CAN PRIVACY BE JUST ANOTHER GOOD?

JOSEPH FARRELL*

I. PRIVACY AS FINAL GOOD AND INTERMEDIATE GOOD 251
II. TREATING PRIVACY AS A FINAL GOOD WITH CONSUMER SOVEREIGNTY: INITIAL AGREEMENTS 253
A. How Initial Agreements Can Work Well 254
B. How Initial Agreements Can Fail 256
III. MODIFYING POLICIES LATER 259
A. Communication and Negotiation 260
B. Direction of Payments and the Role of Content and Advertising 260
CONCLUSION 261

I. PRIVACY AS FINAL GOOD AND INTERMEDIATE GOOD

Consumers care about privacy in part for its own sake: many of us at least sometimes feel it's just icky to be watched and tracked. Some consumers care about this more than others do, and it's ickier in some contexts than in others. Some consumers, and most consumers some of the time, don't care at all; others care a lot. In its preliminary report on protecting consumer privacy, the FTC noted that for some consumers, privacy related harms include "the fear of being monitored or simply having private information 'out there.'"[1] The FTC noted that "consumer

* Professor of Economics, University of California, Berkeley. This article is based on a December 2011 speech at the Silicon Flatirons "The Economics of Privacy" Symposium, during which I was Director of the Federal Trade Commission's Bureau of Economics (as of the publication of this article no longer serve in that role). This article was also based, to a lesser extent, on a March 2012 speech at the payments conference at the Federal Reserve Bank of Kansas City. It represents my views, and not those of the Federal Trade Commission or any individual Commissioner. I thank Ana Jaroszewicz, Paul Pautler, Louis Silversin, Douglas Smith, and especially Alessandro Acquisti and Daniel O'Brien, for helpful comments, but they too are not responsible for my views.

1. FED. TRADE COMM'N, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE: A PROPOSED FRAMEWORK FOR BUSINESSES AND POLICYMAKERS, PRELIMINARY FTC STAFF REPORT 20 (2010) [hereinafter FTC PROPOSED FRAMEWORK], *available at* http://www ftc.gov/os/2010/12/101201privacyreport.pdf.

surveys have shown that a majority of consumers are uncomfortable with being tracked online, although the surveys provide little or no information about the degree of such discomfort . . . ."[2]

Consumers also care about privacy in a more instrumental way. For instance, loss of privacy could identify a consumer as having a high willingness to pay for something, which can lead to being charged higher prices if the competitive and other conditions for price discrimination are present. Some kinds of information leakage could affect the availability or price of employment, credit, or insurance; or could contribute to risks of identity theft or fraud. The FTC discussed these and other harms that may result from "the unexpected revelation of previously private information" in its recently released recommendations for protecting consumer privacy.[3]

Economists sometimes refer to a good valued for its own sake as a *final good*, and to the instrumental kind as an *intermediate good*. Mainstream economics, and mainstream public policy in market economies, often treat these somewhat differently. I am suggesting that privacy has elements of both; so we should expect to draw on both kinds of policies.[4]

When it takes expertise to judge the impact of an intermediate good on final consumer goods, the former are often provided in a relatively regulatory way. Consider for instance airline safety. Most economists would agree that it should be consumer preferences that underlie the tradeoffs made among cost, convenience, and safety; but specific concrete choices, such as whether to take off in a thunderstorm or how much training a commercial pilot must have, are not left to unguided consumer choice. Rather, experts bring together: (a) their knowledge of how the intermediate goods such as the training of the pilot affect the final goods such as safety that consumers care about, with (b) reasonable estimates of the relative value that consumers place on safety versus convenience versus cost. This is essentially the field of benefit-cost analysis. In formulating rules or policies, or in assessing harm for liability, this approach may push toward quantifying how much consumers care and why. In privacy policy, there is a good deal of debate about those questions.

---

2. *Id.* at 29.

3. FED. TRADE COMM'N, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE: RECOMMENDATIONS FOR BUSINESSES AND POLICYMAKERS, REPORT 8 (2012) [hereinafter FTC RECOMMENDATIONS), *available at* http://www ftc.gov/os/2012/03/120326privacyreport.pdf.

4. A third class, of which privacy also partakes, is intermediate goods where the link to final goods is not usefully illuminated by available expertise. If nobody knows how likely it is that revealing your birthday on Facebook might lead to identity theft, it is not clear whether the analogy to expert-mediated intermediate goods or to final goods is more helpful.

For final goods, in contrast, we seldom focus on those questions. Rather, mainstream economics firmly declines to try to explain consumer tastes: the famous Latin tag is *de gustibus non est disputandum* (don't debate tastes). Thus it's jarring for an (or this) economist to hear the notion that economics pushes public policy on privacy towards focusing on quantifiable, tangible, and verifiable specific harms from the loss of privacy, a notion that is also reflected in some court cases.[5] Economics sometimes views intermediate goods that way, but for final goods, it normally takes tastes as given and asks how well a market or an economic system satisfies those tastes. In this article, I very briefly explore some issues in the market provision of privacy as a pure final good.

II. TREATING PRIVACY AS A FINAL GOOD WITH CONSUMER SOVEREIGNTY: INITIAL AGREEMENTS

As with any contract in a complex and shifting environment, good performance requires attention both to initial design and to setting up a framework for subsequent improvements as circumstances change or become known. This section discusses the first of these: agreement on the initial design of a relationship between a consumer and a firm that raises privacy questions. By that phrase I don't mean anything necessarily sinister, but only that the firm will face opportunities to re-purpose information that it's acquired from the consumer; and that it's not obvious which such opportunities should be pursued. In other words, there is a privacy policy choice worth thinking about.

Although there could be a variety of setups, let's think about a firm selling a book to a consumer, and analyze choice of the firm's privacy policy governing later re-purposing of the consumer's information. Consider two possible such privacy policies, A and B. Is the firm's most profitable choice of its privacy policy consistent with consumer preferences and total welfare?[6]

Think of the firm's potential privacy policy A as a benchmark, and introduce two quantities, V and H, that describe its alternative policy B

5. *See, e.g.*, Sasha Romanosky & Alessandro Acquisti, *Privacy Costs and Personal Data Protection: Economic and Legal Perspectives*, 24 BERKELEY TECH. L.J. 1061 (2009).

6. Isolated direct contracting on the use of consumer information can also work well in broadly similar circumstances, a fact that I have found is more transparent for many observers than the problem discussed in the text. *See* Hal R. Varian, *Economic Aspects of Personal Privacy*, INTERNET POLICY AND ECONOMICS 101 (William H. Lehr & Lorenzo Maria Pupillo eds., 2009), *available at* http://people.ischool.berkeley.edu/~hal/Papers/privacy/; *see also* Eli M. Noam, *Privacy and Self-Regulation: Markets for Electronic Privacy*, *in*, U.S. DEP'T OF COMMERCE, PRIVACY AND SELF-REGULATION IN THE INFORMATION AGE (1997), *available at* http://www ntia.doc.gov/reports/privacy/selfreg1 htm#1B.

relative to the benchmark of policy A. The first quantity, V, is the *additional* profits, per customer, that the firm can derive from using or selling the re-purposed information if policy B is adopted instead of policy A. The second quantity, H, is the incremental "harm" to the consumer if policy B were substituted for policy A and if the price charged for the book remained the same. Either or both of these quantities could be negative, but for concreteness let's talk as if both are positive: that is, policy B is less restrictive (it gives the firm more rights to re-purpose the consumer's information) than policy A.

If V is positive, why wouldn't the firm always insist on policy B? The answer is that H may affect consumer demand for the book. A book, bundled with privacy policy B, is a less attractive good than the book bundled with privacy policy A. If consumers pay attention to that difference, each consumer will be willing to pay H less for the former than for the latter. This demand shift based on H works against the direct profit effect based on V. How do they balance out? Remarkably, we can give a strikingly simple and optimistic answer to that question—but one that rests on strong assumptions. The argument is illustrated in Figures 1 and 2.

A. How Initial Agreements Can Work Well

Suppose, to illustrate, that V is $3 and H is $4. For a profit-seeking firm, the prospect of $3 per customer in follow-on revenues has the same impact on business decisions as a $3 per customer reduction in costs.[7] Thus if, for instance, the marginal cost of the book is $20, the firm's profits are given by (p - $20)*Q + $3*Q, where p is the price of the book and Q is demand; this is arithmetically the same as (p - $20 + $3)*Q, which is the profit function that would arise if costs fell by $3, replacing the $20 marginal cost by $17.[8]

Moreover, if customers see clearly before buying the book that privacy policy B is going to apply rather than policy A, then the demand curve shifts uniformly down by $4. That is, for any price p, the number of consumers that value the book, with privacy policy B, at $p or more, is the same as the number of consumers that value the book, with privacy policy A, at $(p + 4) or more. One might write this as $Q_B(p) = Q_A(p+4)$.

We can re-state the $3 downward shift in cost plus the $4 downward

7. It yields the same answers to the key questions such as: Do I want to be in this business? How do I want to price the initial purchase? How many customers do I want to deal with?

8. This simple form of profit function is not necessary for the very general statement in the text to hold.

shift in demand in another way: as a $4 downward shift in both, followed by a $1 upward shift in costs. That may seem an odd re-framing, but we will see it pay analytical dividends.

Consider first, then, the effects of shifting both costs and demand down by $4. Quite generally—and with no assumptions about the state of competition, about the shape of the cost curve or the demand curve—if both the cost curve and the demand curve shift vertically down by $4, the profit maximizing price for the book will go down by $4. To see this very general result, which is well known in the economics of public finance, suppose that the demand curve under privacy policy A is given by the function Q(p), where p is the price of the book. Under policy A, p will be chosen to maximize (p - $20)*Q(p), where $20 is the ordinary marginal cost of the book. Under policy B, the price of the book, p', will be chosen to maximize (p' - $20 + $4)*Q(p' + $4). If we simply rewrite this with a new variable z, defined as p' + $4 (intuitively we can think of z as the full cost to the consumer, including both price p' and non-price cost $4), we get (z - $20)*Q(z). Comparing this to the formula giving profits as a function of price p under policy A, it follows that z will be the same as p. Since z was defined as p' + $4, it follows that p', the price with the $4 lower cost curve and the $4 lower demand curve, is exactly $4 lower than p. Moreover, with the price, the cost, and the demand all shifting down by the $4, it follows that the quantity traded, profits, and consumer welfare, are exactly what they were without the shift.

In short, the $4 downward shifts in both costs and demand together netted out to a complete non-event. So then, the combination of the $4 downward shift in demand, and the $3 downward shift in costs, amounts to just a net $1 (i.e. $4 - $3) increase in costs. Such a $1 cost increase obviously hurts the firm, and almost certainly hurts consumers, too—if a firm's incremental costs go up, the firm will normally raise its price.

Remarkably, then, with the assumptions we made, the firm's and the consumer's preferences about the choice between policies A and B are completely aligned. If, as we assumed to illustrate, H was bigger ($4) than V ($3), then both firm and consumer prefer policy A. Intuitively, one might say that the consumer prefers A because he doesn't suffer the incremental harm H from the less stringent privacy protection; and the firm prefers A because consumers make that preference felt through their demand response. On the other hand, if H were less than V, our analysis would say that a change from policy A to policy B involves a downward shift in both demand and costs equal to H, combined with a further downward shift in costs equal to (V – H); the uniform shifts in both curves affect no quantities or levels of welfare, while the further downward shifts in costs makes both firm and consumer better off.

One might reasonably ask whether this logic depends on the firm

facing a competitive environment. Technically the answer is no: a downward shift in demand hurts a monopolist just as sketched above. This is much the same logic as the "one monopoly rent theorem" in antitrust. However, one might reasonably be more confident that harm is unlikely if there is also choice (beyond take-it-or-leave-it) for consumers. Moreover, more technically, if firms supplying "necessities" price on inelastic portions of their residual demand curve, then the argument doesn't work. However, it is perhaps more interesting to examine the information and contracting assumptions, as I briefly do next.

B. How Initial Agreements Can Fail

Any economic model has assumptions that are not strictly accurate. Such models can nevertheless provide a lot of insight, if one can understand the logic and see how things are working. Here, the key point is that if the firm adopts policy B rather than policy A, and consumers notice that fact, then the demand curve will come down, by an amount exactly related to how much they care (H). That downward shift in the demand curve will be set against whatever other gains (V) the firm might make from reuse of the data. That demand shifting effect can give firms an incentive to act responsibly in their data use. But several things have to go right, and I think there are reasons to believe that they often fail to go right.

First, the firm's choice of policy must be a real commitment; this requires that defections from it be readily observable and provable, and that there be an effective discipline should such defections be found. This is important and not automatic in privacy policy, where it can be hard to tell later how certain information got out—or, in the case of identity theft, even what information it was that went astray. In its recent Privacy Report, the FTC stated that it intends to use law enforcement to support the commitment value of firms' privacy promises to consumers.

> To the extent that strong privacy codes are developed, the Commission will view adherence to such codes favorably in connection with its law enforcement work. The Commission will also continue to enforce the FTC Act to take action against companies that engage in unfair or deceptive practices, including the failure to abide by self-regulatory programs they join.[9]

Second, consumers have to notice which privacy policy is on offer:[10] in economists' language, the policy must be "observable." That is

9. FTC RECOMMENDATIONS, *supra* note 3, at vi.

10. In real-world privacy policy, it is also important that they understand the final-good

much stronger than the "fulfilled expectations" condition that consumers are not persistently fooled. There are two very different ways in which the less protective policy B might turn out to be optimal for the firm. In the first, B is optimal because V exceeds H, so even though the firm would get the benefit of the upward shift (by H) in demand if it switched to policy A, that wouldn't make up for the loss of V. In the second, consumers expect policy B not because they can see that the firm actually adopts it, but because they have learned generally that firms will chase follow-on revenues and won't protect privacy, so they assume it.[11] In the first of these ways, policy B has passed a market test—with the conditions sketched above, a stringent test; in the second, B is expected because consumers have learned to be wary and cynical, and is adopted because the firm knows that the policy it adopts won't affect demand.

While the difference may sound technical, it can make all the difference between the well-functioning market sketched above and a wary, cynical market failure. If consumers have learned that "there is no privacy; get over it," and thus expect policy B *independent of what policy the firm actually were to choose*, consumers' wariness would protect them against being actually deceived and against going into a transaction that would on net be harmful for them. But these stubbornly pessimistic expectations fail to support efficient transactions, leading to a different kind of consumer harm, in which the firm is unable to benefit from making privacy-protective promises; this can make mutually beneficial trades impossible. In technical terms, systematic deception and purchases that consumers actually regret are (in this model) disequilibrium phenomena, but there can be a highly dysfunctional equilibrium, in which consumers' expectations are—regrettably—correct.

In its 2010 Preliminary Report, the FTC Staff recognized the value to consumers and competition from avoiding the cynical equilibrium where firms' privacy promises are viewed by consumers as not credible. In describing possible benefits from the adoption of the privacy framework, the Preliminary Report stated, "The Commission staff believes that such a simplified approach to providing choices will not only help consumers make decisions during particular transactions, but also will facilitate consumers' ability to compare privacy options that different companies offer. Thus, the staff's approach could promote

implications of the intermediate-good aspects of this; I remind the reader that I am focusing only on privacy as a final good.

11. This is the view perhaps captured by the quip, "There is no privacy; get over it." Learning that "there is no privacy" and getting "over it" will protect consumers against actual deception but will not in itself support market incentives for efficient or competitive privacy policies.

meaningful competition on privacy."[12]

Notoriously, privacy policies are often poorly disclosed, seldom read, and seldom understood;[13] and if left untreated, this problem seems apt to get worse as more of the relevant interactions take place with consumers on mobile phones with small screens. If many consumers contemplating buying the book don't notice whether privacy policy A or B is implicated, then the demand curve won't shift as I assumed above. If that's true, then H won't much affect the choice of policy, which will be driven primarily by V: the firm will offer the privacy policy that gives it more follow-on revenues, with little attention to what consumers would want if they knew.[14] This is the dysfunctional equilibrium I'm concerned about.

If firms perceive that few consumers shift their demand in response to actual privacy policies, then the firm's incentives are to make its policy noncommittal and/or non-protective, and to go for the biggest available V—or, perhaps less dramatically, to go for bigger V disproportionately over minimizing H. It would then be tempting to design disclosures so as not to really communicate the choice of policy, if it is possible to obfuscate for the minority of consumers while retaining the ability to claim that the policy was disclosed. Meanwhile, if consumers perceive that firms behave in this kind of way, they will not expect attentive reading of privacy policies to be a rewarding activity.

---

12. FTC PROPOSED FRAMEWORK, *supra* note 1, at 51.

13. For example, in a large dataset of consumer clickstream data, Bakos et al. find that the frequency at which retail shoppers look at End User Licensing Agreements (EULAs) on websites is roughly 1 to 2 times per thousand shopping visits; for various definitions of shopping visits to a website they find frequencies of EULA reading as low as 0.05% and no more than 0.22%. Among shoppers that do look at an EULA, the median time spent on the page is 29 seconds, far less than the approximately 8-10 minutes it takes to read a typical EULA. *See* Yannis Bakos et al., *Does Anyone Read the Fine Print? Testing a Law and Economics Approach to Standard Form Contracts* (N.Y. Univ. Sch. of Law Ctr. for Law, Econ. & Org., Working Paper No. 09-40, 2009), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1443256. McDonald and Cranor estimate the cost if all individuals read all privacy policies on websites they visited, using data on how long it takes on average to read a policy, how many unique websites the average internet user visits (in a year), how many Americans are online, and estimates of the value of time. Their point estimate is $781 billion per year. Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 ISJLP 540, 543 (2008). These numbers are clearly difficult to estimate reliably, and estimates differ a great deal. For instance, Milne and Culnan give much higher numbers for the frequency of reading privacy notices: their paper identifies 83.7% of respondents as readers, but that includes those who indicate they rarely read (33.3%) or sometimes read (31.8%) privacy notices. 17.3% indicated they never read. An alternative breakdown would be 50.6% rarely or never read, and 85.4% read sometimes, rarely or never. George R. Milne & Mary J. Culnan, *Strategies for Reducing Online Privacy Risks: Why Consumers Read (Or Don't Read) Online Privacy Notices*, 18:3 J. OF INTERACTIVE MKTG. 15, 16-29 (2004).

14. This is of course still a simplified model: for instance, reputations might develop over time.

These patterns of conduct and expectations would reinforce each other, which is what makes them a game-theoretic or economic equilibrium.[15]

It is often difficult to escape a dysfunctional equilibrium if there are large numbers of players involved. A consumer can't simply decide to start reading privacy policies—or rather, she can, but it won't do a lot of good, since firms will still expect that few consumers do so, so the consumer is apt either to learn little (noncommittal or vague policies) or be confirmed in her wary cynicism (policies that reveal a lack of protection—the rational choice for the firm when it expects that few consumers' behavior will be affected). A small firm, likewise, can't simply decide to break out of the equilibrium by adopting more protective policies and clearer disclosures, because its demand won't shift by much, so it will mostly just be sacrificing V. Escape from a dysfunctional equilibrium often, and probably here, requires action by large and powerful players, and/or concerted action by groups of players.

Compounding the problem, it's technically hard to effectively disclose a complex or nuanced policy. Teaching is hard, even when both teacher and student are really trying. The technical difficulty of truly effective disclosure interacts with the incentive problems: the difficulty makes it hard for regulators or courts—or perhaps even the reputation mechanism—to insist on fully effective disclosure, and the resulting wiggle room offers scope for adverse incentives to play out. While effective communication is hard even without incentive problems, advertising often seems more effective than many mandated disclosures. Incentives can outperform rules.

Summing up, if consumers know the implications of different up-front privacy policies and the policies are truly effectively disclosed—including drawing consumers' attention—then the demand-shift effect could provide strong incentives for firms to choose responsible policies. Unfortunately, there is also a dysfunctional equilibrium in which few consumers devote much attention to disclosures, disclosures are vague, noncommittal, or even if explicit, mostly ignored; and the privacy policies chosen are inefficiently non-protective. One of the grand challenges for policy is to escape such equilibrium.

III. MODIFYING POLICIES LATER

Since clear disclosures need to be reasonably simple, but the world is full of complex contingencies, it is natural to ask about later modifying

15. A recent paper states that, "70% of people surveyed disagreed with the statement 'privacy policies are easy to understand', and few people make the effort to read them." Janice Y. Tsai et al., *The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study*, 22:2 INFO. SYS. RES. 254, 254-68 (2011) (internal citations omitted).

privacy understandings to take advantage of late-arising opportunities, better market information, etc. Clearly, too liberal an approach to later modification could gut the concept of commitment to a privacy policy. Here I assume that a good modification regime involves ex post consumer consent.[16] This raises some logistical issues, because whereas it is easy to slightly modify the price of a book to take account of small follow-on revenues or consumer impacts of privacy policies, making stand-alone small payments involves a real challenge lest the transaction cost swamp the payment. This is sometimes called the problem of micropayments, although getting transaction costs down is a broader goal in the payments industry.

A. Communication and Negotiation

Part of the challenge is simply for the firm to communicate a proposed modification of policy to the consumer. Sometimes this will be relatively simple, other times not—for instance, in some relationships the consumer is regularly visiting a website controlled by the firm, while in other cases there may be little or no ongoing contact. Presumably the firm can describe the proposed change; but incentives may be needed to get the consumer to focus on it, or to agree to it even if it is in the consumer's and firm's joint interest. The problems of such (broadly construed) micropayments differ somewhat according to the direction in which the payment may need to flow.

B. Direction of Payments and the Role of Content and Advertising

In some cases, the consumer must be compensated to accept the change in policy: this will be the case, for instance, if the firm has belatedly discovered a new way, not previously contemplated, to profit from re-purposing the consumer's information. How can the firm offer micro-compensation to a consumer?

In other cases, though this is less obvious, it will be efficient to modify the privacy policy to strengthen the consumer's rights. The firm could simply refrain from certain actions, but in order to provide full incentives for such ex post improvements in privacy, one would ideally want there to be a mechanism for consumers to micro-compensate the firm for doing so.

16. The FTC's criterion for whether to require consumer consent for new uses of information depends on "the extent to which the [new use] is consistent with the context of the transaction or the consumer's existing relationship with the business, or is required or specifically authorized by law." FTC RECOMMENDATIONS, *supra* note 3, at 38-39.

We don't have a widely used system for money payments of fractions of a cent, in either direction, that doesn't involve high transaction costs. However, consumers can in a sense make such payments to firms by viewing ads, and firms perhaps can make payments to consumers by offering free attractive content. Oddly perhaps, the exchange of content for the viewing of ads, which is one source of modern privacy concerns, could thus also help with their resolution.

CONCLUSION

Privacy is a subtle good, whose economic character varies widely. In this speech I explored some aspects of the scope for market provision of privacy to the extent that it is a final good, valued for its own sake by consumers whose preferences are respected and who broadly know how to evaluate it. I stressed that if standard commitment and information conditions for efficient contracting hold, there seems no obvious reason why privacy should not be efficiently chosen by profit-seeking firms—but that there are particular reasons why those standard conditions may be especially challenging in the privacy context.

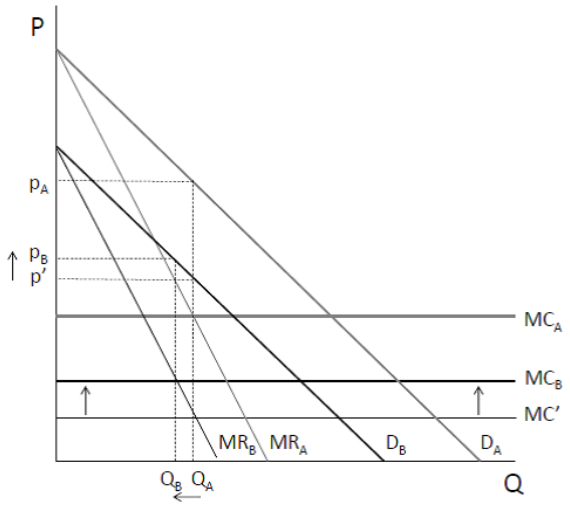

FIGURE 1

Under policy B (when information is used), demand falls by a downward shift of H = $4 from D_A to D_B. Similarly, if V = $4, this can be illustrated as marginal cost falling $4 from MC_A to MC' (we indicate a shift in V equal to the shift in H by the addition of a' to each curve or value). The firm would sell the same quantity as under policy A at p' = p_A - $4. The firm's surplus under policy A is the same as it is under policy B, and welfare does not change.

D_A= Consumer demand for the book under policy A (information is not used)

MR_A = Marginal revenue for the book under policy A

D_B = Consumer demand for the book under policy B (information is used)

MR_B = Marginal revenue for the book under policy B

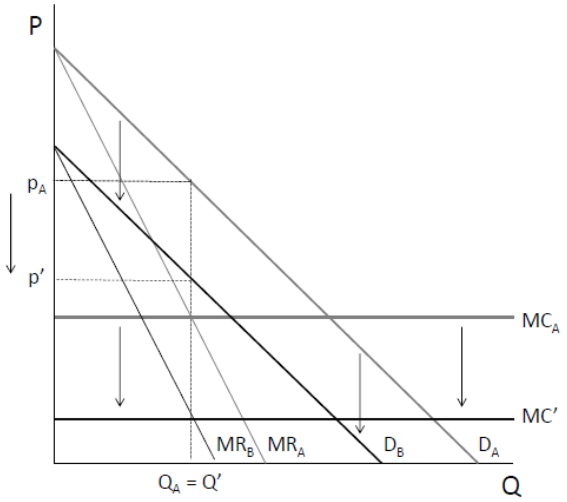

MC_A = Marginal cost of the book under policy A

MC' = Effective marginal cost of the book when H = V

Q_A = Quantity of the book sold under policy A

Q' = Quantity of the book sold when H = V

p_A = Price of the book under policy A

p' = Price of the book when H = V

FIGURE 2

Because V = $3, the marginal cost only falls $3. We can think of this as a $1 increase in MC from MC' to MC_B. Compared to when V = $4, the firm is worse off because the effective marginal cost has risen. Because the firm's optimal price is now higher and quantity sold is lower, consumers are also worse off.

MC_B = Marginal cost of the book under policy B

Q_B = Quantity of the book sold under policy B
p_B = Price of the book under policy B