FHA underwriting works mainly with a cutoff credit score, which the base case satisfies. Perhaps borrowers with credit scores close to the minimum acceptable score are less willing to shop, because they do not realize that their scores are acceptable to all FHA lenders.

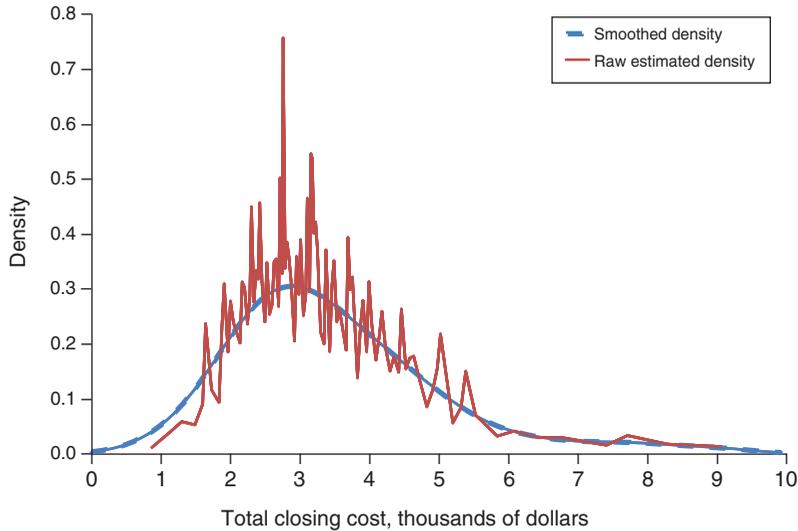


FIGURE 5. RAW AND SMOOTHED DENSITIES FOR THE BASE CASE

III. Consumer Confusion

We believe that our results reveal substantial consumer confusion, in the sense that the process of obtaining a mortgage through a broker results in many borrowers paying far more for the broker's origination services than they would if they better understood how to shop for a mortgage and if they understood that the broker was not doing their shopping for them. Our first approach to supporting this conclusion is to fit a model of optimal shopping to our data and then observe that the payoff to more intensive shopping—specifically, the benefit of getting a quote from an additional broker—is far higher than seems reasonable. Borrowers behave as if it costs over \$1,000 to get a quote from another broker. Of course, we cannot rule out the alternative conclusion that shopping among brokers for the best deal has such a high psychic cost that the choice to limit shopping is a rational response to that high cost.

Not every borrower who gets a good deal from a broker is necessarily free of confusion. Borrowers who are close to the margin for qualifying for a loan may have neither much cash to give the broker nor the ability to pay a higher coupon rate that delivers a handsome YSP, so to gain the borrower's business the broker has to charge little for his services.

A. Modeling Market Equilibrium with Rational Shopping

A recent body of research considers the econometric problem of inferring shopping costs from market data. Hong and Shum (2006) and Hortaçsu and Syverson (2004) are notable contributions, and Moraga-González (2006) is a recent survey. This literature considers posted prices—consumers visit stores and buy from the one with the best posted offering. We have not found any empirical model in that literature that would apply to the problem we consider here, even if we accepted the assumption of

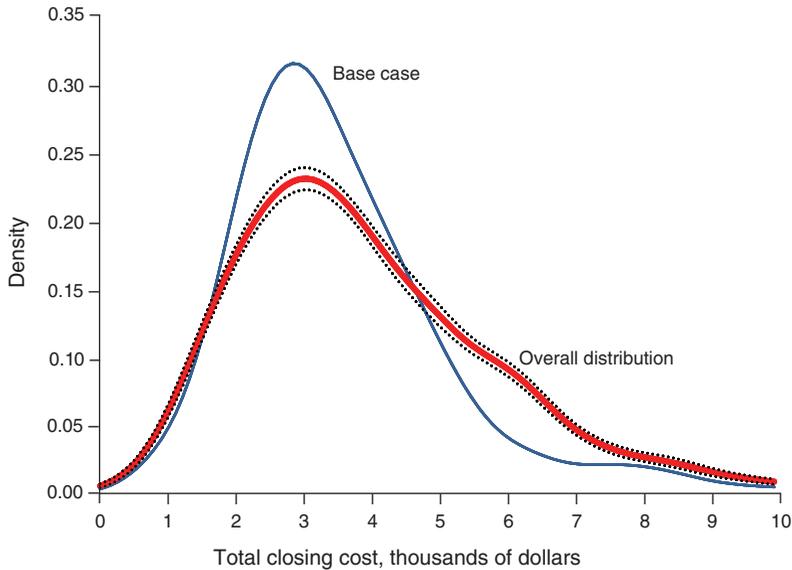


FIGURE 6. SMOOTHED DENSITIES FOR THE BASE CASE AND FOR THE ACTUAL ORIGINATION CHARGE

posted prices. Hong and Shum (2006) follow Burdett and Judd (1983) in assuming that all sellers have the same cost. That assumption is untenable in our data, because it implies that the common value of cost must be lower than the smallest observed price. We find a tiny fraction of prices that are only a few hundred dollars. We can only make sense out of the data under the assumption of fairly heterogeneous cost.

On the other hand, Hortaçsu and Syverson's (2004) study of mutual funds assumes heterogeneity in both seller product cost and buyer cost of search, the case we believe describes the mortgage market. Their results support the conclusion that consumers suffer confusion in their choices of mutual funds by failing to shop hard enough for a low expense ratio. But they use the market shares of sellers to identify the two distributions, along with the observed distribution of posted prices (in their case, of mutual funds). Their approach is suitable to a market where each seller posts the same price for all buyers. Mortgage brokers tailor their proposals to individual borrowers, so the assumption fails in our case. In addition, we lack data on origination volume by broker.

The assumption of posted prices seems out of place in the mortgage origination market. The essence of a posted price is a commitment not to consider a counteroffer from a customer. Posted prices make sense in Internet commerce, studied by Hong and Shum (2006) and many others, and in regulated mutual funds, studied by Hortaçsu and Syverson (2004). Mortgage brokers, mostly working as individuals or in small groups, have no technology to commit to a stated price. If a borrower turns down an offer and makes a plausible case in favor of a lower origination charge, the broker may well make a better offer rather than risk losing the business. For a discussion of this issue in the labor market, see Hall and Krueger (2012).

Cebul et al. (2011) study employer-provided health insurance, another brokered market with negotiated prices. They find that search frictions increase insurance

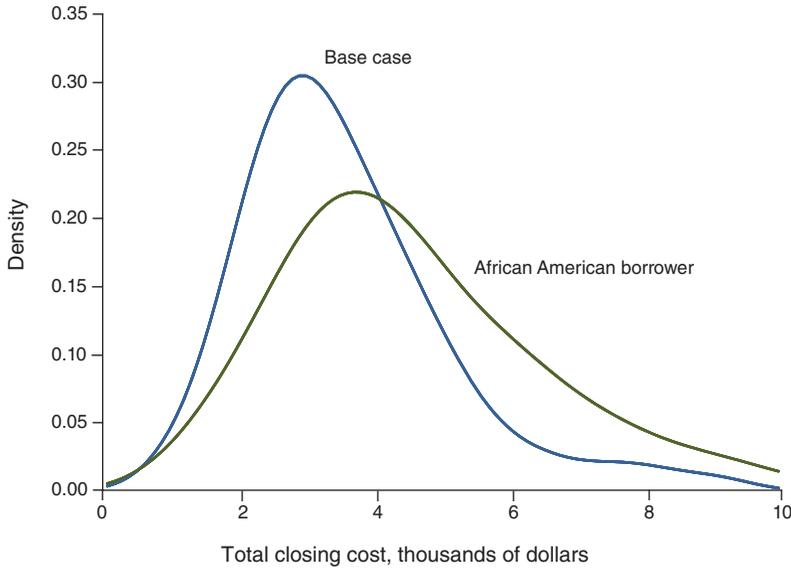


FIGURE 7. SMOOTHED DENSITIES FOR THE BASE CASE AND FOR THE CASE OF AN OTHERWISE SIMILAR AFRICAN AMERICAN BORROWER

premiums, increase policy turnover, and potentially undermine incentives for preventive care. Frictions may arise from buyer confusion, state-level limitations on entry, or adverse selection. The theoretical study closest to the problem that a borrower faces in the mortgage origination market is McAfee and McMillan (1988). That paper considers the optimal strategy for commercial procurement from suppliers with heterogeneous costs. The optimal mechanism induces potential suppliers to reveal their costs; the buyer continues to play the revelation game with additional potential suppliers until the benefit from adding another falls short of the cost of playing the game again.

We take a view that is a simplification of McAfee and McMillan (1988) suited to the case where the buyer is a borrower seeking origination services from a broker. Rather than inducing the revelation of cost, which would require a strong commitment to a sealed-bid second-price auction setup, we suppose that the borrower seeks initial quotes from N brokers, uses the quotes to extract better proposals, until the process ends with one quote that no other broker is willing to beat. This process is an English or Vickrey or second-price auction. Modeling of markets based on these auctions is made easier by the fact that bidding behavior is non-strategic. By standard auction principles, the best ultimate bid is the reservation value (cost) of the second-lowest-cost bidder. Because the winning bidder does not reveal cost, the bidder is protected from opportunistic action by the borrower based on a departure from the McAfee-McMillan commitment, by making a counteroffer to that bidder at a charge lower than the second-lowest but above the winner's cost. We also believe that borrowers who get the best deals on mortgage origination actually do their shopping the way that we model the process.

Notice that we are making an assumption of commitment by the borrower to decline to engage in further bargaining with the sole remaining bidder once the

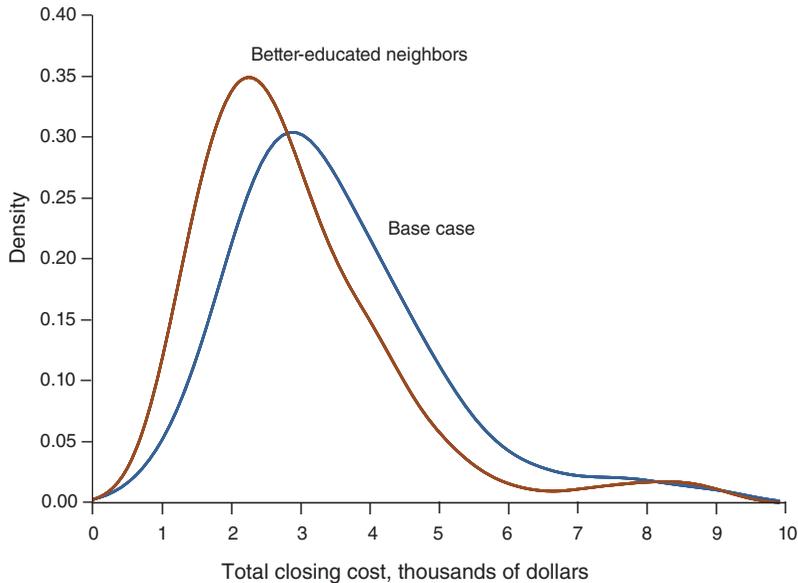


FIGURE 8. SMOOTHED DENSITIES FOR THE BASE CASE AND FOR THE CASE OF MORE EDUCATED NEIGHBORS

second-place broker has dropped out. The borrower accepts the last offer from the winner rather than making a counteroffer. We believe this assumption is generally realistic and it has the further advantage of leading to a simple and clean way to interpret the observed distribution of origination charges.

B. *Alternative Shopping Models*

Some surveys suggest that a fraction of borrowers consider only a single broker. Our analysis below confirms that the data on origination charges—especially the upper tail of remarkably high charges—show that borrowers forego opportunities to pay substantially lower charges by shopping more intensively. Thus one model that deserves attention has the borrower consulting only a single broker.

There is all the difference in the world between shopping from a single broker and from two. If the borrower thinks that the only way to buy a house is to finance it through the broker at hand, that broker has the power to prevent the borrower from buying the house and thereby can capture up to the total benefit that the borrower stands to gain from buying this house. For that matter, if the borrower continues to believe that there is, in effect, only one broker in the universe, the broker can capture up to the total benefit to the borrower from buying *any* house. The distribution of origination charges in that setting would depend on the benefit associated with the underlying house purchase. It would not reveal anything useful about the benefit of shopping from additional brokers. Though we believe that a fraction of borrowers shop from a single broker, we are unable to integrate that belief into our empirical work.

Another model of a borrower who may wind up shopping from only a single broker is that the borrower has some idea about the best terms available in the origination market and tries to make a bilateral bargain with the broker at hand, based on the

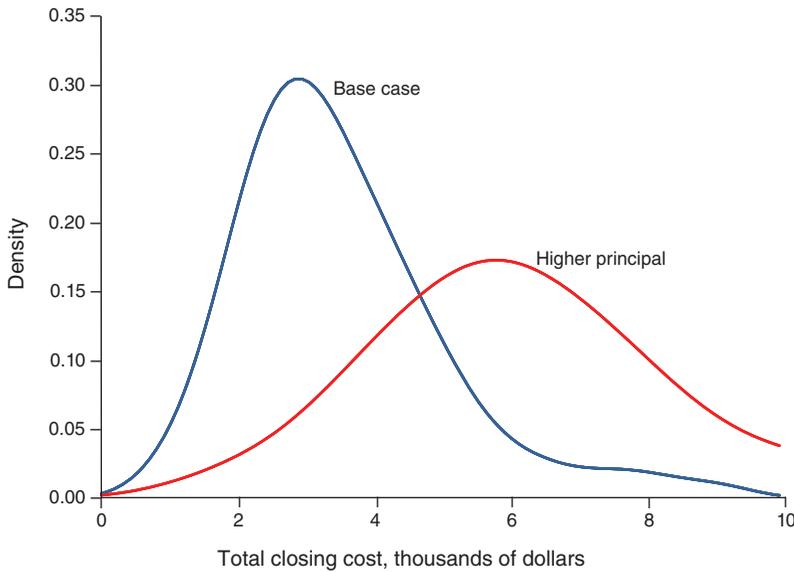


FIGURE 9. SMOOTHED DENSITIES FOR THE BASE CASE AND FOR THE CASE OF A HIGHER PRINCIPAL

threat to move on to another broker if this one fails to result in acceptable terms. This model of sequential bargaining would resemble the Diamond-Mortensen-Pissarides model of the labor market (Mortensen and Pissarides 1994). The model would imply a reservation value of the broker's cost. If the broker at hand had a cost below the reservation value, the broker and borrower would have a surplus—they would make a deal and the resulting price would split the surplus. If the cost was above the reservation cost, the borrower would move on to the next broker and repeat the process. The reservation value would balance the cost of search, a decreasing function of the reservation cost, against the disadvantage of a higher charge, an increasing function of the reservation cost. The observed distribution of the origination charge would be truncated above by the reservation cost. No broker with a cost above the reservation cost would get any business. But the data show no sign of any truncation. A small fraction of borrowers pay enormous origination charges. The data do not fit the sequential bargaining model. In addition, based on our conclusions about heterogeneity in broker's costs, most borrowers would shop from more than one broker unless they attributed extreme costs to moving on to another broker after encountering one with high cost.

C. Finding the Distribution of Broker Cost when Borrowers Shop from N Brokers

We proceed with our model in which the borrower invites N brokers to enter a bidding process that eventually gives the business to the broker with the lowest cost, with the borrower paying an origination charge equal to the cost of the second-lowest-cost bidder.

Our empirical analysis within this framework rests on a fundamental hypothesis—that the heterogeneity in the origination charge paid by borrowers with a given

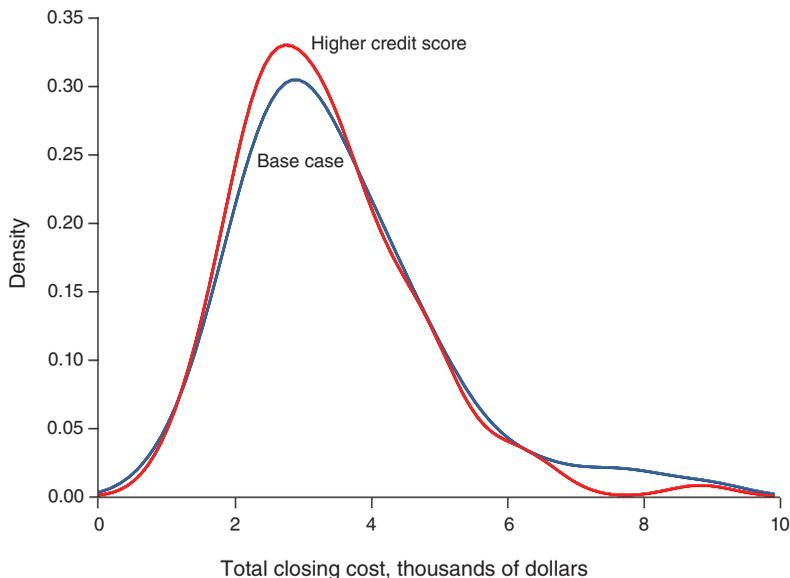


FIGURE 10. SMOOTHED DENSITIES FOR THE BASE CASE AND FOR THE CASE OF A HIGHER CREDIT SCORE

set of observed characteristics arises from the heterogeneity of costs of brokers and not from heterogeneity of the origination cost for that type of borrower. That is, we assume that the random variable $k_{i,j}$, broker j 's cost for originating a loan for borrower i , within a group defined by the observed characteristics, is independently and identically distributed across brokers. This assumption rules out a random effect among borrowers with the same observed characteristics that affects the costs of all brokers serving that type of borrower, because in that case, $k_{i,j}$ would be correlated across brokers j for a particular borrower i . To see the importance of this assumption, consider its opposite, where all heterogeneity among borrowers with a given set of characteristics arises from their idiosyncratic origination costs and all brokers have the same cost for any particular borrower. Then the borrower could run the auction described above between just two brokers. The outcome would be an origination charge equal to the common value of the two brokers' cost. If all borrowers shopped from two brokers, the market would be perfectly competitive. Under that assumption, the observed distribution of origination charges would be the same as the distribution of costs.

We believe that our assumption of no heterogeneity of origination cost among borrowers with the same observed characteristics is a reasonable approximation. First, we observe all of the major characteristics that appear likely to be determinants of cost. Second, what matters is the cost that the broker perceives at the time that he makes the deal with the borrower, not the actual realization of cost. It is true that brokers occasionally incur high costs lining up a loan for a borrower, but these costs are largely unpredictable. At the time that the broker makes the deal with the borrower, the broker knows little more than we do about the borrower's origination cost.

For brevity, our notation omits the dependence of our calculations on the observed characteristics, but we include a full set in our empirical work. We let $B(k)$ be the cumulative distribution of loan origination cost among brokers—the fraction of brokers whose cost is no greater than k . The origination charge τ is the cost level of the runner-up in the bidding for the business of a borrower. We let $H(\tau)$ be the cumulative distribution of the total origination charge τ among borrowers.

The relation between the two distributions is easy to derive. The probability that a random draw of broker's cost is greater than τ is $1 - B(\tau)$. The probability that none of N draws is as low as τ is $(1 - B(\tau))^N$. The probability that one of N draws is as low as τ is $NB(\tau)(1 - B(\tau))^{N-1}$. The probability that two draws or more are as low as τ (that is, the second-lowest draw is not greater than τ) is one minus these two probabilities:

$$(4) \quad H(\tau) = 1 - (1 - B(\tau))^N - NB(\tau)(1 - B(\tau))^{N-1}.$$

We find the B distribution in the following way: On a grid of values of B in $[0, 1]$, we evaluate the right-hand side of equation (4) and associate it with the value of τ where $H(\tau)$ is closest to that value. We take that value of B and $k = \tau$ to be a point on the distribution of broker cost, $B(k)$. The set of values of k that emerges from this process is a resampling of the original set of values of τ such that the values appear according to the distribution of broker cost rather than the original distribution of origination charges. Every borrower pays a charge that is equal to some broker's cost—that broker was the runner-up in the negotiation. That is why we relabel τ as k . We use a standard kernel-smoothing density estimator on the resample to find the smoothed density of broker cost. Because our earlier work resulted in a complete model of the full distribution of origination charges as a function of the observed characteristics, we can calculate the underlying distribution of broker cost for any combination of those characteristics.

Figure 11 shows the distributions of broker cost for the cases where borrowers shop from $N = 2, 3,$ and 4 brokers. Because more intensive shopping results in charges that are closer to cost, the distributions for higher numbers of brokers involve generally higher costs, so as to rationalize the same observed distribution of total charges. All of the distributions, including the one for the case of least intensive shopping from only two brokers, suggest that the implied level of cost is quite high for most mortgages and that the upper tail of cost includes implausible costs.

Figure 12 shows the implied distributions for the five cases in the case of shopping from only two brokers. The distribution for the high-principal loans lies far to the right—if it is true that borrowers only shop from two brokers, the cost of originating loans with principals of \$200,000 is far higher than that of smaller loans, a conclusion we find implausible. The alternative interpretation is that borrowers seeking bigger loans often shop from only one broker, who is able to capture some part of the surplus that the borrower enjoys from buying the house, which will generally be larger for more expensive houses. A borrower dealing with only one broker faces an outside option in the bargain of not buying the house at all, while the shopper who knows to go to at least one other broker has the outside option of dealing with that broker. The tendency for consumers to leave more on the table for large transactions has been widely reported in the literature on consumer behavior.

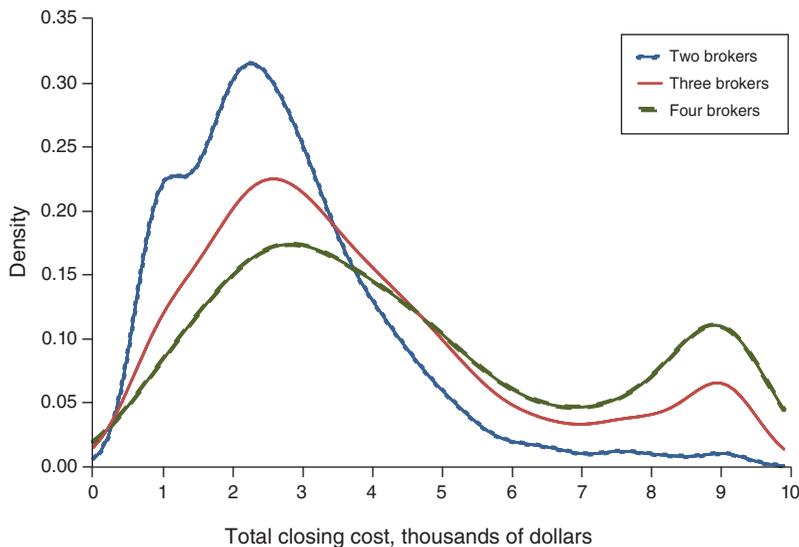


FIGURE 11. DISTRIBUTIONS OF BROKER COST INFERRED FROM THE HYPOTHESES THAT BORROWERS SHOP AMONG TWO, THREE, AND FOUR MORTGAGE BROKERS, BASE CASE

D. *Benefits that Borrowers Failed to Gain from more Intensive Shopping*

Our next step provides the main message of the paper—mortgage borrowers could save really a lot of money by shopping harder, from more brokers. Based on our results above on the distribution of cost among brokers, and on the survey evidence cited earlier in the paper, we conclude that few borrowers shop from more than two brokers and there are signs, such as the huge and implausible implied extra broker cost of high-principal mortgages, that borrowers often shop from only a single broker.

We should be candid about what is possible given the data. The most important limitation is that we have no empirical handle on the case where borrowers seek a quote from only a single broker. As we suggested above, the presence of these borrowers in our data appears to cause our estimates of the broker cost distribution to assign implausibly high costs to a small fraction of brokers.

If we had outside information about the distribution of costs among brokers, we could calculate the benefit of shopping from two, three, or four brokers (or even more), by solving for the distribution of second-lowest costs. Lacking such outside information, what we *can* do is to take the implied broker cost distribution for two-broker shopping and calculate the better distributions of total origination charges that a borrower would pay if she shopped from three or four brokers. Although this approach shows obvious signs of overstating the fraction of brokers with very high cost—thanks to the inclusion of some borrowers who shopped from only a single broker—most of the distribution is reasonable in our view. Of course, we do not believe that a large majority of borrowers literally shop by working a pair of brokers against each other. Rather, we believe that borrowers are sufficiently well-informed about mortgage pricing that they gain an outcome roughly comparable to shopping from two brokers.

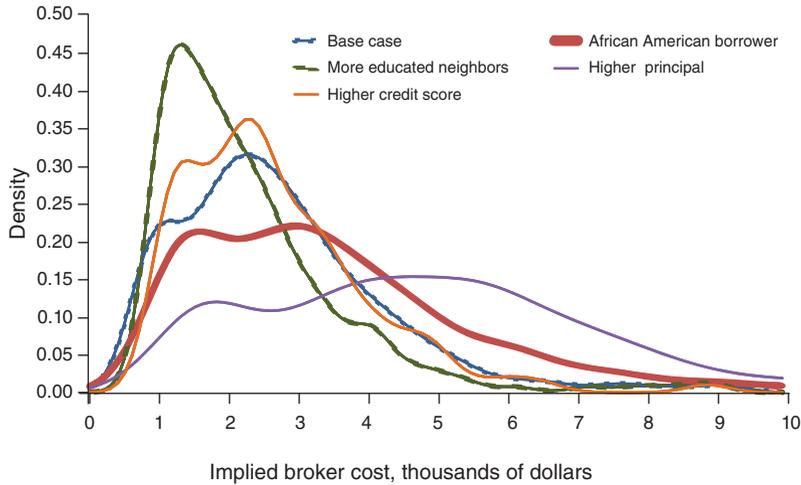


FIGURE 12. DISTRIBUTIONS OF BROKER COST INFERRED FROM THE HYPOTHESES THAT BORROWERS SHOP FROM TWO BROKERS, FOR THE FIVE CASES

E. Results

Figure 13 displays the distributions of the origination charges that a base-case borrower would pay, given the broker cost distribution inferred on the hypothesis that borrowers shop from only two brokers, for alternative shopping strategies involving shopping from two, three, and four brokers. The distributions shift substantially to the left for each added broker.

Table 5 shows the median gains that would be achieved from more intensive shopping, for the cases we have been considering. The gain from going to one additional broker ranges from \$836 for the white borrowers with smaller mortgages and high credit scores (who seem to shop more effectively than other groups, or who perhaps seem more savvy to brokers and consequently get better offers from them) to a colossal \$1,866 for the borrowers seeking a high-principal loan. These gains rise to \$1,197 and \$2,664 for adding a fourth broker. We conclude that borrowers are failing to use a simple method to obtain better deals on their origination charges. While it's possible that most people find the bargaining process so unpleasant that they knowingly overpay their brokers, we suspect that a lack of awareness of the advantage of more intensive shopping is a bigger part of the story. Brokers seem to have mastered the art of dissuading their customers from doing the kind of shopping that comes naturally for other expensive purchases.

Our approach to testing the hypothesis of consumer rationality has some of the character of a *reductio ad absurdum*. We posit a model that captures a natural and realistic solution to finding the best price in a market with negotiated prices. Then we estimate the underlying distribution of brokers' costs that would produce the observed distribution of prices borrowers pay. Given this distribution, we ask: How much would borrowers save by shopping from one more broker? The benefit is much too large to make sense. Some assumption fails. We find it implausible that the failing assumption is that buyers visit only a pair of brokers, because the implied distribution of broker cost in the presence of more intensive shopping has much too

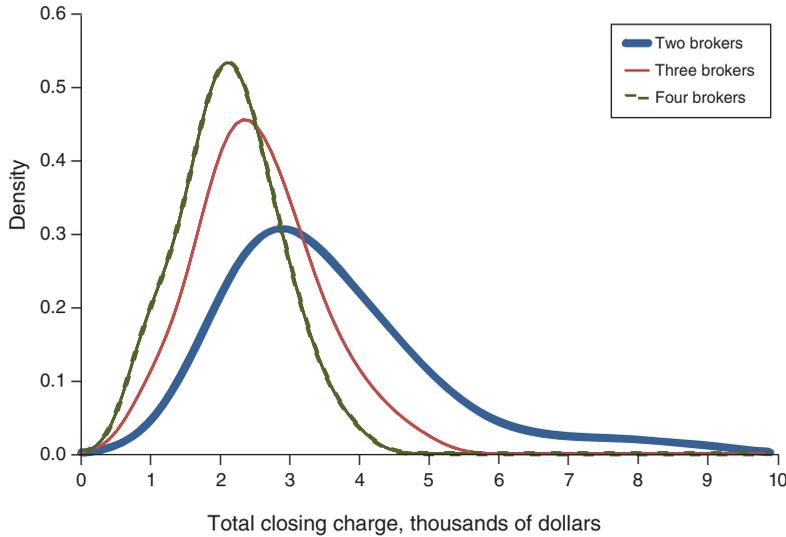


FIGURE 13. DISTRIBUTIONS OF BROKER CHARGES PAID BY A BASE-CASE BORROWER SHOPPING FROM TWO, THREE, AND FOUR BROKERS

large a fraction of brokers with huge costs. We conclude that the failing assumption is borrower understanding of the complex features of the transaction with a broker.

Many readers expressed discomfort with this approach, because we wind up rejecting the model that we use to reach the conclusion. We believe that the basic logic of rejection is sound because it is a proper application of the principle of *reductio ad absurdum*. Of course, our rejection of full-information rationality is less than dispositive. We have earlier emphasized the importance of our assumption that costs do not vary across borrowers with given observed characteristics. Our rejection also reflects our judgments about what distributions of broker cost are implausible and what levels of foregone gains from more intensive shopping are implausible.

We do not believe that borrowers are irrational in the sense that they are failing to shop harder even though they know how the market works, but rather that they do not understand the market institutions they encounter. Some borrowers probably think that the price for origination services is essentially fixed. Others believe that brokers are shopping for them. They likely do not know about and understand the incentives offered to brokers by wholesale lenders.

IV. The Division of the Origination Charge between Cash and the YSP

Earlier we noted that the function governing the YSP, as revealed in our data from actual rate sheets, has a shape that makes some borrower-broker pairs essentially indifferent between using the borrower's cash to pay the broker and borrowing the funds to pay the broker. Given that conclusion, we have studied the total payment. In this section, we explore the division of the total payment between cash and YSP.

We construct the variable $s = \frac{L}{L+Y}$ for each loan, the fraction of the total closing cost paid in cash. Thus a "no-cost" loan, where the broker receives only the YSP, has $s = 0$, while a loan at the par coupon rate, with zero YSP, has $s = 1$. Figure 14

TABLE 5—MEDIAN REDUCTION IN TOTAL ORIGATION CHARGE

	Case				
	Base	African American borrower	Highly educated neighbors	High principal	High credit score
Mean gain from shopping from one additional broker, dollars	981 (26)	1,270 (74)	901 (63)	1,866 (93)	836 (47)
Mean gain from shopping from two additional brokers, dollars	1,393 (36)	1,839 (107)	1,232 (85)	2,664 (125)	1,197 (68)

Note: Bootstrap standard errors in parentheses.

shows the distribution of s among the loans in the FHA sample. The modal loan's total closing cost is around half cash and half YSP. Loans with closing costs paid mostly in cash—the right side of the distribution—are rare. Loans with closing costs paid mostly through the YSP are not common, but constitute an important minority.

Our earlier discussion of the choice of the division between cash and YSP (amplified in the online Appendix) suggests that there is a mapping between two characteristics of the borrower—her discount rate and the time she expects to keep the loan—and the division. Those with higher discount rates and lower keeping times should opt for mostly YSP and vice versa. In that case, we can treat the observed value of the division, s , as a personal characteristic of the borrower. To pursue this idea, we divide the range of s (the unit interval, aside from a few with negative cash or negative YSP, which are included in the first and last bins) into ten equal bins and introduce the corresponding ten dummy variables into our earlier estimation framework. We remove the constant, which is the sum of the dummy variables.

Table 6 shows the results for four estimating equations, stated as the implied values of the total closing cost at the tenth, fiftieth, and ninetieth percentile points and at the mean. The bottom estimate is for the mean and is estimated by regression. In general, borrowers who opt to fund their total closing costs mainly with the YSP, at the left edge of the table, pay less in total closing costs. Those who opt to pay with cash alone, at the right edge, pay less than those in the middle if they are below or at the median, but not if they are high payers, in the ninetieth percentile.

The estimates for the tenth percentile describe the relationship between s and the total origination charge τ among borrowers whose total origination charges are quite low compared to the majority—that is, 90 percent of the borrowers pay more than this group. That point occurs at the very low level of \$1,366 in total origination charges for those who relied almost entirely on the YSP to pay those charges. The tenth percentile occurs at about double that level of total payment in the more popular case where the borrower pays 30 to 40 percent of the total in cash and the rest from the YSP. The tenth percentile falls back to lower levels among those who rely mostly on cash, though observations in that category are sparse and make the decline statistically ambiguous.

The results for the typical borrower, viewed as the median (second set of estimates) or mean (fourth set of estimates), also show a substantial advantage for those choosing not to pay in cash or to pay only a small fraction in cash. The advantage of paying less than 10 percent in cash relative to the total origination charge paid by

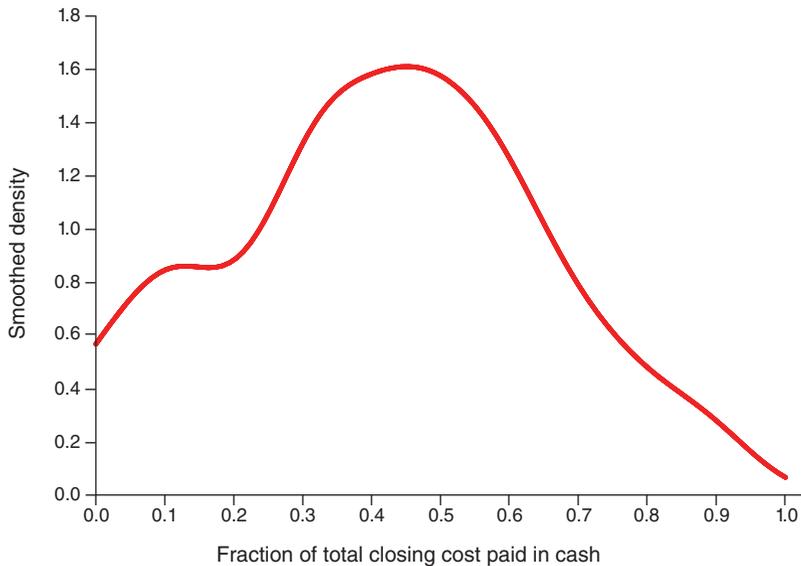


FIGURE 14. SMOOTHED DENSITY OF THE FRACTION OF TOTAL CLOSING COST PAID IN CASH

those using 30 to 40 percent cash is \$1,444 in the median and \$1,493 in the mean. But these borrowers with typical shopping prowess do not gain a similar advantage if they use all cash—at the median, they save around \$1,010 compared to borrowers who use a fairly even mix of cash and YSP.

The results for the least-successful shoppers, those at the ninetieth percentile for the total origination charge, show a huge advantage for the low-cash shopping strategy, though even the borrowers in the lowest-cash category pay a total origination charge at the high level of \$3,500, almost triple the level of the tenth percentile group. Unlike the other percentiles and the mean, in this group, the total closing payment keeps rising with the cash share up to 70 percent cash, where the total payment is \$8,460, an astronomical amount to pay for origination services for a \$100,000 mortgage.

We conclude that all borrowers who opt to rely on the YSP to fund their origination costs rather than paying a mixture of YSP and cash are likely to make better deals.

To interpret the findings in, we consider a number of hypotheses:

- *Baseline*: The borrower has a given s . The borrower and broker negotiate total origination charge τ that is uncorrelated with s . The two components of the payment to the broker are the cash charge $L = s\tau$ and the YSP $Y = (1 - s)\tau$.
- *High-discount disadvantage*: High-discount borrowers are more costly for the broker to serve and place lower value on the deferred payments to finance the yield-spread premium Y , so on both accounts the total origination charge τ is higher for the low- s borrowers who rely on the yield-spread premium to save themselves scarce current cash.
- *Marginal home purchasers*: Some borrowers are only barely able to manage the purchase of the house they are trying to finance. Not only are they short of cash to pay the broker, but they have hardly enough income to qualify for the loan. They need a low s because of their cash shortage, but cannot qualify for much of a yield-spread premium. Because they have a credible threat not to

TABLE 6—DIVISION BETWEEN CASH CHARGE AND THE YSP

Estimation	<i>Total Closing Cost in the Base Case, by Division between Cash and YSP, Dollars</i>									
	<i>Fraction of Total Closing Cost Paid in Cash</i>									
	0.0 to 0.1	0.1 to 0.2	0.2 to 0.3	0.3 to 0.4	0.4 to 0.5	0.5 to 0.6	0.6 to 0.7	0.7 to 0.8	0.8 to 0.9	0.9 to 1.0
Quantile, 10th percentile	1,366 (158)	1,825 (229)	2,116 (90)	2,636 (127)	2,369 (157)	2,120 (126)	2,082 (103)	1,884 (135)	1,682 (323)	1,063 (463)
Quantile, median	2,393 (86)	2,594 (76)	3,186 (95)	3,849 (86)	3,515 (85)	3,436 (117)	3,747 (354)	3,895 (387)	3,375 (353)	2,827 (424)
Quantile, 90th percentile	3,500 (163)	3,971 (286)	4,701 (266)	4,813 (103)	5,613 (329)	6,245 (229)	8,460 (621)	8,227 (879)	7,167 (723)	6,491 (1,051)
Regression, mean	2,237 (143)	2,501 (170)	3,257 (156)	3,730 (129)	3,786 (122)	3,827 (132)	4,432 (158)	4,595 (202)	3,787 (255)	3,392 (395)

Note: Bootstrap standard errors in parentheses.

complete the purchase, the broker is forced to offer a total origination charge τ close to his cost.

- *Single-dimension shopping advantage:* Borrowers who elect no-cost loans ($s = 0$) or those with low cash charges can shop for the lowest coupon rate without having to balance the cash charge against the coupon rate; they have no trouble determining the best deal among a set of proposed mortgages, as they can shop on rate alone. Similarly, those who elect no yield-spread premium and thus to pay the par coupon rate can shop for the lowest cash charge (if they can truly detect a YSP, a big if). A third possibility for no-cost loans is that setting the broker's expectations that no cash will be paid by the borrower at closing precludes the broker from nibbling by adding on additional fees, such as fax or courier fees or charges for preparing an amortization table. And a fourth possibility is that brokers are willing to accept lower origination charges for no-cost loans because they have higher coupon interest rates and thus are more likely to bring the borrower back to the broker to refinance.

The results refute the baseline hypothesis. Under that hypothesis, conditioning on the cash share s would not change the distribution of the total origination charge τ . Table 6 shows a strong tendency for the borrowers who contribute less cash, with low values of s , to pay substantially less by all four measures. The results plainly establish an interaction between s and the total origination charge.

The results also refute the hypothesis of a high-discount disadvantage. The view seems plausible that these borrowers are more costly to serve, because the same factors that cause them to have high discounts also means that they have trouble finding the needed documents, filling out applications, and performing the other steps in the application process. Brokers probably do more hand-holding for high-discount borrowers. By definition, these borrowers place a lower discounted value on the future coupon payments that fund the yield-spread premium, another factor that would help the broker extract a high total origination charge. Our earlier analysis showed that high-discount borrowers will opt for low- s loans. High-discount borrowers will more often lack the cash to pay the broker up front. All of these plausible

considerations point toward *higher* total broker charges for the low- s borrowers. But the evidence in Table 6 is that these borrowers pay thousands of dollars *less* than those with s around 0.5.

Table 6 supports the hypothesis that marginal home purchasers get better deals from their brokers. The marginal home purchaser can't pay the broker much in cash and can't qualify for a mortgage with a high enough coupon to support much of a yield-spread premium, so the broker has the choice between doing origination for only a small profit or not doing it at all. The marginal home purchaser gets a low-price origination without shopping hard.

Finally, the results in Table 6 support the hypothesis that consumers shop better when the price has a single dimension. This hypothesis helps explain why the borrowers with high values of s do somewhat better, as well as explaining the striking advantage of the more common single-dimension shopping strategy based on the no-cost loan. The borrowers with intermediate values of s seem to respond to the cash and YSP elements of the total charge as if they had nothing to do with each other—at the tenth percentile, they pay \$1,185 for each element at $s = 0.5$. Those who choose to pay only one component—either Y or L but not both—incur only half the total cost of those who pay with both. Borrowers behave as if they are unaware that incurring a higher YSP should entitle them to a lower cash payment.

We conclude that, among the four hypotheses, the data plainly support those relating to marginal house purchasers and single-dimension shopping strategies, while the data refute the baseline hypothesis (where s is irrelevant) and the hypothesis of a disadvantage for high-discount borrowers.

V. Concluding Remarks

Untrained, inexperienced borrowers interact with specialist mortgage brokers in the mortgage origination market. Brokers earn two kinds of compensation, explicit charges the borrower pays in cash and a commission the lender pays based on the spread between the coupon rate the borrower agrees to and the par mortgage interest rate. Both types of broker compensation seem to confuse borrowers. The wholesale lender's commission is determined by financial dynamics understood by a tiny group of professionals, and the rate sheet that summarizes the possible payments is never shown to borrowers. Even the cash fees are not called "broker's fees" or anything like that. Rather, it is a bewildering assortment of fees, each seeming to cover some aspect of origination, often including a charge for "points." Almost everybody, including regulators and many economists, seem to believe that points have a mechanical relation to a reduced interest rate, when in fact the term "points" is nothing more than another category for the broker's compensation.

When analyzed in the context of a minimal amount of shopping—shopping from a pair of brokers—the data suggest that most borrowers would benefit substantially by shopping from additional brokers. The data also suggest that simplifying shopping by considering only loans where the broker receives all his compensation from the lender and the borrower pays no cash results in substantially better terms for borrowers. The borrowers who receive the worst terms tend to pay both large cash charges and to agree to high interest rates that give the broker nearly equal additional compensation in the form of the yield-spread premium.

Although brokers tend to make large margins over cost in each origination, we do not believe that they earn incomes above those available from similar lines of work. Rather, they dissipate the margin with efforts to find borrowers. Equilibrium in the broker origination market appears to be inefficient. If borrowers spent more effort seeking offers from more brokers—which is easy—brokers would earn lower margins and would dissipate less effort trying to locate shy borrowers.

With respect to policy changes that might help achieve a more efficient equilibrium, we believe in evidence-based design. Disclosure law has historically been in the hands of lawyers, who designed dense forms that may help absolve their clients of blame for consumer error, but which did little to help consumers find better deals. A new movement to design disclosures that are proven to be helpful, through field experiments, may result in some progress. Whether these forms can overwhelm the persuasion of skilled expert salesmen remains to be seen. We are inclined to believe that simple admonitions, such as “mortgage brokers are salesmen and the only way to get a good deal is to shop and bargain” and “you are more likely to get a good deal if you shop for no-cost loans” are more likely to yield improvements than, for example, trying to teach borrowers enough financial economics to understand the tradeoff between cash and the interest rate.

REFERENCES

- Berndt, Antje, Burton Hollifield, and Patrik Sandås.** 2010. “The Role of Mortgage Brokers in the Subprime Crisis.” National Bureau of Economic Research Working Paper 16175.
- Burdett, Kenneth, and Kenneth L. Judd.** 1983. “Equilibrium Price Dispersion.” *Econometrica* 51 (4): 955–69.
- Cebul, Randall D., James B. Rebitzer, Lowell J. Taylor, and Mark E. Votruba.** 2011. “Unhealthy Insurance Markets: Search Frictions and the Cost and Quality of Health Insurance.” *American Economic Review* 101 (5): 1842–71.
- Federal Reserve Board.** 2009. “Design and Testing of Truth-in-Lending Disclosures for Closed-End Mortgages.” <http://www.federalreserve.gov/newsevents/press/bcreg/bcreg20081218a7.pdf>.
- Hall, Robert E., and Alan B. Krueger.** 2012. “Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search.” *American Economic Journal: Macroeconomics* 4 (4): 56–67.
- Hong, Han, and Matthew Shum.** 2006. “Using Price Distributions to Estimate Search Costs.” *RAND Journal of Economics* 37 (2): 257–75.
- Hortaçsu, Ali, and Chad Syverson.** 2004. “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds.” *Quarterly Journal of Economics* 119 (2): 403–56.
- Koenker, Roger, and Kevin F. Hallock.** 2001. “Quantile Regression.” *Journal of Economic Perspectives* 15 (4): 143–56.
- Lacko, James, and Janis Pappalardo.** 2007. “Improving Consumer Mortgage Disclosures.” Federal Trade Commission Bureau of Economics Staff Report. <http://www.ftc.gov/os/2007/06/P025505MortgageDisclosureReport.pdf>.
- McAfee, R. Preston, and John McMillan.** 1988. “Search Mechanisms.” *Journal of Economic Theory* 44 (1): 99–123.
- Moraga-González, José Luis.** 2006. “Estimation of Search Costs.” Unpublished.
- Mortensen, Dale T., and Christopher A. Pissarides.** 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies* 61 (3): 397–415.
- US Census Bureau.** 2001. “American Housing Survey.” http://www.census.gov/hhes/www/housing/ahs/ahs01_2000wts/tab315.html.
- Woodward, Susan.** 2008. A Study of Closing Costs for FHA Mortgages, (prepared for HUD by the Urban Institute). http://www.huduser.org/Publications/pdf/FHA_closing_cost.pdf.
- Woodward, Susan E., and Robert E. Hall.** 2012. “Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence: Dataset.” *American Economic Review*. <http://dx.doi.org/10.1257/aer.102.7.3249>.