A Simple Framework for Estimating Consumer Benefits from Regulating Hidden Fees

Sumit Agarwal, Souphala Chomsisengphet, Neale Mahoney, and

Johannes Stroebel

ABSTRACT

Policy makers are increasingly turning to regulation to reduce hidden or nonsalient fees. Yet the overall consumer benefits from these policies are uncertain because firms may increase other prices to offset lost fee revenue. We show that the extent to which firms offset reduced hidden-fee revenue is determined by a simple equation that combines two sufficient statistics, which can be estimated or calibrated in a wide range of settings: a parameter that captures the degree of market competitiveness and a parameter that captures the salience of the hidden fee. We provide corroborating evidence for this approach by drawing upon evidence on the effect of fee regulation under the 2009 Credit Card Accountability Responsibility and Disclosure Act. We also illustrate the applicability of our approach by using the framework to assess a hypothetical regulation of airline baggage fees.

1. INTRODUCTION

A growing body of research emphasizes the distortionary effects of hidden fees (Ellison 2005; Gabaix and Laibson 2006; Heidhues and Kőszegi

2010).[1] At the same time, policy makers are increasingly using regulation to reduce or eliminate these hidden or nonsalient charges. One focus of this regulation has been the market for consumer financial products. For example, in 2009, the Federal Reserve limited the ability of banks to charge overdraft fees at ATMs and on debit card transactions by requiring consumers to opt in to overdraft fee programs (74 Fed. Reg. 59033 [November 17, 2009]). In 2012, the U.S. Department of Labor (2012) issued regulation requiring more transparent disclosure of fees and expenses associated with 401(k)-type retirement plans. And, in 2011, a bill to ban fees on airline passengers' first checked bag was introduced in the U.S. Senate (*Times-Picayune* 2012). Yet policy efforts face skepticism that fee reductions will benefit consumers because firms will increase other prices to offset lost fee revenue.

This paper presents a simple framework for estimating the overall consumer cost savings from regulating hidden fees. We first present a simple model, adapted from Agarwal et al. (forthcoming), to achieve two goals. The model provides intuition for when limits to hidden fees will bring about an across-the-board reduction in costs to consumers and presents a simple equation that can be used to estimate the consumer benefits from a proposed fee regulation that combines two sufficient statistics, which can be estimated or calibrated in a wide range of settings.

This simple equation shows that the degree to which firms will offset a reduction in fee revenue through increasing other prices is determined by a parameter that captures the degree of competitiveness in the market and a parameter that captures the salience of the hidden fee. The degree of competitiveness can be measured by the pass-through rate—the fraction of a cost shock that firms pass through into prices. The salience parameter can be either measured using a survey or experiment or calibrated on the basis of an understanding of the institutional setting and estimates of salience in comparable settings in the literature.

We provide corroborating evidence for our proposed approach by analyzing a regulation of hidden fees by the 2009 Credit Card Account-

1. Ellison (2005) presents a model in which profits from difficult-to-observe add-ons, such as credit cards' late fees or unnecessary rental car insurance, are not competed away by lower prices on the base good or service, which leads to higher consumer costs. Gabaix and Laibson (2006, p. 505) argue that a "curse of debiasing" can explain why firms may not engage in advertising to inform consumers about their competitors' use of hidden fees.

ability Responsibility and Disclosure (CARD) Act (Pub. L. No. 111–24, 123 Stat. 1734). This exercise compares ex ante predictions that could have been constructed using our framework with the ex post estimates of the actual effects presented in Agarwal et al. (forthcoming), where we show that the CARD Act reduced credit card fees by 1.6 percent of average daily balances (or an aggregate $11.9 billion annually), without any offsetting increase in interest charges or a reduction in access to credit. Here we show that an analyst using our framework would have produced an accurate forecast of the size of the offset.

We also illustrate the applicability of our approach by using the framework to assess a hypothetical fee regulation that has not yet been implemented. Specifically, we assess the extent to which a regulation that limits airline baggage fees would reduce total travel costs to consumers rather than simply lead to higher airline ticket prices. The goal is to provide an example of how a researcher might put our framework into practice, without explicitly advocating for or against this type of regulation.

In addition to the literature on hidden fees, our modeling approach is related to the literature on pricing in aftermarkets, which analyzes markets in which the sale of a durable foremarket product, such as a printer, requires consumers to purchase supplies, such as ink cartridges, in an aftermarket (for example, Shapiro 1994; Carlton and Waldman 2010). Our model is most closely related to that of Farrell (2008), who considers the welfare impacts of making aftermarkets more competitive under different assumptions of the awareness of consumers regarding the costs in these aftermarkets.

2. A SIMPLE EQUATION FOR ESTIMATING CONSUMER BENEFITS

In this section we provide a simple model, adapted from Agarwal et al. (forthcoming), to illustrate that for limits on hidden fees to be offset it is sufficient for either markets to be perfectly competitive or fees to be perfectly salient. If markets are perfectly competitive, then aggregate prices inclusive of hidden fees will be forced down to marginal costs, and any regulation that reduces a certain fee will be offset with a similarly sized increase on another pricing dimension. If hidden fees are perfectly salient, then demand is responsive only to the total price. As a result, demand will not respond to an equal-sized reduction in one cost component and increase in another. If, however, markets are not perfectly

competitive and the fee is partially nonsalient, then regulators can be successful in lowering aggregate consumer costs.

2.1. Model Setup

There are n identical firms with identical constant marginal costs c. These firms compete to offer a product with an observable price p_1 (for example, an annual fee for a checking account) and a potentially hidden fee p_2 (for example, an overdraft fee). Since firms are identical, they charge the same prices in equilibrium. Aggregate demand is given by $q(p_1 + \psi p_2)$, where $\psi \in [0, 1]$ parameterizes the degree of salience of p_2. A value of $\psi = 1$ indicates that the price is perfectly observable; a value of $\psi = 0$ indicates that consumers are completely oblivious to the price. Following Heidhues, Kőszegi, and Murooka (2012), we assume that there is a maximum $\bar{p}_2$ that is determined by regulation or some other factor. Because demand is weakly less elastic in p_2, it is optimal for firms to set the potentially nonsalient price p_2 to the maximum allowable amount $\bar{p}_2$.[2]

The first-order condition for the observable price p_1 is given by

$$p_1 + p_2 - c = \theta \mu (p_1 + \psi p_2), \tag{1}$$

in which the markup of price over marginal cost is set equal to the product of a market competitiveness parameter $\theta \in [0, 1]$ and an absolute markup term $\mu(p_1 + \psi p_2)$. In particular, the degree of competition is decreasing in θ with perfect competition given by $\theta = 0$ and monopoly given by $\theta = 1$, and the absolute markup term $\mu(p_1 + \psi p_2)$ is equal to p_1 times the inverse elasticity of aggregate demand: $\mu(p_1 + \psi p_2) \equiv -(q/q') = p_1/\varepsilon_{p_1}$. The specification is quite general and nests a Cournot model of competition when $\theta = 1/n$, where n is the number of firms. See Mahoney and Weyl (2013), Weyl and Fabinger (2013), and Bresnahan (1989) for extended discussions of the microfoundations of this specification.

2.2. Pricing Offset

Next consider a regulation that decreases the maximum allowable hidden fee $\bar{p}_2$. We want to know how much of the decline in p_2 is offset by an increase in p_1. For small changes in p_2, this offset is given by $\omega \equiv$

---

2. To see this, suppose a firm sets a p_2 less than $\bar{p}_2$. For $\psi < 1$, the firm can increase profits by decreasing the salient price by ψdp_2 and increasing the nonsalient price by dp_2. This pricing change has no effect on demand because $q[p_1 - \psi dp_2 + \psi(p_2 + dp_2)] = q(p_1 + \psi p_2)$ but raises total profits by $(1 - \psi)dp_2 q(p_1 + \psi p_2) > 0$. This means that $p_2 < \bar{p}_2$ cannot be an equilibrium.

$-(dp_1/dp_2)$. We will say that there is a full offset if $\omega = 1$ and no offset if $\omega = 0$. In principle, the offset can be greater than full, with $\omega > 1$. Assume that θ and ψ are invariant to the price. Totally differentiating the first-order conditions (equation [1]) with respect to p_2 and rearranging yields

$$\omega = \frac{1 - \psi\theta\mu'}{1 - \theta\mu'}, \tag{2}$$

where we have suppressed the arguments of μ for notational simplicity. To gain intuition, consider some extreme cases. The offset is full ($\omega = 1$) when there is perfect competition ($\theta = 0$). Since competition drives price to marginal cost, any decrease in p_2 must be fully offset by an increase in p_1 to maintain zero markup in equilibrium. Similarly, the offset is full ($\omega = 1$) if p_2 is perfectly salient ($\psi = 1$). In this case, consumers view both prices as equivalent, and firms can maintain their desired level of demand by increasing p_1 one for one with the decline in p_2.

   Intuitively, the offset can be less than full when there is both imperfect competition and imperfect salience. In Agarwal et al. (forthcoming), we prove that, under reasonable conditions, the offset is indeed declining as the market becomes less competitive and as p_2 becomes less salient.

2.3. Offset Equation

The model also provides a simple equation that can be used to estimate the consumer benefits from a reduction in hidden fees in a wide range of settings. The degree of competitiveness can be measured by the pass-through rate—the fraction of a cost shock that firms pass through to higher prices. The salience parameter can be either measured using a survey or experiment or calibrated on the basis of an understanding of the institutional setting and estimates of salience in comparable settings in the literature.

   Let $\rho \equiv dp_1/dc$ denote the pass-through of a increase in marginal costs. Differentiating the first-order conditions with respect to c yields the well-known pass-through formula

$$\rho = \frac{1}{1 - \theta\mu'}, \tag{3}$$

where θ is the competition parameter and μ' is the derivative of the absolute markup term. There is one-for-one pass-through of a change in marginal costs ($\rho = 1$) if the market is perfectly competitive ($\theta = 0$).

The pass-through rate can be less than full when there is imperfect competition.

Simple algebraic manipulation of equation (2) and equation (3) allows us to rewrite the offset as a function of the pass-through rate and the salience parameter (for details, see Agarwal et al., forthcoming):

$$\omega = \rho + (1 - \rho)\psi. \tag{4}$$

When p_2 is fully hidden ($\psi = 0$), the offset is equal to the pass-through rate ($\omega = \rho$). Since consumers are completely oblivious to p_2, a reduction in revenue from this fee is exactly the same as an increase in marginal costs. When p_2 is fully observed ($\psi = 1$), the offset is 1. When p_2 is partially hidden ($0 < \psi < 1$), the offset is between ρ and 1. Since consumers, in a sense, observe part of the reduction in the hidden fee, the firm passes through some of this reduction like a cost shock and the other part like a reduction in a fully observable price.

3. CORROBORATING EVIDENCE FROM THE 2009 CREDIT CARD ACCOUNTABILITY RESPONSIBILITY AND DISCLOSURE ACT

The CARD Act placed significant restrictions on hidden fees, and critics of the bill argued that the reduction in fee revenue would be offset with higher interest charges and other fees.[3] Here we examine how an analyst might have used our framework to predict the impact of the CARD Act prior to its implementation. We discuss results from Agarwal et al. (forthcoming) and confirm that the analyst's prediction would have been correct.

3.1. An Ex Ante Estimate

The CARD Act had two main provisions aimed at reducing hidden fees. First, it required account holders to explicitly opt in to having their credit card company process rather than decline overlimit transactions and also imposed limits on the magnitude of overlimit fees that could be charged each billing cycle. Second, the CARD Act made it more difficult for banks to charge late fees by regulating the circumstances and the magnitude of permissible late fees. The overlimit provisions came

3. Response to Notice and Request for Information Regarding the Credit Card Market, letter from Nessa Eileen Feddis, senior vice president and deputy chief counsel for consumer protection and payments, American Bankers Association's Center for Regulatory Compliance, to the Consumer Financial Protection Bureau, February 19, 2013 (http://www.aba .com/Advocacy/commentletters/Documents/clCardAct2013Feb.pdf).

into effect February 22, 2010. The late-fee provisions came into effect February 22, 2010, and August 22, 2010.

Suppose an analyst was given the task of predicting the degree of offset of CARD Act fee restrictions as the bill was being debated in Congress in 2008. In Section 2, we showed that the size of the offset is determined by a parameter that captures the degree of competitiveness in the market and a parameter that captures the salience of the hidden fee. Specifically, the offset ω is given by the equation

$$\omega = \rho + (1 - \rho)\psi, \tag{5}$$

where ρ is the pass-through rate (the fraction of a cost shock that firms pass through in higher prices) and ψ is the salience of the hidden fee.

In a classic paper, Ausubel (1991) examines the time-series correlation between costs and interest rates in the credit card market. He finds that interest rates are extremely sticky, with credit card issuers passing through essentially none of the large changes in the cost of funds he observes over the time period he analyzes. While a more detailed analysis would want to consider whether changes in the cost of funds were passed through into other prices, such as interchange fees charged to merchants, choosing $\rho \approx 0$ seems like a reasonable estimate for this industry.

Similarly, there was significant evidence that late fees and overlimit fees were not salient to consumers (Sunstein 2006; Bar-Gill and Warren 2008; Mullainathan, Barr, and Shafir 2009). For example, Peter Davidson, executive vice president at Speer & Associates, argued that penalty fees are a "good source of revenue" because the industry perceives that "there [are] very few cardholders that switch cards because the late fee is too high" (Lazarony 1998). On the basis of this evidence, the analyst might have chosen $\psi \approx .1$ as a sensible approximation.

Combining these estimates in equation (6), the analyst would estimate that for every dollar in fee reduction, credit card issuers will increase prices by about 10 cents:

$$\omega = \rho + (1 - \rho)\psi = 0 + (1 - 0).1 = .1. \tag{6}$$

3.2. Corroborating Evidence

To examine whether this prediction would have been correct, we turn to an ex post analysis of the CARD Act fee restrictions from Agarwal et al. (forthcoming), which examines the effect of these fee regulations and the potential offsetting response using the Credit Card Metrics data set assembled by the U.S. Office of the Comptroller of the Currency

Figure 1. Credit card fees by FICO score

(OCC). The data set has account-level information on a near universe of credit card accounts at the eight largest U.S. banks.

Key results from the paper are conveyed in Figures 1 and 2.[4] Figure 1 plots monthly average fee payments for consumer credit cards as an annualized percentage of average daily balances (ADB) separately by the FICO score of the account when the card was originated. Consumers with FICO scores below 620 (roughly the bottom 20 percent of the distribution) experience an overall reduction in fees from 23 percent to about 9 percent of ADB. Figure 2 shows annualized interest charges as a percentage of ADB for account holders of different FICO scores at origination. Consistent with the findings in Ausubel (1991), credit card interest rates are very sticky and do not increase noticeably around the two CARD Act implementation dates or at all over the sample period. More surprisingly, we see no differential trend in the interest charges paid by low-FICO-score borrowers (for whom banks saw a very significant decline in fee revenue) relative to high-FICO-score borrowers, who

4. Values are for all reported accounts and are weighted by average daily balance. Vertical lines are plotted for February 2010 and August 2010, the two key implementation dates of the Credit Card Accountability Responsibility and Disclosure Act. See Agarwal et al. (forthcoming) for more details.
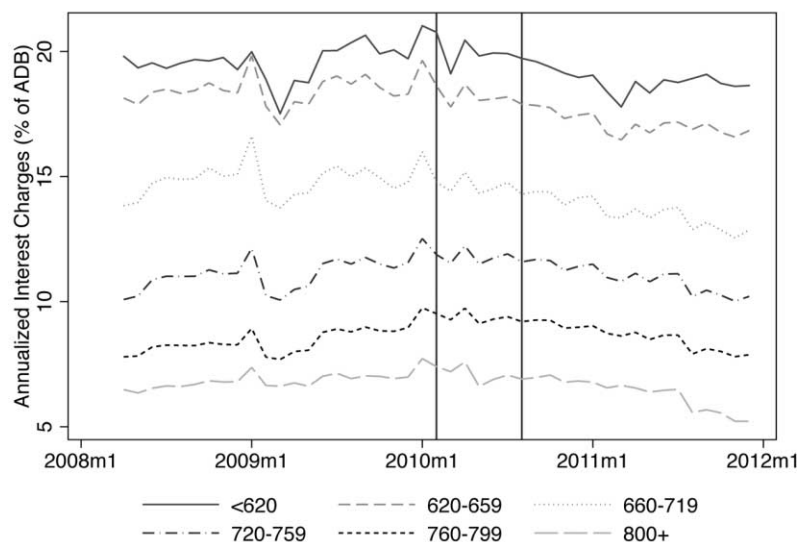
Figure 2. Credit card interest charges by FICO score

did not pay high fees in the first place. Agarwal et al. (forthcoming) presents a formal difference-in-differences analysis that confirms this finding of no offset. In that paper, we compare outcomes for fees, interest charges, and credit volume for consumer credit cards, which were affected by the CARD Act, with small business credit cards, which were not regulated by the CARD Act. The point estimate for the offset is $-.08$, and we can rule out an offset of greater than .57 with 95 percent confidence.

The evidence from the CARD Act corroborates our approach to estimating offsets. Combining the simple equation with estimates from the literature, an analyst would have predicted an offset of $\omega = .1$. And an ex post analysis from Agarwal et al. (forthcoming) confirms that the offset from the CARD Act fee reductions was indeed very small.

4. A HYPOTHETICAL EXAMPLE OF REGULATING AIRLINE BAGGAGE FEES

To illustrate the applicability of our approach, we use our framework to assess the consumer benefits from a hypothetical regulation of airline baggage fees. In 2008 and 2009, airlines collected $3.9 billion in checked-bag fees (*Times-Picayune* 2011). Policy makers are concerned that these fees are not salient to consumers at the point of sale, and in

Figure 3. Pass-through of jet fuel costs

2011 a bill was introduced in the U.S. Senate to prevent airlines from charging fees on the first checked bag. However, observers of the Senate bill were skeptical as to its potential for success, since "[t]he airlines are just going to find some other way to make it up by charging us for something else" (Halsey 2011, p. A14).

To assess the consumer benefits from such a policy using our framework, we require estimates of the pass-through rate of a cost shock (ρ) and the salience of baggage fees at the point of sale (ψ). To estimate ρ, we examine the extent to which changes in the price of jet fuel are passed through into ticket prices. Figure 3 shows quarterly average route-adjusted ticket prices (left axis) and quarterly average jet fuel costs per ticket (right axis), where both the left and right axes are adjusted to have a range of $100 for comparability.[5] Ticket prices are constructed using the Bureau of Transportation Statistics' Full-Scope Air Travel Price Index, which measures changes in prices for identical routes and classes of services relative to the first-quarter 1995 base period. Average jet fuel

5. The price index is multiplied by the inflation-adjusted average airfare in 1995 to convert the index to dollar values. Values are inflation adjusted to 2012 using the Bureau of Labor Statistics' Consumer Price Index for All Urban Consumers.

Table 1. Regressions of Pass-Through of Jet Fuel Costs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Average jet fuel cost								
per ticket	1.49	.95	.84	.99	1.49	.95	.84	1.00
	(.11)	(.15)	(.14)	(.15)	(.12)	(.15)	(.14)	(.15)
Time trend:								
Linear	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Square	No	No	Yes	Yes	No	No	No	Yes
Cubic	No	No	No	Yes	No	No	No	Yes
Quarter fixed effects	No	No	No	No	Yes	Yes	Yes	Yes
R^2	.73	.80	.84	.85	.73	.81	.84	.86

Note. Values are estimates from ordinary least squares regressions of quarterly average route-adjusted ticket prices on and quarterly average jet fuel costs per ticket. The dependent variable is average route-adjusted ticket price. Standard errors are in parentheses. $N = 65$.

costs per ticket are calculated by multiplying the average price of jet fuel per gallon (from the U.S. Gulf Coast Kerosene-Type Jet Fuel Spot Price data series from the U.S. Energy Information Administration) and the average number of gallons per passenger (using data from the Bureau of Transportation Statistics on total jet fuel consumption and number of passengers for U.S. carriers with no less than $20 million in annual revenue). The plot shows a strong time-series correlation between ticket prices and jet fuel costs.

Table 1 shows the corresponding regressions of ticket prices on jet fuel costs that allow us to recover the pass-through coefficient ρ. We control for trends in ticket prices (for example, from industry consolidation) with a polynomial in time and for seasonality (for example, greater demand during the winter holidays) with quarter-of-the-year fixed effects. In the specifications without controls for time, we estimate a coefficient on average jet fuel greater than $\rho = 1$. This is likely due to broad increases in ticket and fuel prices over the time period. In specifications with controls for time trends, the coefficient on fuel costs is approximately $\rho \approx 1$ and reasonably stable across specifications.

The estimate that airlines fully pass through fuel costs implies, using our framework, that regulation of baggage fees will also be fully offset by higher ticket prices. Plugging $\rho = 1$ into the offset equation (4) yields $\omega = 1$ for all values of ψ. However, our estimates do not allow us to rule out modest cost savings for consumers. In our preferred specification with the full set of controls, we can only rule out pass-through below $\rho = .71$ with a 95 percent confidence interval, and more sophisticated empirical strategies, which account for macroeconomic variables or

changes in industry structure, could in principle lead to even lower estimates.

Moreover, even if regulation of higher baggage fees is fully passed through into ticket prices, regulation of baggage fees might still have important benefits. Regulation of baggage fees might transfer surplus from consumers who carry on their baggage to consumers who check bags, and a reduction in carry-on bags might have positive externalities by freeing up overhead space for consumers who have a higher demand to carry on luggage and by reducing the number of bags going through the carry-on line at airport security screenings.

5. CONCLUSION

In many settings, policy makers are contemplating the regulation of hidden fees. Any benefit-cost analysis of such regulation requires an understanding of the extent to which firms offset the decline in fee revenue on other dimensions of pricing.[6] In this paper, we provide a simple framework to forecast the potential offset that combines an estimate of firms' pass-through of cost shocks and an estimate of the salience of the regulated hidden fee.

We provide corroborating evidence for our approach by showing that an ex ante analysis of the CARD Act's regulation of late fees and over-limit fees would have predicted the low degree of offset found in Agarwal et al. (forthcoming). We also illustrate the applicability of our approach by examining the extent to which limits on airline baggage fees would reduce total travel costs for consumers rather than simply lead to higher airline ticket prices. While more work is needed to tie down the exact pass-through and salience parameters, our estimates suggest a full offset of baggage fee regulation.

REFERENCES

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. Forthcoming. Regulating Consumer Financial Products: Evidence from Credit Cards. *Quarterly Journal of Economics*.

6. Other important ingredients for benefit-cost analyses include the relative valuation of producer and consumer surplus, in particular for the analysis of regulations such as the CARD Act, which precipitated a large shift of surplus from producers to consumers. In this paper we remain agnostic about how these should be traded off.

Ausubel, Lawrence M. 1991. The Failure of Competition in the Credit Card Market. *American Economic Review* 81(1):50–81.

Bar-Gill, Oren, and Elizabeth Warren. 2008. Making Credit Safer. *University of Pennsylvania Law Review* 157(1):1–101.

Bresnahan, Timothy F. 1989. Empirical Studies of Industries with Market Power. Pp. 1011–57 in *Handbook of Industrial Organization,* vol. 2, edited by Richard Schmalensee and Robert D. Willig. Amsterdam: Elsevier.

Carlton, Dennis W., and Michael Waldman. 2010. Competition, Monopoly, and Aftermarkets. *Journal of Law, Economics, and Organization* 26(1):54–91.

Ellison, Glenn. 2005. A Model of Add-On Pricing. *Quarterly Journal of Economics* 120:585–637.

Farrell, Joseph. 2008. Some Welfare Analytics of Aftermarkets. Working paper. University of California, Department of Economics, Berkeley.

Gabaix, Xavier, and David Laibson. 2006. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quarterly Journal of Economics* 121:505–40.

Halsey, Ashley, III. 2011. Bill Targets Airline Fees for Checked Luggage. *Washington Post,* November 21. http://articles.washingtonpost.com/2011–11–21/local/35281151_1_check-one-bag-bag-fees-carry-on-bags.

Heidhues, Paul, and Botond Kőszegi. 2010. Exploiting Naivete about Self-Control in the Credit Market. *American Economic Review* 100:2279–2303.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka. 2012. Exploitive Innovation. Working paper. University of California, Department of Economics, Berkeley.

Lazarony, Lucy. 1998. Penalty Fees Are Rising as Credit Card Companies Fight to Keep Their Profits. *Bankrate.com,* August 17. http://www.bankrate.com/brm/news/cc/19980817.asp.

Mahoney, Neale, and E. Glen Weyl. 2013. Imperfect Competition in Selection Markets. Working paper. University of Chicago, Booth School of Business, Chicago.

Mullainathan, Sendhil, Michael Barr, and Eldar Shafir. 2009. The Case for Behaviorally Informed Regulation. Pp. 25–62 in *New Perspectives on Regulation,* edited by David Moss and John Cisternino. Cambridge, Mass.: The Tobin Project.

Shapiro, Carl. 1995. Aftermarkets and Consumer Welfare: Making Sense of *Kodak*. *Antitrust Law Journal* 63:483–511.

Sunstein, Cass R. 2006. Boundedly Rational Borrowing. *University of Chicago Law Review* 73:249–70.

Times-Picayune. 2012. Sen. Mary Landrieu Wants to Ban Fees for Fliers' First Checked Bag. November 26. http://www.nola.com/politics/index.ssf/2011/11/sen_mary_landrieu_wants_to_ban.html.

U.S. Department of Labor. 2012. Final Rule to Improve Transparency of Fees

and Expenses to Workers in 401(k)-Type Retirement Plans. Fact sheet. February. http://www.dol.gov/ebsa/pdf/fsparticipantfeerule.pdf.

Weyl, E. Glen, and Michal Fabinger. 2013. Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy* 121(3):528–83.