

A Model of Biased Intermediation*

Alexandre de Cornière[†] and Greg Taylor[‡]

Accepted for publication at *the RAND Journal of Economics*
2019

Abstract

We study situations in which consumers rely on a biased intermediary's advice when choosing among sellers. We introduce the notion that sellers' and consumers' payoffs can be *congruent* or *conflicting*, and show that this has important implications for the effects of bias. Under congruence, the firm benefiting from bias has an incentive to offer a better deal than its rival and consumers can be better-off than under no bias. Under conflict, the favored firm offers lower utility and bias harms consumers. We study various policies for dealing with bias and show that their efficacy also depends on whether the payoffs exhibit congruence or conflict.

Keywords: intermediary, bias, regulation.

JEL Classification: D21, L15, L40.

1 Introduction

Consumers often turn to an intermediary for advice when choosing among firms. The intermediary's job is to provide information about which are the best products and which products best fit each consumer's needs. However, the intermediary does not necessarily have the consumer's best interests in mind. Particular concerns have repeatedly been raised about the fact that a vertically integrated intermediary has a clear incentive to

*This paper previously circulated under the title "Quality Provision in the Presence of a Biased Intermediary". We are grateful to the Editor, David Myatt, and two anonymous referees, as well as Diane Coyle, Jacques Crémer, Serena Druifuca, Mikhail Drugov, Anthony Dukes, Péter Esö, Renaud Foucart, Doh-Shin Jeon, Ben Hermalin, Paul Klemperer, Laurent Lamy, Jon Levin, Meg Meyer, Ines Moreno de Barreda, Martin Obradovits, Sander Onderstal, Régis Renault, Andrew Rhodes, Michael Riordan, Andre Veiga, Chengsi Wang, Chris Wilson, Peyton Young, Jidong Zhou, and participants at various seminars and conferences for useful comments and discussions. The authors gratefully acknowledge the financial support of this research by the NET Institute, <http://www.NETinst.org>.

[†]Toulouse School of Economics, University of Toulouse Capitole; alexandre.de-corniere@tse-fr.eu; <https://sites.google.com/site/adecorniere>

[‡]Oxford Internet Institute, University of Oxford; greg.taylor@oii.ox.ac.uk; <http://www.greg-taylor.co.uk>

bias its advice in favour of its own product offerings rather than those of rivals. These concerns have given rise to high-profile investigations in many markets by regulators and competition authorities, who worry about potentially harmful effects of bias.

For example, in the current debate about platform regulation a major issue is that of “own-content bias”, i.e. when a multi-sided platform that is also present on one side of the market behaves in such a way as to steer consumers towards its own products or services. Google, for instance, has been investigated in many jurisdictions for promoting its own products in search results.¹ In 2017 the European Commission imposed a €2.42bn fine after Google was found guilty of favoring its own comparison shopping website. Before that, the U.S. Federal Trade Commission and the High Court of England and Wales had ruled against fining Google for related practices.² Other companies may soon face similar charges: the European Commission recently launched an investigation to determine whether Amazon uses its dominant position in the e-commerce market to favor its own, private-label sellers.³ Some of Apple’s practices also tend to favor its own services at the expense of competitors (see Krämer and Schnurr, 2018, for a thorough discussion). The European Commission, particularly active when it comes to the issue of platform regulation, recently put forward a proposal to require platforms to, among other things, be transparent regarding “how they treat their own goods or services compared to those offered by their professional users”.⁴ This issue may become more prevalent in the not-too-distant future with the growth of the market for virtual assistants, currently dominated by vertically integrated firms like Amazon, Google, and Apple. Forbes already reports that “Eighty-five percent of Amazon customers select the recommended Amazon product when voice shopping”.⁵

The mechanisms of own-content bias may also apply to situations where a leading application or hardware provider (the “intermediary”) chooses one of its own complementary applications as the default option for consumers: Safari as the default browser for iPhones, Google as the default search engine for its browser Chrome, etc. Being the default option is a powerful way to foster usage of a product or service, as the default bias is well-documented (Sunstein and Thaler, 2008). Companies sometimes choose long-term contracts as an alternative to integration: For example, the technology press reports that

¹Edelman and Lai (2013) provide evidence that Google’s own-content bias drives more traffic to its affiliate flight-search service than it would get otherwise.

²See http://ec.europa.eu/competition/elojade/isef/case_details.cfm?proc_code=1_39740 for the EC’s decision, https://www.ftc.gov/sites/default/files/documents/public_statements/statement-commission-regarding-googles-search-practices/130103brillgooglesearchstmt.pdf for the FTC’s one, and <http://www.bailii.org/ew/cases/EWHC/Ch/2016/253.html> for the High Court’s decision in *Streetmap.EU Ltd v Google Inc. & Ors* (2016); all pages accessed 12 October 2018.

³See <https://www.ft.com/content/a8c78888-bc0f-11e8-8274-55b72926558f>, accessed 12 October 2018.

⁴A summary of the proposal is available at http://europa.eu/rapid/press-release_IP-18-3372_en.htm, accessed 12 October 2018.

⁵See <https://www.forbes.com/sites/kirimasters/2018/06/30/amazon-voice-commerce-a-huge-opportunity-for-brands-or-too-early-to-tell/#2080f0a83d43>, accessed 16 October 2018.

heavy bidding forced Google to pay \$300 million for the right to be the default search engine in Mozilla Firefox and \$1 billion for similar rights across Apple’s suite of products.⁶

The idea that an intermediary might provide biased recommendations because of financial incentives is, of course, not restricted to the digital economy. Similar concerns have been raised in the pharmaceutical and financial industries, where doctors and financial advisers sometimes receive commissions potentially leading them to recommend one treatment/investment over another which would be better suited to their patient/client (Engelberg, Parsons, and Tefft, 2013; Cookson et al., 2017; Egan, 2017). Default option bias may also apply there, as when a bank offers its own insurance alongside a loan.

Biased intermediation raises two main concerns: First, are consumers directed towards inferior or higher priced products or services? Second, does bias distort firms’ incentives in a way that harms consumers? Even though these concerns are common to the situations described above, a certain degree of heterogeneity remains as to the institutional or strategic environments. A key concern in the price comparison website industry, for example, has been the effect of paid endorsements on final consumer prices. Search engines and websites, by contrast, are mostly free to access so that attention in the search bias debate has focused on innovation and investment decisions rather than on pricing. It is therefore doubtful that the analysis of a particular market could be transposed directly to another one. We show that, in order to understand the effects of bias, it is often possible to reduce complex strategic heterogeneity to a simpler problem of identifying the relationship between firms’ and consumers’ marginal payoffs. This insight allows us to nest a variety of strategic environments and study how they drive both the consequences of bias and the efficacy of various policy interventions.

More concretely, we consider a market composed of two sellers, a unit mass of consumers, and one intermediary. Sellers are differentiated, and at least some of the consumers cannot observe the sellers’ characteristics and offers. This creates a role for the intermediary, whose technology allows it to identify the best match for each consumer. However, the intermediary is integrated with one of the sellers and therefore has incentives to produce biased recommendations, i.e., to recommend its own seller more often than would be objectively justified. We measure the extent of bias as the share of uninformed consumers who are wrongly directed towards the integrated seller.

In order to accommodate the strategic heterogeneity mentioned above, we abstract from the details of firm’s product design choices and use a modelling framework where sellers make two decisions: (i) the level of utility they offer to consumers, and (ii) their per-unit mark-up. This framework, which encompasses the Armstrong and Vickers (2001)

⁶See <http://allthingsd.com/20111222/google-will-pay-mozilla-almost-300m-per-year-in-search-deal-besting-microsoft-and-yahoo/> and <https://www.recode.net/2014/4/16/11625704/marissa-mayers-secret-plan-to-get-apple-to-dump-google-and-default-to->, accessed 12 October 2018.

competition-in-utility model, allows us to nest a variety of strategic environments: pure price competition; quality competition among free, ad-funded products; competition in price and quality; competitive investments in cost-reducing technology; and others.

Within this model, we introduce the notions of *conflict* and *congruence* between a seller's and its customers' payoffs. Conflict arises when the most efficient way for a seller to increase the utility offered to consumers is to reduce its per-unit mark-up (as in standard price competition, where such a reduction corresponds to a price decrease). Congruence, on the other hand, arises when higher utility levels are also associated with higher mark-ups. This is the case, for instance, when quality is the most important dimension for competition (in a sense made precise below).

Our main contribution is to show that the notions of congruence and conflict are critical in determining (i) the effect of bias on competition between sellers, (ii) the implications for consumer surplus, and (iii) the efficacy of various policy interventions.

Intuitively, in our model, bias acts as a demand-shifter by increasing the number of consumers directed to the integrated seller and by reducing the demand for its competitor. This provides an incentive for the integrated seller to select a strategy yielding a higher per-unit mark-up—corresponding to a higher utility under congruence, and to a lower utility under conflict. The opposite effects are at play for the non-integrated seller.

In terms of consumer surplus, bias is always harmful under conflict: consumers are mis-matched more often, and the favored seller offers a lower utility than its rival. Under congruence, on the other hand, bias is sometimes good for consumers because it provides the favored seller with stronger incentives to offer higher levels of utility, which may offset the mis-match costs. Thus, suppose a regulator is considering reducing bias (e.g., by imposing a fiduciary duty on the intermediary). Such an intervention will generally improve consumer outcomes under conflict, but may not do so under congruence. Besides the effects of direct regulation of bias, we study several other policy interventions that have been mooted in recent bias cases: neutrality (forcing the intermediary to grant equal prominence to the two sellers), transparency policies that alert consumers to the presence of bias, and the break-up (or divestiture) of the integrated firm. The neutrality policy is always ineffective, while the efficacy of the other interventions extends only to situations of conflict. Thus, whether the environment exhibits congruence or conflict not only drives the effects of bias, but also influences the range of effective tools available to policy makers. As part of our analysis of transparency, we endogenize the intermediary's choice of the level of bias and find that bias tends to be higher under congruence than conflict.

1.1 Related literature

This paper contributes to the literature on information intermediaries. Several recent papers study intermediaries whose role is to help consumers choose among products, and focus on how commissions or other contracts affect both the quality of the recommendation and sellers' behavior. Inderst and Ottaviani (2012a) study the effects of mandatory disclosure of commissions, while Inderst and Ottaviani (2012b) focus on the intermediary's equilibrium compensation structure and associated policies. In Teh and Wright (2018) firms pay commissions to influence the intermediary's advice and pass the cost on to consumers; but the intermediary is unbiased in equilibrium because firms offer symmetric commissions. De Cornière and Taylor (2014) and Burguet, Caminal, and Ellman (2015) study the determinants of search engine bias and its effect on websites' strategies. Armstrong and Zhou (2011) look at several models where firms can become prominent, one of which involves the payment of commissions to an intermediary. These papers only consider environments with, using our terminology, *conflicting payoffs*. An important contribution of our work is to consider a richer strategic environment in which payoffs may be congruent rather than conflicting. Indeed, we will see that congruence arises endogenously from fairly natural models of competition. Moreover, we show that both the implications of bias and the efficacy of various policy responses are quite different depending on whether the environment is one of congruence or conflict.

A related literature studies certification intermediaries, whose role is to disclose to consumers the quality of the products that are offered. In contrast with the papers mentioned above and with ours, this literature has mostly abstracted away from competition between sellers.⁷ In the absence of commitment power by the intermediary,⁸ an important concern is the threat of collusion between the intermediary and the firms it is supposed to certify. Several papers examine the role of reputational incentives as a disciplining device for the intermediary (Biglaiser, 1993; Strausz, 2005; Peyrache and Quesada, 2011; Durbin and Iyer, 2009). Rather than studying *conditions* under which collusion can or cannot occur when qualities are exogenous, we study the *effects* of collusion ("bias" in our model) on the equilibrium behavior of firms (e.g., choice of quality and price).⁹

Some papers cover related themes in the context of intermediation in online markets. One aspect of the net neutrality debate concerns agreements between Internet Service Providers (ISPs) and Content Providers (CPs) whereby some CPs can ensure preferential treatment for themselves by paying the ISP (a "fast lane access"). Choi and Kim (2010) and Economides and Hermalin (2012), for instance, study how such agreements affect the ISP's incentives to invest. In contrast to this, we focus on *sellers'* investment incentives.¹⁰

⁷An exception is Buehler and Schuett (2014).

⁸See Lizzeri (1999) and Albano and Lizzeri (2001) for models with perfect commitment power, respectively with exogenous and endogenous quality.

⁹Biglaiser and Friedman (1994) deal with reputational incentives in a setup with endogenous qualities.

¹⁰A few papers look at downstream investment, but from a different perspective to that taken here.

Another series of papers study the effect of news aggregators on competition among content providers (Dellarocas, Katona, and Rand, 2013; Jeon and Nasr, 2016; Rutt, 2011). News aggregators help consumers identify quality content, and the above papers study how their presence affects content providers’ incentives to invest in quality. Unlike the present paper, this literature has not investigated cases where the intermediary is biased towards a subset of content providers.

A closely related literature is that on the tying of complementary products (e.g., Whinston, 1990; Nalebuff, 2004; Carlton and Waldman, 2002; Carbajo, De Meza, and Seidmann, 1990; Chen, 1997). In particular, Choi (2004) considers how tying may affect both the pricing and the incentives to invest in quality for the tied product. Some of our results are reminiscent of his (in particular that bundling has a positive effect on the tied product’s quality and a negative one on the rival’s product). However, his focus is quite different from ours, as he is mostly interested in the profitability of bundling in a setup with strategic substitutes whereas we emphasize the importance of the congruence/conflict dichotomy and consider a wider set of policy interventions.

2 The model

The market we consider is composed of two sellers ($i \in \{1, 2\}$), a unit mass of consumers, and one intermediary that is integrated with seller 1.

2.1 Competition between sellers: the (u, r) -model

In modelling competition between sellers, our objective is to have a parsimonious framework that still allows us to capture a variety of strategic environments beyond simple price competition. To do so, we follow Armstrong and Vickers (2001) by assuming that sellers compete in utilities. More precisely, suppose that seller i can design an offer that generates a base level of utility u_i to its customers. Suppose that the associated demand function is $D_i(u_1, u_2)$, increasing in u_i and decreasing in u_j (we adopt a Hotelling specification for demand, as described in the next subsection).

In general, a base level of utility u_i is achieved through various instruments (price, quality, number of advertisements displayed alongside content, etc.). However these instruments are not equivalent. Increasing quality may require a seller to pay larger fixed costs,¹¹ while the “cost” of lowering the price consists in a lower mark-up, and is

Choi and Kim (2010) assume that such investments do not benefit consumers, unlike in our setup. Choi, Jeon, and Kim (2015) and Peitz and Schütt (2016) study investments that have externalities for other sellers (to reflect the issue of congestion, which is an important technical aspect of the net neutrality debate), which is not a feature of our environment.

¹¹For example, when the product in question is a website’s content, the costs of improving quality (hiring more and better writers, substituting original for syndicated content, etc.) are almost entirely fixed.

thus mostly borne on a per-consumer basis.¹² To allow for a discussion of these various instruments, we write profits as a function of both the utility offered and the per-consumer mark-up, r_i . The combination of u_i and r_i generates a fixed cost $C(u_i, r_i)$. This cost is such that $\frac{\partial^2 C}{\partial u^2} \geq 0$, $\frac{\partial^2 C}{\partial r^2} \geq 0$, and $\frac{\partial^2 C}{\partial u \partial r} \geq 0$.¹³ Seller i 's profit is therefore

$$\pi_i(r_i, u_i, u_j) = r_i D_i(u_1, u_2) - C(u_i, r_i). \quad (1)$$

For a given u_j , we assume throughout the paper that π_i is concave.

Let us discuss some examples of environments that can be subsumed into the (u, r) framework. These fall into two categories: two-dimensional cases where sellers choose both u and r ; and one-dimensional cases where sellers choose (only) u , with $r(u)$ determined as a function of that choice.

Two dimensions: Price and quality competition with a fixed cost The most straightforward application of the (u, r) -framework is a situation where sellers have control over a price p_i and a quality q_i . If the cost of quality is $K(q_i)$, sellers' marginal cost is γ , and if utility is $u_i = q_i - p_i$, we can write $r_i = p_i - \gamma$ and $C(u_i, r_i) = K(u_i + r_i + \gamma)$. The seller's problem is then transformed into $\max_{u_i, r_i} \pi_i(r_i, u_i, u_j)$, as in (1).

A very similar logic would apply if sellers could choose a price and invest in a reduction of marginal cost.

Two dimensions: A seller with three decision variables This (u, r) -framework may also be applied to some situations where sellers' vector of decisions is of dimension larger than 2. For instance, suppose that sellers simultaneously choose a price p_i , a quality q_i and a cost-reduction level Δ_i so that the marginal cost is $\gamma - \Delta_i$ (the cost of this investment is $G(\Delta_i)$, with $G'' > 0$).

We then have $u_i = q_i - p_i$ and $r_i = p_i - \gamma + \Delta_i$. We can write profit as $r_i D_i(u_1, u_2) - K(u_i + p_i) - G(r_i - p_i + \gamma)$. For a given (u_i, r_i) , let $p(u_i, r_i)$ be the value of p_i that minimizes $K(u_i + p_i) + G(r_i - p_i + \gamma)$ (a convex function of p_i). We can write $C(u_i, r_i) \equiv K(u_i + p(u_i, r_i)) + G(r_i - p(u_i, r_i) + \gamma)$, so that seller i 's problem can be written in the form $\max_{u_i, r_i} \pi_i(r_i, u_i, u_j)$.

One dimension: Pure pricing and pure quality models In other setups, sellers may not have the ability to choose both r_i and u_i , due to a lack of available instruments. For instance, if sellers' only available strategic tool is their price then u_i and r_i are simultaneously determined by the price level:¹⁴ $r_i = p_i$ and $u_i = v - p_i$ (v is the exogenous

¹²Some quality improvements may only entail higher marginal costs (e.g. if they consist in using higher quality inputs), in which case they are formally closer to a price reduction.

¹³In most of our applications, the cost is a function of $u_i + r_i$.

¹⁴Here we normalize the marginal cost to 0.

willingness to pay). The cost associated to any level of utility is $C(u_i, r(u_i)) = 0$, which corresponds to the model of Armstrong and Vickers (2001). In this case, we can write $r_i = v - u_i \equiv r(u_i)$, and the seller's profit is simply a function of u_i : $\pi_i = r_i(u_i)D_i(u_1, u_2)$.

For another uni-dimensional example, suppose that 1 and 2 are websites whose revenues solely come from advertising. They can invest $K(q_i)$ in the quality q_i of their content, which simultaneously determines the utility $u_i = u(q_i)$ that consumers get, and advertisers' willingness to pay to be displayed to a consumer alongside the content, $r_i = r(q_i)$. Given $u'(q_i) > 0$, we can think of the seller as choosing u_i , resulting in quality $q(u_i) = u^{-1}(u_i)$ and revenue $r(q(u_i))$. Again, we can write the seller's problem as $\max_{u_i} \pi_i = \max_{u_i} \{r(q(u_i))D_i(u_1, u_2) - K(q(u_i))\}$.

One dimension: Price-quality competition with per-unit quality costs Another possibility is that sellers choose price and quality, (p_i, q_i) , but with $K(q_i)$ being a marginal rather than fixed cost. We then have $u_i = q_i - p_i$ and $r_i = p_i - K(q_i)$. The seller's problem can be reformulated as $\max_{u_i} r_i(u_i)D_i(u_1, u_2)$, where $r_i(u_i) = \max_{q_i} \{q_i - u_i - K(q_i)\}$ is the most profitable way of providing u_i . Thus, the problem reduces to the one-dimensional case where the seller chooses only u_i .

2.2 Demand side: preferences, information, and intermediation

We adopt the Hotelling formulation to model demand:¹⁵ sellers are located at the ends of a unit length segment, over which consumers are uniformly distributed. The utility that a consumer derives from consumption of product i is $u_i - td_i$, where d_i is the distance separating the consumer from seller i and t is the differentiation parameter. A consumer's utility is maximized by choosing seller 1 if his location, x , satisfies $x < x^* \equiv \frac{1}{2} + \frac{u_1 - u_2}{2t}$, and seller 2 if $x > x^*$.¹⁶ A fraction $1 - \mu$ of consumers are "informed" and choose the seller that maximizes their utility. The remaining μ consumers are "uninformed" and rely on the intermediary to help them choose a seller.

The intermediary observes both sellers' offers and each consumer's location. One interpretation is that consumers submit a precise query to the intermediary, which can then infer how suited to the consumers' needs the two products are. The intermediary then provides advice to the consumer.

The intermediary is integrated with seller 1. Therefore, while it has the ability to match each consumer to their best option it does not necessarily have the incentive to do so. Indeed, everything else equal, the intermediary would want to direct as many consumers as possible towards seller 1. Thus, the intermediary biases its recommendation to steer

¹⁵In an earlier version of this paper we found similar results in a model without horizontal differentiation. In Appendix B we show that the main result holds with various alternative discrete choice demand specifications.

¹⁶Throughout the paper we focus on cases where the market is covered, and ignore the possibility of large deviations that result in the market being uncovered.




Figure 1: Allocation of uninformed consumers

consumers who are close to indifferent towards the integrated seller (for example, by featuring seller 1 more prominently or exaggerating seller 1's advantages). More formally, we assume that the intermediary is able to steer uninformed consumers with $x < x^* + b$ to seller 1, leaving only those with $x > x^* + b$ to choose seller 2 (see Figure 1).¹⁷ Thus, b is a measure of the intermediary's ability to influence consumers through bias. We endogenize the level of bias, b , in section 4.

The total demands (from informed and uninformed consumers) are therefore given by

$$D_1(u_1, u_2, b) = \left(\frac{1}{2} + \frac{u_1 - u_2}{2t} \right) + \mu b, \quad (2)$$

$$D_2(u_1, u_2, b) = \left(\frac{1}{2} + \frac{u_2 - u_1}{2t} \right) - \mu b. \quad (3)$$

We solve for the (Nash) equilibrium choices of u and r by the sellers.

3 Equilibrium Analysis

3.1 Conflict and congruence

We can decompose a seller's problem into two parts. The first step is to identify the value of r_i , denoted $\hat{r}_i(u_i, u_j, b)$, that would accompany a given choice of u_i . In one-dimensional cases such as pure price or pure quality competition, we exogenously have $\hat{r}_i(u_i, u_j, b) = r(u_i)$. When the seller can choose both u_i and r_i , on the other hand, $\hat{r}_i(u_i, u_j, b)$ is endogenously chosen as the most profitable way to provide a target level of utility. The first-order condition giving \hat{r}_i , namely $\frac{\partial \pi_i(r_i, u_i, u_j, b)}{\partial r_i} = 0$, is then

$$D_i(u_1, u_2, b) = \frac{\partial C(u_i, \hat{r}_i(u_i, u_j, b))}{\partial r_i}. \quad (4)$$

¹⁷Formally, it is as though the intermediary were able to increase consumers' perceived value of seller 1 by $2tb$. Indeed, we can write

$$x^* + b = \frac{\tilde{u}_1 - u_2 + t}{2t},$$

where $\tilde{u}_1 = u_1 + 2tb$ is consumers' *perceived* value of good 1. This would lend bias a similar interpretation to persuasive advertising à la Dixit and Norman (1978).

The optimal way to provide a utility level u_i depends on u_j , through its effect on demand. For example, if firm i faces a choice between reducing its price or increasing its quality, the former option is more attractive when there are fewer consumers, i.e. when u_j is larger.

Given \hat{r}_i , i 's decision is reduced to a one-dimensional problem of choosing u_i to maximize $\pi_i(\hat{r}_i(u_i, u_j^*, b), u_i, u_j^*)$, where u_j^* is the expected equilibrium play of j .

We now introduce the notions of conflict and congruence, which will play a key role in our analysis.

Definition 1. For a given u_j , we say that seller i 's and its customers' payoffs are *conflicting* if $\frac{\partial \hat{r}_i(u_i, u_j, b)}{\partial u_i} < 0$. They are *congruent* if $\frac{\partial \hat{r}_i(u_i, u_j, b)}{\partial u_i} > 0$.

In the uni-dimensional case, congruence or conflict is an intrinsic property determined entirely by the (exogenous) sign of $r'(u)$. When sellers choose both u and r , on the other hand, congruence or conflict arise endogenously as a strategic consequence of firms' product design decisions. Whether payoffs are congruent or conflicting then depends on the sign of the cross-derivative $\frac{\partial^2 \pi_i}{\partial r_i \partial u_i}$: there is congruence if it is positive, and conflict otherwise.

Notice that the definitions of conflict and congruence apply to the mark-up \hat{r}_i , not to the entire profit π_i . In other words, congruence does not imply that seller i 's profit increases as it offers a higher level of utility.

Most previous work has focused on environments that exhibit conflict. The possibility of congruence has received much less attention, despite the fact that it emerges quite naturally in plausible models of competition. Let us apply Definition 1 to some of the examples discussed above:

Price and quality model with fixed costs Suppose sellers choose a price and a quality, with positive fixed cost and marginal cost normalized to zero. We saw above that $\pi_i(r_i, u_i, u_j) = r_i D_i(u_1, u_2, b) - C(u_i + r_i)$. We therefore have

$$\text{congruence} \iff \frac{\partial^2 \pi_i(r_i, u_i, u_j)}{\partial r_i \partial u_i} > 0 \iff \frac{\partial D_i(u_1, u_2, b)}{\partial u_i} - C''(u_i + r_i) > 0.$$

Payoffs are more likely to be congruent when demand is more sensitive to the utility offered, and when the cost of providing quality is not too convex. Intuitively, if seller i finds it optimal to increase utility mostly by reducing its price then payoffs are conflicting. If, on the other hand, it finds it optimal to increase both its quality and, to a smaller extent, its price, then payoffs are congruent.

An example to which we will return is the price-quality model where costs are quadratic:¹⁸

$$u_i = v + q_i - p_i, \quad r_i = p_i, \quad \text{fixed cost } K(q_i) = \frac{C}{2} q_i^2 = \frac{C}{2} (u_i + r_i - v)^2. \quad (5)$$

¹⁸We abuse notation by using C to denote the constant cost parameter in this example. Parameter v , which may equal zero, is the value of the product when a firm does not invest.

From (5), $q_i = u_i - v + r_i$ so we can substitute q_i out of the profit expression. Solving the resulting first-order condition $\frac{\partial \pi_i}{\partial r_i} = 0$ yields \hat{r}_i . The optimal (p, q) to provide a given utility is then:

$$p_i(u_i, u_j, b) \equiv \hat{r}_i(u_i, u_j, b) = \frac{1 - 2Ct}{2Ct}u_i + \left[v + \frac{2tb\mu + t - u_j}{2Ct} \right], \quad (6)$$

$$q_i(u_i, u_j, b) \equiv u_i + p_i(u_i, u_j, b) - v = \frac{t + u_i - u_j + 2bt\mu}{2Ct}.$$

From (6), we have congruence if $Ct < 1/2$ and conflict if $Ct > 1/2$. Notice that the intensity of competition, as measured by the parameter t (a lower t meaning that competition is more intense), determines whether payoffs are congruent or conflicting: intense competition goes hand in hand with congruent payoffs. To see why, notice first that the cost of provided a higher utility is either fixed if this is achieved through an increase in quality, or variable if achieved through a lower price. When t is small, demand increases a lot following an increase in u_i , which leads firm i to prefer to use a mix of higher quality and not-so-higher price (i.e., congruence) rather than to use a decrease in the price.

Pure price and pure quality competition models In the one-dimensional model, congruence or conflict is given by the sign of $r'(u_i)$. In the model where sellers only compete in prices (Armstrong and Vickers, 2001), we have $r_i \equiv p_i = v - u_i$, so $r'(u_i) < 0$. In words: payoffs are conflicting.

In the model where websites compete purely in content quality, there is congruence if $r'(u_i) > 0$. This is true if $u'(q_i)$ and $r'(q_i)$ have the same sign—that is, if advertisers are willing to pay more to be exposed alongside content of higher quality, a reasonable assumption in practice.

An example parameterization that accommodates both the pure price and pure quality cases is

$$r(u_i) = v - \psi u_i, \quad \text{fixed cost } C(u_i) = \frac{C}{2}u_i^2, \quad (7)$$

which yields conflict if $\psi > 0$ and congruence when $\psi < 0$.

Price and quality competition with per-unit cost Suppose sellers choose price and quality, with the cost of quality provision incurred on a per-unit basis. We saw in Section 2.1 that $r(u_i) = \max_{q_i} \{q_i - u_i - K(q_i)\}$. This is decreasing in u_i , so the market is one with conflict.

3.2 Equilibrium effects of bias

Recall that x^* is the location of the consumer who is indifferent between firms 1 and 2. Suppose that b is small enough that $x^* + b \leq 1$ in equilibrium. Demands are given by (2)

and (3).

Define $\Pi_i(u_i, u_j, b) \equiv \pi_i(\hat{r}_i(u_i, u_j, b), u_i, u_j, b)$ as the profit evaluated at the optimal r_i , and let $\hat{u}_i(u_j) \equiv \operatorname{argmax}_{u_i} \Pi_i(u_i, u_j, b)$ be seller i 's best-response to the utility provided by seller j . We make the following assumption:

Assumption 1 (Stability). For $i = 1, 2$, $|\hat{u}'_i(u_j)| < 1$.

Assumption 1 ensures that the equilibrium, which is such that $\hat{u}_i(u_j^*) = u_i^*$, is stable.¹⁹ For example, Assumption 1 requires $Ct > 1/3$ in the price-quality specification of (5), or $Ct > -3\psi/2$ in the one-dimensional specification given in (7).

To obtain the equilibrium effect of bias on u_1 and u_2 , one needs to take into account both the direct and the strategic effects: bias affects the best-response function of each seller (direct effect), but also what seller i expects seller j to play (strategic effect). In other words, we must consider both the shift in sellers' reaction functions, and their slopes. The next Lemma, which we obtain by totally differentiating the first-order condition $\frac{\partial \Pi_i}{\partial u_i} = 0$, characterizes these effects. Omitted proofs are in the Appendix.

Lemma 1. *Consider the best-response functions $\hat{u}_i(u_j)$.*

(i) *An increase in b causes \hat{u}_1 to shift up when 1's payoffs are congruent, and to shift down when 1's payoffs are conflicting.*

(ii) *An increase in b causes \hat{u}_2 to shift down when 2's payoffs are congruent, and to shift up when 2's payoffs are conflicting.*

(iii) *u_1 and u_2 are strategic substitutes when payoffs are congruent, and strategic complements when payoffs are conflicting.*

The intuition for the direct effect in parts (i) and (ii) is the following: an increase in bias b means that seller 1 faces relatively more inframarginal, or "captive", consumers, and that seller 2 faces relatively fewer of them. The incentive to choose a higher per-consumer revenue r_i is therefore increased for seller 1, and decreased for seller 2. Under congruence, a higher r_1 corresponds to a higher u_1 , and a lower r_2 to a lower u_2 . The reverse holds under conflict.

Notice that an increase in b can be interpreted both as a decrease in seller 1's demand elasticity (thereby suggesting a higher price) and as an increase in its scale (usually associated with higher quality). Whether payoffs are congruent or conflicting determines which of these two forces dominates.

Regarding the issue of strategic substitutability or complementarity, the intuition is more subtle. Suppose that payoffs are congruent. The marginal effect on π_i of an increase in seller j 's offer is $\hat{r}_i \frac{\partial D_i}{\partial u_j}$. In the Hotelling model, the term $\frac{\partial D_i}{\partial u_j}$ is independent of u_i (see

¹⁹ Assumption 1 is a sufficient condition for stability. The necessary condition is that $|\hat{u}'_1(u_2)| |\hat{u}'_2(u_1)| < 1$. However, if Assumption 1 was violated, the equilibrium would have the implausible property that $\frac{dD_1(u_1^*, u_2^*, b)}{db} < 0$ when payoffs are conflicting.




Figure 2: Congruent payoffs: u_1 increases in b , u_2 decreases.

below for a discussion). Therefore, under congruence, the impact of a rise in u_j is more important for high values of u_i (which correspond to high values of \hat{r}_i). This means that $\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j} < 0$, i.e. $\hat{u}_i(u_j)$ is decreasing. The reverse is true under conflict.

Given Lemma 1, we can represent the equilibrium effect of bias graphically. Figure 2 corresponds to the case of congruence. There, the direct and strategic effects go in the same direction: more bias leads u_1 to go up and u_2 to go down (direct effect). By strategic substitutability, the increase in u_1 is reinforced by the decrease in u_2 , and reciprocally. The overall effect is therefore clear: bias leads u_1 to increase, and u_2 to decrease.

In the case of conflicting payoffs, things are less clear a priori (see Figure 3). If the direct effect implies that u_1 goes down while u_2 goes up, strategic complementarity goes in the opposite direction. We show in Proposition 1 that, under our assumptions, the direct effect always dominates.

Proposition 1. *Let $u_i^*(b)$ be the equilibrium utility offered by seller i .*

1. *If seller 1's payoffs are congruent then $u_1^{*'}(b) > 0$; if 1's payoffs are conflicting then $u_1^{*'}(b) < 0$.*
2. *If seller 2's payoffs are congruent then $u_2^{*'}(b) < 0$; if 2's payoffs are conflicting then $u_2^{*'}(b) > 0$.*

Proposition 1 implies that when both sellers' payoffs are conflicting (e.g., when prices are the most important dimension for competition), we have $u_1 < u_2$. When sellers' payoffs are congruent the opposite is true: $u_1 > u_2$. In both cases, $|u_1 - u_2|$ increases in b .




Figure 3: Conflicting payoffs: u_1 decreases in b , u_2 increases.

3.3 Welfare analysis

In order to obtain results about the effects of the policies on consumer surplus and welfare, we restrict attention to the one-dimensional and the two-dimensional setups with quadratic costs, described in (7) and (5) respectively. Recall that we have conflict in the one-dimensional case when $\psi > 0$, and in the two dimensional case when $Ct > 1/2$ (and congruence otherwise).

As a preliminary result, the following Proposition describes the equilibrium in both the one- and two-dimensional cases.

Proposition 2. *In the two-dimension case (where payoffs are given by (5)), if the level of bias is b , we have :*

$$u_1^* = v + \frac{(1 - 2Ct)(1 - Ct(2b\mu + 3))}{2C(1 - 3Ct)}, \quad u_2^* = v + \frac{(1 - 2Ct)(1 - Ct(3 - 2b\mu))}{2C(1 - 3Ct)}. \quad (8)$$

In the one-dimension case (where payoffs are given by (7)), if the level of bias is b , we have :

$$u_1^* = \frac{2Ct^2 + v}{2Ct + \psi} - t - \frac{2bt\mu\psi}{2Ct + 3\psi}, \quad u_2^* = \frac{2Ct^2 + v}{2Ct + \psi} - t + \frac{2bt\mu\psi}{2Ct + 3\psi}. \quad (9)$$

Consumer surplus is then

$$\begin{aligned} \text{CS} = (1 - \mu) \left\{ \int_0^{x^*} [u_1^* - tx] dx + \int_{x^*}^1 [u_2^* - t(1 - x)] dx \right\} + \\ \mu \left\{ \int_0^{x^*+b} [u_1^* - tx] dx + \int_{x^*+b}^1 [u_2^* - t(1 - x)] dx \right\}. \end{aligned} \quad (10)$$

Comparative statics on b lead to the following:

Proposition 3. *There exists values \overline{Ct}_S and \overline{Ct}_W (given in the proof) such that:*

1. *Under conflict, consumer surplus always decreases as b increases. Under congruence, consumer surplus is increasing in b if $Ct < \overline{Ct}_S$, decreasing otherwise.*
2. *Total welfare increases in b if $Ct < \overline{Ct}_W$, and decreases otherwise.*
3. *Industry profit is always increasing in b , so that $\overline{Ct}_S \leq \overline{Ct}_W$.*

The one-dimensional model has the advantage that the existence of congruence/conflict is determined solely by ψ , independent of the model's demand and cost parameters. It therefore allows us to study in isolation the effect of conflict and congruence. We find that the two regimes have quite different implications for consumer surplus, with bias being unambiguously harmful under conflict but not under congruence. The same pattern is repeated for the two-dimensional case.

Under congruence, bias causes seller 1 to endogenously increase its utility offer in response to a larger captive audience. Thus, more consumers are steered towards a better seller as the level of bias increases. This gain from improved investment must be weighed against the preference mis-matching implied by bias. If Ct is not too large then the positive effects of bias dominate because (i) low transport costs mean that the mismatching is not very harmful, and (ii) low C amplifies the increase in seller 1's investment in its product.

Under conflict, an increase in bias results in more consumers being directed to a seller at the same time as that seller's utility offer becomes endogenously *worse*. Combined with the mismatching implied by bias, this must leave consumers worse-off in aggregate.

Regarding industry profit, an increase in b moves the situation nearer to a monopoly one, and therefore increases total profit (even though firm 2 is worse-off). This implies that, even if firm 2 could offer side payments to the integrated intermediary, this would not be enough to deter it from offering biased results to consumers. Given that industry profit increases with bias, total welfare is more likely than consumer surplus to also increase with b , which explains that $\overline{Ct}_S < \overline{Ct}_W$.

Figure 4 illustrates Proposition 3.

3.4 Discussion

Robustness to alternative model specifications Our baseline model makes two functional form assumptions: consumer preferences are taken to be à la Hotelling, and bias is taken to be additive. We begin by discussing the robustness of Proposition 1 to these assumptions.

Firstly, the additive bias specification ($D_1 = x^* + \mu b$, $D_2 = 1 - x^* - \mu b$) reflects the idea that bias—perhaps achieved by exaggerating seller 1's qualities, making it prominent, or otherwise giving it a slight advantage—is most likely to influence consumers who




Figure 4: Welfare effects of bias

were close to indifferent to begin with. An alternative specification would see consumers affected by biased advice uniformly at random, yielding demand $D_1 = x^*(1 - b) + b$ and $D_2 = (1 - x^*)(1 - b)$.²⁰ The main practical difference is that an increase in b now reduces the sensitivity of demand for both firms: $\partial^2 D_i / \partial u_i \partial b < 0$. This implies that seller 2's best response function sometimes shifts down with bias under conflict and seller 1's best response sometimes decreases under congruence, which can make the sign of du_i / db ambiguous for one of the sellers. Nevertheless, a slightly weaker version of Proposition 1 continues to hold. Indeed, in Appendix B we prove that bias still induces $u_1^* > u_2^*$ under congruence and $u_2^* > u_1^*$ under conflict, with the difference increasing in b .

Secondly, the Hotelling demand specification, $x^* = \frac{1}{2t}(t + u_1 - u_2)$, is such that $\partial^2 D_i / \partial u_i \partial u_j = 0$. Under alternative demand specifications this cross derivative is non-zero, with the consequence that congruence no longer implies strategic substitutability, nor does conflict imply strategic complementarity.²¹ We study the robustness of our model to such alternative demand specifications in Appendix B. There, we compute the equilibrium for two popular discrete choice demand models (logit and probit), and show that, in both cases, Proposition 1 continues to hold. In particular, even where the correspondence between congruence/conflict and substitutability/complementarity breaks down, u_1^* increases with b under congruence and decreases under conflict (with u_2^* moving

²⁰We take $\mu = 1$ here for the sake of brevity.

²¹Taking the two-dimensional case as an example, the slope of the reaction functions has the same sign as

$$\underbrace{\frac{\partial \hat{r}_i}{\partial u_j}}_{<0} \underbrace{\frac{\partial^2 \pi_i}{\partial u_i \partial r_i}}_{(a)} + \hat{r}_i \frac{\partial^2 D_i}{\partial u_i \partial u_j},$$

where term (a) is positive when payoffs are congruent and negative when they are conflicting.

in the opposite direction).

Lastly, in Appendix B we also study the case with both the alternative multiplicative specification for bias and the alternative discrete choice demand specifications. Again, we find that the basic message of the model is robust.²²

A note on corner solutions Proposition 1 holds when $x^* + b \leq 1$, i.e. at an interior equilibrium. In such an equilibrium seller 2 can attract additional uninformed consumers by slightly increasing u_2 . If bias is very strong (b very large) then we would instead obtain a corner solution in which seller 2 does not serve any of the uninformed consumers in a neighbourhood around u_2^* . A similar result to Proposition 1 obtains in this case.

Corollary 1. *In a corner equilibrium $u_1^* > u_2^*$ if payoffs are congruent, and $u_2^* > u_1^*$ if payoffs are conflicting.*

In order to avoid having to deal with different regimes and streamline the exposition, from now on we focus on cases where the solution is interior:

Assumption 2 (Interior Equilibrium). $b \leq \bar{b}$, where \bar{b} satisfies $\frac{1}{2} + \frac{u_1^*(\bar{b}) - u_2^*(\bar{b})}{2t} + \bar{b} = 1$.²³

4 Policy analysis

The existence of bias in equilibrium is a priori problematic, as it means that at least some consumers are misled by the intermediary. In this section we study several policy interventions that have been contemplated or implemented in markets where biased intermediation is a concern.

We consider several behavioral remedies (imposition of a fiduciary duty, neutrality obligation, transparency regulation) as well as a structural one, namely the breaking-up of the integrated firm, an intervention we refer to as divestiture.

In order to obtain clean results about the effects of the policies on consumer surplus, we restrict attention to the two-dimensional model with quadratic costs, described in (5). Recall that we have conflict when $Ct > 1/2$ (and congruence otherwise).

4.1 Regulating bias: fiduciary duty

The most natural intervention to consider is one where the intermediary would be required to provide objective advice. Such a fiduciary duty is, for instance, commonly imposed on financial advisers.

²²Another variation would be to stick with the Hotelling setup, but assume consumers are non-uniformly located along the line. For example, if we let consumers' location follow distribution Φ then demand for firm 1 would be $\Phi(x^* + b)$ (combining non-linearity with non-separability in bias). We also check our results in this case (with Φ being truncated normal). Details are in the web appednix.

²³A solution to this equation exists. In cases where there is more than one solution, we take the smallest one.

Formally, we model this policy as a requirement to lower the level of bias b . Ideally, fiduciary duty would mean $b = 0$, but the policy may be difficult to implement because of monitoring costs or imperfect compliance by the intermediary. The effects of this policy can be obtained as a direct corollary to Proposition 3:

Corollary 2. *Under conflict, imposing a fiduciary duty increases consumer surplus. Under congruence, it increases consumer surplus if $Ct > \overline{Ct}_S$, and reduces it otherwise.*

This analysis implies that imposing a fiduciary duty on the intermediary benefits consumers when payoffs are conflicting, but may be misguided under congruence when the cost of providing utility C is small and when competition between sellers (as inversely measured by t) is strong.

4.2 Neutrality

When the level of bias is not observable, or not verifiable by a court, the previous policy may be difficult to implement. A crude way to approximate $b = 0$ without having to construct a measure of bias is to force the intermediary to send fraction $1/2$ of consumers to seller 1. For instance, the intermediary might be required to recommend seller 1 and 2 equally often or otherwise extend equal treatment to the two sellers. We refer to this kind of intervention as a neutrality policy. This kind of remedy was considered (though ultimately dismissed) when, in 2016, the UK's high court ruled on *Streetmap.EU Ltd v Google Inc.* Specifically, Streetmap argued that Google should be required to randomize the provider of the map shown to users who conduct a location-based query (instead of always showing them a Google map). On one hand, the policy was mooted to provide sellers with a level playing field and stimulate competition between them. On the other hand, the court was concerned that neutrality requirements limit the intermediary's ability to favor sellers with very high quality products. Here we investigate the merits of these claims.

Formally, we adapt the baseline model such that the intermediary is constrained to send uninformed consumers with $x \leq 1/2$ to seller 1, while the remainder are sent to seller 2.

Under neutrality, sellers' profits are

$$\pi_i = r_i \left[(1 - \mu) \left(\frac{1}{2} + \frac{1}{2t}(u_i - u_j) \right) + \frac{\mu}{2} \right] - \frac{C}{2}(u_i + r_i - v)^2. \quad (11)$$

We observe two changes from the baseline model. Firstly, profits are symmetric. Thus, sellers choose symmetric equilibrium strategies, $u_1^* = u_2^*$, resulting in $D_1 = D_2 = 1/2$. Because half of the consumers *should* go to each seller when they play symmetric strategies, the intermediary ends up sending all consumers to their preferred alternative. In other words, neutrality eliminates bias as intended. However, neutrality also affects competition

between sellers. Given that each seller has a guaranteed market share of $1/2$ over the uninformed consumers, sellers only compete to attract the informed ones. Indeed, under neutrality we have $\partial D_i / \partial u_i = (1 - \mu) / 2t$, whereas in the baseline case $\partial D_i / \partial u_i = 1 / 2t$. This negative competition-softening effect dominates the no-bias result under both congruence and conflict, leaving consumers worse-off overall.

Proposition 4. *A neutrality policy*

1. *results in ex post optimal matching of consumers and sellers,*
2. *reduces the average utility offered by sellers (weighted by market share),*
3. *reduces overall consumer surplus.*

Our analysis thus vindicates one of the main concerns about this policy raised in the Streetmap v Google case. Neutrality is also subject to some practical concerns. An industry may have many sellers and frequent entry by new sellers (some of whom are of dubious quality). Policymakers must specify which of these sellers are eligible to benefit from the neutrality policy. Thus, instead of the intermediary vetting sellers (a task in which it is specialized), neutrality imposes this burden on the policy authority.

4.3 Transparency

In the baseline model we treat the level of bias as given and investigate its effect on sellers' equilibrium choices. The underlying assumption is that uninformed consumers are unaware of the existence of bias or unable to do anything about it. This could be the case for several reasons. Firstly, it might be that consumers naïvely expect the intermediary to act in their best interest (or are simply unaware of the bias). This frames bias in the same intellectual tradition as persuasive advertising, which is taken to distort consumers' perception of their own preferences without them noticing (see, e.g., Dixit and Norman, 1978).²⁴ Secondly, it may be that the intermediary serves many different markets but is integrated with a seller in just a few. For instance, even though Google is integrated in a few markets (videos, maps, shopping, etc.), the majority of queries do not fall into these categories. Therefore the “average” level of own-content bias is probably quite small, which may lead consumers to follow the recommendation even if they know that it will be biased in a few instances.²⁵

In response to concerns that uninformed consumers might be vulnerable to deception, policymakers often consider transparency policies that force intermediaries to disclose the

²⁴In Footnote 17 we describe how bias in the Hotelling model can be formally framed in this light.

²⁵This informal reasoning can be made rigorous. Below, when consumers observe b , we show that all consumers will follow the intermediary's recommendation for b sufficiently low. Thus, if a small fraction of queries are subject to bias but consumers don't know which ones (so that the expected level of bias on each given query is positive but low) then consumers would follow the advice even if they are fully rational.

extent to which they are biased. Our assumption here is that such a disclosure would make uninformed consumers aware of the existence (and magnitude) of bias, and lead them to rationally choose whether to follow the intermediary's recommendation.²⁶ In turn, given the new awareness of consumers, the policy could lead the intermediary to be less biased towards its own affiliate.

To model the effect of a transparency policy, let us introduce some more structure on consumers' information. Informed consumers are perfectly informed and always go directly to the seller offering them the highest utility. Uninformed consumers are unable to observe the utilities offered by each firm. Additionally, while uninformed consumers may know their own tastes, they do not know how those tastes map onto the available products.²⁷ In other words, they view their own position, x , on the Hotelling line as a uniform draw from $[0, 1]$. As an illustrative example, a tourist might know their own food preferences, but not know which restaurants in a town best cater to these tastes or have the best food/prices. The intermediary, on the other hand, can observe the utilities and determine each consumer's x (e.g. via a matching algorithm), thereby being able to identify the best match for each consumer.

We use the following timing: at $\tau = 1$, the intermediary selects a level of bias b . At $\tau = 2$, sellers 1 and 2 observe b and choose their actions (u_i, r_i) . At $\tau = 3$, uninformed consumers consult the intermediary, which recommends seller 1 to those located to the left of $\frac{t+u_1-u_2}{2t} + b$, and seller 2 to the others. At $\tau = 4$, informed consumers select their preferred seller, while uninformed consumers choose whether to follow the recommendation or buy from the non-recommended seller.²⁸ We consider two cases. In the post-transparency case, uninformed consumers correctly observe the magnitude and direction of b , form correct beliefs about the equilibrium u_i s that result, and rationally choose whether to follow the intermediary's advice or not. The pre-transparency case is the same, except we assume that uninformed consumers form beliefs and behave at $\tau = 4$ as if they had observed $b = 0$.

Thus, suppose the intermediary recommends firm 1 to a consumer who perceives a level of bias $\hat{b} \in \{0, b\}$ and expects firms to offer utility u_i^e . Upon receiving the recommendation, and using Bayes' rule, the consumer believes that his position is $x \in [0, x^e + \hat{b}]$, where $x^e = \frac{t+u_1^e-u_2^e}{2t}$. The consumer finds it optimal to follow the intermediary's advice if

$$\int_0^{x^e+\hat{b}} \frac{1}{x^e+\hat{b}}(u_1^e - tx) dx \geq \int_0^{x^e+\hat{b}} \frac{1}{x^e+\hat{b}}(u_2^e - t(1-x)) dx. \quad (12)$$

²⁶An alternative interpretation of this extension is as a reduced-form long-run model in which consumers have become aware of bias via the intermediary's reputation and choose whether to follow its advice or not.

²⁷Formally, this could be modelled as follows: each consumer has a type $\theta \in \Theta$ that encapsulates his or her preferences. Each taste, θ , corresponds to a location $x = f(\theta)$ on the Hotelling line. However, uninformed consumers do not know the mapping f , they know only that the resulting x are uniform. The intermediary, on the other hand, knows $f(\cdot)$ and can therefore determine each consumer's x .

²⁸Similar results would obtain if the consumers' outside option were to forego the intermediary's advice altogether and choose a seller at random.

For any $b \geq 0$, a consumer to whom the intermediary recommends firm 2 always wants to follow the recommendation because firm 2 is never recommended to a consumer with $x < x^*$.

Firms' optimal choice of u depends on how many consumers they expect to follow the intermediary's recommendation, while consumers' choice to follow the recommendation depends on the u s they expect firms to offer. Firms' and consumers' expectations should be consistent in equilibrium, which can result in multiple equilibria. As a selection device, we focus on equilibria in which uninformed consumers follow the recommendation when such equilibria exist.

As a preliminary result, we verify that the intermediary would choose to increase bias as much as possible if it could be sure that consumers will always follow the recommendation.²⁹ Recall that we have defined \bar{b} as the maximal level of bias, i.e. such that $\frac{t+u_1(\bar{b})-u_2(\bar{b})}{2t} + \bar{b} = 1$ (and its value is thus different under congruence and conflict).

Lemma 2. *Suppose all uninformed consumers follow the intermediary's recommendation. For any $b \leq \bar{b}$, firm 1 benefits from an increase in b .*

Given Lemma 2, we now investigate how biased the intermediary can be in equilibrium. Prior to the transparency intervention, uninformed consumer behave as if $b = 0$ and thus find it optimal to follow the intermediary's advice ((12) is always satisfied for $\hat{b} = 0$). It is immediate that the intermediary can increase b to \bar{b} in this case and will choose to do so.

Now suppose that the intermediary is forced to publicly disclose b , resulting in $\hat{b} = b$. We ask whether this affects its ability to influence consumer behavior with biased results.

Proposition 5. *Prior to transparency, the intermediary is maximally biased: $b^* = \bar{b}$. Under transparency:*

1. *When payoffs are congruent, there exists an equilibrium in which the intermediary is maximally biased, $b^* = \bar{b}$.*
2. *When payoffs are conflicting, we must have positive but less than maximal bias, $0 < b^* < \bar{b}$, in equilibrium.*

The proof is instructive, so we include it here:

Proof. The pre-transparency case follows immediately from Lemma 2.

Suppose that payoffs are congruent, and that $b = \bar{b}$. If all consumers follow the recommendation, we have $u_1(b) > u_2(b)$ (Proposition 1). When $b = \bar{b}$ the recommendation gives no information regarding the horizontal match (everybody is directed towards seller

²⁹We should not take this for granted: bias causes seller 2 to become a tougher competitor under conflict (u_2 increases in b), which could conceivably leave the integrated firm worse off. Lemma 2 establishes that this does not happen. The proof of the lemma holds for any convex cost function, not only the quadratic case.

1), so following the advice or going against it both yield the same distribution of transport costs. Therefore, each consumer is better-off following the recommendation.

When payoffs are conflicting and $b = \bar{b}$, we have $u_1(\bar{b}) < u_2(\bar{b})$. Given that the recommendation is non-informative about the horizontal match, consumers would be better-off going against the recommendation and each choosing seller 2. Thus we must have $b < \bar{b}$ in equilibrium.

If $b = 0$, consumers strictly prefer following the intermediary: both sellers provide the same level of utility $u_i(0)$, but the expected transportation cost is smaller if they follow the (unbiased) recommendation. By continuity, for a small enough level of bias $b = \epsilon > 0$ consumers still prefer to follow the recommendation. ■

Discussion The congruence/conflict dichotomy plays a key role in determining the efficacy of the transparency requirement. Under congruence, the requirement does not prevent the intermediary from being maximally biased because bias leads seller 1 to offer a higher utility than seller 2. When payoffs are conflicting, the intermediary cannot be maximally biased because consumers would go against its advice. However, the intermediary can still achieve a positive level of bias in equilibrium because, for small levels of bias, the informativeness regarding the horizontal match outweighs the fact that the favored seller is of lower intrinsic quality.

Transparency policies are frequently suggested or implemented as a solution to bias. For example, in 2018 the European Commission issued new rules requiring “providers of online intermediation services [to] formulate and publish general policies on [...] how they treat their own goods or services compared to those offered by their professional users”.³⁰

Price comparison websites (PCWs) offer an interesting case in point. In this industry, the presence of vertical integration (between PCWs and sellers)³¹ and of commissions paid by sellers, often in return for favorable positioning, has led to regulatory scrutiny. For instance, OfGem, the UK energy regulator, provides PCWs with an accreditation provided they prominently list the energy companies from which they receive commission on sales, as well as make it clear that they earn commission on certain tariffs.³² In France, recent legislation³³ requires PCWs across all sectors to be more transparent regarding (i) the criteria they use to rank the offers they display, (ii) the existence of contracts (or capitalistic links) with firms, and (iii) the final price paid by the consumer and other “essential characteristics” of the offers. We argue that, in the context of PCWs, competition

³⁰See http://europa.eu/rapid/press-release_IP-18-3372_en.htm, accessed 12 October 2018.

³¹For instance, the insurance PCW Confused.com is part of the Admiral group, a motor insurance company.

³²See <https://www.ofgem.gov.uk/information-consumers/domestic-consumers/switching-your-energy-supplier/confidence-code>, accessed 18 July 2016.

³³<https://www.legifrance.gouv.fr/eli/decret/2016/4/22/EINC1517258D/jo>, accessed 18 July 2016.

is mostly in prices and the interests of sellers and of consumers are in conflict. Therefore we view such interventions as likely to lead to less bias, lower prices, and increased consumer surplus.

In other markets, our assessment of the benefits of transparency is more cautious. For instance, there have been calls for Google to be more transparent over its ranking algorithms, and in particular about the way results are biased to favor Google’s own services. We argue that several of the markets in question are characterized by payoff congruence: maps, browsers, and many other services are free and advertising-supported, and competition is mostly on the quality dimension. Our model indicates that, in cases where bias is harmful, transparency would do little to alleviate the issue.

Another interpretation of our results could be that transparency, even if it is sometimes useless, does no harm, and that hence it should be encouraged. Here we would like to draw attention to the fact that transparency has a specific meaning in our model, namely that the intermediary reveals the extent of its bias to consumers. We do not say anything about requiring intermediaries to fully disclose the inner workings of their algorithms, which might reduce incentives to invest in algorithm design and leave the algorithms vulnerable to manipulation.

4.4 Divestiture

The three interventions discussed above represent behavioral remedies against bias and, as such, would require some form of continued monitoring of the intermediary’s conduct to have any effect. An alternative approach would consist in imposing structural remedies, such as a separation of seller 1 and the intermediary. Such a policy was, for example, proposed during the European Commission’s antitrust investigations against Google.³⁴ The Indian government recently proposed a regulation preventing e-commerce intermediaries from selling products from companies they have an equity stake in,³⁵ and Elizabeth Warren made a similar proposal part of her platform for the 2020 U.S. presidential election.³⁶ To model divestiture, we assume that the intermediary is still able to sell its recommendation to one of the two sellers, but that the contract can only be a short-term one, i.e. bargaining occurs after sellers have chosen their actions (u_i, r_i) .³⁷

If seller i is recommended, the number of consumer it serves is $\frac{t+u_i-u_j}{2t} + \mu b$, versus $\frac{t+u_i-u_j}{2t} - \mu b$ if seller j is recommended. It is therefore willing to pay up to $2\mu b r_i$ for the recommendation.

³⁴See, e.g., <https://www.ft.com/content/617568ea-71a1-11e4-9048-00144feabdc0>, accessed 25 March 2019.

³⁵<https://www.cnn.com/2018/12/26/indias-tightens-e-commerce-rules-likely-to-hit-amazon-flipkart.html>

³⁶<https://medium.com/@teamwarren/heres-how-we-can-break-up-big-tech-9ad9e0da324c>

³⁷If bargaining occurred before the choice of u and r , the firms could write contracts that replicate integration.

In equilibrium, efficiency dictates that the intermediary should grant prominence to the firm with the largest r_i . The equilibrium price then depends on the specific bargaining process. In a first-price auction, if we restrict attention to equilibria in undominated strategies, firm i pays $2\mu br_j$ to the intermediary. If, on the other hand, the intermediary has all the bargaining power,³⁸ it can charge $2\mu br_i$. We allow for a range of outcomes in between these two extreme cases by assuming that the intermediary has bargaining power $\alpha \in [0, 1)$, resulting in a price $2\mu b(\alpha r_i + (1 - \alpha)r_j)$.

We start by describing the equilibrium of the game under divestiture. If the two sellers choose the same value of r , each one is made prominent with probability $1/2$. The profit of seller i is then

$$\pi_i = r_i \left[\frac{1}{2} + \frac{1}{2t}(u_i - u_j) \right] - \mu br_i - \frac{C}{2}(u_i + r_i - v)^2.$$

Even though the intermediary is always biased, the expected demand for seller i is the same as if there was no bias. The second term is the expected payment to the intermediary: with probability $1/2$, seller i pays the intermediary $2\mu b(\alpha r_i + (1 - \alpha)r_j)$, with $r_j = r_i$.

If $r_i > r_j$, the profit of seller i is

$$\pi_i = r_i \left[\frac{1}{2} + \frac{1}{2t}(u_i - u_j) + \mu b \right] - 2\mu b[\alpha r_i + (1 - \alpha)r_j] - \frac{C}{2}(u_i + r_i - v)^2, \quad (13)$$

and the profit of seller j is

$$\pi_j = r_j \left[\frac{1}{2} + \frac{1}{2t}(u_j - u_i) - \mu b \right] - \frac{C}{2}(u_j + r_j - v)^2. \quad (14)$$

Lemma 3. *Under divestiture, the equilibrium is asymmetric. One seller chooses a higher r_i than its rival, and is recommended by the intermediary with probability 1.*

Even though divestiture puts sellers 1 and 2 in an ex ante symmetric situation by breaking the contractual ties between seller 1 and the intermediary, the equilibrium is necessarily asymmetric: one seller (say i) anticipates that it will be recommended and chooses a high r_i , while seller j anticipates that it will not be recommended and chooses a low r_j . However, this asymmetric equilibrium is not equivalent to the one under integration because the intermediary can extract part of the value of seller i 's investment. Comparing divestiture and integration, we obtain the following result:

Proposition 6. *Suppose the intermediary divests its interest in seller 1. Then,*

1. *under conflict, consumer surplus increases compared to integration;*
2. *under congruence, consumer surplus falls.*

³⁸If for instance it can make sequential take it or leave it offers.




Figure 5: The effect of divestiture.

To understand the intuition, it is useful to compare the profit functions. Suppose without loss of generality that firm 1 is the one choosing a high value of r in equilibrium. Seller 2's profit function is the same under both integration and divestiture, therefore its best-response $\hat{u}_2(u_1)$ is unchanged. On the other hand seller 1's profit function is not the same under both regimes: under divestiture, it has to pay an amount $2\mu b(\alpha r_1 + (1-\alpha)r_2)$ to the intermediary. The marginal payoff from increasing r_1 is thus lower. Under congruence, this means that the payoff of choosing a higher u_1 (and the associated $\hat{r}_1(u_1, u_2)$) is also lower: seller 1's best-reply $\hat{u}_1(u_2)$ shifts down. The reverse is true under conflict. Figure 5 then illustrates that both sellers end up providing higher utility levels under conflict (which benefits consumers), while only seller 2 does so under congruence. In the latter case, because seller 2 serves fewer consumers, consumer surplus goes down compared to integration.

5 Conclusion

Because information intermediaries play an influential role in shaping consumers' choices, and because they often have a stake in the outcome of those choices, it is not surprising that intermediaries have been accused of biasing their advice in response to economic incentives. But information intermediaries exist in many different industries, where firms employ many different business models. This presents a challenge for policy makers who are asked to consider the effects of bias in a wide variety of strategic environments. To what extent can lessons from one market context (e.g., search engine bias) be transferred to another (e.g., paid promotion on price comparison websites)?

To cast light on this question, we introduce the notions of congruence and conflict, and show that they play an important role in shaping the effects of bias and the efficacy of various policy interventions. Environments exhibiting *conflict* are those where higher

revenues are obtained by extracting more surplus at consumers' expense (for example, when firms compete mostly in prices). In such an environment, intermediary bias leads the favored firm to endogenously offer lower utility. Thus, consumers are systematically mismatched in favour of a firm they like less, which tends to harm them. Although bias is harmful, we find that a range of policy responses (direct regulation of bias, divestiture, and transparency) are all at least somewhat effective at improving consumer outcomes.

We contrast this case with environments exhibiting *congruence*. Congruence arises when strategies that increase firms' per-consumer revenues also increase consumers' utility (such as when firms improve the quality of their product and share the resulting surplus with consumers). Here, an increase in bias leads the beneficiary firm to invest in improving its utility offer. Thus, while bias implies that some consumers are still systematically mismatched, the mismatching now happens in favor of an endogenously better product. This softens the effects of bias for consumers and, indeed, means that an increase in bias can leave consumers better-off overall. Although bias can be less problematic under conditions of congruence, any problems that do arise are likely to be more difficult to deal with. This is because the same factors that allow bias to be beneficial (in particular, the favored firm's enhanced investment incentive) tend to make policy interventions less effective under congruence than conflict. Divestiture, neutrality, and transparency interventions fail entirely to improve consumer outcomes, with only direct imposition of a fiduciary duty remaining effective.

References

- Albano, Gian Luigi and Alessandro Lizzeri (2001). "Strategic certification and provision of quality". *International Economic Review* 42.1, pp. 267–283.
- Armstrong, Mark and John Vickers (2001). "Competitive price discrimination". *RAND Journal of Economics*, pp. 579–605.
- Armstrong, Mark and Jidong Zhou (2011). "Paying for Prominence". *Economic Journal* 121.556, F368–F395.
- Biglaiser, Gary (1993). "Middlemen as Experts". *RAND Journal of Economics* 24.2, pp. 212–223.
- Biglaiser, Gary and James W Friedman (1994). "Middlemen as guarantors of quality". *International Journal of Industrial Organization* 12.4, pp. 509–531.
- Buehler, Benno and Florian Schuett (2014). "Certification and minimum quality standards when some consumers are uninformed". *European Economic Review* 70, pp. 493–511.
- Burguet, Roberto, Ramon Caminal, and Matthew Ellman (2015). "In Google we trust?" *International Journal of Industrial Organization* 39, pp. 44 –55. ISSN: 0167-7187. DOI: <https://doi.org/10.1016/j.ijindorg.2015.02.003>. URL: <http://www.sciencedirect.com/science/article/pii/S016771871500017X>.

- Carbajo, Jose, David De Meza, and Daniel J Seidmann (1990). “A strategic motivation for commodity bundling”. *The Journal of Industrial Economics* 38.3, pp. 283–298.
- Carlton, Dennis W and Michael Waldman (2002). “The strategic use of tying to preserve and create market power in evolving industries”. *RAND Journal of Economics* 33.2, pp. 194–220.
- Chen, Yongmin (1997). “Equilibrium product bundling”. *Journal of Business* 70.1, pp. 85–103.
- Choi, Jay Pil (2004). “Tying and Innovation: A Dynamic Analysis of Tying Arrangements”. *Economic Journal* 114, pp. 83–101.
- Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim (2015). “Net Neutrality, Network Capacity, and Innovation at the Edges”. Working Paper.
- Choi, Jay Pil and Byung-Cheol Kim (2010). “Net Neutrality and Investment Incentives”. *RAND Journal of Economics* 41.3, pp. 446–471.
- Cookson, Gordon, Tim Jenkinson, Howard Jones, and Jose Vicente Martinez (2017). “Best Buys and Own Brands: Investment Platforms Recommendations of Mutual Funds”. Working Paper.
- De Cornière, Alexandre and Greg Taylor (2014). “Integration and Search Engine Bias”. *RAND Journal of Economics* 45.3, pp. 576–597.
- Dellarocas, Chrysanthos, Zsolt Katona, and William Rand (2013). “Media, Aggregators, and the Link Economy: Strategic Hyperlink Formation in Content Networks”. *Management Science* 59.10, pp. 2360–2379.
- Dixit, Avinash and Victor Norman (1978). “Advertising and Welfare”. *The Bell Journal of Economics* 9.1, pp. 1–17.
- Durbin, Erik and Ganesh Iyer (2009). “Corruptible advice”. *American Economic Journal: Microeconomics* 1.2, pp. 220–242.
- Economides, Nicholas and Benjamin E Hermalin (2012). “The economics of network neutrality”. *The RAND Journal of Economics* 43.4, pp. 602–629.
- Edelman, Benjamin and Zhenyu Lai (2013). “Exclusive Preferential Placement as Search Diversion: Evidence from Flight Search”. Working Paper.
- Egan, Mark (2017). “Brokers vs. Retail Investors: Conflicting Interests and Dominated Products”. Working Paper.
- Engelberg, Joseph, Christopher A. Parsons, and Nathan Tefft (2013). “First, Do No Harm: Financial Conflicts in Medicine”. Working Paper.
- Inderst, Roman and Marco Ottaviani (2012a). “Competition Through Commissions and Kickbacks”. *American Economic Review* 102.2, pp. 780–809.
- (2012b). “How (not) to pay for advice: A framework for consumer financial protection”. *Journal of Financial Economics* 105.2, pp. 393–411.

- Jeon, Doh-Shin and Nikrooz Nasr (2016). “News Aggregators and Competition among Newspapers in the Internet”. *American Economic Journal: Microeconomics* 8.4, pp. 91–114.
- Krämer, Jan and Daniel Schnurr (2018). “Is there a need for platform neutrality regulation in the EU?” *Telecommunications Policy*.
- Lizzeri, Alessandro (1999). “Information Revelation and Certification Intermediaries”. *RAND Journal of Economics* 30.2, pp. 214–231.
- Nalebuff, Barry (2004). “Bundling as an Entry Barrier”. *The Quarterly Journal of Economics* 119.1, pp. 159–187.
- Peitz, Martin and Florian Schütt (2016). “Net Neutrality and Inflation of Traffic”. *International Journal of Industrial Organization* 46.1, pp. 16–62.
- Peyrache, Eloic and Lucia Quesada (2011). “Intermediaries, credibility and incentives to collude”. *Journal of Economics & Management Strategy* 20.4, pp. 1099–1133.
- Rutt, James (2011). “Aggregators and the News Industry: Charging for Access to Content”. Working Paper.
- Strausz, Roland (2005). “Honest certification and the threat of capture”. *International Journal of Industrial Organization* 23.1, pp. 45–62.
- Sunstein, Cass and Richard Thaler (2008). *Nudge: Improving Decisions about Health, Wealth and Happiness*. Yale University Press.
- Teh, Tat-How and Julian Wright (2018). “Steering by information intermediaries”. Working Paper.
- Whinston, Michael D. (1990). “Tying, Foreclosure, and Exclusion”. *American Economic Review* 80, pp. 837–860.

A Omitted Proofs

A.1 Proofs from Section 3

Proof of Lemma 1. The proof proceeds slightly differently in the one- and two-dimensional cases. Begin with the case where sellers choose both u and r . Recall that we define $\Pi_i(u_i, u_j, b) \equiv \pi_i(\hat{r}_i(u_i, u_j, b), u_i, u_j, b)$. In order to prove the results, we totally differentiate the first-order condition

$$\frac{\partial \Pi_i(u_i, u_j, b)}{\partial u_i} = 0.$$

We thus obtain

$$\frac{du_i}{db} \Big|_{du_j=0} = - \frac{\frac{\partial^2 \Pi_i}{\partial u_i \partial b}}{\frac{\partial^2 \Pi_i}{\partial u_i^2}}. \quad (15)$$

By the second-order condition the denominator is negative. By the envelope theorem, we have $\frac{\partial \Pi_i}{\partial u_i} = \frac{\partial \pi_i}{\partial u_i}$, so that

$$\frac{\partial^2 \Pi_i}{\partial u_i \partial b} = \frac{\partial \hat{r}_i}{\partial b} \frac{\partial^2 \pi_i}{\partial u_i \partial r_i}$$

(we are also using the fact that $\frac{\partial^2 \pi_i}{\partial u_i \partial b} = \frac{\partial^2 D_i}{\partial u_i \partial b} = 0$ in the Hotelling model with additive bias, see the discussion in section 3.4).

The term $\frac{\partial^2 \pi_i}{\partial u_i \partial r_i}$ is positive under congruence, and negative under conflict. The direction of the shift in firm i 's best-response thus depends on the sign of $\frac{\partial \hat{r}_i}{\partial b}$.

Using the implicit function theorem over (4), we get

$$\frac{\partial \hat{r}_1}{\partial b} = \frac{\frac{\partial D_1}{\partial b}}{\frac{\partial^2 C}{\partial r_1^2}} > 0 \quad \text{and} \quad \frac{\partial \hat{r}_2}{\partial b} = \frac{\frac{\partial D_2}{\partial b}}{\frac{\partial^2 C}{\partial r_2^2}} < 0. \quad (16)$$

Parts (i) and (ii) directly follow. For part (iii), we use a similar reasoning for $\frac{du_i}{du_j} \Big|_{db=0}$, using the fact that

$$\frac{\partial \hat{r}_i}{\partial u_j} = \frac{\frac{\partial D_i}{\partial u_j}}{\frac{\partial^2 C}{\partial r_i^2}} < 0. \quad (17)$$

In the one-dimensional case where $r_i = r(u_i)$, we have

$$\frac{\partial \Pi_i}{\partial u_i} = r'(u_i) \frac{\partial \pi_i}{\partial r_i} + \frac{\partial \pi_i}{\partial u_i} = r'(u_i) D_i + \frac{\partial \pi_i}{\partial u_i}.$$

Using the property that $\frac{\partial^2 \pi_i}{\partial u_i \partial b} = 0$, we find that $\frac{\partial^2 \Pi_i}{\partial u_i \partial b}$ is of the same sign as $\frac{\partial D_i}{\partial b}$ under congruence, and of the opposite sign under conflict, thereby proving parts (i) and (ii). A similar reasoning gives the slope of the best-response function (part (iii)). ■

Proof of Proposition 1 (two-dimension case). We start by proving the result for the case where sellers choose both u and r . The proof for the one-dimensional case where

sellers choose (only) u follows the same steps and is included below. Let $u_1^*(b)$ and $u_2^*(b)$ be the equilibrium utility levels when bias is b . The first-order condition for seller i is

$$\frac{\partial \Pi_i(u_1^*(b), u_2^*(b), b)}{\partial u_i} = 0.$$

By the implicit function theorem, we therefore have, for $i = 1, 2$:

$$\frac{\partial^2 \Pi_i(u_1^*(b), u_2^*(b), b)}{\partial u_i^2} u_i^{*'}(b) + \frac{\partial^2 \Pi_i(u_1^*(b), u_2^*(b), b)}{\partial u_i \partial u_j} u_j^{*'}(b) + \frac{\partial^2 \Pi_i(u_1^*(b), u_2^*(b), b)}{\partial u_i \partial b} = 0.$$

Solving this two-equation system gives

$$u_i^{*'}(b) = \frac{\frac{\partial^2 \Pi_i}{\partial u_i \partial b} \frac{\partial^2 \Pi_j}{\partial u_j^2} - \frac{\partial^2 \Pi_i}{\partial u_i \partial u_j} \frac{\partial^2 \Pi_j}{\partial u_j \partial b}}{\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j} \frac{\partial^2 \Pi_j}{\partial u_j \partial u_i} - \frac{\partial^2 \Pi_i}{\partial u_i^2} \frac{\partial^2 \Pi_j}{\partial u_j^2}}. \quad (18)$$

Our task is to sign this expression.

Preliminaries: Using similar steps as in the proof of Lemma 1, we have

$$\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j} = \frac{\partial \hat{r}_i}{\partial u_j} \frac{\partial^2 \pi_i}{\partial u_i \partial r_i} = \frac{\partial \hat{r}_i}{\partial u_j} \left(\frac{\partial D_i}{\partial u_i} - \frac{\partial^2 C}{\partial u_i \partial r_i} \right). \quad (19)$$

We have already shown that $\frac{\partial \hat{r}_i}{\partial u_j} < 0$ (proof of Lemma 1) and we know that $\frac{\partial^2 \pi_i}{\partial u_i \partial r_i}$ is, by definition, positive under congruence and negative under conflict.

Similarly, we have

$$\frac{\partial^2 \Pi_i}{\partial u_i \partial b} = \frac{\partial \hat{r}_i}{\partial b} \frac{\partial^2 \pi_i}{\partial u_i \partial r_i} = \frac{\partial \hat{r}_i}{\partial b} \left(\frac{\partial D_i}{\partial u_i} - \frac{\partial^2 C}{\partial u_i \partial r_i} \right), \quad (20)$$

where $\frac{\partial \hat{r}_i}{\partial b}$ is positive for $i = 1$ and negative for $i = 2$.

Using the implicit-function theorem on seller i 's first-order condition, $\frac{\partial \Pi_i}{\partial u_i} = 0$, we have

$$\hat{u}_i'(u_j) = - \frac{\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j}}{\frac{\partial^2 \Pi_i}{\partial u_i^2}}. \quad (21)$$

Denominator: Equation 21 implies that the denominator of (18) can be rewritten as $\frac{\partial^2 \Pi_i}{\partial u_i^2} \frac{\partial^2 \Pi_j}{\partial u_j^2} (\hat{u}_i'(u_j^*) \hat{u}_j'(u_i^*) - 1)$. By the second-order condition, $\frac{\partial^2 \Pi_i}{\partial u_i^2} < 0$ for both sellers. By the stability condition (Assumption 1), the term in brackets is negative. Therefore, the denominator in (18) is negative.

Numerator: Using (19)–(21), the numerator in (18) for $i = 1$ is equal to

$$\text{NUM}_1 = \frac{\partial^2 \pi_1}{\partial u_1 \partial r_1} \left\{ \frac{\partial^2 \pi_2}{\partial u_2 \partial r_2} \right\} \left[\frac{\partial \hat{r}_1}{\partial b} \left(- \frac{\frac{\partial \hat{r}_2}{\partial u_1}}{\hat{u}_2'(u_1)} \right) - \frac{\partial \hat{r}_1}{\partial u_2} \frac{\partial \hat{r}_2}{\partial b} \right]. \quad (22)$$

Using expressions (17) and (16), and the fact that $\frac{\partial D_i}{\partial b} = \mu$ for $i = 1$ and $-\mu$ for $i = 2$, the term in square brackets is of the same sign as

$$\left(-\frac{\frac{\partial D_2}{\partial u_1}}{\hat{u}'_2(u_1)} + \frac{\partial D_1}{\partial u_2} \right) = \left(\frac{1}{2\hat{u}'_2(u_1)t} - \frac{1}{2t} \right), \quad (23)$$

which is positive when seller 2's payoffs are in conflict and negative when they are congruent. The term in curly brackets in (22) is positive when 2's payoffs are congruent and negative when they are conflicting. Because the terms in curly and square brackets have opposite signs, and because the denominator of (18) is negative, (18) has the same sign as the first term of (22). This is positive if seller 1's payoffs are congruent and negative otherwise. The proof for seller 2 is identical except that the analog of (23) has the opposite sign. ■

Proof of Proposition 1 (one-dimension case). We now prove Proposition 1 for the uni-dimensional case where sellers choose only u . The proof follows similar steps to the two-dimension case. Using $\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j} = \frac{\partial r_i}{\partial u_i} \frac{\partial D_i}{\partial u_j}$, $\frac{\partial^2 \Pi_i}{\partial u_i \partial b} = \frac{\partial r_i}{\partial u_i} \frac{\partial D_i}{\partial b}$, and $\hat{u}'_i(u_j) = -\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j} / \frac{\partial^2 \Pi_i}{\partial u_i^2}$, we can write the numerator of (18) for seller 1 as

$$NUM_1 = -\frac{\partial r_1}{\partial u_1} \frac{\partial r_2}{\partial u_2} \left(\frac{\partial D_1}{\partial b} \frac{\partial D_2}{\partial u_2} \frac{1}{\hat{u}'_2(u_1)} + \frac{\partial D_1}{\partial u_2} \frac{\partial D_2}{\partial b} \right).$$

Substitute $\frac{\partial D_1}{\partial b} \frac{\partial D_2}{\partial u_2} = -\frac{\partial D_2}{\partial b} \frac{\partial D_1}{\partial u_2} = -\mu/2t$ to yield

$$NUM_1 = \frac{\partial r_1}{\partial u_1} \frac{\partial r_2}{\partial u_2} \left(\frac{1}{\hat{u}'_2(u_1)} - 1 \right) \frac{\mu}{2t}.$$

When 2's payoffs are congruent we have $\frac{\partial r_2}{\partial u_2} > 0 > \hat{u}'_2(u_1) > -1$. Given that the denominator of (18) is negative, (18) must therefore have the same sign as $\frac{\partial r_1}{\partial u_1}$: positive when seller 1's payoffs are congruent and negative when they are conflicting.

When 2's payoffs are conflicting, $\frac{\partial r_2}{\partial u_2} < 0$ and $\left(\frac{1}{\hat{u}'_2(u_1)} - 1 \right) > 0$, so that, again, the sign of $u_1^{*'}(b)$ is positive if 1's payoffs are congruent and negative otherwise.

We can apply a symmetric reasoning for seller 2. ■

Proof of Proposition 2. Part 1: two-dimensional case. Sellers' profits can be written as

$$\pi_1 = r_1 \left[\frac{1}{2} + \frac{1}{2t}(u_1 - u_2) + \mu b \right] - \frac{C}{2}(u_1 + r_1 - v)^2, \quad (24)$$

$$\pi_2 = r_2 \left[\frac{1}{2} + \frac{1}{2t}(u_2 - u_1) - \mu b \right] - \frac{C}{2}(u_2 + r_2 - v)^2. \quad (25)$$

Computing the first-order condition $\partial \pi_i / \partial r_i = 0$ yields (6). The condition for congruence is $\frac{\partial \hat{r}_i(u_i, u_j)}{\partial u_i} > 0$, i.e. $Ct < 1/2$.

Substituting these \hat{r}_i s into the profit functions and differentiating π_i with respect to u_i yields a system of best responses that is solved to give the equilibrium utility offers: (8).

The consumer indifferent between 1 and 2 is interior if

$$\frac{1}{2} + \frac{1}{2t}(u_1 - u_2) + b \leq 1 \iff b < \frac{1 - 3Ct}{4C\mu t - 6Ct - 2\mu + 2}. \quad (26)$$

Part 2: one-dimensional case. when the sellers' choice is one-dimensional, profits are

$$\begin{aligned} \pi_1 &= (v - \psi u_1) \left[\frac{1}{2} + \frac{1}{2t}(u_1 - u_2) + \mu b \right] - \frac{C}{2} u_1^2 \\ \pi_2 &= (v - \psi u_2) \left[\frac{1}{2} + \frac{1}{2t}(u_2 - u_1) - \mu b \right] - \frac{C}{2} u_2^2. \end{aligned}$$

Computing $\partial \pi_i / \partial u_i = 0$ and simultaneously solving this system of first-order conditions yields (9). ■

Proof of Proposition 3. Consumer surplus: Consumer surplus is given by (10). In the two-dimensional case (evaluated at (8)), this gives

$$CS|_{u_1^*, u_2^*} = v + \frac{b^2 \mu^2 t (1 - 2Ct)^2}{(1 - 3Ct)^2} - b^2 \mu t + \frac{2 - 5Ct}{4C}. \quad (27)$$

Differentiating,

$$\frac{\partial CS|_{u_1^*, u_2^*}}{\partial b} = 2b\mu t \left(\frac{\mu(1 - 2Ct)^2}{(1 - 3Ct)^2} - 1 \right), \quad (28)$$

which is positive if and only if

$$Ct < \frac{3 - 2\mu + \sqrt{\mu}}{9 - 4\mu} \equiv \overline{Ct}. \quad (29)$$

Because $\overline{Ct} \leq 1/2$, this condition can only be satisfied under congruence.

In the one-dimensional case, substituting u_1^* and u_2^* from (9) into (10) yields

$$CS|_{u_1^*, u_2^*} = \frac{1}{36} \left(\frac{36(2Ct^2 + v)}{2Ct + \psi} + \frac{64b^2 C^2 t^3 \mu^2}{(2Ct + 3\psi)^2} - \frac{64b^2 C t^2 \mu^2}{2Ct + 3\psi} - t(45 + 4b^2 \mu(9 - 4\mu)) \right).$$

The derivative with respect to b is

$$\frac{\partial CS|_{u_1^*, u_2^*}}{\partial b} = - \frac{2bt\mu(4C^2 t^2 + 12Ct\psi + (9 - 4\mu)\psi^2)}{(2Ct + 3\psi)^2}, \quad (30)$$

which is negative when $\psi > 0$. It is positive if $\psi < 0$ and

$$Ct < \overline{Ct} \equiv \frac{-3\psi - 2\sqrt{\mu}\psi}{2t}. \quad \blacksquare$$

Welfare: Start with the two-dimensional case. Consumer surplus is given by (27).

Sellers' profits are found by substitution of $(u_1^*, u_2^*, r_1^*, r_2^*)$ into (24) and (25). Total welfare is then evaluated as

$$W = \text{CS} + \pi_1 + \pi_2 = \frac{1}{4C} - \frac{t}{4} + v - b^2 t \mu + \frac{b^2 t (1 + Ct(-5 + 8Ct)) \mu^2}{(1 - 3Ct)^2}.$$

The derivative with respect to b is positive if

$$\frac{\partial W|_{u_1^*, u_2^*}}{\partial b} = 2bt\mu \left(\frac{(1 + Ct(8Ct - 5))\mu}{(1 - 3Ct)^2} - 1 \right) > 0,$$

i.e., if

$$Ct < \widetilde{Ct} \equiv \frac{6 - 5\mu + \sqrt{\mu(8 - 7\mu)}}{18 - 16\mu}.$$

It is easily checked that $\partial W/\partial b$ is greater than $\partial \text{CS}/\partial b$ (given in (28)).

For the one-dimensional case, we proceed in a similar fashion, taking the expressions for CS , u_1^* , and u_2^* from (9). We have

$$\frac{\partial W|_{u_1^*, u_2^*}}{\partial b} = \frac{2bt\mu [4Ct\psi(\mu\psi - 3) + \psi^2(4\mu(1 + \psi) - 9) - 4C^2t^2]}{(2Ct + 3\psi)^2},$$

which is positive if

$$Ct < \widetilde{Ct} \equiv \frac{1}{2} \left(\mu\psi^2 + \sqrt{4\mu\psi^2 - 2\mu\psi^3 + \mu^2\psi^4} - 3\psi \right).$$

Again, $\partial W/\partial b > \partial \text{CS}/\partial b$, where the latter derivative is given in (30).

Industry profit: We have $\pi_i(u_i, u_j, r_i, b) = r_i D_i(u_i, u_j, b) - C(u_i, r_i)$. By the envelope theorem we have

$$\frac{d\pi_i}{db} = \frac{\partial \pi_i}{\partial u_j} \frac{\partial u_j}{\partial b} + r_i \frac{\partial D_i}{\partial b}$$

Using the fact that $D_1 + D_2 = 1$, and that $\frac{\partial D_i}{\partial u_j} = \frac{\partial D_j}{\partial u_i}$, the effect of b on industry profit is

$$\frac{d(\pi_1 + \pi_2)}{db} = (r_1 - r_2) \frac{\partial D_1}{\partial b} + \left(r_1 \frac{\partial u_2}{\partial b} + r_2 \frac{\partial u_1}{\partial b} \right) \frac{\partial D_1}{\partial u_2}$$

In the 2 dimension model we can further simplify this to

$$\frac{d(\pi_1 + \pi_2)}{db} = (r_1 - r_2) \left(\mu - \frac{(1 - 2Ct)2Ct\mu}{4Ct(1 - 3Ct)} \right) = (r_1 - r_2) \frac{\mu}{2(3Ct - 1)} (4Ct - 1)$$

We always have $r_1 \geq r_2$ in equilibrium, and $Ct > 1/3$ by the stability condition. Thus industry profit increases in b .

In the one-dimensional model, the condition is $\frac{d(\pi_1 + \pi_2)}{db} \geq 0 \iff Ct + \psi \geq 0$, which is true (because $Ct \geq \max\{-\frac{3\psi}{2}, 0\}$ by the stability condition).

Proof of Corollary 1. Profits are

$$\pi_1(r_1, u_1, u_2, b) = r_1 \left[(1 - \mu) \left(\frac{1}{2} + \frac{u_1 - u_2}{2t} \right) + \mu \right] - C(u_1, r_1) \quad (31)$$

$$= (1 - \mu) \left[r_1 \left(\frac{1}{2} + \frac{u_1 - u_2}{2t} + \tilde{b} \right) - \tilde{C}(u_1, r_1) \right], \quad (32)$$

$$\pi_2(r_1, u_1, u_2, b) = r_2(1 - \mu) \left[\frac{1}{2} + \frac{u_2 - u_1}{2t} \right] - C(u_2, r_2) \quad (33)$$

$$= (1 - \mu) \left[r_2 \left(\frac{1}{2} + \frac{u_2 - u_1}{2t} \right) - \tilde{C}(u_2, r_2) \right]. \quad (34)$$

where $\tilde{b} \equiv \frac{\mu}{1-\mu}$ and $\tilde{C}(u, r) \equiv \frac{C(u, r)}{1-\mu}$. Notice that the profit of seller 1 in (32) is of the same form as in the interior equilibrium. Therefore the best-response function shares the same properties: it is decreasing in u_2 under congruence and increasing under conflict. Moreover, an increase in \tilde{b} shifts seller 1's best-response upwards under congruence, and downwards under conflict. A similar reasoning for seller 2 reveals that the best-response function is decreasing under congruence and increasing under conflict. If \tilde{b} was equal to zero, the two best-response functions would intersect on the 45° line (with u_1 on the horizontal axis and u_2 on the vertical axis). With $\tilde{b} > 0$, the equilibrium is therefore below the 45° line under congruence, and above it under conflict. ■

A.2 Proofs from Section 4

Proof of Proposition 4 (Neutrality). Computing, from (11), the first-order condition $\partial \pi_i / \partial r_i = 0$ yields $\hat{r}_i(u_i, u_j)$:

$$\hat{r}_i(u_i, u_j) = v - u_i + \frac{t + (1 - \mu)(u_i - u_j^*)}{2Ct}.$$

Substituting $\hat{r}_i(u_i, u_j)$ into $\pi_i(r_i, u_i, u_j)$ to get $\Pi_i(u_i, u_j)$ and finding the symmetric solution to $\partial \Pi_i / \partial u_i = 0$ yields equilibrium u s:

$$u_i^* = \frac{1}{2C} + v - \frac{t}{1 - \mu}. \quad (35)$$

Consumer surplus is

$$\begin{aligned} \text{CS} = (1 - \mu) \left\{ \int_0^{D_1(u_1, u_2, 0)} [u_1 - tx] dx + \int_{D_1(u_1, u_2, 0)}^1 [u_2 - t(1 - x)] dx \right\} + \\ \mu \left\{ \int_0^{1/2} [u_1 - tx] dx + \int_{1/2}^1 [u_2 - t(1 - x)] dx \right\}, \quad (36) \end{aligned}$$

which evaluates to

$$\text{CS}|_{u_1^*, u_2^*} = \frac{1}{2C} - \left(\frac{1}{1-\mu} + \frac{1}{4} \right) t + v.$$

Comparing this to the surplus from the baseline model (given in (27)) yields part 3.

Part 1 of the result follows immediately because sellers' utility offers are symmetric and sending consumers with $x < 1/2$ to seller 1 therefore yields the optimal match.

To see part 2, we compute the weighted average utility offer in the baseline model as

$$u_1^* D_1(u_1^*, u_2^*, b) + u_2^* D_2(u_1^*, u_2^*, b) = \frac{1}{2C} - t + v - \frac{2b^2 t (1 - 2Ct) \mu}{1 - 3Ct} + \frac{2b^2 t (1 - 2Ct)^2 \mu^2}{(1 - 3Ct)^2}.$$

Subtracting (35) from this yields

$$t\mu \left(\frac{1}{1-\mu} - \frac{2b^2(1-2Ct)(1-\mu-Ct(3-2\mu))}{(1-3Ct)^2} \right).$$

This difference is positive whenever (26) is satisfied. ■

Proof of Lemma 2. Under congruence we have

$$\frac{d\Pi_1}{db} = \frac{\partial \Pi_1}{\partial r_1} \frac{dr_1}{db} + \frac{\partial \Pi_1}{\partial u_1} \frac{du_1}{db} + \frac{\partial \Pi_1}{\partial u_2} \frac{du_2}{db} + \frac{\partial \Pi_1}{\partial b}.$$

The first two terms are zero by the envelope theorem. The last two terms can be rewritten as

$$r_1 \left(\mu - \frac{du_2}{db} \frac{1}{2t} \right) > 0.$$

For conflict, we have

$$\frac{d\Pi_1}{db} = \frac{\partial r_1}{\partial b} \left(D_1 - \frac{\partial C(u_1, r_1)}{\partial r_1} \right) + r \frac{dD_1}{db} - \frac{du_1}{db} \frac{\partial C(u_1, r_1)}{\partial u_1}.$$

The first term is zero by the envelope theorem and we know that $du_1/db < 0$ under conflict (Proposition 1). It remains to show that $dD_1/db > 0$. From (18), and using the observation (from the proof of Proposition 1) that $\hat{u}'_i(u_j) = -\frac{\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j}}{\frac{\partial^2 \Pi_i}{\partial u_i^2}}$, we have

$$u_1^{*'}(b) - u_2^{*'}(b) = \frac{[1 - \hat{u}'_2(u_1^*)] \frac{\frac{\partial^2 \Pi_1}{\partial u_1 \partial b}}{\frac{\partial^2 \Pi_1}{\partial u_1^2}} - [1 - \hat{u}'_1(u_2^*)] \frac{\frac{\partial^2 \Pi_2}{\partial u_2 \partial b}}{\frac{\partial^2 \Pi_2}{\partial u_2^2}}}{\hat{u}'_1(u_2^*) \hat{u}'_2(u_1^*) - 1}. \quad (37)$$

From (17), (16), (19), and (20):

$$\frac{\frac{\partial^2 \Pi_1}{\partial u_1 \partial b}}{\frac{\partial^2 \Pi_1}{\partial u_1^2}} = -2t\mu \frac{\frac{\partial^2 \Pi_1}{\partial u_1 \partial u_2}}{\frac{\partial^2 \Pi_1}{\partial u_1^2}} = -2t\mu \hat{u}'_1(u_2), \quad \frac{\frac{\partial^2 \Pi_2}{\partial u_2 \partial b}}{\frac{\partial^2 \Pi_2}{\partial u_2^2}} = 2t\mu \frac{\frac{\partial^2 \Pi_2}{\partial u_1 \partial u_2}}{\frac{\partial^2 \Pi_2}{\partial u_2^2}} = 2t\mu \hat{u}'_2(u_1)$$

Substituting this into (37) yields

$$u_1^*(b) - u_2^*(b) = -2t\mu \frac{[1 - \hat{u}_2'(u_1^*)] \hat{u}_1'(u_2^*) + [1 - \hat{u}_1'(u_2^*)] \hat{u}_2'(u_1^*)}{\hat{u}_1'(u_2^*) \hat{u}_2'(u_1^*) - 1}.$$

Starting from (2), we have

$$\frac{dD_1}{db} = \mu + \frac{1}{2t} [u_1^*(b) - u_2^*(b)] = \mu \frac{1 + \hat{u}_1'(u_2^*) + \hat{u}_2'(u_1^*) - 3\hat{u}_1'(u_2^*) \hat{u}_2'(u_1^*)}{1 - \hat{u}_1'(u_2^*) \hat{u}_2'(u_1^*)}.$$

Because there is conflict, and using Assumption 1, $0 < \hat{u}_i'(u_j^*) < 1$. Thus, $\frac{dD_1}{db} > 0$, completing the proof that the integrated seller's profit is increasing in b . ■

Proof of Lemma 3. Given the payoff structure, it is slightly more convenient to work with r as the main choice variable. That is, for given values of r_i, r_j and u_j , we first find the optimal u_i , denoted \hat{u}_i . It turns out that \hat{u}_i only depends on r_i , and does not depend on whether a seller expects to win or lose the contest: $\hat{u}(r_i) = v + r_i(\frac{1}{2Ct} - 1)$. Notice that the condition for payoffs to be congruent or conflicting is the same as under integration.

Suppose that sellers play a symmetric strategy profile $(r, \hat{u}(r))$. The right-hand derivative of the profit is $\frac{t-r}{2t} + \mu b(1 - 2\alpha)$. The left-hand derivative of the profit is $\frac{t-r}{2t} - \mu b$, which is always smaller than the right-hand derivative. Therefore we cannot have a symmetric equilibrium.

In an asymmetric equilibrium, one seller (say 1) maximizes (13) and the other maximizes (14). The solution is

$$r_1 = t + \frac{2bt(\alpha + Ct(1 - 4\alpha))}{3Ct - 1} \quad \text{and} \quad r_2 = t + \frac{2bt(\alpha - Ct(1 + 2\alpha))}{3Ct - 1} \quad (38)$$

■

Proof of Proposition 6 (Divestiture). Consumer surplus is given by (10). Using (38), this is equal to

$$CS|_{u_1^*, u_2^*} = \frac{1}{2C} - \frac{5t}{4} + v - \frac{b(Ct(b - 2\alpha) + \alpha)\mu}{C} + \frac{b^2t(1 - 2Ct)^2(1 - \alpha)^2\mu^2}{(1 - 3Ct)^2}.$$

This is less than (27) under congruence ($Ct < 1/2$). It is greater than (27) under conflict whenever $b \leq \frac{1-3Ct}{2(1-3Ct-(1-2Ct)(1-\alpha)\mu)}$ (which is the condition for demand to be interior). ■

B Alternative model specifications

B.1 Model with multiplicative bias

Here we consider an alternative specification for the effect of bias on demand. In our baseline specification, bias steers b consumers who are close to indifferent towards seller 1. Here, we instead take bias as directing a random and uniformly chosen sample of b consumers towards the integrated seller. In particular, we take demands to be $D_1 = \frac{1}{2t}(u_1 - u_2)(1 - b) + b$, while $D_2 = \frac{1}{2t}(u_2 - u_1)(1 - b)$. To be concise, we take $\mu = 1$ and focus on the multi-dimensional case where sellers choose both u and r .

We have the following result:

Proposition 7. *When $D_1 = x^*(1 - b) + b$ and $D_2 = (1 - x^*)(1 - b)$:*

1. *If both sellers' payoffs are congruent then $u_1^* - u_2^*$ is positive and increases with b .*
2. *If both sellers' payoffs are conflicting then $u_1^* - u_2^*$ is negative and decreases with b .*

Proof of Proposition 7. *Preliminaries:* The first-order condition for firm i , $\frac{\partial \Pi_i(u_i, u_j, b)}{\partial u_i} = 0$, can be written as follows:

$$\hat{r}_i(u_i, u_j, b) \frac{\partial D_i(u_i, u_j, b)}{\partial u_i} + \frac{\partial \hat{r}_i(u_i, u_j, b)}{\partial u_i} D_i(u_i, u_j, b) - \frac{\partial \hat{r}_i(u_i, u_j, b)}{\partial u_i} \frac{\partial C(u_i, \hat{r}_i(u_i, u_j, b))}{\partial r_i} - \frac{\partial C(u_i, \hat{r}_i(u_i, u_j, b))}{\partial u_i} = 0. \quad (39)$$

Plugging (4) into (39), and dropping the arguments, we find that the best-reply $\hat{u}_i(u_j, b)$ is implicitly given by

$$\hat{r}_i \frac{\partial D_i}{\partial u_i} - \frac{\partial C}{\partial u_i} = 0. \quad (40)$$

Totally differentiating (40) yields

$$\frac{\partial^2 \Pi_i}{\partial u_i^2} du_i + \left[\frac{\partial D_i}{\partial u_i} - \frac{\partial^2 C}{\partial u_i \partial r_i} \right] \left(\frac{\partial \hat{r}_i}{\partial u_j} du_j + \frac{\partial \hat{r}_i}{\partial b} db \right) + \hat{r}_i \frac{\partial^2 D_i}{\partial u_j \partial b} db = 0,$$

which is the same as in the additive bias case except for the new final term. This implies that, as in the additive bias case, we have strategic complements under conflict and strategic substitutes under congruence. Expressions (17), (16), (19), and (20) are unchanged, as is the condition that $\hat{u}'_i(u_j) = -\frac{\frac{\partial^2 \Pi_i}{\partial u_i \partial u_j}}{\frac{\partial^2 \Pi_i}{\partial u_i^2}}$ from Proposition 1.

Main proof: Following the analysis in the proofs of Proposition 1 and Lemma 2 carries us again to (37). Using (17) and (16) we can write

$$\frac{\partial \hat{r}_1}{\partial b} = \frac{\partial \hat{r}_1}{\partial u_2} \frac{\partial D_1}{\partial u_2} = -\frac{\partial \hat{r}_1}{\partial u_2} \frac{2t}{1 - b} (1 - x^*).$$

By (19) and (20), we can therefore replace the $\frac{\partial^2 \Pi_1}{\partial u_1 \partial b} / \frac{\partial^2 \Pi_1}{\partial u_1^2}$ term in (37) with

$$\frac{\frac{\partial^2 \Pi_1}{\partial u_1 \partial b}}{\frac{\partial^2 \Pi_1}{\partial u_1^2}} = -\frac{\frac{\partial^2 \Pi_1}{\partial u_1 \partial u_2}}{\frac{\partial^2 \Pi_1}{\partial u_1^2}} \frac{2t}{1-b} (1-x^*) = -\frac{2t}{1-b} (1-x^*) \hat{u}'_1(u_2). \quad (41)$$

By the same token, we can rewrite $\frac{\partial^2 \Pi_2}{\partial u_2 \partial b} / \frac{\partial^2 \Pi_2}{\partial u_2^2}$ as

$$\frac{\frac{\partial^2 \Pi_2}{\partial u_2 \partial b}}{\frac{\partial^2 \Pi_2}{\partial u_2^2}} = -\frac{\frac{\partial^2 \Pi_1}{\partial u_1 \partial u_2}}{\frac{\partial^2 \Pi_1}{\partial u_1^2}} \frac{\partial D_2}{\partial b} = \frac{2t}{1-b} x^* \hat{u}'_2(u_1). \quad (42)$$

Substituting (41) and (42) into (37) yields

$$u_1'(b) - u_2'(b) = -\frac{2t}{1-b} \frac{[1 - \hat{u}'_2(u_1^*)] (1-x^*) \hat{u}'_1(u_2) + [1 - \hat{u}'_1(u_2^*)] x^* \hat{u}'_2(u_1)}{1 - \hat{u}'_1(u_2^*) \hat{u}'_2(u_1^*)}.$$

Given Assumption 1, this is positive if $\hat{u}'_1(u_2) < 0$ and $\hat{u}'_2(u_1) < 0$, which we know to be true under congruence. It is negative if $\hat{u}'_1(u_2) > 0$ and $\hat{u}'_2(u_1) > 0$, which holds under conflict. ■

B.2 Other Discrete Choice Models

The results above have been obtained in a Hotelling model where the demand for seller 1 is $D_1 = \frac{1}{2t}(u_1 - u_2) + b$.³⁹ But our main observation about the importance of congruence and conflict is robust to other common discrete choice demand specifications.

Suppose that a consumer obtains utility $u_i + \epsilon_i$ from seller i 's product, where ϵ_i is an i.i.d. seller-specific taste shock.⁴⁰ We consider two cases. In the first case, the shock is normally distributed, yielding demand for seller 1 of

$$D_1 = \Phi\left(\frac{u_1 - u_2}{\sqrt{2}}\right) + b,$$

where $\Phi(\cdot)$ is the CDF of the standard normal distribution. The second case is the logit demand model (where the shocks follow a type-I extreme value distribution). Demand is then

$$D_1 = \frac{e^{u_1}}{e^{u_1} + e^{u_2}} + b.$$

In either case, seller 2's demand is $1 - D_1$.

To allow us to compute the equilibrium, we focus on the one-dimensional case where sellers choose u_i , resulting in per-consumer revenue $r(u_i) = 1 - \psi u_i$ and cost $C(u_i) = u_i^2$. Given this model setup, profits are as in (1), and the equilibrium is found where $\frac{\partial \pi_1}{\partial u_1} =$

³⁹Here, we focus on $\mu = 1$ for brevity.

⁴⁰In our baseline Hotelling model, $\epsilon_2 = -t - \epsilon_1$, with ϵ_1 uniformly distributed on $[-t, 0]$.




Figure 6: The effect of an increase in b on u_1^* and u_2^* . Figures (a) and (b) show the case for normally distributed shocks, while (c) and (d) show the type-I extreme value distribution (logit) case.

$\frac{\partial \pi_2}{\partial u_2} = 0$. Solving for and plotting the equilibrium values of u_1^* and u_2^* yields Figure 6.

The key parameter of interest is ψ : payoffs are congruent when $\psi < 0$ and conflicting when $\psi > 0$. Thus, according to Proposition 1, we would expect that u_1 increases with b when $\psi < 0$ and decreases with b when $\psi > 0$. Conversely, we would expect u_2 to decrease with b when $\psi < 0$ and increase with b when $\psi > 0$. This is indeed what we observe for both the logit and probit cases.⁴¹

⁴¹We have computed the equilibrium for every combination of $b \in \{0, \frac{1}{100}, \frac{2}{100}, \frac{3}{100}, \dots, 1\}$ and $\psi \in \{-1, -\frac{49}{50}, -\frac{48}{50}, \dots, \frac{49}{50}, 1\}$ and verified that u_1 is increasing (and u_2 decreasing) in b if $\psi < 0$. This is reversed if $\psi > 0$.




Figure 7: The effect of an increase in b on $u_1^* - u_2^*$. Figures (a) shows the case for normally distributed shocks, while (b) treats the logit case.

B.3 Discrete choice specifications with multiplicative bias

We can repeat the analysis from the previous subsection for the case with multiplicative bias. Thus, sellers' demands are

$$D_1 = \Phi\left(\frac{u_1 - u_2}{\sqrt{2}}\right)(1 - b) + b, \quad D_2 = \Phi\left(\frac{u_2 - u_1}{\sqrt{2}}\right)(1 - b),$$

in the case of normally distributed preference shocks and

$$D_1 = \frac{e^{u_1}}{e^{u_1} + e^{u_2}}(1 - b) + b, \quad D_1 = \frac{e^{u_2}}{e^{u_1} + e^{u_2}}(1 - b)$$

for the logit case. As above, we take $r(u_i) = 1 - \psi u_i$ and $C(u_i) = u_i^2$. In light of Proposition 7, it makes sense to consider the effect of bias on $u_1^* - u_2^*$ and we would expect to see this increase under congruence and decrease under conflict. This is indeed the case. A plot can be found in Figure 7.

B.4 Non-uniform consumer locations in the Hotelling model

Lastly, we repeat the analysis by returning to the baseline Hotelling specification, but assuming that, for any $x \in (0, 1)$, there are $\tilde{\Phi}(x)$ consumers located to the left of x and $1 - \tilde{\Phi}(x)$ located to the right, where $\tilde{\Phi}(\cdot)$ is the CDF of the normal distribution with mean $1/2$ and variance 1 .⁴² Thus, firm 1's demand is $\tilde{\Phi}(x^* + b)$ and firm 2's demand is $1 - \tilde{\Phi}(x^* + b)$.

As above, we take $r(u_i) = 1 - \psi u_i$ and $C(u_i) = u_i^2$. Bias is no longer additively separable in firms' demand and, similarly to Proposition 7, we therefore focus on the effect

⁴²For convenience, we truncate the distribution by locating the residual mass of consumers at $x \in \{0, 1\}$ rather than distributing them along the $[0, 1]$ interval.




Figure 8: The effect of an increase in b on $u_1^* - u_2^*$ when consumers are distributed according to a truncated normal distribution on the Hotelling line.

of bias on $u_1^* - u_2^*$. We would expect to see $u_1^* - u_2^*$ increase in b under congruence and decrease under conflict. This is indeed the case. A plot can be found in Figure 8.