

Search Costs, Hassle Costs, and Drip Pricing: Equilibria with Rational Consumers and Firms*

Michael R. Baye

John Morgan

Indiana University

University of California, Berkeley

March 2019

Abstract

This paper examines drip pricing related to compulsory charges—a situation where firms intentionally make it costly for consumers to discover mandatory fees or surcharges that “drip” into the full (total) price, which is only revealed after incurring the hassle cost of completing a purchase. We show that drip pricing can arise as an equilibrium phenomenon with fully rational consumers and profit-maximizing firms. We also show that, when consumers and firms are rational, (a) situations where drip pricing raises prices and harms consumers are unlikely to arise from unilateral business decisions, and that (b) the most likely avenue by which drip pricing harms consumers is through the coordinated adoption of drip pricing.

*We thank Nathan Blalock, Joe Farrell, Rick Harbaugh, Jan Pappalardo, Mary Sullivan, two anonymous referees, and participants of the Consumer Protection Economics Symposium organized by the Federal Trade Commission and *Economic Inquiry* for helpful comments.

1 Introduction

In an influential paper, Sullivan (2017) defines drip pricing as “the practice of advertising only part of a product’s price up-front and revealing additional charges later as consumers go through the buying process.” For example, a consumer might visit a hotel’s website, expend time and effort navigating through various pages, enter her address and credit card information and so on before ultimately reaching the checkout page—only then discovering a higher total price owing to a mandatory resort fee. Recently, such strategies have come under scrutiny by the Federal Trade Commission (FTC), who sent the following warning to one or more hotels:

“We reviewed your website...and found that in at least some instances mandatory resort fees are not included in the reservation rate quoted to consumers. We strongly encourage you to review your company’s website to ensure you are not misrepresenting the total price consumers can expect to pay when making a reservation to stay in your hotel. Please be advised that the FTC may take action to enforce and seek redress for any violations of the FTC Act as the public interest may require.”¹

Similar concerns have been raised by the Department of Transportation (DOT), the former Office of Fair Trading (OFT) in the UK, and other agencies charged with protecting consumers from unfair business practices.²

When consumers suffer from behavioral biases, policies that prevent drip pricing may be necessary to protect consumers.³ Largely for this reason, the policy debate is almost

¹<https://www.ftc.gov/sites/default/files/attachments/press-releases/ftc-warns-hotel-operators-price-quotes-exclude-resort-fees-other-mandatory-surcharges-may-be/121128hoteloperatorsletter.pdf>

²Much of this concern derives from the observation that drip pricing may permit firms to exploit consumer irrationality to increase prices and, consequently, harm consumers. See, for instance, Sullivan (2017), DOT (2011, 2012), and OFT (2010a,b).

³See, for instance, the surveys by Ahmetoglu (2010), Sullivan (2017) and the OFT (2010a,b).

exclusively based on analyses that assume consumers are irrational.⁴ The interested reader is encouraged to read Ahmetoglu, *et al.* (2010) and Sullivan (2017) for a discussion of the theoretical and empirical literature on behavioral economics and its relevance to drip pricing.

The present paper focuses on drip pricing related to compulsory surcharges (rather than optional add-on charges) in an environment where consumers and firms are fully rational. In this setting, consumers have full information about everything but price, and hence do not gain any new information about product quality or add-on options during the buying process. Instead, consumers discover information about mandatory fees and surcharges that “drip” into the total price—the actual price paid—which is only revealed at the end of the process.⁵ We show that the conventional wisdom—that drip pricing regarding compulsory charges cannot harm rational consumers—is incomplete. Harm from drip pricing depends on the “hassle costs” (e.g., the time and effort required to discover the total price) imposed on consumers. Indeed, we show that with fully rational consumers and firms, a continuum of drip pricing equilibria can arise in which consumers pay higher average prices as a result of firms’ injection of frictions into the price discovery process. The magnitude of the hassle costs orders the levels of prices and consumer welfare in these equilibria.

Our analysis complements several recent papers on obfuscation that demonstrate firms may not have incentives to disclose all relevant information to rational consumers; see Wilson (2010) and Ellison-Wolitzky (2012) for analyses of obfuscation in oligopoly settings and Petrikaitė (2018) for an analysis of obfuscation by a monopolist. While these models provide useful insights into obfuscation, we focus on a drip pricing environment where a large number of firms each choose a sequence of prices and hassle costs to maximize profits. Specifically,

⁴Jovanovic (1982, p. 42) identifies conditions in which “. . . the free market offers ample incentives for disclosure. . .” and, in his model, there is “. . . no support for a policy that makes business disclosure mandatory.” Brown, *et al.* (2010) and Sullivan (2017) correctly note that these and related models (e.g., Milgrom, 1981) are based on the assumption that consumers are fully rational.

⁵An important difference between discovering the price of optional features versus compulsory fees is that the former entails changes in the nature/quality of the product or “bundle” purchased during the buying process (resulting in additional revenues to the firm only if a consumer opts for add-ons), while the latter does not.

a firm engaging in drip pricing discloses (a) an initial price ($p_0 \geq 0$) and subsequent “drips” ($p_t > 0$) that reveal the total price $p = p_0 + \sum_{t=1}^T p_t$ only at checkout, and (b) potentially requires a consumer to incur a hassle cost κ_t to discover drip t . We model the disclosure and discovery of this price information as endogenous decisions of rational firms and consumers. Our analysis is further differentiated from the obfuscation literature in that we do not assume that obfuscation results in consumer fatigue (as in Ellison and Wolitzky) or serves as a commitment mechanism (as in Wilson). Unlike Petrikaitė who abstracts from competition among firms to focus on screening incentives, our focus is on competition among firms selling similar products.

Furthermore, to focus purely on drip pricing, we rule out reputational considerations that might also influence firms’ incentives to engage in drip pricing by assuming that consumers engage in optimal non-directed sequential search to visit sellers’ websites and possibly purchase a product of known quality. That is, we consider the stark environment where consumers cannot target a specific firm, firms produce homogeneous products, and the search/purchase decision is one-shot. The cost of navigating to a particular firm’s website—the traditional search cost—is positive and exogenous, but potentially negligible. Unlike traditional search models, landing on a firm’s website reveals an initial price that may not coincide with the final, all-in price. Instead, a consumer only discovers the total price by incurring the hassle costs of navigating the firm’s website to discover each price drip until ultimately reaching the checkout page. Each profit-maximizing firm unilaterally decides on the initial price as well as the sequence of drips and hassle costs. A firm’s profit-maximizing price depends on consumer demand, as well as its marginal cost (which is private information and varies across firms).

We focus on the weak perfect Bayesian equilibria (wPBE) of this environment. On the surface this is a difficult task because firm choices are multidimensional; that is, they consist of sequences of endogenous price drips and hassle costs. Our strategy is to examine equilib-

rium in simple (constrained) models that limit the number of elements in these sequences (Propositions 1 and 2), and then show that any equilibrium in the general (unconstrained) model has an economically equivalent analog in the constrained model (Proposition 3). These results allow us to readily analyze equilibrium when firms endogenously choose hassle costs (Propositions 4 and 5). We first show that, in any wPBE, no firm chooses an initial price, sequence of drips, or sequence of hassle costs that causes a consumer to abandon the firm before discovering its total price. This implies that many different initial prices and sequences of drips (including ones that provide differing degrees of information about a given firm’s total price) are consistent with non-abandonment. Despite this indeterminacy, all of these sequences give rise to the same (unique) distribution of total prices and purchase behavior. Indeed, we establish uniqueness in economic outcomes: for any equilibrium level of total hassle costs, the resulting distribution of total prices and purchase behavior is unique. In terms of policy, hassle costs associated with drip pricing harm consumers through higher (average) transactions prices and, as a result, more deadweight loss.

Our results contribute to the literature on obfuscation by identifying a distinct avenue through which drip pricing may harm consumers—the imposition of needless hassle costs. Harm in our setting derives from the imposition of hassle costs by *other* firms. In existing models, a firm has a unilateral incentive to engage in obfuscation even if other firms do not. Such unilateral incentives may stem from behavioral factors or by denying information to rational consumers. These are not the sources of harm stemming from our model of drip pricing; while initial prices and subsequent drips may only provide partial information to consumers, firms in our model do not benefit from unilaterally concealing information through drip pricing. Indeed, each firm has an incentive to ensure that every visitor discovers its total price; competition—and the ability of high-cost firms to mimic the initial prices and drips of low-cost firms—constrains firms’ abilities to harm consumers by concealing information.

The harm we identify stems from the impact of hassle costs on consumers’ costs of gathering price information at other firms. No firm has a unilateral incentive to impose hassle costs, but when all firms do so it raises the reservation prices of consumers, softens price competition, and increases industry profits. However, if it is costly for firms to impose hassle costs, no firm does so in any wPBE. Thus, consumer harm from drip pricing stems from coordinated (or collusive) efforts by firms—not through inadequate disclosures or the concealing of information through obfuscation.

The next section presents a formal definition of drip pricing and our model. Section 3 presents technical results, and then we provide a policy-oriented discussion of our results (and their limitations) in Section 4. We conclude in Section 5 by providing a broader context for how our analysis complements other analyses and differing world views regarding the appropriate role of consumer protection policies. Formal proofs of all of the propositions in the text are provided in the Appendix.

2 Model

Drip pricing is the practice of (a) dividing the total price into two or more parts (e.g., product price, shipping & handling, surcharge, taxes), (b) disclosing these parts sequentially rather than simultaneously, and (c) where this practice possibly makes it costly for consumers to discover the total price. Notice that there are two main components to this definition: the labeling of various portions of the total price, and the timing in which parts of the total price are revealed. We consider a model in which consumers are fully rational, and therefore are not susceptible to framing effects arising from labeling. Hence, we focus purely on the timing in which the total price is revealed. Letting p be the total price, we say that a firm engages in drip pricing with T “drips” when, after observing an initial price (p_0), T additional components of the total price $\{p_t\}_{t=1}^T$ are listed sequentially, producing the total price $p = p_0 + \sum_{t=1}^T p_t$. For example, a consumer wishing to book a room at a resort might

visit a firm’s website to discover the room rate (p_0), and only after navigating the site (or showing up at the resort) discover additional mandatory charges such as taxes, cleaning fees, and/or resort fees. Formally:

Definition 1 *Drip pricing is the practice of sequentially disclosing the total price (p) through an initial up-front price (p_0) and $T \geq 1$ subsequent price “drips” containing additional compulsory charges ($p_t > 0$). The total price is thus*

$$p = p_0 + \sum_{t=1}^T p_t > p_0.$$

We examine drip pricing in an environment where a large number of firms engage in price competition for a large number of identical consumers. We normalize the number of consumers per firm to be unity.⁶ A given firm privately learns its marginal cost, m , which we view as an *iid* realization from a continuous distribution with cdf

$$G(m) \text{ on } [\underline{m}, \bar{m}], \text{ where } \underline{m} \geq 0 \text{ and } \bar{m} < \infty. \quad (1)$$

When there is no ambiguity, we also refer to m as a firm’s type. Let p_m denote the *total* price charged by a firm of type m (e.g., a firm with marginal cost, m), and let κ denote a generic level of hassle costs. Otherwise, we index variables (e.g., $\kappa_{t,m}$, and $p_{t,m}$) to identify variables associated with drip t by firm m .

Multiple steps are required to purchase a product. A consumer must incur (an exogenous) search cost, $c > 0$, to visit a firm’s website to observe p_0 on its landing page. Search is *non-directed*; that is, consumers visit a firm selected at random and are unable to target a specific firm. After arriving at a firm’s landing page, a consumer must incur the (potentially endogenous) hassle costs, $\kappa_m \equiv \sum_{t=1}^T \kappa_{t,m} \geq 0$, of navigating through firm m ’s website to its checkout page in order to discover its total price, $p_m \equiv p_{0,m} + \sum_{t=1}^T p_{t,m}$.⁷ At any point (t)

⁶Technically, let $\lambda = a/b$, where a is the (large) number of consumers and b is the (large) number of firms. Letting a and b approach infinity while holding λ constant provides a measure of the number of consumers per firm (*cf.*, Reinganum (1978, p. 6)). None of our results depend on the assumption that $\lambda = 1$.

⁷Notice that one can allow different firms to utilize different numbers of steps by interpreting T to be the largest possible number of steps and setting $\kappa_{t,m} = 0$ (and $p_{t,m} = 0$) for firms employing fewer steps.

in the shopping process, a consumer may abandon firm m to search at another firm or quit altogether. Hassle costs may vary across firms; different values of $\kappa_{t,m}$ represent different possible consumer experiences from arrival to checkout. Low values represent circumstances where price discovery is relatively straightforward; high values of $\kappa_{t,m}$ represent situations where price discovery is time-consuming and arduous. In practice, firms may achieve the latter by requiring consumers to navigate to additional pages to learn the total price. A consumer only learns $\kappa_{t,m}$ after incurring it; that is, hassle costs are not observable *ex ante*. Finally, upon discovering a firm's total price, a consumer determines whether (and how many units) to purchase. We assume that there is *free recall*; that is, if at any point during the shopping process a consumer decides not to gather additional price information, she may costlessly return to a previously visited firm to purchase at the total price discovered there or, alternatively, resume shopping at any firm previously visited without having to incur c or any of the hassle costs *previously* incurred.

A consumer who visits n firms, learns the total price of each, and purchases $q(p_m)$ units from a firm charging a total price, p_m , obtains (indirect) utility

$$V = v(p_m) + Y - \left(cn + \sum_{i \in N} \kappa_i \right) \quad (2)$$

where N denotes the set of firms visited and Y is income. By Roy's identity, a consumer's demand is $q(p) \equiv -v'(p)$. Following Reinganum (1979), we assume (a) consumers have identical isoelastic demands given by $-v'(p) = q(p) = p^\varepsilon$, where $\varepsilon < -1$; thus, firm m 's monopoly profits, $\pi^M(p_m) \equiv (p_m - m)q(p_m)$, are strictly increasing up to its monopoly price, ρ_m :

$$\rho_m \equiv \left(\frac{\varepsilon}{1 + \varepsilon} \right) m \quad (3)$$

(b) consumers engage in optimal non-directed sequential search with free recall, and (c) consumers find it optimal to search at least once. A sufficient condition for (c) is that hassle costs have a finite upper bound, $\bar{\kappa} > 0$, such that $\kappa_m \leq \bar{\kappa} \equiv v(\rho_m) - c$ for all m . That is, a consumer would be willing to incur search cost c for the privilege of purchasing at the

monopoly price of the firm with the highest marginal cost, even if doing so entails incurring the highest possible hassle cost. We will show that the Reinganum model is nested as the special case where $T = 0$ or where $T \geq 1$ and $\kappa_m = 0$ for all m . Finally, let $\bar{\rho} \equiv \rho_{\bar{m}}$ and $\underline{\rho} \equiv \rho_{\underline{m}}$ denote, respectively, the monopoly price of a firm with the highest and lowest marginal cost.

The timing of decisions is as follows. Nature first endows firms with heterogenous marginal costs; m is private information but G is common knowledge. Firms then choose their sequences of $p_{t,m}$ and $\kappa_{t,m}$ simultaneously but independently. Consumers do not observe firms' prices or hassle costs; rather, they discover them through non-directed sequential search. Search results in a visit to a firm's landing page that only reveals $p_{0,m}$. A consumer may then choose to visit another firm's landing page or incur the cost $\kappa_{1,m}$ to discover $p_{1,m}$, and so on. A consumer cannot purchase the product until reaching a checkout page—that is, observing the total price charged by at least one firm.

Our equilibrium concept is weak perfect Bayesian equilibrium (wPBE) in pure strategies. This requires that (a) each firm maximizes its expected profits given beliefs about consumer search and buying strategies, (b) each consumer maximizes her expected surplus given beliefs about firms' prices and hassle costs, (c) beliefs must be correct in equilibrium, and (d) at every decision node in or out of equilibrium, consumers decisions are optimal given observed firm behavior.

3 Analysis and Equilibrium

Notice that firm choices are multidimensional; that is, they consist of a sequence of prices $\{p_{0,m}, p_{1,m}, \dots, p_{T,m}\}$ and hassle costs, $\{\kappa_{1,m}, \dots, \kappa_{T,m}\}$. Our strategy is to first examine equilibrium in simple (constrained) models that limit the number of elements in these sequences, and then show that any equilibrium in the general (unconstrained) model has an economically equivalent analog in the constrained model. Initially, we consider two simple environments. In one firms are constrained not to engage in drip pricing ($T = 0$); in the other, firms are

constrained to impose a single drip ($T = 1$), disclose an initial price of zero ($p_{0,m} = 0$) and to impose a common level of hassle costs ($\kappa \geq 0$).

Subsection 3.1 examines the first constrained model. Proposition 1 shows that the $T = 0$ case corresponds to the analysis in Reinganum (1979), where the only informational friction is the exogenous search cost $c > 0$ of sampling a randomly selected firm. Subsection 3.2 examines equilibrium in the second constrained model; Proposition 2 shows that the $T = 1$ and $p_{0,m} = 0$ case induces a wPBE corresponding to a Reinganum equilibrium where search costs are $c + \kappa$. Among other things, this implies that (a) when $\kappa = 0$ the wPBE in the constrained model with $T = 1$ and $p_{0,m} = 0$ is identical to that arising when firms cannot engage in drip pricing ($T = 0$), and (b) when $\kappa > 0$ the equilibrium arising in the constrained model with $T = 1$ and $p_{0,m} = 0$ is identical to an environment where firms impose (exogenous) hassle costs of $\kappa > 0$ through a single (exogenous) drip.

Finally, Proposition 3 shows that the equilibrium arising in the constrained model with $T = 1$ and $p_{0,m} = 0$ corresponds to that arising when the number of drips ($T > 1$) is arbitrary and firms may freely choose sequences $\{p_{0,m}, p_{1,m}, \dots, p_{T,m}\}$ of initial prices and drips. This permits us to use the second constrained model to determine firms' unilateral incentives to impose hassle costs in the more complex environment where both $\{p_{0,m}, p_{1,m}, \dots, p_{T,m}\}$ and $\{\kappa_{1,m}, \dots, \kappa_{T,m}\}$ are endogenous. This analysis is conducted in subsection 3.3.

3.1 The No Drip Pricing Benchmark ($T = 0$)

As a benchmark, consider the situation where drip pricing is not possible ($T = 0$, and hence $\kappa_m = 0$ for all m). In this case, each firm makes a single decision (the total price to post on the landing page of its website) and the only cost to consumers of discovering a firm's total price is the exogenous search cost, c , of visiting a firm. In this case, (a) the industry fundamentals in our model are identical to those in Reinganum (firm costs, consumer preferences, and monopoly prices are as described in equations (1), (2), and (3), respectively), (b) consumers

engage in optimal non-directed sequential search with free recall, and (c) consumers find it optimal to search at least once. Thus, the model and equilibrium exactly correspond to that in Reinganum.

Proposition 1 *When $T = 0$, the model corresponds to the economic setting in Reinganum.*

In the Reinganum equilibrium:

(1) *Consumers engage in sequential search and purchase $q(p_m) > 0$ units from firm m if and only if its total price does not exceed the reservation price, r_c ;*

(2) *A firm of type m charges a total price of $p_m = \min\{\rho_m, r_c\}$. This strategy induces a distribution of total prices given by*

$$F_c(p) = \begin{cases} G\left(\frac{1+\varepsilon}{\varepsilon}p\right) & \text{if } p < r_c \\ 1 & \text{if } p = r_c \end{cases}$$

where the reservation price, r_c , is implicitly defined by

$$\int_{\underline{p}}^{r_c} (v(p) - v(r_c)) dF_c(p) = c$$

(3) *The expected profits of a firm of type m are given by*

$$\pi_m = \begin{cases} (\rho_m - m) q(\rho_m) > 0 & \text{if } \rho_m \leq r_c \\ (r_c - m) q(r_c) > 0 & \text{if } \rho_m > r_c \end{cases}$$

Furthermore, the Reinganum equilibrium is a weak perfect Bayesian equilibrium (wPBE) in pure strategies.

Notice that search in the Reinganum equilibrium is socially efficient in that each consumer visits only a single firm. Thus, while the costs consumers incur searching reduce social welfare, these costs are minimized. Despite this, consumers' options to visit other stores constrains the prices charged by high-cost firms. Firms with low costs, however, charge their monopoly prices. Indeed, because demand is downward sloping rather than rectangular, low-cost firms earn more by charging their (lower) monopoly price than pricing at the reservation price of consumers.

3.2 Drip Pricing ($T \geq 1$) with Exogenous Hassle Costs

Next, consider a simplified environment where firms are constrained to engage in a form of drip pricing with $T = 1$ and that hassle costs κ as well as the initial price $p_0 = 0$ are exogenous and common across firms. That is, each firm’s landing page contains price information that is completely uninformative, such that consumers must navigate to the checkout page to discover a firm’s total price. This might represent a situation where, for reasons outside the model, all the firms’ websites have the same “look and feel,” and the initial price displayed by firms is so low that rational consumers know they will pay a higher total price, but are unsure of the level of the final price. Later, we shall relax these assumptions and permit firms to endogenously choose initial prices, as well as entire sequences of price drips and hassle costs.

In this simple environment, consumers learn nothing about *any* firm’s total price after incurring the cost c of visiting a landing page. And regardless of the firm visited, a consumer anticipates that she must incur an additional cost of κ to discover that firm’s total price. Thus, the effective search cost—the prospective cost of discovering any firm’s total price—is $c + \kappa$. Since there are many firms with marginal costs that are *iid* draws from G and search is non-directed,⁸ a consumer’s beliefs about the distribution of total prices from an additional search remain the same regardless of her search history. It follows that a stationary stopping rule is optimal for consumers and that an optimizing consumer never abandons her shopping cart, in equilibrium.⁹ Indeed, as shown below, one may view the resulting equilibrium as the wPBE arising in the Reinganum model when the reservation price is $r_{c+\kappa}$ and the distribution of prices is $F_{c+\kappa}$ rather than r_c and F_c .

To see this, note that an optimizing consumer compares the expected benefits of searching again to the prospective costs (including hassle costs), $c + \kappa$. Thus, analogous to Reinganum,

⁸By non-directed search we mean that consumers choose a firm at random rather than choosing to visit specific firm (or type of firm).

⁹The formal arguments are similar to those presented in the Appendix to Reinganum (1978) and thus omitted. More generally, see Morgan and Manning (1985).

search costs $c+\kappa$ induce an equilibrium with a stationary reservation price (denoted $r_{c+\kappa}$) and a distribution of prices (denoted $F_{c+\kappa}$); neither consumers nor firms can gain by deviating. In such an equilibrium, each firm charges the minimum of its monopoly price, $\rho_m = \left(\frac{\varepsilon}{1+\varepsilon}\right) m$, and the reservation price, $r_{c+\kappa}$. Since marginal costs vary across firms according to $G(m)$, this induces a distribution of prices with support $[\underline{\rho}, r_{c+\kappa}]$ and cdf

$$F_{c+\kappa}(p) = \begin{cases} G\left(\frac{1+\varepsilon}{\varepsilon}p\right) & \text{if } p < r_{c+\kappa} \\ 1 & \text{if } p = r_{c+\kappa} \end{cases}$$

where the reservation price satisfies¹⁰

$$\int_{\underline{\rho}}^{r_{c+\kappa}} (v(p) - v(r_{c+\kappa})) dF_{c+\kappa}(p) = c + \kappa \quad (4)$$

The *reservation price*, $r_{c+\kappa}$, represents the highest acceptable price to a consumer. At higher prices, a consumer strictly gains from visiting another firm to discover its price; at lower prices, she strictly gains from buying. Hence, consumer behavior is optimal given the price distribution $F_{c+\kappa}$. Likewise, no firm can gain by altering its price. Clearly, a firm whose monopoly price lies below $r_{c+\kappa}$ can do no better by deviating. A firm whose monopoly price lies above the reservation price optimally charges $r_{c+\kappa}$ since it earns less by charging a lower price and, since all prices $p > r_{c+\kappa}$ are rejected, earns nothing by charging a higher price. Hence, firms of all types are behaving optimally. In such an equilibrium, a firm of type m earns profits of

$$\pi_m = \begin{cases} (\rho_m - m) q(\rho_m) & \text{if } \rho_m \leq r_{c+\kappa} \\ (r_{c+\kappa} - m) q(r_{c+\kappa}) & \text{if } \rho_m > r_{c+\kappa} \end{cases}$$

Note that $\pi_m > 0$ for all m . For future reference, we refer to these pricing, search, and buying strategies as a “Reinganum $c + \kappa$ equilibrium.” We formalize these observations as:

Proposition 2 *Fix $\kappa \in [0, \bar{\kappa}]$, $T = 1$, and $p_{0,m} \equiv p_0 = 0$. Then the distribution of total prices and consumer purchase behavior arising in a wPBE with drip pricing and hassle costs*

¹⁰Like Reinganum, we assume an interior solution for the reservation price. Baye, *et al.* (2006) also characterize corner cases.

of κ is identical to that in a setting without drip pricing (e.g., $T = 0$) where exogenous search costs are $c + \kappa$.

It is clear that equilibrium consumer and firm behavior in Proposition 2 when $T = 1$ does not critically depend upon the assumption that each firm lists the same initial price, $p_{0,m} \equiv p_0 = 0$, on its landing page. For example, if each firm chooses $p_{0,m}$ randomly from any common distribution with support $[0, \underline{m}]$, then firms' initial prices remain uninformative, and the distribution of total prices and the purchase behavior identified Proposition 2 remains an equilibrium.

In fact, the implications of the results for the $T = 1$ and $p_0 = 1$ case are even more general. As our next proposition shows, consumer purchasing behavior and the distribution of total prices (and hence, consumer welfare) in the constrained model are identical to that arising in the unconstrained model where firms impose an arbitrary ($T > 1$) number of drips and endogenously choose sequences $\{p_{t,m}\}_{t=0}^T$ of initial prices and drips that determine their total prices, $p_m = p_{0,m} + \sum_{t=1}^T p_{t,m}$.

Proposition 3 *Fix $\kappa \in [0, \bar{\kappa}]$. Then the equilibrium distribution of total prices and consumer purchase behavior arising in the constrained model with $T = 1$ and $p_{0,m} = 0$ is identical to that arising in the general model with $T > 1$ drips, unconstrained sequences of initial prices and drips, $\{p_{t,m}\}_{t=0}^T$, and total hassle costs $\kappa_m = \sum_{t=1}^T \kappa_{t,m} = \kappa$.*

Among other things, Proposition 3 implies that the equilibrium distribution of total prices and purchases in the general model is identical to the constrained model where $T = 1$ and $p_{0,m} = 0$ for all m . Importantly, the proposition implies that real economic outcomes (total prices and consumer purchase behavior) do not depend on the timing or amount of information disclosed prior to checkout; for a given κ , a firm's optimal price is the same regardless of whether it discloses 95% of its total price up front and drips the rest or discloses nothing up front. Hence, the failure of firms to disclose the total price on the landing page

is not sufficient for consumer harm. Consumer harm depends solely on the *total* hassle cost consumers incur to arrive at checkout.¹¹

In the equilibrium identified in Proposition 2 and Proposition 3, each firm’s *total* hassle cost is exogenously given by $\kappa_m = \kappa \geq 0$. Rational consumers optimally respond by setting a higher reservation price ($r_{c+\kappa}$) than they would in a situation where hassle costs are absent (r_c). While this does not impact the prices charged by firms with low marginal costs ($\rho_m \leq r_c$), firms with higher marginal costs increase their prices. In particular, firms whose monopoly prices lie in the interval $[r_c, r_{c+\kappa}]$ increase prices to their monopoly levels, ρ_m , while firms whose monopoly prices are above $r_{c+\kappa}$ raise their prices to match the new reservation price. In short, adding exogenous hassle costs $\kappa > 0$ raises industry profits, transactions prices, deadweight loss, and therefore lowers consumer welfare.

This implication of Proposition 2 and Proposition 3—that hassle costs associated with price discovery lead to higher prices and lower consumer welfare—might lead one to conclude that individual firms have incentives to engage in drip pricing. In the next section, we examine this more formally.

3.3 Drip Pricing ($T \geq 1$) with Endogenous Hassle Costs

We are now in a position to examine equilibrium in the general model where hassle costs are strategic decisions of individual firms. That is, prior to consumer search, firms simultaneously (but independently) choose their initial prices and drips $\{p_{t,m}\}_{t=0}^T$ as well as their hassle costs, $\{\kappa_{t,m}\}_{t=1}^T$.

Proposition 4 *When it is costless for firms to impose endogenous hassle costs, there is a continuum of wPBEs with drip pricing. In each of these equilibria, the distribution of total prices and consumer purchase behavior corresponds to a Reinganum $c + \kappa$ equilibrium where each firm imposes total hassle costs $\kappa_m = \sum_{t=1}^T \kappa_{t,m} = \kappa \in [0, \bar{\kappa}]$.*

¹¹Note that Proposition 3 does not require firms’ sequences of hassle costs to be identical.

Proposition 4 implies that, even when firms choose different initial prices, different sequences of drips, and different sequences of hassle costs, each of the equilibria in the continuum identified in Proposition 2 with $\kappa \in [0, \bar{\kappa}]$ can arise when hassle costs are endogenous. This highlights that the observed level of drip pricing in the market entails a coordination problem. Proposition 4 understates the degree of the coordination problem. For example, for any non-degenerate distribution of hassle costs with mean κ and support $S \subseteq [0, \bar{\kappa}]$, Proposition 4 still holds. Regardless of the realized hassle cost of discovering a particular firm’s total price, a consumer’s purchase decision depends on her prospective evaluation of the (expected) cost of obtaining an additional price quote, which is $c + \kappa$.

Note that the continuum of equilibria is Pareto ordered; equilibria in which the industry coordinates on higher levels of hassle costs result in greater industry profits (and lower consumer welfare). We summarize this observation as:

Corollary 1 *When firms act in concert to raise the costs of price discovery, industry profits rise and consumer welfare falls.*

Depending on the nature of industry coordination, Corollary 1 has potential ramifications for whether Section I of the Sherman Act or Section 5 of the FTC Act is the most appropriate avenue to pursue enforcement actions regarding drip pricing.¹²

It is important to stress that, in our framework, the reduction in consumer welfare associated with equilibria with positive hassle costs does *not* stem from an individual firm “surprising” consumers with an unexpectedly high hassle cost or price. Faced with such a situation, the consumer evaluates the firm’s offer against the value of an additional search, which is unchanged by encountering unexpectedly high hassle costs. Hence, even though a firm can surprise rational consumers with a high hassle cost, it gains nothing by doing so.

To summarize:

¹²Section I of the Sherman Act prohibits agreements by firms to coordinate on higher prices, while Section 5 of the FTC Act prevents “unfair” business practices.

Corollary 2 *Acting alone, a firm cannot gain by surprising consumers or raising the cost of discovering its price.*

We conclude this section by noting that our finding that drip pricing can harm fully rational consumers is fragile. It turns out that, if firms must incur any expense whatsoever to raise their hassle costs above some status quo level, denoted $\hat{\kappa} \in [0, \bar{\kappa})$, they will never do so. This conclusion holds regardless of whether one views $\hat{\kappa}$ as the lowest feasible level of hassle costs, the socially efficient level of hassle costs, or merely an arbitrary level.

Proposition 5 *Suppose hassle costs are endogenous and that it is costly for firms to increase them above some status quo, $\hat{\kappa}$. Then there does not exist an equilibrium in which firms impose hassle costs above the status quo.*

Proposition 5 illustrates that, while the industry as a whole may benefit (at the expense of consumers) from hassle costs that soften price competition by raising consumers' costs of discovering prices, there is a free-rider problem when it is costly to create price discovery barriers. While all firms would like their brethren to construct such barriers, they are not keen on doing it themselves.

Corollary 3 *If it is costly for firms to impose hassle costs, they never do so unilaterally.*

4 Discussion

This section connects our findings to practice and policy. We primarily focus on drip pricing related to compulsory charges, i.e., where the consumer desires a specific item and must undergo some inconvenience (hassle) to learn the full price a given seller charges. As we have already mentioned, our formal results do not pertain to add-on pricing or more exotic pricing practices that may (or may not) have efficiency rationales.

Central to all of our results is this: Even when consumers and firms are fully rational, industry-wide drip pricing can raise average prices, increase industry profits, and harm consumers. The larger is the hassle to consumers of discovering firms’ all-in prices, the greater is the magnitude of the potential effect.

One might be tempted to view the role of drip pricing and, more broadly, the practice of unnecessarily inconveniencing consumers for strategic advantage, merely as a disguised version of the Diamond Paradox, where search frictions provide the seller with maximal leverage in negotiating with a “captive” consumer, allowing it to charge its monopoly price. This is not the potential source of harm from drip pricing that we identify. Instead, firms derive leverage as a consequence of the consumer’s cost to searching *elsewhere*—no seller gains by employing drip pricing to inconvenience its own customers.

Even acting in concert, firms with the lowest marginal costs gain *nothing* from drip pricing, as they already have maximal leverage and thus charge monopoly prices. Competition from these firms (with the lowest monopoly prices) constrains the prices of firms with moderate and high marginal costs. The imposition of industry-wide hassle costs relaxes this constraint. Firms with moderate marginal costs begin charging monopoly prices. Those with the highest costs raise their price to the (now higher) reservation level.

Industry Coordination

Since a seller gains nothing by inconveniencing its own customers, moving away from transparent pricing would require some degree of industry coordination. In markets with numerous sellers, such coordination would seem difficult. Regardless, our results may cast drip pricing policies in a new light: Rather than focusing on the practice of drip pricing *per se*, the more relevant consideration might be to focus on industry coordination.

Importantly, however, evidence that all firms in an industry impose similar levels of “hassle costs” need not imply that drip pricing stems from coordinated behavior. Firms in an industry often adapt and react to common features of the market in similar ways, for

efficiency and other reasons. This, in turn, leads to common industry practices even absent coordination motives. Indeed, such practices may be efficient depending on the nature of the product or the evolution of the industry. Asymmetric benefits to participating firms further impede coordination; only the less efficient (higher marginal cost) firms benefit from industry-wide drip pricing in our model.

Industry Lock-In

It would of course be a mistake to conclude that an industry will automatically evolve to minimize hassle costs. Since low cost firms in our model gain nothing by adopting drip pricing, they also gain nothing by abandoning it. The reason is that, with non-directed search, a seller cannot *ex ante* signal its degree of price transparency, and hence cannot attract more customers through transparent pricing. Moreover, because hassle costs are sunk at the time of purchase, a seller cannot charge a price premium for offering a more streamlined shopping experience. Thus, firms with low marginal costs gain nothing from abandoning drip pricing. Those with high costs have even less incentive to abandon this practice.

In principle, centralized policy interventions, such as the DOT’s requirement that “...all mandatory taxes and fees must be included together in the advertised fare” of airline tickets, might remedy such *industry* lock-in.¹³ The key challenge is determining what practices to require and the degree of specificity. On the other hand, such regulation may be unnecessary. Southwest Airlines, for instance, successfully advertised a transparent pricing strategy to induce customers to engage in *directed search* to its website. More broadly, if firms can influence consumer choice through marketing or reputation, this may mitigate the need for regulation. We return to these issues when we discuss *mandatory disclosures*.

The Free-rider “Problem”

Most of regulatory attention has, rightly, focused on the costs to *consumers* of drip pricing and similar practices. But decisions about how to best organize a store’s website or

¹³<https://www.transportation.gov/affairs/2012/dot0812.html>.

floor space, and the cost to *firms* of implementing these decisions, also bear consideration. While our results indicate that competition may not lead to more transparent pricing, a countervailing force—free-riding—might. When firms must incur costs to intentionally add frictions into the buying experience, such as redesigning websites or reconfiguring floor space, they are no longer indifferent between price transparency and opacity. Firms exercise market power because *others* engage in drip pricing—their own disclosure practices do not figure.

Thus, firms have an incentive to free-ride: They will not invest in drip pricing when doing so is costly. Moreover, firms’ costs of adding frictions to price discovery need not be direct. Indirect costs, such as legal risk, bad publicity, or reputational damage, might also matter, and such considerations may even dwarf direct costs. Accounting for indirect costs amplifies firms’ incentives to free-ride, thereby mitigating the risk of deleterious drip pricing.

On the other hand, if informational frictions arise from acts of omission rather than acts of commission, free-rider incentives work against price transparency. Consider an industry where firms have settled on a given level of frictions, perhaps even the socially optimal level, given the available technology. A new innovation appears that reduces frictions, but requires firm investments. The migration of internet stores from Web 1.0 to Web 2.0 offers a good, real world, example. In this case, free-riding incentives hurt rather than help—no firm has a unilateral incentive to embrace Web 2.0, regardless of how cheap or how good it is. Even acting collectively, firms have no incentives to invest.

Mandatory Standards

In principle, a regulatory body or enforcement agency might attempt to solve the problems alluded to above by mandating standards or imposing transparency rules. But what level of transparency is appropriate? For instance, it may be technologically feasible to create a buying process that is virtually frictionless, but at a prohibitively high cost to firms. Determining the “right” level of transparency is no easy task. Merely observing the existence of hassle costs—that price disclosure is not perfect—implies nothing about whether more or

less disclosure is socially desirable. Likewise, survey evidence indicating that an alternative presentation of information would result in enhanced price discovery does not resolve the question either. Depending on the particulars of the environment, observing positive hassle costs can be consistent with too much or too little disclosure friction relative to the social optimum.

This observation readily generalizes. In many regulatory environments, a “full information” benchmark is inappropriate for evaluating various disclosure policies. Even if one discounts the cost to firms of achieving a given standard, the full disclosure benchmark can still be misleading as it ignores consumers’ costs of processing more detailed disclosures. When presented with “full disclosure,” effort is required for a consumer to plow through a mass of data, and these costs must be taken into account. For instance, even in a simple case where there are L binary choices that determine the final product and price, full disclosure would require a consumer to evaluate and compare 2^L different prices—and that for only a single firm. When L is large, is hard to imagine any consumer benefiting from this form of “transparency.” The broader point is that regulators must employ considerable care in evaluating market transparency or imposing disclosure rules.

5 Conclusion

We have shown that consumer harm from drip pricing derives not from the timing and extent of information disclosure, but rather from the hassle costs imposed by firms. We stress that drip pricing comes in many flavors, and not all are harmful to consumers or competition. For instance, a firm may engage in drip pricing because disclosing the total price up front is costly and/or would confuse consumers. Product complexity may make it difficult for a firm to initially disclose the total price when the total price depends on the optional or add-on features selected by particular consumers. For these sorts of reasons, many custom-home builders do not disclose the total price up front, as this price depends on a plethora

of customer-specific options including the structure (brick or vinyl) and other features that vary in price, grade and/or quality (e.g., paint/wall paper, appliances, molding, energy efficiency, and so on). Our analysis does not incorporate these considerations although, in some settings, they may be important.

Optimal consumer protection policy regarding drip pricing or obfuscation depends on the source of market imperfections, such as whether consumers are fully rational or suffer from behavioral biases. The degree to which behavioral factors affect consumer choice is hotly debated by academics and policy makers with markedly different world views. Our paper contributes to this debate. Like Wilson (2010), Ellison-Wolitzky (2012) and others, we demonstrate that there is scope for consumer protection policies even when consumers are fully rational. We find it unhelpful to promulgate policy based on ideological beliefs concerning rationality—appropriate policy should be guided by facts and evidence rather than dogma. Depending on the facts, imperfections in market outcomes may stem from either (a) consumer biases and departures from full rationality, (b) from market imperfections that permit firms to exploit rational consumers, or both. Effective consumer protection policies must (1) target the pathology producing the harm, and (2) minimize the prospect of causing unintended harm to competition or consumers. If, as in our model, drip pricing harms consumers via coordinated hassle costs, rather than through inadequate disclosure or behavioral factors, then policies requiring the cessation of obfuscation (e.g., full up-front disclosure of the total price) will fail on both counts: They target the wrong pathology and could, unintentionally, cause harm.

Over the past 40 years, demands for more and better disclosure have increasingly been viewed as sound consumer protection policy. On the surface, this makes sense; perfect information theoretically enhances competition and improves consumer decisions.¹⁴ A broad implication of our analysis is that an idealized benchmark, like perfect information through a

¹⁴Beales et al (1981) provide a useful summary of the rationale for promoting disclosure, as well as important caveats.

full up-front disclosure, is an inappropriate counterfactual for measuring policy effectiveness. Indeed, Stigler (1961) anticipated this almost 60 years ago when he noted that information acquisition and disclosure is inherently costly and that it is not generally optimal for any decision-maker to pursue or obtain perfect information. Our model of drip pricing illustrates this starkly—a policy requiring full up-front disclosure of the total price provides no benefits to consumers or to competition so long as hassle costs remain the same. Thus, if a policy requiring greater disclosure entails any costs whatsoever, it reduces social welfare. While we view this, too, as a mere benchmark, it highlights an important lesson for contemporary policy makers: Even information disclosure suffers from diminishing (or possibly negative) returns to consumers and to competition.

A Appendix

Proof of Proposition 1

Since $\kappa_m \leq \bar{\kappa} \equiv v(\bar{p}) - c$ for all m , each consumer is willing to visit a randomly selected firm when holding beliefs consistent with the putative equilibrium. Given equilibrium beliefs about consumer behavior (including the consumer's stationary reservation price, r_c), each firm maximizes its expected profits by charging $p_m = \min\{\rho_m, r_c\}$. This induces the distribution of prices, $F_c(p)$. This distribution of prices is consistent with consumer beliefs, and given F_c each consumer maximizes her expected payoff by visiting a randomly selected firm and purchasing if and only if $p_m \leq r_c$. Hence, neither consumers nor firms can gain by pursuing different strategies; see Reinganum (1978, 1979) and Baye et. al (2006) for details. This behavior also constitutes a wPBE in which each consumer believes that each firm charges $p_m = \min\{\rho_m, r_c\}$, and every firm m believes that any consumer visiting its website will purchase $q(p_m) > 0$ units if and only if p_m does not exceed r_c . In particular, (a) each firm is maximizing its expected profits given these beliefs, (b) each consumer maximizes her expected surplus given these beliefs, (c) beliefs are correct in equilibrium, and (d) consumer decisions are optimal at every decision node reached in equilibrium.

Given the timing of decisions, all that remains is to specify out-of-equilibrium consumer beliefs that are consistent with the putative equilibrium behavior. Suppose that, when encountering an out-of-equilibrium action by firm m , consumer beliefs regarding other firms remain unchanged.¹⁵ Then it remains optimal for a consumer to purchase $q(p_m) > 0$ units from firm m if and only if $p_m \leq r_c$. Given out-of-equilibrium consumer beliefs and actions, it follows that firm m finds it optimal to charge $p_m = \min\{\rho_m, r_c\}$. Thus, we have shown that the Reinganum equilibrium is a wPBE. ■

Proof of Proposition 2

Since $p_0 \leq \underline{\rho}$ and consumers know the distribution of firm costs, consumers rationally

¹⁵Essentially, we are assuming consumers hold passive beliefs to permit a direct comparison with Reinganum.

believe that any firm visited will charge a price p_m in excess of $p_0 = 0$ with probability one. Since the initial price is completely uninformative, a consumer must expend an additional κ to discover any firm's total price. It follows that, prospectively, consumers anticipate that the cost of discovering the price charged by any randomly selected firm is $c + \kappa \leq \bar{\kappa}$, which is identical to that in the Reinganum model where exogenous search costs are $c + \kappa$. Since $v(\bar{p}) \geq c + \kappa$ by assumption, an optimizing consumer will sample at least one firm. It follows from arguments akin to those in the proof of the previous proposition that the equilibrium is a wPBE in which consumers visit a randomly selected firm and purchase $q(p_m) > 0$ units if and only if $p_m \leq r_{c+\kappa}$, each firm charges $p_m = \min\{\rho_m, r_{c+\kappa}\}$, and firms' prices are distributed according to $F_{c+\kappa}(p)$. ■

Proof of Proposition 3

We first show that any consumer who reaches a firm's checkout page purchases the product if and only if p_m does not exceed some stationary reservation value, r_σ , and that anticipating this, firms optimally charge p_m . Fix sequences of hassle costs, $\{\kappa_{t,m}\}_{t=1}^T$, such that $\kappa_m = \sum_{t=1}^T \kappa_{t,m} = \kappa \in [0, \bar{\kappa}]$. Note that a consumer who reaches firm m 's checkout page to discover its total price, $p_m = p_{0,m} + \sum_{t=1}^T p_{t,m}$ has expended c to visit the firm's website and another κ in (sunk) hassle costs. Furthermore, she expects to incur an incremental cost of $c + \kappa$ to discover another firm's total price if she does not buy from this firm. It follows that, if consumers believe that firms' total prices are distributed by $F_{c+\kappa}$, then upon observing firm m 's final drip ($p_{T,m}$) and discovering its total price p_m , perfection requires that a consumer buy $q(p_m) > 0$ units from firm m if and only if $p_m \leq r_{c+\kappa}$. Anticipating this, each firm finds it optimal to charge a total price of $p_m = \min\{\rho_m, r_{c+\kappa}\}$ on its checkout page, so consumer beliefs are correct. This implies that in any wPBE where consumers reach the checkout page of any firm visited, the distribution of total prices and consumer purchase behavior in the general model (with $T > 1$ drips and unconstrained initial prices) must coincide with that in the constrained model (with $T = 1$ and $p_0 = 0$).

We next show that, in any wPBE, every firm m in fact chooses an initial price and sequence of drips (e.g., $\{p_{t,m}\}_{t=0}^T$), and elements of the sequence of hassle costs satisfying $\sum_{t=1}^T \kappa_{t,m} = \kappa$ to ensure that every optimizing consumer visiting its website navigates to its checkout page. While each firm's *total* hassle costs are constrained by $\kappa_m = \kappa$, the elements of the sequences are not. But no firm has an incentive to choose a sequence that would induce a consumer to abandon it; indeed, each firm can readily prevent their sequence of hassle costs from triggering abandonment by mimicking other firms (e.g., choosing any $\{\kappa_{t,m}\}_{t=1}^T = \{\kappa_t\}_{t=1}^T$, where $\sum_{t=1}^T \kappa_t \equiv \kappa$). At first blush, it might also seem obvious that no firm has an incentive disclose drips that would induce a consumer to abandon its shopping cart before reaching checkout; the only subtle issue is the ability of low-cost firms to potentially signal their types (e.g., a low marginal cost and hence a low total price) through initial prices. Here, the ability to learn something about firm m 's total price in advance of incurring hassle costs is not relevant unless this prospect induces consumers to abandon some before discovering its total price. A simple argument demonstrates that, while initial prices and drips may convey information, the information can never trigger abandonment in any wPBE.

Consider, for example, firm m 's decision regarding its initial price, $p_{0,m}$. Since the definition of drip pricing requires a firm's initial price to be lower than its total price—and firm m 's optimal total price is $p_m = \min\{\rho_m, r_{c+\kappa}\}$ —optimality requires firm m to choose $p_{0,m} < \min\{\rho_m, r_{c+\kappa}\}$. Importantly, this constraint on the initial price is more stringent for low-cost firms than high-cost firms (owing to the monotonicity of monopoly prices in unit costs); that is, any higher-cost firm can mimic the initial price of any lower-cost firm. Additionally, high-cost firms cannot gain by revealing their type to consumers; nor can low-cost firms gain by attempting to mimic the behavior of a high-cost firm. Among other things, this implies that in any wPBE—including where a low-cost firm might signal that its total price is at the lower end of the distribution—the information revealed through initial prices cannot induce consumers to abandon any firm's shopping cart. A similar argument reveals

that this is also true for every element of the sequence, $\{p_{t,m}\}_{t=1}^T$. Thus, in any wPBE, firms will endogenously choose initial prices and drips to ensure that no consumer abandons its shopping cart prior to checkout. Thus, while there is a unique optimal total price for each firm (and distribution of prices), numerous sequences of drips are generally consistent with wPBE.

Next, we establish that the distribution of total prices and consumer buying behavior are unique. By way of contradiction, suppose not. Since optimality requires that firms set initial prices and drips in a manner that ensures that no consumer abandons it prior to checkout, consumers must employ a stationary reservation price, $r_\sigma \neq r_{c+\kappa}$ and firms must charge $p_m = \min\{\rho_m, r_\sigma\}$ in the putative equilibrium. Let F_σ denote the distribution of total prices induced by $G(m)$ and $p_m = \min\{\rho_m, r_\sigma\}$. A consumer who purchases at r_σ enjoys utility of

$$v(r_\sigma) - c - \kappa$$

Alternatively, if she rejects this total price and searches again, her expected utility is

$$\int_{\underline{\rho}}^{r_\sigma} v(p) dF_\sigma(p) - 2c - 2\kappa$$

Thus, the expected gain to rejecting r_σ and searching again is

$$\begin{aligned} \Delta U &= \int_{\underline{\rho}}^{r_\sigma} v(p) dF_\sigma(p) - 2c - 2\kappa - [v(r_\sigma) - c - \kappa] \\ &= \int_{\underline{\rho}}^{r_\sigma} [v(p) - v(r_\sigma)] dF_\sigma(p) - c - \kappa \\ &= \sigma - (c + \kappa) \end{aligned}$$

where the last equality follows from the definition of r_σ . If $\sigma < c + \kappa$, then $\Delta U < 0$ and hence a consumer strictly gains by accepting prices slightly higher than r_σ , a contradiction. If $\sigma > c + \kappa$, then $\Delta U > 0$ and hence a consumer strictly gains by rejecting a price slightly below r_σ , a contradiction. We conclude that distribution of total prices and consumer buying behavior are unique.

Finally, we establish that this equilibrium is, in fact, a wPBE. It is immediate that (a) each firm's choice of initial and total prices maximizes its expected profits given its beliefs about consumer behavior, (b) each consumer maximizes her expected surplus given beliefs about firm behavior, (c) beliefs are correct in equilibrium, and (d) consumer decisions are optimal at every decision node reached in equilibrium. Given the timing of decisions, all that remains is to specify out-of-equilibrium consumer beliefs that are consistent with the putative equilibrium behavior. Suppose that, when encountering an out-of-equilibrium action by firm m , consumer beliefs regarding other firms remain unchanged. In this case, a firm cannot gain by displaying an out-of-equilibrium initial price (or sequence of drips) that induces a consumer to visit another firm before observing its total price; by doing so the firm would earn profits of zero instead of its equilibrium profits of $(p_m - m) q(p_m) > 0$ by setting $p_m = \min\{\rho_m, r_{c+\kappa}\}$. Furthermore, given her out-of-equilibrium beliefs, it remains optimal for a consumer reaching firm m 's checkout page to purchase $q(p_m) > 0$ if and only if $p_m \leq r_{c+\kappa}$. And given out-of-equilibrium consumer beliefs and actions, firm m finds it optimal to charge $p_m = \min\{\rho_m, r_{c+\kappa}\}$. ■

Proof of Proposition 4

By Proposition 3, it suffices to examine the $T = 1, p_0 = 0$ case. It follows immediately from Proposition 2 that a consumer cannot gain by deviating from a reservation price $r_{c+\kappa}$, given firm strategies in the putative equilibrium. Likewise, it follows from Proposition 2 that no firm can unilaterally gain by changing its price under in the putative equilibrium. It remains to show that no firm can gain by altering κ , with or without a price deviation.

Suppose a firm unilaterally deviates, such that a consumer observes (p, κ') at checkout. If the consumer stops searching, her utility is

$$v(p) + M - c - \kappa'$$

If the consumer rejects the offer and searches again, her expected utility is

$$\begin{aligned}
& F_{c+\kappa}(p) \left(\int_{\underline{p}}^p v(z) \frac{dF_{c+\kappa}(z)}{F_{c+\kappa}(p)} \right) + [1 - F_{c+\kappa}(p)] v(p) + (M - 2c - \kappa' - \kappa) \\
&= \int_{\underline{p}}^p v(z) dF_{c+\kappa}(z) + [1 - F_{c+\kappa}(p)] v(p) + (M - 2c - \kappa' - \kappa)
\end{aligned}$$

Differencing, the expected gain from rejecting the firm's offer and searching again is

$$\int_{\underline{p}}^p [v(z) - v(p)] dF_{c+\kappa}(z) - c - \kappa$$

Since this expression equals zero at $p = r_{c+\kappa}$, such a deviation has no impact on the consumer's reservation price; $r_{c+\kappa}$ remains the same regardless of a deviation in κ . And hence, a firm cannot gain by altering its price. Therefore, the consumer's strategy remains optimal and so the firm gains nothing by altering its price. ■

Proof of Proposition 5

By Proposition 3, it suffices to examine the $T = 1$, $p_0 = 0$ case. Let $C(\kappa) \geq 0$ denote the cost of increasing hassle costs to $\kappa \geq \hat{\kappa}$, with $C(\kappa) = 0$ if and only if $\kappa = \hat{\kappa}$. Since firms choose strategies simultaneously, independently, and are *ex ante* identical, it follows that a consumer's optimal stopping rule consists of a reservation price, r' . Given this stopping rule, a firm of type m optimally charges $p_m = \min\{\rho_m, r'\}$.

Suppose to the contrary that there exists a wPBE where some firm chooses p_m and sets $\kappa_m > \hat{\kappa}$. Since, by choosing $\kappa_m > \hat{\kappa}$, a firm of type m incurs added costs of $C(\kappa_m) > 0$, it follows that such a firm can gain $C(\kappa_m)$ by charging p_m but deviating to $\kappa_m = \hat{\kappa}$. Thus, no firm will choose $\kappa_m > \hat{\kappa}$ in equilibrium. ■

References

- [1] Ahmetoglu, Gorkan, Simon Fried, John Dawes, and Adrian Furnham (2010), “Pricing Practices: Their Effects on Consumer Behaviour and Welfare,” Report Prepared for the Office of Fair Trading.
- [2] Baye, Michael R., and John Morgan (2001), “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review* 91(3): 454-474.
- [3] Baye, Michael R., John Morgan, and Patrick Scholten (2006), “Information, Search, and Price dispersion,” in Elsevier: *Handbook on Economics and Information Systems* 1: 323-375.
- [4] Beales, Howard, Richard Craswell, and Steven Salop (1981), “Information Remedies for Consumer Protection,” *The American Economic Review* 71 (2): 410-413.
- [5] Brown, Jennifer, Tanjim Hossain, and John Morgan (2010), “Shrouded Attributes and Information Suppression: Evidence from the Field,” *The Quarterly Journal of Economics* 125(2): 859-876.
- [6] Department of Transportation (2011), “Enhancing Airline Passenger Protections,” *Federal Register* 76: 23110-23111.
- [7] Department of Transportation (2012), “New Airline Passenger Protections Take Effect This Week,” available at <https://www.transportation.gov/affairs/2012/dot0812.html>.
- [8] Ellison, Glenn, and Alexander Wolitzky (2012), “A Search Cost Model of Obfuscation.” *The RAND Journal of Economics* 43 (3): 417-441.
- [9] Jovanovic, Boyan (1982), “Truthful Disclosure of Information,” *The Bell Journal of Economics* 13(1): 36-44.

- [10] Morgan, Peter, and Richard Manning (1985), “Optimal Search,” *Econometrica* 53 (4) : 923-944.
- [11] Milgrom, Paul R. (1981), “Good News and Bad News: Representation Theorems and Applications,” *The Bell Journal of Economics*, 12(2), Autumn, 380-391.
- [12] Office of Fair Trading (2010a), “Advertising of Prices,” *OFT Report* 1291.
- [13] Office of Fair Trading (2010b), “The Impact of Price Frames on Consumer Decision Making,” *OFT Report* 1226.
- [14] Reinganum, Jennifer F. (1978), “A Simple Model of Equilibrium Price Dispersion.” Northwestern University Working paper No. 335, available at <https://www.kellogg.northwestern.edu/research/math/papers/335.pdf>
- [15] Reinganum, Jennifer F. (1979), “A Simple Model of Equilibrium Price Dispersion.” *Journal of Political Economy*, 87 (4): 851-858.
- [16] Shelanski, Howard A., Joseph Farrell, Daniel Hannerm, Christopher J. Metcalf, Mary W. Sullivan, and Brett W. Wendling (2012), “Economics at the FTC: Drug and PBM Mergers and Drip Pricing,” *Review of Industrial Organization* 41: 303-319.
- [17] Stigler, George J. (1961), “The Economics of Information,” *The Journal of Political Economy*: 213-225.
- [18] Sullivan, Mary W. (2017), “Economic Analysis of Hotel Resort Fees,” Economic Issues Paper, Bureau of Economics, Federal Trade Commission.
- [19] Petrikaitė, Vaiva (2018), “Consumer Obfuscation by a Multiproduct Firm,” *The RAND Journal of Economics* 49 (1): 206-223.
- [20] Wilson, Chris M. (2010), “Ordered Search and Equilibrium Obfuscation,” *International Journal of Industrial Organization* 28 (5): 496-506.