

Competition for Attention*

Pedro Bordalo

Royal Holloway, University of London,

pedro.bordalo@rhul.ac.uk

Nicola Gennaioli

Università Bocconi, and IGIER,

nicola.gennaioli@unibocconi.it

Andrei Shleifer

Harvard University,

shleifer@fas.harvard.edu

First draft, April 2013. This version: August 2015.

*Royal Holloway, University of London, Università Bocconi and IGIER, Harvard University. We are grateful to Charles Angelucci, Daniel Benjamin, Stefano DellaVigna, Matthew Gentzkow, Daniel Gottlieb, Emir Kamenica, Botond Koszegi, Annie Liang, Jeffrey Picel, Josh Schwartzstein, Jesse Shapiro, Charlie Sprenger and Yanos Zylberberg for extremely helpful comments. Shleifer thanks the Kauffman Foundation for research support.

Abstract

We present a model of market competition in which consumers' attention is drawn to the products' most salient attributes. Firms compete for consumer attention via their choices of quality and price. Strategic positioning of a product affects how all other products are perceived. With this attention externality, depending on the cost of producing quality some markets exhibit "commoditized" price salient equilibria, while others exhibit "de-commoditized" quality salient equilibria. When the costs of quality change, innovation can lead to radical shifts in markets, as in the case of decommodification of the coffee market by Starbucks. In the context of financial innovation, the model generates the phenomenon of "reaching for yield".

1 Introduction

In many markets, consumers' attention to particular attributes of a product seems critical. In fashion goods, business class airline seats, and financial products, consumers focus on quality rather than price. In these markets, firms advertise quality to draw consumers' attention. In fast food, economy air travel, or standard home goods, consumers seem much more attentive to prices. In these markets, firms typically advertise their low prices.

Scholars of strategy and marketing are keenly aware of these distinct modes of market competition, and tirelessly emphasize the importance of having differentiated attributes and drawing consumer attention to them (Levitt 1983, Rangan and Bowman 1992, Mauborgne and Kim 2005). Southwest wants to be known as "the low cost airline;" Singapore as the winner of prizes for luxury and comfort. Walmart touts its everyday low prices, Nordstrom's its service. Successful firms "frame" competition by focusing consumers' attention on their best attribute (quality or price). These mechanisms do not arise naturally in standard economic models, in which consumers attend to all product attributes equally.

This paper seeks to understand these phenomena. We take a standard model in which firms compete on quality and price, and add to it the mechanism of salience we developed elsewhere (Bordalo, Gennaioli, and Shleifer 2012, 2013). According to salience theory, the attention of decision makers is drawn to the most unusual, surprising, or salient attributes of the options they face, leading them to overweight these attributes in their decisions. Salience theory applied to consumer choice can shed light on a host of lab and field evidence on consumers' context dependent behavior. Such context dependence is well established in experiments, including the well known decoy effects (Huber, Payne and Puto 1982) and compromise effects (Simonson 1989). More recently, Hastings and Shapiro (2013) show using field data that, after a parallel increase in the prices of all gas grades, the demand for premium gas drops to an extent that cannot be accounted for by standard income effects. The salience model accounts for this evidence by recognizing that surprising price hikes focus consumer attention on gas prices, rather than quality, favoring the choice of cheaper grades.

In this paper, we show that the influence of prices and qualities on consumer attention has significant implications for market competition. In competitive markets, the salience of

price and quality are endogenously determined by the firms' strategic choices, and create an attention externality that lies at the heart of our model. A high quality good draws attention not only to its own quality, but also to the fact that the competitor product has lower quality, reducing the competitor's relative valuation. A good with a low price draws attention to the competitor's higher price, reducing the competitor's relative valuation. When salience matters, firms compete for consumer attention via the choice of quality and price.

We show that, depending on the cost of producing quality, some markets exhibit price salient equilibria in which consumers are most attentive to prices and less sensitive to quality differences. In these markets firms compete on prices, and quality could be under-provided relative to the efficient level. Because consumers neglect quality upgrades, escaping such "commodity magnets" is difficult. Fast food and budget air travel can be described in this way.

In other markets, equilibria are quality salient in that consumers are attentive to quality and are less sensitive to price differences. Firms compete on quality, which can be over-supplied relative to the efficient level. In these markets, it is again difficult to escape the high quality equilibrium because consumers neglect price cuts. Financial services or fashion can be described in this way.

We investigate how market equilibrium depends on the cost of providing quality. We explore the possibility of radical change in markets when the cost of producing quality changes dramatically. This can take the form of de-commoditization, whereby a firm acquires access to a technology of producing quality at a much lower cost than its competitor, and is able to change the market from a price-salient to a quality-salient equilibrium. Prices can then rise substantially, but quality as perceived by consumers rises more. Market transformation can also take the form of commoditization, which arises when industry costs fall dramatically, so that large price cuts become possible. As price becomes salient, and quality differences are neglected, firms reduce quality in order to cut prices even more.

Some of these effects can also arise in a traditional model, under judicious assumptions about consumer heterogeneity. Section 5 describes similarities and differences between salience and the traditional approach to innovation by using two real world examples. We begin by considering the case of financial innovation in the form of new products with higher

expected return and risk, such as mortgage backed securities (MBS). We show that such innovation is especially attractive in low interest rate environments, and when the innovation offers higher returns at a moderately higher risk. Higher returns are salient to investors when alternative yields are extremely low and the (small) extra risk of the new product is underweighted. The model generates the well documented phenomenon of “reaching for yield” in a psychologically intuitive way, based on the properties of salience (Becker and Ivashina 2014, Greenwood and Hanson 2013). The model is also consistent with the evidence on structured financial products sold to European consumers in the last decade (C  l  rier and Vall  e 2015).

We conclude the analysis by considering the de-commoditization of the coffee market after the entry of Starbucks. We show how innovation led to a radical expansion of the market for specialty coffee, with substantial increases in both quality and price. One key difference between ours and the more standard approach to this market lies in the drivers of change. In standard models, it is typically the marginal consumers who shift in response to changes in quality or price. In our model, in contrast, the attention and thus the price-sensitivity of all consumers changes in response to significant innovation. As a consequence, shifts in demand and market structure can be massive in a short period of time.

Our paper is related to recent work on “behavioral industrial organization” (Ellison 2006, Spiegler 2011). In some models, consumers restrict their attention to a subset of available options, the consideration set, which can be manipulated by firms by expending a marketing cost (Spiegler and Eliaz 2011a,b and Hefti 2012), by setting a salient low price on some products (Ellison and Ellison 2009), or by setting an inconspicuous price (de Clippel, Eliaz and Rozen 2014). In our model, the attention externality operates within a *given* consideration set.

A related literature recognises that firms seek to “frame” competition, for example by exposing the consumer to specific price formats that are favourable to them or that hinder price comparisons (Salant and Rubinstein 2008, Spiegler and Piccione 2012, Spiegler 2013). In our model, consumers’ price sensitivity is determined by firms’ choices of product attributes themselves, so “framing” efforts by firms affect quality provision.

Another strand of the literature considers the working of market competition in settings in which some product attributes are “shrouded”, namely sufficiently obscured that consumers

find it difficult to compare them across products (Gabaix and Laibson 2006, Ellison and Ellison 2009, Armstrong and Chen 2009). This literature takes as given the attributes that consumers pay attention to. In our analysis, instead, consumers’ differential attention to quality and price is shaped the the firms’ choices of attributes, so biases are endogenously determined in market equilibrium.

Azar (2008), Cunningham (2012), and Dahremöller and Fels (2012) explore models in which the relative weight that consumers put on different attributes depends on the choice context, and can thus be manipulated by firms. These papers model consumer attention by using approaches different from salience and explore a different set of issues, such as properties of markups or the monopolist problem. Finally, our analysis builds on recent work relating inattention to consumer demand. Some approaches – such as Gabaix (2014), Matějka and McKay (2012), and Persson (2012) – are grounded in the rational inattention framework, in which attention to different product features is efficiently allocated ex-ante. In our salience model consumer attention to different product attributes is drawn ex-post, depending on which attribute stands out. Koszegi and Szeidl (2013) and Bushong, Rabin and Schwartzstein (2014) follow related approaches.

The paper is organized as follows. In section 2, we describe our basic model of competition and show how salience would influence product valuations by consumers. In section 3, we take qualities as fixed and examine the basic analytics of price competition and of price and quality salient equilibria. Section 4 focuses on the full model of quality competition, and derives our main results for markets for products where attribute salience matters. In section 5, we apply the model to discuss innovation in the markets for financial products and for coffee. Section 6 concludes.

2 The Model

There are two firms, 1 and 2. Each firm $k = 1, 2$ produces one unit of a good having quality q_k at cost $c_k(q_k)$. Cost functions are common knowledge to firms (and consumers) and include a quality-dependent and a quality-independent component. Formally, $c_k(q) = F_k + v_k(q)$, where $v_k(q)$ is an increasing and convex function satisfying $v_k(0) = 0$. Here F_k captures the

cost of producing one unit of good k (not a fixed entry cost), so we refer to it as a unit cost. To obtain closed form solutions, we sometimes use the quadratic form:

$$c_k(q) = F_k + \frac{v_k}{2} \cdot q^2, \text{ for } k = 1, 2. \quad (1)$$

We assume that firm 1 has weakly lower total and marginal costs of quality than firm 2, namely $c_1(q) \leq c_2(q)$ and $c'_1(q) \leq c'_2(q)$ for all qualities q . In the quadratic case, this implies $F_1 \leq F_2$ and $v_1 \leq v_2$.

There is a measure one of identical consumers, each of whom chooses one unit of one good from the choice set $C \equiv \{(q_1, p_1), (q_2, p_2)\}$, where (q_k, p_k) stand for the quality and price of the good produced by firm k .¹ Both qualities and prices are measured in dollars and assumed to be known to the consumer. Absent salience distortions, each consumer values good $k = 1, 2$ at:

$$u(q_k, p_k) = q_k - p_k. \quad (2)$$

A salient thinker departs from (2) by inflating the weight attached to the attribute that he perceives to be more salient in the choice set $C \equiv \{(q_1, p_1), (q_2, p_2)\}$.

For each good k , its salient attributes are those whose levels are unusual or surprising, in the sense of being furthest from the reference attribute levels in the choice set C . Following Bordalo, Gennaioli and Shleifer (BGS 2013), we take the reference attribute levels to be the average levels in the choice set; thus, the reference good (\bar{q}, \bar{p}) , has the average price $\bar{p} = (p_1 + p_2)/2$ and the average quality $\bar{q} = (q_1 + q_2)/2$ in C .²

We model salience using a salience function $\sigma(x, y)$ that satisfies two main properties: ordering and homogeneity of degree zero. According to ordering, if an interval $[x, y]$ is contained in a larger interval $[x', y']$, then $\sigma(x, y) < \sigma(x', y')$. According to homogeneity of

¹Several of our results continue to hold with more than two firms in the market; see footnotes 12 and 14. To apply the salience framework to a more general model of market competition, the relevant market should be taken as the definition of the choice set.

²BGS (2013) defined reference attributes as the average over the expected and realised attribute levels. Here, attributes are deterministic so reference attributes coincide with the average over the choice set. This specification is the simplest measure of context, and we refer the interested reader to BGS (2013) for a more detailed discussion of this assumption. In terms of robustness, our results are invariant to replacing the average with any strict convex combination of the two goods in the choice set, see footnote 7. This holds for our baseline 2-goods setting, but also for the analysis of the *symmetric* equilibrium with more than 2 goods (generally, the specification matters only once there are three or more different quality price profiles).

degree zero, $\sigma(\alpha x, \alpha y) = \sigma(x, y)$ for any $\alpha > 0$, with $\sigma(0, 0) = 0$. In the choice set C , the salience of price for good k is $\sigma(p_k, \bar{p})$ while the salience of quality for good k is $\sigma(q_k, \bar{q})$. Good k 's quality is more salient than its price – or, for short, quality is salient – if and only if $\sigma(q_k, \bar{q}) > \sigma(p_k, \bar{p})$. Ordering and homogeneity of degree zero of the salience function imply that the salience of a good's quality is an increasing function of the percentage difference between the good's quality and the average quality in the choice set, and similarly for price.³

In our main analysis, we assume that salience distorts consumer valuation by attaching a higher and fixed weight to the most salient attribute of a good. This “rank based weighting” allows a stark and intuitive characterisation of the central implications of salience. In the Online Appendix B we analyze the model under a continuous weighting formulation, and show that our main results continue to hold. To keep the model tractable under both formulations, we impose an additional intuitive condition on the salience of attributes when the choice set has two goods, namely symmetry: $\sigma(a_1, \bar{a}) = \sigma(a_2, \bar{a})$ for $a = p, q$. In words, any attribute is equally salient for the two goods. As an example, the function $\sigma(a, \bar{a}) = |a - \bar{a}| / \bar{a}$ laid out in BGS (2012), which measures attribute salience as the proportional difference from the average value of the attribute, satisfies this symmetry property in our two goods context.

Under rank based weighting, the salient thinker's perceived utility from (q_k, p_k) is:

$$u^{ST}(q_k, p_k) = \begin{cases} q_k - \delta p_k & \text{if quality is salient} \\ \delta q_k - p_k & \text{if price is salient} \\ q_k - p_k & \text{if equal salience} \end{cases}, \quad (3)$$

where $\delta \in [0, 1]$ captures the extent to which valuation is distorted by salience (the above expression omits for simplicity the normalization factor $2/(1+\delta)$, see BGS (2013)). When $\delta = 1$, valuation coincides with (2) and the salient thinker behaves like a rational consumer. When $\delta < 1$, the salient thinker overweights the salient attribute. The competitive equilibrium then depends on δ , allowing us to study how salience affects market competition.

³In particular, consumers have diminishing sensitivity to attribute differences: increasing the prices of both goods by a uniform amount ϵ makes prices weakly less salient, $\sigma(p_k + \epsilon, \bar{p} + \epsilon) \leq \sigma(p_k, \bar{p})$ for $k = 1, 2$, and strictly so when $p_k \neq \bar{p}$. This property is consistent with Weber's law of sensory perception.

We assume consumer homogeneity for simplicity. In Online Appendix C we show that our main results extend to the case in which salience weighting varies across otherwise identical consumers. Allowing for heterogeneity in consumer tastes is an important topic for future research, particularly with regards to testing empirically the effect of salience on consumer demand.

Firms compete in two stages. In the first stage, each firm makes a costless commitment to produce quality $q_k \in [0, +\infty)$, and quality choices are observed by both firms. In the second stage, each firm competitively sets an optimal price p_k given the quality-cost bundle (q_k, c_k) it committed to, where $c_k \equiv c_k(q_k)$.⁴ To account for consumers' exogenous budget constraints, we assume that possible prices are bounded above, i.e. $p_k \leq p_{\max} < \infty$. With qualities and prices chosen, demand materialises and firms produce. Firm k 's payoff or profit is $\pi_k = d_k \cdot [p_k - c_k(q_k)]$ where $d_k \equiv d_k(q_k, q_{-k}, p_k, p_{-k})$ is the demand for good k at stage 2, such that $d_k = 1 - d_{-k}$.

To map the consumer preferences in (3) into demand functions d_k, d_{-k} , a “sharing” rule is required that specifies demand when ties arise in salience ranking or in valuation. Suppose that, at the action vector $(q_k, q_{-k}, p_k, p_{-k})$, good $-k$ is weakly preferred to good k , namely $u^{ST}(q_k, p_k) \leq u^{ST}(q_{-k}, p_{-k})$, but that good k is strictly preferred if its price is slightly reduced: formally, there exists $\bar{\epsilon}$ such that $u^{ST}(q_k, p_k - \epsilon) > u^{ST}(q_{-k}, p_{-k})$ for any $\epsilon \in (0, \bar{\epsilon}]$.⁵ We then specify the following sharing rule at vector $(q_k, q_{-k}, p_k, p_{-k})$: if firm k prices above cost, $p_k > c_k$, while firm $-k$ prices at cost, $p_{-k} = c_{-k}$, we set $d_k = 1, d_{-k} = 0$; in all other cases we set $d_k = d_{-k} = 1/2$. This sharing rule captures the idea that, at $(q_k, q_{-k}, p_k, c_{-k})$, firm k can infinitesimally reduce its price and capture the market with a profit.⁶

⁴This game is similar to the one in Shaked and Sutton (1982), except that we abstract from the initial stage in which firms decide whether to enter the market. In our game, firms always post their quality choices at stage 1 and their prices at stage 2, and they only incur costs if there is demand for their products.

⁵Such a shift in utility rankings can occur in three cases: i) the salience ranking is constant around p_k , and $u^{ST}(q_k, p_k) = u^{ST}(q_{-k}, p_{-k})$; ii) the salience ranking changes at p_k but the ranking of valuations (keeping salience fixed) does not; and iii) both the salience and the utility rankings change at p_k . In case i) it must be that $u^{ST}(q_k, p_k) = u^{ST}(q_{-k}, p_{-k})$ so when firm k lowers its price, the consumers goes from indifference to a strict preference for good k . In case ii), valuation jumps discontinuously at the salience bound. Because valuation is slack, it must be that $u^{ST}(q_k, p_k) < u^{ST}(q_{-k}, p_{-k})$. By lowering its price, firm k renders its advantage salient and strictly reverses the consumer's preference ranking. Finally, case iii) occurs when the goods are identical.

⁶This sharing rule is determined jointly with strategy selection in equilibrium. As in Reny (1999), we adopt this endogenous sharing rule to deal with discontinuities in the firms' payoff functions. Another way to avoid discontinuities arising from ties is to discretize the set of prices firms can set. In this case, the firm

We solve this game by finding subgame perfect equilibria. We restrict our attention to equilibria in pure strategies, but Online Appendix A proves that mixed strategy equilibria do not exist in our model. We describe equilibria in two steps. First, in Section 3 we take each firm’s quality and cost (q_k, c_k) as given and study price competition among firms. This price setting stage is of independent interest from endogenous quality choice because in the short run firms often take quality as given, and react to cost shocks only by adjusting their prices (in some settings firms may be unable to adjust quality, due to regulatory or technological constraints). The pricing game generally admits multiple equilibria, since the losing firm is indifferent between different strategy choices that yield zero market share (and thus zero profits). We restrict the equilibrium set by using the standard refinement that excludes equilibria in weakly dominated strategies. This refinement constrains firms to price weakly above cost so that, in equilibrium, the losing firm prices at cost.

In Section 4 we investigate the full game, endogenizing quality choice (given the refinement of the pricing game). Our main analysis deals with the case in which firms are symmetric, in the sense of having the same cost of quality $c(q)$. In this setting, under some conditions the game admits multiple equilibria that vary in the quality provided by the losing firm. However, we show that quality provision by the firm that captures the market is uniquely determined in equilibrium. To ease exposition, in the main text we focus on the symmetric equilibria of the game, which are unique. We characterise the full set of equilibria in the Appendix. We also consider what happens to this symmetric equilibrium when a shock occurs that reduces the cost of one of the two firms.

3 Price Competition

We begin with an analysis of price competition between firms 1 and 2, assuming that qualities q_1, q_2 and costs c_1, c_2 are fixed, and only prices are set by firms. Suppose that firm 1 chooses weakly higher quality than firm 2 and, as a consequence, incurs a weakly higher production cost, namely $q_1 \geq q_2$ and $c_1 \geq c_2$. In Section 4 we show that this is indeed the relevant case

delivering higher perceived surplus would set the highest price consistent with the consumer choosing its product, but this price will generally not leave the consumer indifferent between the two products.

when quality and costs are determined endogenously. Before characterizing the outcome under salience, consider the rational benchmark that obtains when $\delta = 1$.

Lemma 1 *When $\delta = 1$, the price competition subgame admits a unique pure strategy equilibrium under refinement, which satisfies:*

- i) If $q_1 - c_1 > q_2 - c_2$, then equilibrium prices are $p_1 = c_2 + (q_1 - q_2)$ and $p_2 = c_2$. Demand satisfies $d_1 = 1$ and firm 1 makes positive profits $\pi_1 = (q_1 - q_2) - (c_1 - c_2)$.*
- ii) If $q_1 - c_1 < q_2 - c_2$, then equilibrium prices are $p_1 = c_1$ and $p_2 = c_1 - (q_1 - q_2)$. Demand satisfies $d_2 = 1$ and firm 2 makes positive profits $\pi_2 = (c_1 - c_2) - (q_1 - q_2)$.*
- iii) If $q_1 - c_1 = q_2 - c_2$, then equilibrium prices are $p_1 = c_1$ and $p_2 = c_2$. Demand satisfies $d_1 = d_2 = 1/2$ and both firms make zero profits.*

All proofs are in the Appendix. In the rational benchmark, the firm creating greater surplus $q_k - c_k$ captures the entire market and makes a profit equal to the differential surplus created. When, as in case *iii*), the two goods yield the same surplus, firms share the market and make zero profits, as in standard Bertrand competition. The benchmark of fully homogeneous goods and zero profits corresponds to the special case $q_1 = q_2 = q$, and $c_1 = c_2 = c$.

To see how salience affects price competition, suppose that the firm producing the lower quality product 2 sets a lower price $p_2 \leq p_1$. The Appendix proves that this always holds in equilibrium. In particular, good 1 has higher (and good 2 has lower) quality and price than the reference levels, $\bar{q} = \frac{q_1 + q_2}{2}$ and $\bar{p} = \frac{p_1 + p_2}{2}$. Then, homogeneity of degree zero of the salience function implies that the same attribute – either quality or price – is salient for both goods. To see this, note that quality is salient (that is, quality is more salient than price for both goods) provided $\sigma(q_k, \bar{q}) > \sigma(p_k, \bar{p})$ for $k = 1, 2$, which holds if and only if the proportional difference in quality across goods is greater than the proportional difference in prices:

$$\frac{q_1}{q_2} > \frac{p_1}{p_2}. \quad (4)$$

Equivalently, quality is salient when the high quality good has a higher quality to price ratio than the low quality good (i.e., $q_1/p_1 > q_2/p_2$), while price is salient if and only if the reverse

inequality holds.⁷ Because the good that fares better along the salient attribute is overvalued relative to the other good, equation (3), salience tilts preferences in favor of the good that has the highest ratio of quality to price (BGS 2013).

This logic implies that the valuation of a good depends on the entire competitive context. In particular, by changing its price a firm imposes an “attention externality” on the competing good. To see this, suppose that $q_1 > q_2$ and $p_1 > p_2$, and the high quality firm reduces its price p_1 . This improves the consumer’s valuation of good 1 but, by making prices *less* salient, it also draws the consumer’s attention to the low quality of good 2. Suppose alternatively that the low quality firm reduces its price p_2 . This improves the consumer’s valuation of good 2, but by making prices *more* salient, it also draws his attention to the high price of good 1. Thus, by reducing price a firm draws the consumer’s attention to the attribute along which it fares better. This attention externality can either strengthen or dampen competitive forces, depending on the situation.

3.1 Salience and Competitive Pricing

When a firm sells to salient thinkers, it sets its price to render salient the advantage of its product relative to its competitor. To see how this affects competitive pricing, we examine price setting in two opposite situations, one in which quality is salient and firm 1 wins the market, another in which price is salient and firm 2 wins the market.

Consider first the optimal price set by the high quality firm 1 in order to win a quality-salient market (when it offers a higher perceived surplus to consumers). Suppose that firm 2 has set a price p_2 for q_2 . The maximal price p_1 at which firm 1 attracts consumers into buying its product while keeping quality salient solves:

$$\begin{aligned} \max_{p_1 \geq p_2} \quad & p_1 - c_1 \\ \text{s.t.} \quad & q_1 - \delta p_1 \geq q_2 - \delta p_2, \end{aligned} \tag{5}$$

$$q_1/p_1 \geq q_2/p_2. \tag{6}$$

⁷The condition (4) that determines the salience ranking is invariant to the specification of the reference attributes as any strict convex combination of the attributes in the choice set, namely $\bar{q} = \alpha q_1 + (1 - \alpha)q_2$ and $\bar{p} = \alpha p_1 + (1 - \alpha)p_2$ with $\alpha \in (0, 1)$.

The “valuation constraint” in (5) ensures that the consumer prefers good 1 when quality is salient. The “salience constraint” in (6) ensures that quality is indeed salient. At this point, it is useful to illustrate the sharing rule (and the salience tie-breaking rule) assumed above. Firm 1 captures the entire market, $d_1(q_1, q_2, p_1, p_2) = 1$, when both (5) and (6) hold – even if one constraint holds with equality – as long as $p_1 > c_1$ and $p_2 = c_2$. This is because only firm 1 can lower its price p_1 and satisfy both the salience constraint and the valuation constraint strictly; it thus captures the market with a salient advantage and positive profits.

The optimisation problem above presents two departures from the rational case. On the one hand, firm 1 now has an additional reason to cut its price: by setting p_1 low enough, it makes quality salient in (6), inducing the consumer to buy its high quality product. On the other hand, when quality is salient the high quality good is over-valued, which may allow firm 1 to hike its price p_1 above the rational equilibrium level. This effect of salience is captured by Equation (5).

Consider next the optimal price set by the low cost firm 2 to win a price salient market when firm 2 offers a higher perceived surplus to consumers. The maximal price p_2 at which firm 2 attracts consumers while keeping prices salient solves:

$$\begin{aligned} \max_{p_2 \leq p_1} \quad & p_2 - c_2 \\ \text{s.t.} \quad & \delta q_2 - p_2 \geq \delta q_1 - p_1, \end{aligned} \tag{7}$$

$$q_2/p_2 \geq q_1/p_1. \tag{8}$$

Once again, price setting is constrained by consumer valuation and salience. On the one hand, salience provides firm 2 with an additional incentive to cut its price, as doing so makes its lower price salient, inducing the consumer to buy its cheaper product. On the other hand, by causing an over-valuation of the cheap good, salience can allow firm 2 to charge a higher price than in the rational case.

This analysis suggests that, depending on the balance between the salience and valuation constraints, salient thinking may boost or dampen prices relative to a rational world. We now characterise equilibrium prices under salient thinking. We focus for simplicity on parameter configurations satisfying the restriction:

$$\text{A.1: } \delta(c_1 - c_2) < q_1 - q_2 < \frac{1}{\delta}(c_1 - c_2).$$

Assumption A.1 ensures that salience fully determines the preference of consumers among goods when prices are equal to production costs. If quality is salient, consumers prefer the high quality good 1; if price is salient they prefer the cheap good 2. This is akin to assuming that the two firms produce sufficiently similar surpluses $q_k - c_k$ that changes in salience change the consumer's preference ranking. In the Appendix, we extend the characterisation of equilibria to the full parameter space (not restricted to Assumption A.1).

Proposition 1 *For any parameter values $\delta \in [0, 1]$ and $q_1, q_2, c_1, c_2 \in R_+$ such that $q_1 \geq q_2$ and $c_1 \geq c_2$, the price competition subgame has a unique pure strategy equilibrium under refinement. Under A.1, this equilibrium satisfies:*

i) if $\frac{q_1}{c_1} > \frac{q_2}{c_2}$, prices are $p_1 = \min\{q_1 \cdot \frac{c_2}{q_2}, c_2 + \frac{1}{\delta}(q_1 - q_2)\}$ and $p_2 = c_2$. Quality is salient, demand satisfies $d_1 = 1$ and firm 1 makes positive profits.

ii) if $\frac{q_1}{c_1} < \frac{q_2}{c_2}$, prices are $p_1 = c_1$ and $p_2 = \min\{q_2 \cdot \frac{c_1}{q_1}, c_1 - \delta(q_1 - q_2)\}$. Price is salient, demand satisfies $d_2 = 1$ and firm 2 makes positive profits.

iii) if $\frac{q_1}{c_1} = \frac{q_2}{c_2}$, prices are $p_1 = c_1$ and $p_2 = c_2$. Quality and price are equally salient. Demand satisfies $d_k = 1$ if $q_k - c_k > q_{-k} - c_{-k}$ and $d_k = 1/2$ if $q_k - c_k = q_{-k} - c_{-k}$. Both firms make zero profits.

Under salience, the market equilibrium critically depends on the quality to cost ratios q_k/c_k of different products. A firm with a higher ratio q_k/c_k monopolizes the market and makes positive profits. When the two firms have identical quality to cost ratios, they earn zero profits in the competitive equilibrium.⁸

Proposition 1 holds because the firm having the highest quality to cost ratio can always engineer a price cut turning salience in its favor. When $q_1/c_1 > q_2/c_2$, the high quality firm can set a sufficiently low price that quality becomes salient, monopolizing the market. The low quality firm is unable to reverse this outcome: in fact, doing so would require it to cut price below cost. When instead $q_1/c_1 < q_2/c_2$, the low quality firm can set price sufficiently

⁸As we show in the Appendix, in the full parameter space equilibria in pure strategies of the pricing game exist, are unique, and can also be characterized by the quality-cost ratios of the firms. When A.1 does not hold, a qualitatively new type of equilibrium in pure strategies arises in which a firm may win the market at equilibrium prices for which its advantage is not salient.

low so that price is salient, and it monopolizes the market. The high quality firm is unable to reverse this outcome: once again, doing so would require it to cut price below cost. Finally, consider the case in which $q_1/c_1 = q_2/c_2$. In this case, as soon as a firm tries to extract some consumer surplus by setting a price above cost, its disadvantage becomes salient and the price hike becomes self defeating. The equilibrium outcome is zero profits for both firms.

The central role of the quality to cost ratio is economically appealing because it pins down salience distortions in terms of average costs of quality c_k/q_k . As we show when we endogenize quality, this feature allows our model to make tight predictions about how changes in cost structure affect salience and market outcomes. Before turning to that analysis, it is useful to look more closely at some implications of Proposition 1.

3.2 Price salient vs. Quality salient equilibria

Depending on the quality and cost parameters, salience leads to two types of equilibria: price salient and quality salient. In quality salient equilibria (case *i* of Proposition 1), consumers focus on quality for both goods. This resembles de-commoditized markets described in the marketing literature. In contrast, in price salient equilibria (case *ii*), consumers focus on prices but neglect quality differences among goods. This resembles the canonical description of commoditised markets (Rangan and Bowman 1992).

According to Proposition 1, in both types of equilibria the profits of the winning firm can be either lower or higher than in the rational benchmark. To see this, note that - due to the salience constraint - the equilibrium profits of the winning firm k (the one with lowest average cost) must satisfy:

$$\pi_k^S \leq q_k \cdot \frac{c_{-k}}{q_{-k}} - c_k = q_k \left[\frac{c_{-k}}{q_{-k}} - \frac{c_k}{q_k} \right], \quad (9)$$

where equality holds when the salience constraint binds. Equation (9) shows that equilibrium profits increase in the difference between the average cost of quality of the different firms. Consider the following special cases:

- The two goods yield different surpluses $q_1 - c_1 \neq q_2 - c_2$ but exhibit identical average costs of quality. Under rationality, the high surplus firm would make positive profits.

Under salient thinking, in contrast, industry profits are zero. When average costs of quality are identical (similar), a firm can always undercut its competitor and render its advantage salient. Price cuts are very effective and profits are lower than under rationality.

- The two goods yield the same surplus $q_1 - c_1 = q_2 - c_2$, but differ in their average costs of quality. Here profits are zero under rationality but positive under salient thinking. The reason is that the firm with the lower average cost of quality can set a price above cost and still be perceived as offering a better deal than its competitor. Price cuts by the losing firm are ineffective, and salience dampens competitive forces.

Salience can create abnormal profits in both quality and price salient equilibria (this result extends to industry profits as a whole). In quality salient equilibria, consumers overvalue the high quality good. The high quality firm is then able to hike prices and earn high profits. Financial services and fashion may be examples of this type of competition. In price salient equilibria, consumers are attentive to prices and under-appreciate quality differences among products. This grants an extra advantage to the cheap (and low quality) firm, allowing it to raise the price above cost. Fast-food industry and low-cost airlines may be examples of this type of competition.

4 Optimal Quality Choice

We now examine endogenous quality choice. In the first stage of the game, each firm $k = 1, 2$ makes a costless commitment to produce quality $q_k \in [0, +\infty)$, taking into account the price competition stage. In the second stage, firms compete in prices given the quality-cost attributes $(q_k, c_k(q_k))$, for $k = 1, 2$. The critical question is whether firm 1, which has lower costs, will choose to produce higher or lower quality than firm 2, and what this implies for the equilibrium market outcome.

The bulk of our analysis focuses on the symmetric case, in which firms have the same cost of producing quality, $c_1(q) = c_2(q) \equiv c(q)$. We view this case as capturing the long run outcome arising when all firms, through imitation or entry, adopt the best available

technology. Section 4.2 then considers how this equilibrium changes when one firm is hit by an asymmetric shock reducing its variable cost component.

4.1 Quality Choice in the Symmetric Cost Case

To fix ideas, consider the rational benchmark. Following Lemma 1, In stage 2 the market is monopolized by firm k producing the highest surplus $q_k - c(q_k)$. Anticipating this, in stage 1 the two firms set their qualities as follows.

Lemma 2 *When $\delta = 1$ and firms have identical cost functions, the full game admits a unique symmetric subgame perfect equilibrium in pure strategies: both firms set quality $q^* = \operatorname{argmax}_q [q - c(q)]$ (i.e. such that $c'(q^*) = 1$), price at cost, $p_1 = p_2 = c(q^*)$, and share the market, $d_1 = d_2 = 1/2$.*

In the rational benchmark (where $\delta = 1$), the quality provided in equilibrium maximizes total surplus. Quality provision decreases with the marginal cost of quality $v(q)$ but is independent of the unit cost F . Under the quadratic cost function of Equation (1), firms set:

$$q_1^* = q_2^* \equiv q^* = \frac{1}{v},$$

where v parameterizes the common marginal cost.

Consider now how salience affects quality choice. To build intuition, suppose that firms are at the “rational” quality level q^* . If consumers are salient thinkers, would firm 1 have an incentive to deviate to a different quality $q' \neq q^*$?

Consider the incentive of firm 1 to choose a marginally lower quality, cheaper, product. The new product has quality $q' = q^* - \Delta q$ and cost $c(q') = c(q^*) - \Delta c$. Whether this new product is successful or not against q^* critically relies on salience. If the lower quality q' is salient, the new product fails. If instead the lower price is salient, the new product may be successful. By Proposition 1, price is salient if and only if the quality to cost ratio of q' is higher than that of product q^* :

$$\frac{q^* - \Delta q}{c(q^*) - \Delta c} > \frac{q^*}{c(q^*)} \Leftrightarrow \frac{\Delta c}{\Delta q} > \frac{c(q^*)}{q^*}. \quad (10)$$

A cost cutting deviation works if the marginal cost of quality $\Delta c/\Delta q$ is higher than the average cost $c(q^*)/q^*$ at the rational equilibrium. This is intuitive: when the marginal cost is high, a small quality reduction greatly reduces the cost of firm 1. This allows firm 1 to set a salient low price, and to win the market.

The attention externality plays a key role here. As prices become salient, consumers pay less attention to quality, which reduces consumer valuation of the quality q' offered by the deviating firm. This effect may undermine the profitability of the new product. However, because price is now salient for *both* firms, the valuation by consumers of the competing product q^* drops even more! This externality allows the quality reduction to be profitable for firm 1.

Consider the alternative move whereby firm 1 deviates to a marginally higher quality product $q' = q^* + \Delta q$, which entails a higher cost $c(q') = c(q^*) + \Delta c$. If the higher price of q' is salient, the deviation fails. If however its higher quality is salient, the new product may be successful. This scenario occurs provided:

$$\frac{q^* + \Delta q}{c(q^*) + \Delta c} > \frac{q^*}{c(q^*)} \Leftrightarrow \frac{\Delta c}{\Delta q} < \frac{c(q^*)}{q^*}. \quad (11)$$

A quality improving deviation can work provided the marginal cost of quality is below the average cost at the rational equilibrium. Intuitively, if the marginal cost is low, a large quality improvement entails only a small price hike, making quality salient. Once again, the attention externality is at work. The salience of quality boosts consumer valuation of the new product, but it also draws the consumer's attention to the low quality q^* of the competing product. These effects cause a relative over-valuation of the high quality product q' , allowing the deviating firm to make profits.

This discussion delivers two messages. First, salience creates incentives to deviate away from the rational equilibrium. Second, the deviation can be toward higher or lower quality depending on the relationship between marginal and average costs of quality. This suggests that, if an equilibrium exists, it is likely to entail inefficient quality provision.

Another way to see this is to note that, according to the salience constraint in (6) and (8), the maximum price per unit of quality that firm k can extract (while still having its

advantage salient) is equal to the average cost $c_j(q_j)/q_j$ of the competing firm j . As a consequence, firm k has an incentive to raise quality when its marginal cost $c'_k(q_j)$ is lower than the marginal benefit $c_j(q_j)/q_j$, and to lower quality when the reverse is true. When the average cost of quality is high, the consumer pays a high price while still perceiving quality as salient. The equilibrium may feature quality over-provision. When the average cost of quality is low, the consumer notices even a slight price increase. Firm k now benefits from cutting both quality and price, so that quality under-provision may occur. The analysis of the model confirms that these conjectures are correct.

Proposition 2 *When $\delta < 1$ and firms have identical cost functions, there is a unique subgame perfect symmetric equilibrium in pure strategies. Denote by \bar{q} and \underline{q} the quality levels such that $c'(\bar{q}) = 1/\delta$ and $c'(\underline{q}) = \delta$, and by $\hat{q}(F)$ the quality level minimizing average cost, namely $\hat{q}(F) \equiv \arg \min c(q)/q$. Then, in the unique symmetric equilibrium price and quality are equally salient, quality provision is given by:*

$$q_1^S = q_2^S = q^S \equiv \begin{cases} \bar{q} & \text{if } F > \bar{F} \equiv \bar{q}/\delta - v(\bar{q}) \\ \hat{q}(F) & \text{if } F \in [\underline{F}, \bar{F}] \\ \underline{q} & \text{if } F < \underline{F} \equiv \underline{q}\delta - v(\underline{q}) \end{cases}, \quad (12)$$

(so q^S is weakly increasing in F), firms price at cost, $p_1 = p_2 = c(q^S)$, and share the market, $d_1 = d_2 = 1/2$.

This equilibrium has three main features. First, because costs are identical, firms produce the same quality, face the same production costs, and charge the same price. But then, because firms sell identical products, price and quality are equally salient in equilibrium, so consumers value the products that are offered correctly (as in the case where $\delta = 1$), and firms make zero profits.

Second, although *in equilibrium* consumers correctly value the goods produced, there is inefficient provision of quality (and therefore lower consumer surplus) relative to the rational case. The reason is that salience makes the firms unwilling to deviate towards the socially efficient quality q^* . When quality is over-provided ($q^S > q^*$), reducing quality and price backfires because consumers' attention is drawn to the quality reduction, rather than to

the price cut. This sustains an equilibrium with high quality and high prices. Similarly, when quality is under-provided ($q^S < q^*$), increases in quality and price backfire because consumers focus on the price rather than the quality hike. This sustains an equilibrium with low quality and low prices. Although in equilibrium both attributes are equally salient, we refer to the equilibrium with quality over-provision as quality-salient and to the equilibrium with under-provision as price-salient. This terminology underscores which salience ranking constrains firms from deviating towards the efficient quality level.

The third key feature of the equilibrium is that – unlike in the rational case – quality provision weakly increases in the unit cost F .⁹ Intuitively, F affects average costs and thus the firms’ best responses. When F is high, costs and thus prices are high. By the diminishing sensitivity property, the salience of prices is low. The firm has an incentive to boost quality because any small extra cost can be “hidden” behind the already high price. As a consequence, the small extra price is not salient and quality is over-provided. When in contrast F is low, costs and thus prices are low. By diminishing sensitivity, prices are now very salient. In this case, any price cut is immediately noticed, encouraging firms to cut costs to an extent that quality is under-provided.

The influence of F on quality is a distinctive prediction of the salience model in settings where the composition of demand stays constant (as in Proposition 2). If instead the composition of demand is allowed to change with changes in F , then the rational model can also predict that quality provision changes.¹⁰ Importantly, however, even with consumer heterogeneity the salience model has a distinctive prediction: the price sensitivity for *all* consumers goes down as F increases. The rational model - in which preferences are exogenous - does not share this prediction, which is empirically testable with individual level data.

To see these effects clearly, consider the case of the quadratic cost function.

Corollary 1 *When $\delta < 1$ and firms have identical quadratic costs $c(q) = F + v \cdot q^2/2$, quality*

⁹In fact, $\hat{q}(F)$ satisfies $v'(\hat{q}) \cdot \hat{q} - v(\hat{q}) = F$ and the left hand side increases in q because $v(\cdot)$ is convex.

¹⁰For instance, if taste heterogeneity is large, an increase in F might cause low valuation consumers to drop out of the market, and induce firms to optimally increase quality to attract the remaining high valuation consumers. In Online Appendix C we study a version of the model with consumer heterogeneity.

provision in the symmetric equilibrium is given by:

$$q_2^S = q_1^S = q^S \equiv \begin{cases} \frac{1}{\delta v} & \text{if } F \cdot v > \frac{1}{2\delta^2} \\ \sqrt{\frac{2F}{v}} & \text{if } \frac{1}{2\delta^2} \leq F \cdot v \leq \frac{\delta^2}{2} \\ \frac{\delta}{v} & \text{if } F \cdot v < \frac{\delta^2}{2} \end{cases} . \quad (13)$$

Figure 1 below plots q^S as a function of the unit cost F , and compares it to the surplus maximizing quality, given by $q^* = 1/v$. As evident from the figure, salience causes quality

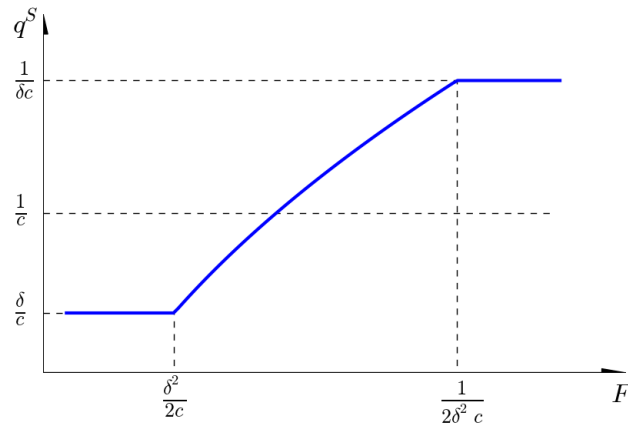


Figure 1: Quality provision in the symmetric equilibrium (quadratic cost).

to be over-provided when the unit cost F is sufficiently high and under-provided otherwise. Recall that for $\delta = 1$, we have $q^* = 1/v$ and quality provision does not depend on F .

This analysis may help explain why sellers of expensive goods such as fancy hotel rooms or business class airplane seats compete mostly on the quality dimension, often providing customers with visible quality add-ons such as champagne, airport lounges, or treats. These visible quality add-ons help make overall product quality salient, and the profit margin associated with them can be hidden behind the high cost of the baseline good. In contrast, sellers of cheap goods such as low quality clothes or fast food compete on the price dimension. These firms cut product quality because it allows them to offer substantially lower prices. These cuts are proportionally larger in the price dimension, draw consumers' attention to prices, and thus enable firms that supply these cheap goods to make abnormal profits. In both cases, equilibrium profits disappear as competing firms adopt the same add-on or quality

cutting strategies, despite the fact that they are providing inefficient levels of quality.^{11,12}

4.2 Innovation as a Cost Shock

We now use our model to explore the implications of salience for product innovation. We view innovation as a change in product characteristics and market equilibrium triggered by a cost shock. The shock hits a market in a long run symmetric equilibrium of Proposition 2. We distinguish between industry-wide cost shocks, such as those caused by deregulation or changes in input prices, and firm-specific shocks such as those stemming from the development of a new technology by an individual firm. This taxonomy illustrates the separate effects of the two key forces driving salience: diminishing sensitivity and ordering. Industry-wide shocks work mainly through diminishing sensitivity because they alter the average value of different attributes in the market. Firm-specific shocks instead work mostly through ordering: they allow one firm’s product to stand out against those of its competitors.

Real world innovation episodes often combine firm-specific and industry-wide factors. Initially only some firms discover new technologies or change their strategies in response to common shocks, so that the initial phase is effectively firm-specific. Subsequently, the new technologies or strategies spread to other firms, becoming industry-wide phenomena. One could view our analysis here as providing snapshots of short and long-run market adjustments to shocks. We leave the modelling of industry dynamics under salience to future research.

In what follows, we restrict attention to the case of quadratic costs, in which $c_k(q_k) =$

¹¹The diminishing sensitivity property is also present in Prospect Theory (reviewed in Tversky and Kahneman, 1981). The distinctive feature of our model is the attention externality, namely the fact that changing attributes of one product alter the valuation of the competing product. This ingredient is important to generate strong reactions to price or quality changes. The benefit for a firm of increasing quality (and price) is particularly large when it induces the consumer to focus more on the full quality provided and on the lower quality of the competing product. In fact, this mechanism implies that there is a complementarity between an add-on quality and the baseline quality level.

¹²These examples illustrate how the results of this Section can be used to study markets with $N > 2$ firms. The symmetric equilibrium that arises when firms are identical, described in Proposition 2, continues to hold for $N > 2$ identical firms. The intuition is simple: consider firm 1’s incentives to deviate from the symmetric equilibrium q^S to quality q_1 . If the deviating firm prices at cost, the reference attributes are now $\bar{q} = \frac{q_1 + (N-1)q^S}{N}$ and $\bar{p} = \frac{c(q_1) + (N-1)c(q^S)}{N}$. Good 1’s advantage relative to the reference good – higher quality, or lower cost – is salient if and only if it has a higher quality cost ratio. This holds if and only if $q_1/c(q_1) > q^S/c(q^S)$, which is exactly the same expression as in the 2 firm case. In this equilibrium analysis, the same attribute is again salient for all goods because there are effectively only two qualities and prices (in more general cases, salience ranking is good-specific, see BGS 2013). As a result, extending the analysis to N identical firms does not change the *symmetric* equilibrium of the model (see Proof of Proposition 2).

$F_k + \frac{v_k}{2} \cdot q_k^2$, for $k = 1, 2$. We begin our analysis by considering industry-wide shocks to an industry in symmetric equilibrium.

Proposition 3 *Suppose that the market is in the equilibrium described by Equation (13).*

We then have:

i) A marginal increase (decrease) in the unit cost F of all firms weakly increases (decreases) equilibrium quality provision under salient thinking ($\delta < 1$) while it leaves quality unaffected under rationality ($\delta = 1$).

ii) A marginal increase (decrease) in the marginal cost of producing quality v of all firms strictly decreases (increases) equilibrium quality provision. Under salient thinking, the decrease (increase) in quality is larger than under rationality ($\delta = 1$) if and only if in the original equilibrium quality is sufficiently over-provided.

With rational consumers, changes in the unit cost F do not affect quality provision. With salient thinkers, they do. The logic is identical to that of Proposition 2: when unit costs, and thus price levels, are higher, given price differences are less salient (by diminishing sensitivity of the salience function). This reduces consumers' price sensitivity, and makes it attractive for firms to upgrade their quality. As an example, the transportation costs involved in exporting German cars to the United States (akin to a rise in F relative to the home market) may cause the car manufacturers to compete on quality provision in the US market, more than in the domestic market, by adding quality add-ons to their cars. Similarly, truffles are served in omelettes in Provence, while truffle "shavings" are added to elegant dishes in the United States, where truffles are in relative terms much more expensive. Lobster is more likely to be served boiled in Boston than in Chicago. Conversely, a reduction in the tariffs on textile imports from China (akin to a drop in F) may induce clothing manufacturers in Europe to intensify price competition relative to the situation with higher tariffs.

The effect of a drop in the marginal cost of producing quality v is more standard. As in the rational case, this shock increases quality provision. However, salience modulates the strength of this effect. The boost in quality provision is amplified at very high cost levels, when there is over-provision of quality, while it is dampened in all other cases. This effect is again due to diminishing sensitivity: by reducing the level of prices, reductions in v render

consumers more attentive to price differences, reducing firms' incentive to increase quality.

Consider next the effect of a firm-specific shock. Suppose that, starting from a symmetric equilibrium, firm 1 acquires a cost advantage that enables it to monopolize the market. One complication is that in this asymmetric case there is typically a multiplicity of equilibria (both under rationality and salience): like in the price-competition subgame of Section 3, the losing firm is indifferent between choosing among quality levels leading to zero profits. To derive comparative statics, and compare the predictions of the salience model to those of the rational model, we introduce an intuitive equilibrium selection rule: we keep the quality of a firm fixed at the pre-innovation, "symmetric play", unless it is strictly profitable for the firm to deviate from it (given the other firm's best response to the original symmetric equilibrium). As we now show, this rule uniquely pins down the equilibrium both in the rational and the salience cases.¹³

For brevity, we report only the effects of reductions in the variable cost of quality. In the rational model, we find:

Lemma 3 *Suppose that, starting from the symmetric equilibrium of Equation (13), the variable cost of firm 1 drops to $v_1 < v_2 = v$. Then, when $\delta = 1$, in equilibrium firm 1 captures the market, $d_1 = 1$, and makes positive profits, $\pi_1 > 0$. Under the "symmetric play" selection rule, there is a unique subgame perfect equilibrium in pure strategies characterized by quality choices q_1^*, q_2^* , such that:*

$$q_1^* = \frac{1}{v_1} > q_2^* = \frac{1}{v}. \quad (14)$$

As a consequence, firm 1 increases quality provision, it wins the market ($d_1 = 1$), and makes positive profits. Equilibrium prices are $p_1^ = p_2^* + (q_1^* - q_2^*)$ and $p_2^* = c_2(q^*)$.*

In the equilibrium pinned down by our selection rule, both firms choose the quality level that - given their own costs - maximizes social surplus. Relative to the symmetric benchmark

¹³In this sense, equilibria are characterized by firms' best response to each other's "symmetric play". This equilibrium selection rule is based on the idea that firms face some inertia in adjusting their quality level, and so they keep quality constant unless it is strictly beneficial for them to unilaterally deviate from the pre-shock symmetric play. A more detailed characterization of asymmetric equilibria that includes equilibria not satisfying this equilibrium selection rule is available upon request.

in which both firms have marginal cost $v = v_2$, firm 1 increases quality provision, wins the market, and makes positive profits.

Consider now the case of salient thinking (namely $\delta < 1$):

Proposition 4 *Suppose that, starting from the symmetric equilibrium of Equation (13), the variable cost of firm 1 drops to $v_1 < v_2 = v$. Then, when $\delta < 1$ firm 1 monopolizes the market, $d_1 = 1$, and makes positive profits. Under the “symmetric play” selection rule, there are two cases:*

- i) The cost shock is large, $v_1 < v/2$. Then, firm 1 boosts both its quality and its price.*
- ii) The cost shock is small, $v_1 > v/2$. Then, there is a threshold $\widehat{F} > 0$ such that firm 1 boosts its quality and price if and only if $F \geq \widehat{F}$. If $F < \widehat{F}$, firm 1 keeps its quality constant at the competitor’s level δ/v and wins the market.*

The size of the cost shock plays a critical role. If the variable cost reduction is drastic, or if the unit cost is high (i.e. $F \geq \widehat{F}$), firm 1 can win the market by boosting quality provision. In this case, prices tend not to be salient, because average costs are high, and therefore quality differences can be large. In this configuration, a substantial quality upgrading alters the market outcome, changing the equilibrium from price- to quality- salient: as firm 1 provides extra quality, the overall quality of its product becomes salient, and consumers’ willingness to pay rises even for infra-marginal quality units. In this sense, the quality add on acts as a complement to baseline quality, greatly increasing the price that firm 1 can charge for its product. This logic provides the testable predictions: i) quality improving innovations regularly occur for goods that are already of high quality (and expensive), and ii) the level of such quality add-ons should respond positively to increases to the unit cost F , and to reductions of the marginal cost of quality.

Matters are different when the cost shock is small, $v_1 > v/2$, and the unit cost is low (i.e., $F < \widehat{F}$). Now prices tend to be salient because of low average cost of quality, and the small cost advantage also makes it very costly for firm 1 to engineer a drastic increase in quality. In this case, quality upgrades make the associated price hikes salient, and thus backfire. As a consequence, it is optimal for firm 1 to keep its quality and price constant at the symmetric equilibrium level, since given the sharing rule firm 1 is then guaranteed

to capture the market. This outcome, while puzzling in a rational model, is natural with salience: in a price-salient equilibrium quality upgradings are neglected, and firms exploit lower costs to cut prices.

An important implication of this analysis is that price-salient equilibria are very stable, particularly for low cost industries, and that in these industries quality upgrades are very hard to materialize. To escape a commoditized market, an individual firm must develop a drastic innovation that allows it to provide sufficiently higher quality than its competitors, and at such reasonable prices that quality becomes salient. Small cost reducing innovations neither beat the “commodity magnet” nor lead to marginal quality improvements. They just translate into lower prices.¹⁴

This result more generally illustrates the working of our model when costs are asymmetric. The low cost firm wins the market, but whether it does so by setting higher quality or lower price depends on the extent of its cost advantage. If it has a large cost advantage, the low cost firm captures the market by setting a salient high quality. If the cost advantage is small, the low cost firm captures the market by setting a salient low price.

5 Applications

We now investigate in greater depth how the effects of innovation in our model differ from those of a standard model. We organise the discussion around two applications. In Section 5.1 we show that our model can capture some features of financial innovation. In Section 5.2 we discuss innovation in the coffee market in the US.

5.1 Financial Innovation

Our model can shed light on financial innovation, and the phenomenon of “reaching for yield” whereby investors are more likely to take risk to earn a higher return in low interest rate environments (Greenwood and Hanson 2013, Becker and Ivashina 2014). We describe innovations that occurred in the safe (AAA) asset market, involving the creation of mortgage

¹⁴This result extends to the case of $N > 2$ firms, where *one* firm receives an idiosyncratic shock to variable costs of quality, while the remaining $N - 1$ firms stay in the long run symmetric equilibrium quality q^S (by the “symmetric play” selection rule). See the proof of Proposition 4.

backed securities (MBS). A security i is characterized by the expected return R_i it yields to investors (net of intermediation fees), and its risk ρ_i . The investor’s “rational” valuation of asset (R_i, ρ_i) is mean-variance, namely:

$$u_i(R_i, \rho_i) = R_i - \rho_i. \tag{15}$$

The assumption of mean-variance preferences is standard in financial economics, and seems particularly appropriate to study the “reach for yield” phenomenon.¹⁵

Under salient thinking, the investor overweights the more salient attribute, which can be either risk or return. We assume salience is determined by comparing only assets in the same AAA risk class.¹⁶ Suppose that the investor chooses between two assets $i = 1, 2$ and the salience function is $\sigma(\cdot, \cdot)$. The following cases can occur. If $\sigma(R_1, R_2) > \sigma(\rho_1, \rho_2)$, returns are salient and the investor values asset i at $R_i - \delta \cdot \rho_i$. If $\sigma(R_1, R_2) < \sigma(\rho_1, \rho_2)$, risk is salient and the investor values asset i at $\delta \cdot R_i - \rho_i$. Finally, if $\sigma(R_1, R_2) = \sigma(\rho_1, \rho_2)$, risk and return are equally salient and the investor’s valuation is rational.

Initially, there are two financial intermediaries, or brokers, $i = 1, 2$, each offering to the investor an identical asset, characterized by a gross expected return \bar{R} and risk ρ .¹⁷ In this pre-innovation benchmark, both intermediaries offer the “standard” asset, such as government bonds. We assume that intermediaries offer this asset to investors at some fee, which can be thought of as brokerage or management fee.

Intermediaries compete by offering investors assets with net expected return $R_i \leq \bar{R}$ and risk ρ . Thus, $\bar{R} - R_i$ is the brokerage or management fee of intermediary i . Investors decide with which intermediary to invest. Competition then works as in Section 2, where quality and

¹⁵Mean-variance utility allows us to map financial assets directly into the previous quality-price model. Moreover, this formalism corresponds to the standard framing of financial products in terms of risk and expected return (at least for retail investors), and may thus capture important psychological aspects of this phenomenon. In turn, this implies that a mean-variance formulation greatly facilitates any empirical analysis of our results (as for example in C el er and Vall e 2015, see below). Studying the model under non-separable preferences such as Expected Utility calls for the use of the salience formalism of Bordalo, Gennaioli and Shleifer (2012).

¹⁶This setting represents investors (possibly including fund managers) choosing which AAA securities to hold, rather than optimizing over the whole range of assets of different risk categories.

¹⁷The model works identically in the case where there are $N > 2$ intermediaries, as long as, as assumed here, only one intermediary has access to an innovation while the remaining $N - 1$ keep their traditional strategy. In this sense, our model captures the initial phases of innovation in which one firm introduces a new product, and studies the conditions that lead to the innovation’s success or failure.

cost are fixed:¹⁸ each intermediary offers a net of fee expected return R_i , which is analogous to product quality, at the cost to the investor of bearing risk ρ , which is analogous to price. As a consequence, the upside of the asset with the highest ratio of expected return to risk is salient, causing that asset to be overvalued relative to its competitor's. Because firms are identical and expected returns and risk are given exogenously, the following equilibrium benchmark holds both in the rational case and with salient thinkers.

Lemma 4 *With no innovation, intermediaries charge zero fees, $R_1 = R_2 = \bar{R}$, and make zero profits, and the investor is indifferent between the two firms.*

As in standard Bertrand competition, the two firms selling the same asset make zero profits, offering the full expected return \bar{R} to the investor (under salient thinking, the logic is the same as that of Proposition 1 point *iii*)).

Against this benchmark, we model financial innovation as the creation by one intermediary of a technology to generate excess return at only a moderate extra risk. The innovator, say intermediary 1, may for example find a way to better diversify the risks from the securities it already manages and thus offer a different asset to investors. Formally, intermediary 1 creates a new asset in the same asset class, with the gross expected return:

$$\bar{R} + \alpha,$$

where α is the new asset's excess expected return. The asset's risk then increases to:

$$\rho + \frac{v}{2} \cdot \alpha^2,$$

where v captures the marginal cost – in terms of added risk – of creating excess expected return α . Intermediary 2 continues to offer the standard product with gross expected return \bar{R} . The no-innovation benchmark can be viewed as the extreme case where v is prohibitively high for both firms.

¹⁸The only difference is that in the setting of Section 2, firms' pricing strategies determine the cost for consumers to buy the good, while here the firms' pricing strategies determine the "quality" of the asset for the investor (namely the investor's return), while cost is exogenously given by the asset's risk.

With fully rational investors, the working of innovation is straightforward. In the spirit of Lemma 3, the innovator: i) captures the entire market by offering the investor a net return of $\bar{R} + (v/2) \cdot \alpha^2$ (which compensates the investor for bearing the extra risk), and ii) sets α to maximize its profit:

$$\max_{\alpha} \alpha - (v/2) \cdot \alpha^2 \quad (16)$$

which implies $\alpha^* = 1/v$. The lower is the extra risk v , the greater is the excess return promised by the new financial product. As intermediary 1 manufactures an asset with a better return/risk combination, its profit and thus social welfare rise (the investor is left indifferent).

In the case of salient thinking, the critical question is whether, compared to the standard asset, the new asset's risk or expected return is salient. Depending on which attribute is salient, the innovator will have an incentive to create a particular return vs. risk profile. The reason is that under salience the investor's risk appetite endogenously depends on the salient features of the new asset. The new equilibrium is as follows.

Proposition 5 *The innovating broker 1 captures the market and makes positive profits. The optimal excess expected return satisfies:*

$$\alpha^* = \begin{cases} \frac{1}{\delta \cdot v} & \text{for } \bar{R} < \delta \cdot \rho \\ \frac{\rho}{\bar{R}} \cdot \frac{1}{v} & \text{for } \bar{R} \geq \delta \cdot \rho \end{cases} \quad (17)$$

Relative to the rational benchmark, under salient thinking there is excessive risk taking if $\bar{R} < \rho$ and too little risk taking if $\bar{R} > \rho$.

The innovation is particularly successful when investors focus on the extra return offered by the new asset and underweight the extra risk that comes with it. As Proposition 5 illustrates, this is the case precisely when the net expected return \bar{R} of the standard asset is low. Diminishing sensitivity generates a “reach for yield” at low interest rates: an excess return of, say, 0.5% is much more salient when the baseline return is 1% than when the baseline return is 6%.¹⁹ Proposition 5 shows that in this case financial intermediaries have an incen-

¹⁹Proposition 5 also shows that financial innovations geared at creating excess returns are much less

tive to offer excessively risky products. When investors focus on return, they underweight risk, enabling the broker to charge high fees. While our analysis focuses on fixed income markets, C  lerier and Vall  e (2015) present striking evidence on financial products offered to retail customers by European financial institutions. These products are characterized by high excess promised returns (as well as high but in part hidden risks) particularly when the benchmark interest rate is low. The authors interpret their findings as support for our model.

An important implication of this analysis is that, when investors’ attention is drawn to expected returns, risks are relatively speaking neglected, and investors end up disappointed when bad returns materialize. Gennaioli, Shleifer, and Vishny (2012, 2013) modeled this neglect of risk as investors’ disregard of tail events, and presented some evidence consistent with the prediction that downside risks were neglected in the period preceding the 2007 – 2008 financial crisis. The salience approach makes a similar point in a perhaps subtler way. During the “reach for yield” episodes, interest rates are low and investors are prone to be inattentive to risks. When investors underweight risks, they engage in too much risk taking. When bad states of the world materialize, these investors wish they had paid more attention.

5.2 Starbucks

In the mid 20th century, coffee was known as “America’s favourite drink” with over half of US adults being daily consumers (Pendergrast 1999, Koehn 2005).²⁰ Yet, most people drank low quality blends brewed at home or drip coffee at restaurants. Up until the mid 1990s, the US coffee industry was dominated by a handful of roasters (Maxwell House, Folgers, Nestle) which sold coffee beans in supermarkets. The market was characterized by low quality and fierce price competition. The major roasters “clashed in frequent price wars, using coupons, discounts and other promotions” and “sought ways to cut costs” (Koehn

successful when net returns are already high. In this case, the investor is much less sensitive to a given increase in return, and the innovating firm must keep the risks of the new asset very low, lest the investors focus on them. In this case, there is too little risk taking, in the sense that the intermediary selects an excess return in (17) below its rational counterpart in (16). Here the intermediary may find it profitable to reduce excess returns and risks relative to the standard asset.

²⁰Coffee was the second most valuable commodity exported by developing countries from 1970 to 2000 (after crude oil), and the seventh largest agricultural export by value in 2005. The retail value of the US coffee market is currently estimated at \$30bn (<http://www.scaa.org/PDF/resources/facts-and-figures.pdf>).

2005). Product innovations effectively translated into reductions in quality that allowed for price cuts: for example, roasters progressively increased the share of the cheaper, lower quality robusta beans in their existing coffee blends.²¹ Overall consumption of coffee declined slightly through the 1970's and 1980's, a trend the National Coffee Association attributed in part to the “price focused” position of the industry's leading producers (Kachra, 1997). When dominant roasters attempted to revive the market and introduce higher quality coffee, these attempts failed (Slywotzky, 1995).

In this regime, high quality whole coffee beans – known as specialty coffee – were a niche market. Starting in the late 1960s, a small number of firms such as Peets Coffee & Tea in San Francisco, and later Starbucks in Seattle, offered high quality roasted beans at high prices to a small devoted clientele.²² This market experienced some growth, particularly in the US Pacific Northwest. As a share of the overall market, however, specialty coffee remained small (less than 10% by 1989, see Figure 2). At this point, most people – including future Starbucks CEO Howard Schultz – had never tasted high quality, specialty coffee, let alone an espresso.²³

Starting in the late 1980s, the coffee market experienced a drastic change. In 1987, under Howard Schultz' direction, Starbucks introduced the Italian coffee shop model, bringing ready to drink, high quality, espresso drinks to the mass market, and providing a “cafe” experience through a comfortable in-shop environment. Starbucks' innovation was to find a profitable way to sell espresso drinks for the mass market, by providing consistently high quality delivered by trained baristas. This innovation revolutionized the market: from a few dozen stores in the late 1980s, Starbucks expanded to over 3,500 stores in 2000, and over

²¹Roasters also progressively increased the share of lower quality Arabica beans, started under-roasting beans (which increased bitterness), and packing them in bricks instead of cans (which required using stale beans, Andrews 1992, Pendergrast 1999). This price competition regime is also illustrated by several quality-discount practices: i) companies “started packing coffee in 14-ounce cans and selling them at prices that previously had applied to one-pound containers” (Koehn, 2005), ii) in the restaurant market, the “bottomless cups” or free coffee refills (essentially a price reduction) were the norm.

²²Specialty coffee is typically understood as made from high quality Arabica coffee beans, and sold at a significant premium over value supermarket brands. In the seventies, the retail price of specialty coffee beans averaged \$5 to \$7 per pound, twice the price of the traditional variety (Koehn, 2005).

²³Schultz relates his first encounter with specialty coffee upon visiting Starbucks, and writes “By comparison, I realised, the coffee I had been drinking was swill” (Schultz and Yang, 1997). Espresso was invented in Milan in the early 1900s. It quickly took over the Italian market, and spread to the rest of Europe, with especially fast growth in the 1950s (Pendergrast, 1999).

11,000 stores in 2010 (and over 17,000 worldwide). By the mid-1990s, Starbucks accounted for over one third of US coffee shops. Fueled by the skyrocketing demand for espresso drinks, specialty coffee expanded from under 10% of sales in 1989 to 40% in 1997 (Vishwanath and Harding 2000), and over 50% today (Figure 2, solid line). Because total coffee consumption stayed approximately constant (Figure 2, dashed line), this rise reflected a substitution away from traditional towards specialty coffee, despite the latter’s significantly higher price (Koehn 2005).²⁴

These changes also transformed the retail market for coffee beans: by the late 1990s, traditional brewers had started to invest in higher quality beans, successfully introducing new premium brands in supermarkets (Pendergrast 1999). Marketing analysts dubbed this the “Starbucks effect”, meaning that Starbucks increased the perceived “premiumness” of the coffee category (Vishwanath and Harding 2000). Competition had shifted to quality across the board, not only in the coffee shop market. Today, even McDonalds advertises “100% Arabica coffee, freshly brewed every 30 minutes”.

5.2.1 Saliency and the “Starbucks effect”

The coffee market’s rapid switch from price to quality competition has a natural interpretation in light of the saliency model. Consider the model of Section 4, with two firms competing on quality and price.²⁵ We distinguish between factors affecting the unit cost F of producing coffee and the marginal cost v of producing higher quality.

The pre-Starbucks era can be characterized by a commoditized coffee market, in which consumers are focused on prices. Selling “coffee in a can” in supermarkets rendered the unit cost F very low, compared for instance with that of offering freshly ground coffee in a shop. Furthermore, real coffee prices exhibited a long gradual decline between the late

²⁴To this point, the National Coffee Association recently estimated that out of the 60% of American adults who drink coffee daily, more than half consume specialty coffee daily. See www.ncausa.org/i4a/pages/index.cfm?pageID=924.

²⁵We map the coffee market into our benchmark model with $N = 2$ firms as follows. Prior to Starbucks, both firms produce filter coffee, converging to an equilibrium (q_{filter}, p_{filter}) . Then one firm introduces espresso technology, providing quality $q_{espresso} > q_{filter}$. As in the previous section, results do not change if we assume the innovator is a third (incoming) firm. With this mapping, the attention externality between the filter coffee and espresso coffee markets is explicit. Naturally, this simple model misses other real world features of our case study, such as heterogeneous demand and the fact that several firms have non-zero market share.

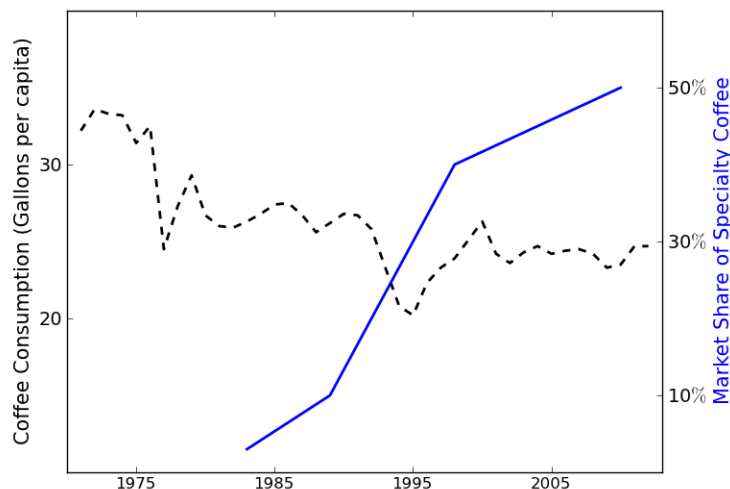


Figure 2: Overall consumption of coffee 1970 – 2010 (dashed line, source: USDA Economic Research Service), and market share of specialty coffee by value 1983 – 2012 (solid line, sources: Kachra (1997), Vishwanath and Harding (2000), and <http://www.scaa.org/PDF/resources/facts-and-figures.pdf>.)

1970s and the 1990s, which also maps into a low F .²⁶ At low unit costs, and consistent with the evidence, our model predicts that firms compete on price. Small quality upgrades by traditional roasters fail, and firms innovate by cutting prices. We are in the commodity magnet of Proposition 4 (case ii).

The “Starbucks effect” then results from Starbucks’ introduction of a different technology that allowed it to offer much higher and salient quality. We view this as a drastic and unilateral reduction of the cost of quality v .²⁷ This innovation decommo- ditized the coffee market (formally, Starbucks became the low cost producer in Proposition 4, case i), causing a reduction in the price sensitivity of *all* consumers, and a drastic increase in quality and price. Decomm- oditization also facilitated a wave of further innovations (Starbucks and other coffee shops now serves several dozen different types of drinks) and induced players

²⁶With respect to the role of supermarkets, technologies for packaging and storing are naturally cheaper in these outlets than in small coffee shops, and nationwide distribution can additionally take advantage of economies of scale. On the drop of the price of coffee, see <http://databank.worldbank.org/data/databases/commodity-price-data>.

²⁷This technology, drawn in part from Schultz’ Italian experience, was the know-how to sell high quality espresso drinks to the mass public in cafes (Pendergrast 1999, Koehn 2005). This required high quality beans, but also trained baristas.

in the traditional coffee market – including roasters selling beans at supermarkets, but also other outlets such as restaurants – to move to high quality, specialty coffee. There was an externality from the coffee shop market to the broader coffee market.²⁸

5.2.2 Conventional Alternatives

Compare the salience account to conventional supply and demand explanations of the same events. The most intuitive account is demand based. If taste for quality increases for many consumers, driven for instance by income growth, the equilibrium quality and price of coffee should rise.

This explanation faces two difficulties. First, it cannot account for the initial commoditization of the coffee market, which witnesses price wars and, if anything, decreasing quality. The US experienced strong economic growth during the 1970s and 1980s, yet this by itself did not lead to decommoditization of the coffee market (specialty coffee niche notwithstanding). Second, it cannot account for the timing of decommoditization, which coincides with Starbucks' innovation and expansion. All the historical analyses we have found emphasize that it was only after Starbucks' introduction of espresso drinks in coffee shops that large numbers of consumers converted to specialty coffee (Vishwanath and Harding 2000, see also Figure 2). Growth of such magnitude can only be explained by a sudden and drastic increase in the taste for coffee, a shock we find implausible.²⁹

If demand alone cannot jointly explain commoditization and decommoditization, a combination of demand and supply shocks may seem promising. For instance, Schultz might have been the first to discover the preference of US consumers for Espresso. Discovery, and not salience, may thus be responsible for decommoditization. This discovery channel, how-

²⁸The externality is even stronger if there is an outside option of not buying. In this case, entry by the high quality firm reduces WTP for the low quality firm, forcing the latter to innovate in order to survive.

²⁹A related possibility, that of an expansion of the coffee market to new consumers, seems unlikely. First, the sheer size of the market suggests there is little room on the extensive margin: according to a 2014 survey by the National Coffee Association, about 61% of adults consume coffee daily (<http://www.ncausa.org/i4a/pages/index.cfm?pageID=924>). These numbers have been broadly stable over time, with the share of daily drinkers not falling below 50% (Koehn 2005). Second, to account for the increase in average taste for quality, the new consumer base would have stronger demand for coffee, and so it is unlikely that it would not have purchased coffee before. We have found no evidence that cohort effects play a significant role, given that a majority of Americans already drink coffee daily and they are at present approximately uniformly distributed across cohorts.

ever, can neither explain why earlier incremental attempts to improve quality failed, nor why decommoditization also extended to filter coffee. In the salience model the role of Espresso is clear: it represented a drastic improvement of quality while keeping costs relatively low, i.e. it rendered quality salient. Consumers' focus shifted to coffee quality across the board.

More complex combinations of shocks may account for the observed patterns of quality and prices in the coffee market.³⁰ We do not attempt here to evaluate these possibilities, but stress that the salience and the rational model yield different predictions that can be tested empirically. For instance, salience accounts for commoditization via a drop in unit costs, while the conventional model requires a hike in marginal costs of quality (yet it may find it hard to reproduce the contemporaneous drop in prices). Additionally, in our model, but not in the conventional one, individual consumers' taste for quality changed upon the introduction of Starbucks, which can be tested by using individual level data.

6 Conclusion

We have shown how salience changes some of the basic predictions of a standard model of competition with vertical product differentiation. Yet the paper has only begun to explore the consequences of salience for market competition. Rather than summarizing our results, in conclusion we mention some issues we have not addressed, but which may be interesting to investigate. These include dynamics of competition, welfare, horizontal product differentiation, and advertising. We have not solved any of these problems, so the discussion here is strictly conjectural.

In a dynamic setting, the salience of a firm's strategy is not only shaped by the background of its competitors, but also by past market outcomes. As we formalized in BGS (2013), the price of a product is salient not only if the product looks expensive relative to substitute goods available today, but also if it looks expensive relative to yesterday's prices. This result

³⁰For example, the early pattern of market commoditization – in which both quality and prices fall – might be explained by considering a composite supply side shock. Real coffee prices declined starting in the late 1970's. In a conventional model, however, lower input costs imply increased (or at least constant) quality, which is inconsistent with early attempts by producers to shade on coffee quality. However, together with falling (unit) cost of coffee (which put downward pressure on prices), it might be that during the same period the extra cost of the higher quality Arabica coffee increased, which induced producers to substitute it with the cheaper robusta coffee.

has interesting implications for the dynamics of entry and imitation. In particular, these dynamics may be very different depending on whether the original innovation ultimately leads to quality-salient or price-salient long run equilibrium. If an innovator finds a way to escape the commodity magnet and produce higher quality at a higher price, the pace at which this change is implemented, and imitated, might be relatively slow. The reason is that firms need to keep quality rather than price salient, and prevent consumers from becoming focused on price increases. This slows down innovation. As an extreme example, if consumers are used to free education, as they are in Europe, charging for education might be extremely difficult even with significant quality improvements because the focus will be entirely on prices. (Of course, once prices are high enough, the pace of innovation and price increases accelerates.) In contrast, precisely because consumers are focused on prices and neglect quality, innovation that reduces price and quality will be extremely fast. The slide to the commodity magnet will be faster than in a rational model.

We have shown that – under the natural assumption that consumer welfare is measured by the undistorted utility – quality provision is generally inefficient in a duopoly, as a consequence of competition for attention between the two firms. An assessment of the welfare consequences of competition when consumers are salient thinkers would require a deeper understanding of the model with heterogeneous consumers, and in particular of monopoly and free entry.

Our approach might also be used to study horizontal differentiation, and to investigate the marketing dictum of “differentiate in any way you can” (Levitt 1983). If a firm horizontally differentiates its product from competitors, then differences along the differentiated attribute become salient, and will attract consumers’ attention. At the same time, differences in prices, which are similar across alternatives, will become non salient. In fact, firms might differentiate their products precisely to segment the market between consumers attracted to different attributes, and thus earn higher profits. This approach has clear applications to product markets, but it might also shed light on political competition, where it can reverse the median voter result in a plausible way. It would suggest that politicians might perhaps converge to the median voter viewpoint on some positions, but also seek to differentiate their views on dimensions that voters might find salient (and attractive). The two parties in

the United States converge on their views on Social Security, for example, making sure that voters do not pay attention to that issue, but then seek to differentiate on the issues they choose, such as immigration or gay marriage.

Finally, salience may have significant implications for how we think about advertising, which deals precisely with drawing consumer attention to products and their attributes. Economists distinguish two broad approaches to advertising: informative and persuasive. The former focuses on provision of hard information about the product; the latter deals with its more emotional appeal. Salience suggests that in fact the two approaches are intimately related, and usually integrated: a key purpose of advertising is to inform about and thus draw attention to the attributes of the product that the seller wants the consumer to think about, but not others. Gas stations sell *regular* and *super* gasoline, even though the difference in octane content is only about 3%. Advertising of attributes is simultaneously informative (sometimes about prices, sometimes about quality, rarely both) and persuasive in that the salience of the attributes being advertised is enhanced. The purpose of advertising is precisely to let some desirable attributes of the product stand out for the potential customers.

In all these situations, firms compete to attract attention to the attributes they want consumers to attend to, and to distract attention from their less attractive attributes.

References

- Anderson, Simon and Andre de Palma. 1992. “Multiproduct Firms: A Nested Logit Approach.” *Journal of Industrial Economics* 40 (3): 261 – 276.
- Andrews, Margaret. 1992. “Avenues for Growth: a 20-Year Review of the U.S. Specialty Coffee Market.” Specialty Coffee Association of America Report.
- Armstrong, Mark and Yongmin Chen. 2009. “Inattentive Consumers and Product Quality,” *Journal of the European Economic Association* 7 (2): 411 – 422.
- Azar, Ofer. 2008. “The Effect of Relative Thinking on Firm Strategy and Market Outcomes: A Location Differentiation Model with Endogenous Transportation Costs.” *Journal of Economic Psychology* 29 (5): 684 – 697.
- Becker, Bo and Victoria Ivashina. 2014. “Reaching for Yield in the Bond Market.” *Journal of Finance*, forthcoming.
- Bordalo, Pedro, Nicola Gennaioli and Andrei Shleifer. 2012. “Salience Theory of Choice under Risk.” *Quarterly Journal of Economics* 127 (3): 1243 – 1285.
- Bordalo, Pedro, Nicola Gennaioli and Andrei Shleifer. 2013. “Salience and Consumer Choice.” *Journal of Political Economy* 121 (5): 803 – 843.
- Bordalo, Pedro, Nicola Gennaioli and Andrei Shleifer. 2013b. “Salience and Asset Prices.” *American Economic Review, Papers & Proceedings* 103 (3): 623–28.
- Bushong, Ben, Matthew Rabin and Joshua Schwartzstein. 2014. “A Model of Relative Thinking.” Unpublished.
- Célérier, Claire and Boris Vallée. 2015. “The Motives for Financial Complexity: An Empirical Investigation.” HEC Paris Research Paper No. FIN-2013-1013.
- Cunningham, Tom. 2012. “Comparisons and Choice.” Unpublished.
- de Clippel, Geoffroy, Kfir Eliaz and Kareen Rozen. 2014. “Competing for Consumer Inattention.” *Journal of Political Economy*, 122 (6): 1203 – 1234.
- Dahremöller, Carsten, and Markus Fels. 2013. “Product Lines, Product Design and Limited Attention.” SSRN Working Paper 1993128.

- Ellison, Glenn. 2006. "Bounded Rationality in Industrial Organization." In *Advances in Economics and Econometrics: Theory and Applications*, Ninth World Congress, edited by Richard Blundell, Whitney Newey and Torsten Persson. Cambridge University Press.
- Ellison, Glenn and Sara Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77 (2): 427 – 452.
- Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 129 (4): 1661 – 1710.
- Gabaix, Xavier and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics* 121 (2): 505 – 540.
- Gennaioli, Nicola, Andrei Shleifer and Robert Vishny. 2012. "Neglected Risks, Financial Innovation and Financial Fragility." *Journal of Financial Economics* 104 (3): 452 – 468.
- Gennaioli, Nicola, Andrei Shleifer and Robert Vishny. 2013. "A Model of Shadow Banking." *Journal of Finance* 68 (4): 1331 – 1363.
- Greenwood, Robin and Samuel Hanson. 2013. "Issuer Quality and Corporate Bond Returns." *Review of Financial Studies* 26 (6):1483 – 1525.
- Hefti, Andreas. 2012. "Attention, Competition and Obfuscation." University of Zurich Department of Economics Working Paper 28.
- Huber, Joel, John Payne and Christopher Puto. 1982. "Adding Asymmetrically Dominated Alternatives: Violations of the Regularity and Similarity Hypothesis." *Journal of Consumer Research* 9 (1): 90 – 98.
- Kachra, Ariff. 1997. "Starbucks." Richard Ivey School of Business Case Study.
- Koehn, Nancy. 2005. "Howard Schultz and Starbucks Coffee Company." Harvard Business School Case Study.
- Koszegi, Botond and Adam Szeidl. 2013. "A Model of Focusing in Economic Choice." *Quarterly Journal of Economics* 128 (1): 53 – 104.
- Levitt, Theodore. 1983. *The Marketing Imagination*. New York : The Free Press.

- Matějka, Filip and Alisdair McKay. 2012. “Simple Market Equilibria with Rationally Inattentive Consumers.” *American Economic Review Papers & Proceedings* 102 (3): 24 – 29.
- Mauborgne, Renée and W. Chan Kim. 2005. *Blue Ocean Strategy: How to Create Uncontested Market Space and Make Competition Irrelevant*. Cambridge, MA: Harvard Business School Press.
- Pendergrast, Mark. 1999. *Uncommon Grounds: the History of Coffee and How It Transformed our World*. New York, NY: Basic Books.
- Persson, Petra. 2012. “Attention Manipulation and Information Overload.” SSRN Working Paper 2102155.
- Rangan, V. Kasturi and George Bowman. 1992. “Beating the Commodity Magnet.” *Industrial Marketing Management* 21 (3): 215 – 224.
- Reny, Philip. 1999. “On the Existence of Pure and Mixed Strategy Equilibria in Discontinuous Games.” *Econometrica* 67 (5): 1029 – 1056.
- Salant, Yuval and Ariel Rubinstein. 2008. “(A,f): Choice with Frames.” *Review of Economic Studies* 75(4): 1287 – 1296.
- Schultz, Howard and Dori Yang. 1999. *Pour Your Heart into It*. New York, NY: Hyperion Books.
- Shaked, Avner and John Sutton. 1982. “Relaxing Price Competition Through Product Differentiation.” *Review of Economic Studies* 49 (1): 3 – 13.
- Simonson, Itamar. 1989. “Choice Based on Reasons: the Case of the Attraction and Compromise Effects.” *Journal of Consumer Research* 16 (2): 158 – 174.
- Slywotzky, Adrian. 1995. *Value Migration*. Cambridge, MA: Harvard University Press.
- Spiegler, Ran. 2012. “Competitive Framing.” *American Economic Journal: Microeconomics* 6 (3): 35 – 58.
- Spiegler, Ran and Kfir Eliaz, 2011a. “On the Strategic Use of Attention Grabbers.” *Theoretical Economics* 6 (1): 127 – 155.

- Spiegler, Ran and Kfir Eliaz, 2011b. “Consideration Sets and Competitive Marketing.” *Review of Economic Studies* 78 (1): 235 – 262.
- Spiegler, Ran and Michele Piccione. 2012. “Price Competition under Limited Comparability.” *Quarterly Journal of Economics* 127 (1): 97 – 135.
- Spiegler, Ran. 2011. *Bounded Rationality and Industrial Organization*. Oxford University Press.
- Tversky, Amos and Daniel Kahneman. 1981. “The Framing of Decisions and the Psychology of Choice.” *Science* 211 (4481): 453 – 458.
- Vishwanath, Vijay, and David Harding. 2000. “The Starbucks Effect.” *Harvard Business Review*.

A Proofs

Lemma 1 (price competition under rationality). When $\delta = 1$ there are no salience distortions and utility is given by Equation (2).

We analyse each case in turn. If $q_1 - c_1 > q_2 - c_2$ then firm 1 sets its price $p_1 = c_2 + (q_1 - q_2)$ and firm 2 sets its price $p_2 = c_2$. The sharing rule determines that at these prices firm 1 captures all demand (because only firm 1 can profitably reduce its price). Firm 1 has no incentive to increase its price, because consumers would then prefer firm 2's good, nor to decrease its price, which would reduce profits. Firm 2 has no incentive to increase its price, since it cannot capture demand by doing so. Firm 2 can also not reduce its price below cost c_2 , as that would entail negative profits. This demonstrates existence. To show uniqueness, assume by contradiction that firm 2 sets its price $p_2 > c_2$. Firm 1 makes positive profits for any price p_1 in the interval $(c_2 + (q_1 - q_2), p_2 + (q_1 - q_2)]$ (though it shares the market with firm 2 when $p_1 = p_2 + (q_1 - q_2)$, since it is then profitable for both firms to unilaterally reduce price). However, no price p_1 in this interval can be an equilibrium: firm 2 would have an incentive to lower its price below p_1 as doing so would allow firm 2 to capture demand and increase its profits. Assume now that firm 2 sets its price $p_2 \in [c_1 + q_2 - q_1, c_2)$. The best response of firm 1 is to set $p_1 = p_2 + q_1 - q_2$. While all strategies within this interval yield zero profit to firm 2 and non-negative profit to firm 1 (hence the multiplicity of equilibria), all strategies are weakly dominated. Since we exclude equilibria in weakly dominated strategies, this constrains firms to price weakly above costs. As a consequence, in the unique equilibrium $p_1 = c_2 + (q_1 - q_2)$ and $p_2 = c_2$.

If $q_1 - c_1 < q_2 - c_2$, the existence and uniqueness arguments carry through switching firms 1 and 2. Finally, if $q_1 - c_1 = q_2 - c_2$, then both firms price at cost and share the market. Neither firm has an incentive to deviate: increasing price would ensure zero demand and thus would not increase profits; decreasing price would lead to negative profits. Uniqueness follows as before: if one firm sets its price above cost, the other firm has an incentive to also set its price above its cost and capture the market. This cannot be an equilibrium because – at this configuration – the first firm has an incentive to slightly reduce its price. Thus, no equilibrium exists in which either firm prices above cost. ■

Proposition 1 (price competition under salient thinking). When $\delta < 1$, utility is given by Equation (3), where salience determines the relative weight of quality and price. We first characterise the equilibria in the full parameter space of exogenous qualities and costs satisfying $q_1 \geq q_2$ and $c_1 \geq c_2$, showing existence and uniqueness. This is illustrated in Figure 3. We then restrict the analysis to parameters satisfying Assumption A.1.

Let $q_1 \geq q_2$ and $c_1 \geq c_2$ be given. Note that it cannot be the case that in equilibrium $p_1 < p_2$. In fact, by increasing its price p_1 to p_2 the high quality firm 1 still captures all demand and increases its profits. To see this, it is enough to show that when $p_1 \leq p_2$, consumers perceive good 1 as dominating good 2, and therefore choose good 1. That this holds is guaranteed by the symmetry of the salience function, which ensures that both goods always have the same salience ranking (see Section 2). When $p_1 < p_2 \frac{q_2}{q_1}$, price is salient for both goods, while for $p_1 \in (p_2 \frac{q_2}{q_1}, p_2]$ quality is salient for both goods.

When $p_1 \geq p_2$ the difference between the (salience-distorted) valuations of good k and k' always strictly decreases in the price p_k (keeping price p_{-k} fixed). To see this, we first show that $u^S(q_1, p_1 + \Delta p_1) - u^S(q_2, p_2) < u^S(q_1, p_1) - u^S(q_2, p_2)$ for any p_1, p_2 and $\Delta p_1 > 0$. Note that the l.h.s. refers to valuation in the choice set $\{(q_1, p_1 + \Delta p_1), (q_2, p_2)\}$, while in the r.h.s. the choice set is $\{(q_1, p_1), (q_2, p_2)\}$; in particular, the salience ranking of good 2 may be not be the same in both contexts. Suppose first that the salience ranking of quality and price does not change upon the price shift $p_1 \rightarrow p_1 + \Delta p_1$. Then $(u^S(q_1, p_1 + \Delta p_1) - u^S(q_2, p_2)) - (u^S(q_1, p_1) - u^S(q_2, p_2)) \propto -\Delta p_1$, which is negative. Suppose now that the salience ranking does change upon this price shift. This means quality is salient at prices p_1, p_2 while price becomes salient when $p_1 \rightarrow p_1 + \Delta p_1$. Thus $(u^S(q_1, p_1 + \Delta p_1) - u^S(q_2, p_2)) - (u^S(q_1, p_1) - u^S(q_2, p_2)) = -\Delta p_1 - (1 - \delta)[(q_1 - q_2) + (p_1 - p_2)] < 0$. For shifts in p_2 , the reasoning is similar: if the salience ranking does not change upon the price shift $p_2 \rightarrow p_2 + \Delta p_2$, we have $(u^S(q_1, p_1) - u^S(q_2, p_2 + \Delta p_2)) - (u^S(q_1, p_1) - u^S(q_2, p_2)) \propto \Delta p_2$, which is positive. If the salience ranking does change, this means price is salient at prices p_1, p_2 while quality becomes salient when $p_2 \rightarrow p_2 + \Delta p_2$. Thus $(u^S(q_1, p_1) - u^S(q_2, p_2 + \Delta p_2)) - (u^S(q_1, p_1) - u^S(q_2, p_2)) = \delta \Delta p_2 + (1 - \delta)[(q_1 - q_2) + (p_1 - p_2)] > 0$.

We now separately consider three cases: one in which firm 1 wins, another in which firms 2 wins, and a final one in which the two firms split the market. As a preliminary

observation, note that – like in the rational case – the losing firm prices at cost. The winning firm maximizes its profit by setting a price that renders either the valuation or the salience constraint binding, given the losing firm’s price (see Equations (5) through (8) in the text). But this can only be an equilibrium if the losing firm prices at cost, because – by the assumed sharing rules – only in this case is the losing firm unable to win the market and make a profit by reducing its price.

- Suppose that in equilibrium firm 1 wins the market, so that $p_1 \geq c_1$ and $p_2 = c_2$ (the latter follows from our restriction to equilibria in non-weakly dominated strategies). At equilibrium prices, either quality or price can be salient. These two types of equilibria arise for different parameter ranges, as described in Equations (18) and (19):

- Firm 1 wins with salient quality. This equilibrium is determined by the following conditions on p_1 :

$$\left\{ \begin{array}{ll} \frac{q_1}{p_1} \geq \frac{q_2}{c_2} & \text{(salience constraint)} \\ p_1 \leq c_2 + \frac{1}{\delta}(q_1 - q_2) & \text{(valuation constraint)} \\ p_1 \geq c_1 & \end{array} \right. \quad (18)$$

The salience constraint ensures that quality is salient at equilibrium prices p_1, c_2 , while the valuation constraint ensures that at these prices good 1 is chosen over good 2.³¹ Both constraints are weak because the sharing rule guarantees that at the price p_1 at which either the salience or valuation constraint binds, firm 1 captures the entire market (formally: at equal salience good 1’s advantage (quality) is overweighted in consumers’ utility, and at equal valuation (under salient quality) consumers choose good 1). This is because, when $p_2 = c_2$, only firm 1 can reduce its price and make a profit.

The above constraints on p_1 are satisfied only within a given range of quality and cost parameters. This range is characterized as follows: i) $\frac{q_1}{c_1} > \delta$ and $\frac{q_2}{c_2} \in \left[\delta, \frac{q_1}{c_1} \right]$,

³¹Recall from footnote 7 that the salience constraint is invariant to specifications of the reference attribute levels that are strict convex combinations of attributes in the choice set. This implies that the price competition stage, described in this Proposition, is invariant to this specification and, as a consequence, so is the full equilibrium analysis of Proposition 2.

so that in equilibrium the salience constraint binds and $p_1 = c_2 \frac{q_1}{q_2}$; ii) $\frac{q_2}{c_2} < \delta$ and $\frac{q_1}{c_1} \geq \delta - \frac{c_1}{c_2} \left(\delta - \frac{q_2}{c_2} \right)$, so that in equilibrium the valuation constraint binds and $p_1 = c_2 + \frac{1}{\delta}(q_1 - q_2)$.

- Firm 1 wins with salient price. This equilibrium is determined by the following conditions on p_1 :

$$\begin{cases} \frac{q_1}{p_1} \leq \frac{q_2}{c_2} & \text{(salience constraint)} \\ p_1 \leq c_2 + \delta(q_1 - q_2) & \text{(valuation constraint)} \\ p_1 \geq c_1 \end{cases} \quad (19)$$

The salience constraint ensures that price is salient at equilibrium prices p_1, c_2 , while the valuation constraint ensures that at these prices good 1 is still chosen over good 2. In this case, only the valuation constraint binds price from above. This constraint is weak because, at the price p_1 at which it binds, only firm 1 can reduce its price and still make a profit, so the sharing rule guarantees that firm 1 captures the market.

The above constraints on p_1 are satisfied only within a given range of quality and cost parameters. This range is characterized as follows: i) the salience constraint binds p_1 from below, $\frac{q_1}{c_1} > \frac{q_2}{c_2}$ and $\frac{q_2}{c_2} > \frac{1}{\delta}$; ii) the cost constraint binds p_1 from below, $\frac{q_1}{c_1} < \frac{q_2}{c_2}$ and $\frac{q_1}{c_1} > \frac{1}{\delta} + \frac{c_2}{c_1} \left(\frac{q_2}{c_2} - \frac{1}{\delta} \right)$, which implies $\frac{q_2}{c_2} > \frac{1}{\delta}$. In both cases, the equilibrium price is $p_1 \leq c_2 + \delta(q_1 - q_2)$.

In particular, when $\frac{q_1}{c_1} > \max \left\{ \frac{1}{\delta}, \frac{q_2}{c_2} \right\}$ firm 1 can win either with salient quality or with salient price. In either case, firm 2 sets its price equal to cost. Here firm 1 has an incentive to choose the price salience configuration as it allows it to set a higher price, and thus obtain a higher profit.

- Suppose that in equilibrium firm 2 wins the market. Then it must be that firm 2 sets its price $p_2 \geq c_2$ while firm 1 prices at cost, $p_1 = c_1$. At equilibrium prices, either quality or price can be salient. These two types of equilibria arise for different parameter ranges, as described below (the analysis is very similar to that of the case where firm 1 wins).

- Firm 2 wins with salient quality, so that $\frac{q_2}{p_2} \leq \frac{q_1}{c_1}$. Salient quality implies that price satisfies $p_2 \leq c_1 + \frac{1}{\delta}(q_1 - q_2)$, $p_2 \geq c_1 \frac{q_2}{q_1}$ and $p_2 \geq c_2$. In this case, only the valuation constraint binds firm 2's price from above. This constraint is weak because, at the price p_2 at which it binds, only firm 2 can reduce its price and still make a profit, so the sharing rule guarantees that firm 2 captures the market.

The set of p_2 satisfying these conditions, if non empty, is bounded above by the valuation constraint. The set is non empty when: i) $\frac{q_1}{c_1} \leq \frac{q_2}{c_2}$ and $\frac{q_1}{c_1} < \delta$, where the salience constraint provides a lower bound for price; or ii) $\frac{q_2}{c_2} < \delta$ and $\frac{q_2}{c_2} \geq \delta + \frac{c_1}{c_2} \left(\frac{q_1}{c_1} - \delta \right)$, where the cost constraint provides a lower bound, implying $\frac{q_1}{c_1} < \delta$.

- Firm 2 wins with salient price, so that $\frac{q_2}{p_2} \geq \frac{q_1}{c_1}$. In this case, price must satisfy $p_2 \leq c_1 \frac{q_2}{q_1}$, $p_2 \leq c_1 - \delta(q_1 - q_2)$, as well as $p_2 \geq c_2$. Both the salience and the valuation constraints are weak because the sharing rule guarantees that at the price p_2 at which either constraint binds, firm 2 captures the entire market (formally: at equal salience good 2's advantage (price) is overweighted in consumers' utility, and at equal valuation (under salient price) consumers choose good 2). This is because, when $p_1 = c_1$, only firm 2 can reduce its price and make a profit.

The set of p_2 satisfying these conditions is non empty when: i) $\frac{q_1}{c_1} > \frac{1}{\delta}$ and $\frac{q_2}{c_2} > \frac{1}{\delta} + \frac{c_1}{c_2} \left(\frac{q_1}{c_1} - \frac{1}{\delta} \right)$. The first condition guarantees that the valuation constraint is binding on p_2 , while the second condition guarantees there exists a $p_2 \geq c_2$ satisfying the valuation constraint. ii) $\frac{q_1}{c_1} < \frac{1}{\delta}$ and $\frac{q_2}{c_2} > \frac{q_1}{c_1}$. In this case, the first condition guarantees that the salience constraint is binding on p_2 , while the second condition guarantees there exists a $p_2 \geq c_2$ satisfying the salience constraint.

- Comparing the two cases above, we find that when $\frac{q_1}{c_1} < \frac{1}{\delta}$ and $\frac{q_2}{c_2} \in \left[\frac{q_1}{c_1}, \frac{1}{\delta} \right]$, firm 2 can win either with salient quality or with salient price, while firm 1 always prices at cost. In this case, in equilibrium firm 2 sets its price such that quality is salient, since it can then obtain a higher profit by doing so.

The analysis above shows equilibria exist for any parameters satisfying $q_1 \geq q_2$ and $c_1 \geq c_2$. Furthermore, the equilibria are unique, since for every choice of quality and cost

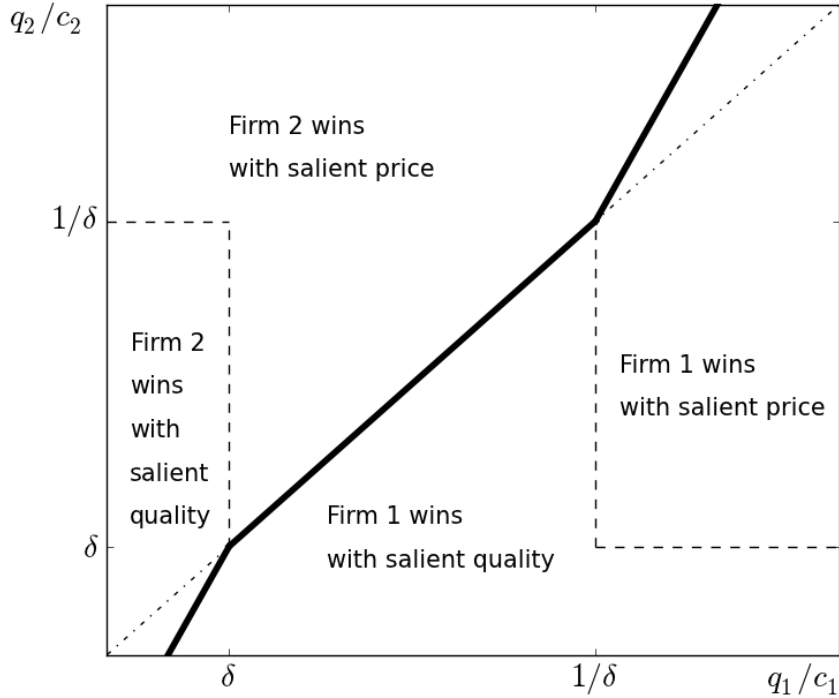


Figure 3: Equilibria of the price competition game when $q_1 > q_2, c_1 > c_2$.

parameters, equilibrium prices are uniquely defined. While in some regimes the firm that wins the market makes its advantage salient (e.g. when $\frac{q_1}{c_1}, \frac{q_2}{c_2} \in [\delta, \frac{1}{\delta}]$), in other regimes – namely when one firm’s quality cost ratio is extreme – a firm might win the market despite having increased its price to the point that its disadvantage (high price or low quality) is salient.

We now restrict the results to the case where Assumption A.1 holds, namely $\delta(c_1 - c_2) < q_1 - q_2 < \frac{1}{\delta}(c_1 - c_2)$. In equilibrium, the firm that wins the price competition sets its price so that its relative advantage is salient, and it captures the market. Thus, if $q_1/c_1 > q_2/c_2$, firm 1 wins the market in equilibrium. Because the salience constraint binds, this corresponds to the region in Figure 3 where firm 1 wins the market with salient quality: firm 1 sets $p_1 = \min \{c_2 \cdot q_1/q_2, c_2 + \frac{1}{\delta}(q_1 - q_2)\}$ and firm 2 sets $p_2 = c_2$. A similar argument shows that, if $q_2/c_2 > q_1/c_1$, firm 2 wins the market, and that equilibrium prices satisfy $p_2 = \min \{c_1 \cdot q_2/q_1, c_1 + \delta(q_2 - q_1)\}$ and $p_1 = c_1$. This corresponds to the region in Figure

3 where firm 2 wins the market with salient price. Finally, if the two firms have the same quality to cost ratio, namely $q_1/c_1 = q_2/c_2$, no firm can raise its price above cost without having its disadvantage salient. Given that, by A.1, consumers do not buy a good whose disadvantage is salient, the only equilibrium is for the two firms to price at cost, setting $p_1 = c_1, q_2 = c_2$. Firms make zero profits, both attributes are equally salient and consumers select the good yielding higher (rational) surplus. This corresponds to the diagonal segment in Figure 3.

■

Lemma 2 (quality competition under rationality with symmetric costs). When $\delta = 1$, the full game admits a set of subgame perfect equilibria in pure strategies. We first characterise this set of equilibria, and show consumers buy the same quality level $q^* = \operatorname{argmax}_q (q - c(q))$ in every equilibrium in this set. Finally, we show the set contains a unique unique symmetric equilibrium, as described in the main text.

To find the subgame perfect equilibria of the game, we identify the best response of firm $j = 1, 2$, which consists of the set $q_j^b(q_{-j})$ of the optimal qualities by firm j as a function of the quality q_{-j} set by firm $-j$, given equilibrium play in the price subgame. Recall that at the quality choice stage, firm j 's optimisation problem is

$$\begin{aligned} \max_{q_j} \quad & ([q_j - c_j(q_j)] - [q_{-j} - c_{-j}(q_{-j})]) \cdot d_j(q_j, q_{-j}) \\ \text{where } d_j(q_j, q_{-j}) = & \begin{cases} 1 & \text{if } q_j - c_j(q_j) > q_{-j} - c_{-j}(q_{-j}) \\ 1/2 & \text{if } q_j - c_j(q_j) = q_{-j} - c_{-j}(q_{-j}) \\ 0 & \text{if } q_j - c_j(q_j) < q_{-j} - c_{-j}(q_{-j}) \end{cases} \end{aligned} \quad (20)$$

Here, we restrict to the case where firms have the same cost function, $c_j(q) = c_{-j}(q) = c(q)$.

We start with two preliminaries. First, we define q^* to be the surplus maximizing quality, i.e., $q^* = \operatorname{argmax}_q (q - c(q))$. Because $c(q)$ is convex and firms have the same technology, q^* is unique and common to j and $-j$. Second, by Lemma 1, firm j wins the market and obtains a strictly positive profit if and only if it yields strictly higher surplus than $-j$, namely $q_j - c(q_j) > q_{-j} - c(q_{-j})$. In this case, firm j 's equilibrium price in the price subgame is equal to $p_j = c(q_{-j}) + q_j - q_{-j}$ and its profit is $\pi_j(q_j, q_{-j}) = q_j - c(q_j) + c(q_{-j}) - q_{-j}$.

As a first step, consider firm j 's best response to quality $q_{-j} \neq q^*$. To win the market, firm j finds it optimal to choose a best response $q_j^b(q_{-j})$ that generates higher surplus than that generated by firm $-j$. Within the set of quality levels yielding higher surplus, the best response of firm j is the quality level that maximizes its profit. The expression for $\pi_j(q_j, q_{-j})$ implies that the best response to any $q_{-j} \neq q^*$ is q^* . This quality level maximizes not only surplus, but also profits.

Consider now firm j 's best response to the surplus maximizing quality $q_{-j} = q^*$. By definition of q^* , there is no feasible quality $q_j \in [0, +\infty)$ at which firm j delivers strictly higher surplus than firm $-j$. It then follows from Lemma 1 that when $q_{-j} = q^*$, firm j 's equilibrium price in the price subgame is $p_j = c(q_j)$, so firm j makes zero profits. This is true regardless of the quality $q_j \in [0, +\infty)$ chosen by the firm. Hence, when $q_{-j} = q^*$ any feasible quality is a best response for j , that is, $q_j^b(q^*) = [0, +\infty)$.

To sum up:

$$q_j^b(q_{-j}) = \begin{cases} q^* & \text{if } q_{-j} \neq q^* \\ [0, +\infty) & \text{if } q_{-j} = q^* \end{cases}. \quad (21)$$

Pure strategy equilibria are then identified at intersections of the best response correspondences of the two firms. Given that firms 1 and 2 have the same technology, they also have the same best response correspondence in (21). This implies that at least one firm must choose the surplus maximizing quality q^* . In turn, the other firm can choose any feasible quality in $[0, +\infty)$. The entire set of subgame perfect equilibria in pure strategies in the rational case is then given by:

$$\{(q_j, p_j), (q_{-j}, p_{-j}) : q_j = q^*, q_{-j} \geq 0, p_j = c(q_{-j}) + q^* - q_{-j}, p_{-j} = c(q_{-j})\}_{j=1,2}.$$

In these equilibria, firm $j = 1, 2$ sets $q_j = q^*$ while firm $-j$ sets $q_{-j} \in [0, +\infty)$. If $q_{-j} \neq q^*$ firm $-j$ makes zero profits, while firm j captures all consumer demand and makes strictly positive profits. There is a unique symmetric equilibrium in which both firms set quality q^* and make zero profits (and the consumer is indifferent between the products of firms 1, 2.)

■

Proposition 2 (symmetric equilibrium under salient thinking). When $\delta < 1$, the

full game admits a set of subgame perfect equilibria in pure strategies. We first characterise this set of equilibria as a function of the fixed cost F , and show that consumers buy the same quality level $q^S \equiv q^S(F)$ in every equilibrium in this set. We then show the set contains a unique symmetric equilibrium, as described in the main text.

As a first step, define the thresholds \underline{q} , \bar{q} and \hat{q} as follows: \underline{q} is the quality that maximizes (total) surplus when price is salient, namely $\underline{q} = \operatorname{argmax}_q (\delta q - c(q))$. Thus \underline{q} is determined by the first order condition $c'(\underline{q}) = \delta$ (recall that the cost function is strictly convex). Quality \bar{q} is the surplus maximizing quality when quality is salient, $\bar{q} = \operatorname{argmax}_q q - \delta c(q)$, which satisfies $c'(\bar{q}) = 1/\delta$. Note that, by convexity of the cost function, $\underline{q} < \bar{q}$. Finally, note that the quality to cost ratio $q/c(q)$ is an inverse-U shaped function with a unique local maximum (this follows from our assumptions on the cost function, namely $c(0) > 0$, $c'(q) > 0$, $c''(q) > 0$). We then denote by \hat{q} the quality level that maximizes the quality to cost ratio (minimizes average cost), namely $\hat{q} = \operatorname{argmax}_q q/c(q)$. This quality level satisfies $c'(\hat{q}) = c(\hat{q})/\hat{q}$.

Recall that a full characterisation of the price subgame equilibria was given in the proof of Proposition 1. As shown there, the firm that wins the market and makes positive profits is the one delivering highest perceived (total) surplus in the price subgame. That is, when quality is salient in the price subgame equilibrium, firm j wins the market and makes a positive profit if and only if $q_j - \delta c(q_j) > q_{-j} - \delta c(q_{-j})$. When price is salient in the price subgame equilibrium, firm j wins the market and makes a positive profit if and only if $\delta q_j - c(q_j) > \delta q_{-j} - c(q_{-j})$. When price and quality are equally salient, firm j wins the market and makes a positive profit if and only if $q_j - c(q_j) > q_{-j} - c(q_{-j})$. In the equilibrium of the price subgame, the losing firm prices at cost while the winning firm extracts all perceived consumer surplus.

The following three cases must then be considered:

1. $\underline{q} < \bar{q} \leq \hat{q}$. It follows from the first order conditions that determine \bar{q} and \hat{q} , and convexity of costs, that this case occurs when $\hat{q}/c(\hat{q}) \leq \delta$.

To characterize the equilibria of the game, the key property when $\underline{q} < \bar{q} \leq \hat{q}$ is that in the equilibrium of any price subgame with quality configuration $q_j = \bar{q}$ and $q_{-j} \neq \bar{q}$, firm j

wins the market and makes positive profits.

To prove this claim, suppose first that $q_{-j} < \bar{q}$. Then, firm j delivers higher quality than firm $-j$ and yields a higher quality to cost ratio, namely $\bar{q}/c(\bar{q}) > q_{-j}/c(q_{-j})$. This property follows from the fact that the quality to cost ratio increases in quality for qualities below \hat{q} . From the proof of Proposition 1, then, in the equilibrium of the price subgame quality is salient, firm j provides strictly higher perceived surplus than firm $-j$ and wins the market.

Suppose next that $q_{-j} \in (\bar{q}, \tilde{q}(\bar{q})]$, where $\tilde{q}(\bar{q})$ denotes the unique quality level matching the quality to cost ratio of \bar{q} , namely $\tilde{q}(\bar{q})/c(\tilde{q}(\bar{q})) = \bar{q}/c(\bar{q})$ and $\tilde{q}(\bar{q}) \neq \bar{q}$. When q_{-j} is in the interior of this range, firm $-j$ provides higher quality than firm j and it also has a weakly higher quality to cost ratio. At the boundary $q_{-j} = \tilde{q}(\bar{q})$, price and quality are equally salient if both firms price at cost. By raising price above cost, firm j renders quality (its disadvantage relative to $\tilde{q}(\bar{q})$) salient. In either case, quality is salient in the price subgame equilibrium. Since firm j provides the surplus maximizing quality level \bar{q} when quality is salient, it wins the market and makes positive profits.

Suppose, finally, that $q_{-j} > \tilde{q}(\bar{q})$. Now firm $-j$ provides higher quality than firm j but it has a strictly lower quality to cost ratio than j . As a result, price is salient if firms price at cost. Note that even with salient price, firm j provides higher perceived surplus than firm $-j$ because $q_j = \bar{q}$ is closer than q_{-j} to the surplus maximizing quality level \underline{q} under salient price. However, the equilibrium of the price subgame still features salient quality. The reason is that firm j can increase price up to $p_j = c(q_{-j}) + \frac{1}{\delta}(\bar{q} - q_{-j})$ and render quality salient. Even though quality is salient and firm $-j$ provides higher quality, firm j wins the market because it provides higher surplus.

We have thus proved that when $\hat{q}/c(\hat{q}) \leq \delta$ the strategy $q_j = \bar{q}$ by firm j beats any other feasible quality $q_{-j} \neq \bar{q}$ set by firm $-j$. This implies that, in this cost range, at least one firm $j = 1, 2$ must play $q_j = \bar{q}$ in equilibrium. In fact, if both firms play $q_j, q_{-j} \neq \bar{q}$, then at least one firm would find it profitable to deviate to \bar{q} . We can characterize equilibria as follows. First, when one firm plays $q_j = \bar{q}$, its opponent $-j$ is sure to make zero profits and is thus willing to play any feasible quality. At the same time, the equilibrium values of q_{-j} are those to which $q_j = \bar{q}$ constitutes a best response. Clearly, $q_j = \bar{q}$ constitutes a best response to all quality levels such that quality is salient in the price subgame: from the above analysis,

this set clearly includes the interval $[0, \tilde{q}(\bar{q})]$ but it also includes those $q_{-j} > \tilde{q}(\bar{q})$ that satisfy $q_{-j}/c(q_{-j}) \geq \bar{q}/p_j(\bar{q}, q_{-j})$ where $p_j(\bar{q}, q_{-j}) = c(q_{-j}) + \frac{1}{\delta}(\bar{q} - q_{-j})$ (the sharing rule ensures salient quality in the case the quality cost ratios are equal). But this condition is satisfied for all $q_{-j} > \bar{q}$ because $q/c(q) \leq \delta$ for all q . As a consequence, $q_j = \bar{q}$ is a best response to any quality level in the set $[0, +\infty)$.

This implies that the entire set of subgame perfect equilibria in pure strategies when $\underline{q} < \bar{q} \leq \hat{q}$ is given by:

$$\left\{ (q_j, p_j), (q_{-j}, p_{-j}) : q_j = \bar{q}, q_{-j} \geq 0, p_j = c(q_{-j}) + \frac{1}{\delta}(\bar{q} - q_{-j}), p_{-j} = c(q_{-j}) \right\}_{j=1,2}.$$

In these equilibria, firm $j = 1, 2$ sets $q_j = \bar{q}$ while firm $-j$ sets $q_{-j} \in [0, +\infty)$. If $q_{-j} \neq \bar{q}$ firm $-j$ makes zero profits, while firm j makes strictly positive profits. There is a unique symmetric equilibrium in which both firms set quality \bar{q} and make zero profits. In either case, the quality level bought by consumers is $q^S(F) = \bar{q}$.

2. $\underline{q} < \hat{q} < \bar{q}$. It is easy to see that this case occurs when $\hat{q}/c(\hat{q}) \in (\delta, 1/\delta)$.

To characterize the equilibria of the game, the key property is that when firm $-j$ sets quality $q_{-j} \neq \hat{q}$, there exists q_j with $q_j/c(q_j) > q_{-j}/c(q_{-j})$ such that firm j wins the market and makes positive profits in the equilibrium of the ensuing price subgame.

To prove this claim, suppose first that firm $-j$ plays $q_{-j} > \hat{q}$. Define $\underline{q}(q_{-j})$ as the lowest quality yielding the same quality to cost ratio as q_{-j} . Formally, $\underline{q}(q_{-j})/c(\underline{q}(q_{-j})) = q_{-j}/c(q_{-j})$ and $\underline{q}(q_{-j}) < \hat{q}$. Then, any quality $q_j \in (\max(\underline{q}, \underline{q}(q_{-j})), \hat{q}]$ beats q_{-j} . By definition, any such q_j has lower quality and a higher quality to cost ratio than q_{-j} . Since $q_j > \underline{q}$, it follows from Proposition 1 that price is salient in the equilibrium of the price subgame. As a consequence, q_j wins the market because it generates larger surplus when price is salient (it is closer than q_{-j} to the quality \underline{q} that maximizes surplus under salient price).

Suppose next that firm $-j$ plays $q_{-j} < \hat{q}$. Then, define $\bar{q}(q_{-j})$ as the highest quality yielding the same quality to cost ratio as q_{-j} . Formally, $\bar{q}(q_{-j})/c(\bar{q}(q_{-j})) = q_{-j}/c(q_{-j})$ and $\bar{q}(q_{-j}) > \hat{q}$. Then, any quality $q_j \in [\hat{q}, \min(\bar{q}, \bar{q}(q_{-j}))]$ beats q_{-j} . By definition, any such q_j has higher quality and a higher quality to cost ratio than q_{-j} . Since $q_j < \bar{q}$, it follows from Proposition 1 that quality is salient in the equilibrium of the price subgame. As a

consequence q_j wins the market because it generates larger surplus when quality is salient (it is closer than q_{-j} to the quality \bar{q} that maximizes surplus under salient price).

Given that, whenever one firm is away from \hat{q} its opponent has the incentive to slightly increase its quality to cost ratio close to the maximum possible level, it follows that the unique equilibrium in this case is for both firms to set the quality level maximizing the quality to cost ratio, namely $q_1 = q_2 = \hat{q}$, and set prices equal to $p_1 = p_2 = c(\hat{q})$. Formally, it is sufficient to consider q_j, q_{-j} within the neighborhood (q, \bar{q}) of \hat{q} . For q_{-j} in this range, the best response by firm j is to set $q_j^b(q_{-j})$ such that $c'(q_j^b(q_{-j})) = c(q_{-j})/q_{-j}$. Firm j 's equilibrium price in the price subgame is then equal to $p_j = c(q_{-j}) \cdot q_j/q_{-j}$.³² Since firm $-j$ has the same cost function, it has the same best response correspondence. The equilibrium of the game is then described by the (unique) fixed point of this correspondence, which is equal to \hat{q} as it satisfies $c'(\hat{q}) = c(\hat{q})/\hat{q}$. In this case, the quality level bought by consumers is $q^S(F) = \hat{q}$.

3. $\hat{q} \leq \underline{q} < \bar{q}$. It is easy to see that this case occurs when $\hat{q}/c(\hat{q}) > 1/\delta$.

To characterize the equilibria of the game, the key property when $\hat{q} \leq \underline{q} < \bar{q}$ is that in the equilibrium of any price subgame with quality configuration $q_j = \underline{q}$ and $q_{-j} \neq \underline{q}$, firm j wins the market and makes positive profits.

To prove this claim, suppose first that $q_{-j} > \underline{q}$. Then, firm j delivers lower quality than firm $-j$ and yields a higher quality to cost ratio, namely $\underline{q}/c(\underline{q}) > q_{-j}/c(q_{-j})$. It follows from the proof of Proposition 1 that in the equilibrium of the price subgame price is salient, firm j provides strictly higher surplus than firm $-j$ and wins the market.

Suppose next that $q_{-j} \in [\tilde{q}(\underline{q}), \underline{q}]$, where $\tilde{q}(\underline{q})$ denotes the unique quality level matching the quality to cost ratio of \underline{q} , namely $\tilde{q}(\underline{q})/c(\tilde{q}(\underline{q})) = \underline{q}/c(\underline{q})$ and $\tilde{q}(\underline{q}) \neq \underline{q}$. When q_{-j} is in the interior of this range, firm $-j$ provides lower quality than firm j and it also has a weakly higher quality to cost ratio. At the boundary $q_{-j} = \tilde{q}(\underline{q})$, price and quality are equally salient if both firms price at cost. By raising price above cost, firm j renders price (its disadvantage relative to $\tilde{q}(\underline{q})$) salient. In either case, price is salient in the price subgame equilibrium.

³²At this price level, quality and price are equally salient. The sharing rule then specifies that firm j captures the entire demand at this price, as it is the only firm that can reduce its price and still make a profit.

Since firm j provides the surplus maximizing quality level \underline{q} when price is salient, it wins the market and makes positive profits.

Suppose finally that $q_{-j} < \tilde{q}(q)$. Now firm $-j$ provides lower quality than firm j but it has a strictly lower quality to cost ratio than j . As a consequence, quality is salient if firms price at cost. Note that even with salient quality, firm j provides higher perceived surplus than firm $-j$ because $q_j = \underline{q}$ is closer than q_{-j} to the surplus maximizing quality level \bar{q} under salient quality. However, by raising price up to $p_j = c(q) + \delta(\underline{q} - q)$, firm j renders price salient. This is because its quality to price ratio is now lower than the quality to cost ratio of $\tilde{q}(q)$ (since p_j decreases in q for $q < \underline{q}$) and therefore lower than the quality cost ratio of any $q < \tilde{q}(q)$. Since the quality level \underline{q} maximizes surplus given salient price, firm j wins the market and makes positive profits.

We have thus proved that when $\hat{q}/c(\hat{q}) \geq 1/\delta$ the strategy $q_j = \underline{q}$ by firm j beats any other feasible quality $q_{-j} \neq \underline{q}$ set by firm $-j$. This implies that, in this range of quality to cost ratios, at least one firm $j = 1, 2$ must play $q_j = \underline{q}$ in equilibrium. In fact, if both firms play $q_j, q_{-j} \neq \underline{q}$, then at least one firm would find it profitable to deviate to \underline{q} . Equilibria are then characterized as follows. First, when one firm plays $q_j = \underline{q}$, its opponent $-j$ is sure to make zero profits and is thus willing to play any feasible quality. At the same time, the equilibrium values of q_{-j} are those to which $q_j = \underline{q}$ constitutes a best response. Clearly, $q_j = \underline{q}$ constitutes a best response to all quality levels q_{-j} such that price is salient in the price subgame; from the above analysis, this set clearly includes the interval $[\tilde{q}(q), +\infty)$ but it also includes those lower qualities $q_{-j} < \tilde{q}(q)$ that have higher quality to cost ratios, namely $q_{-j}/c(q_{-j}) \geq \underline{q}/p_j(\underline{q}, q_{-j})$ where $p_j(\underline{q}, q_{-j}) = c(q_{-j}) + \delta(\underline{q} - q_{-j})$. Because firm j provides higher quality \underline{q} at a lower quality to price ratio, price is salient in the equilibrium of the price subgame. To work through this condition, note that $q_{-j}/c(q_{-j}) = \underline{q}/p_j(\underline{q}, q_{-j})$ if and only if $q_{-j}/c(q_{-j}) = 1/\delta$ (recall that $q_{-j} < \hat{q}$). As a consequence, $q_j = \underline{q}$ is a best response to any quality level in the set $[\check{q}, +\infty)$, where \check{q} is defined by $\check{q}/c(\check{q}) = 1/\delta$ and $\check{q} < \hat{q}$.

This implies that the entire set of subgame perfect equilibria in pure strategies when $\hat{q} \leq \underline{q} < \bar{q}$ is given by:

$$\left\{ (q_j, p_j), (q_{-j}, p_{-j}) : q_j = \underline{q}, q_{-j} \geq \check{q}, p_j = c(q_{-j}) + \delta(\underline{q} - q_{-j}), p_{-j} = c(q_{-j}) \right\}_{j=1,2}.$$

In these equilibria, firm $j = 1, 2$ sets $q_j = \underline{q}$ while firm $-j$ sets $q_{-j} \in [\check{q}, +\infty)$. If $q_{-j} \neq \check{q}$ firm $-j$ makes zero profits, while firm j makes strictly positive profits. There is a unique symmetric equilibrium in which both firms set quality \underline{q} and make zero profits. In either case, the quality level bought by consumers is $q^S(F) = \underline{q}$.

Finally, note that the proof above also implies that, when there are $N > 2$ identical firms, the configuration $q_k = q^S(F)$ for all k is a (symmetric) equilibrium. To see this, consider firm j 's best response q_j when the remaining $N - 1$ firms set quality q_{-j} . This setting is formally equivalent to the 2-firm setting above, except that the reference attribute levels are now given by $\bar{q} = \frac{q_j + (N-1)q_{-j}}{N}$ and $\bar{p} = \frac{c(q_j) + (N-1)c(q_{-j})}{N}$. However, because the salience constraint is invariant to reference attributes as strict convex combinations of attributes in the choice set, the results go through as above. In particular, firm j has no positive incentives to deviate from the symmetric equilibrium where every other firm chooses quality $q^S(F)$ as defined above. ■

Corollary 1 (symmetric equilibrium for quadratic costs). We work out the symmetric equilibrium in Proposition 2 when costs are quadratic, $v(q) = \frac{v}{2}q^2$. In this case, $c'(q) = v \cdot q$ so that $\bar{q} = \frac{1}{\delta v}$, and $\underline{q} = \frac{\delta}{v}$. Moreover, $\bar{F} = \frac{1}{2\delta^2 v}$ and $\underline{F} = \frac{\delta^2}{2v}$. Finally, \hat{q} satisfies $c'(\hat{q}) = c(\hat{q})/\hat{q}$, which yields $\hat{q} = \sqrt{2F/v}$. ■

Lemma 3 (quality competition under rationality with asymmetric costs). We consider an asymmetric shock that reduces the marginal cost v_1 of firm 1, keeping v_2 constant. Because firm 1 can always produce the quality of firm 2 at lower cost, firm 2 cannot win the market in equilibrium. Equation (20) then indicates that to maximize its profit, firm 1 chooses the surplus maximizing quality, that solves $c'_1(q_1^*) = 1$. Under quadratic costs, we have $q_1^* = 1/v_1$. Our “symmetric play” equilibrium selection rule then implies that the losing firm 2 sets quality as in the symmetric equilibrium in which both firms have the same cost function $c_2(q)$. Thus, firm 2 sets $c'_2(q_2^*) = 1$, namely $q_2^* = 1/v$. Equilibrium prices are $p_2 = c_2(q_2^*)$ and $p_1 = c_2(q_2^*) + (q_1^* - q_2^*)$. ■

Proposition 3 (industry wide cost shocks). Under the symmetric equilibrium of Equation (13), consider an increase in the unit cost of all firms, from F_0 to $F_1 > F_0$. If the

interval $[F_0, F_1]$ has a non-empty overlap with the interval $[\underline{F}, \overline{F}]$, then equilibrium quality strictly increases from $\max\{\delta/v, \sqrt{2F_0/v}\}$ to $\min\{\sqrt{2F_1/v}, 1/(\delta v)\}$. Otherwise, equilibrium quality provision does not change, staying at δ/v if $F_1 < \underline{F}$ or at $1/(\delta v)$ if $F_0 > \overline{F}$.

Note that, when $\delta < 1$, the equilibrium quality can be written as $\frac{1}{v} \cdot A(v, F)$, where $A(v, F) = \max\{\delta, \min\{\sqrt{2Fv}, 1/\delta\}\}$. As a consequence, following an increase in the marginal cost of producing quality for all firms, quality provision strictly decreases. Formally, $\partial_v \frac{1}{v} \cdot A(v, F) < \partial_v \sqrt{\frac{2F}{v}} < 0$.

We can also ask when is the change in quality provision in reaction to a marginal increase in v larger than in the rational case? When $\delta = 1$, quality provision equals $1/v$. Therefore, the change in quality provision increases when $\delta < 1$ if and only if $A(v, F) > 1$, namely when quality is over provided to begin with (i.e. if $F > \frac{1}{2v}$). ■

Proposition 4 (firm specific cost shocks). Starting from the symmetric equilibrium of Equation (13), let the marginal cost of firm 1 drop to $v_1 < v_2 = v$. This implies that firm 1 will win the market and firms 2 will lose it, making zero profits. Suppose in fact that this was not the case. Then, firm 1 could adopt the same quality of firm 2, produce it at lower cost, and win the market.

To work out the equilibrium in which the low cost firm 1 wins we proceed in two steps: i) we first compute firm 1's best response from the symmetric equilibrium quality provision q^S of firm 2; ii) we then show that firm 2 has no incentive to deviate; the resulting configuration is thus an equilibrium. Because the losing firm plays the "symmetric" quality strategy, this is the equilibrium selected by our refinement.

When the unit costs are sufficiently high, $F > \frac{\delta^2}{2v}$, the average costs of firm 2 satisfy $c(q^S)/q^S > \delta$. It then follows from the analysis in Proposition 2 that firm 1's best response is to engineer a salient quality increase: i) when $c(q^S)/q^S \in [\delta, 1/\delta]$, firm 1 sets q_1^* satisfying $c'_1(q_1^*) = c(q^S)/q^S$. With quadratic costs, this reads $q_1^* = q^S \cdot \frac{v}{v_1} > q^S$. Firm 2 has no incentive to deviate: in this parameter range, it is already minimizing average cost, so it cannot engineer a quality innovation that gives it a salient advantage. Together with the fact that it has higher costs, this precludes any profitable deviation. The equilibrium prices are then $p_2 = c_2$, $p_1 = c_2 \frac{q_1}{q_2}$. As in the analysis of Proposition 1, the sharing rule ensures

that at these prices quality is salient and firm 1 captures the entire market.

ii) when the average costs of firm 2 exceed $1/\delta$, the quality provision in the symmetric equilibrium satisfies $c'(q^S) = 1/\delta$. In this case, firm 1 boosts quality to q_1^* satisfying $c'_1(q_1^*) = 1/\delta$, so that $q_1^* > q^S$. Firm 2 again has no incentive to deviate, since increasing quality (thereby diminishing average costs, if possible below that of firm 1) is never profitable: if firm 2 engineers a salient quality advantage then it decreases its valuation, while if it creates a salient price advantage it cannot price above cost. The equilibrium prices are then $p_2 = c_2$, $p_1 = c_2 + \frac{1}{\delta}(q_1 - q_2)$. As in the analysis of Proposition 1, the sharing rule ensures that at these prices firm 1 captures the entire market.

Consider now the case where $F < \frac{\delta^2}{2v}$. While firm 2 sets q^S such that $c'(q^S) = \delta$, firm 1's best response is to set q_1^* satisfying $c'_1(q_1^*) = c(q^S)/q^S$, provided $q_1^* > q^S$. With quadratic costs, this reads $q^S = \delta/v$ and $q_1^* = \frac{c(q^S)/q^S}{v_1}$. Thus, $q_1^* > q^S$ requires $F > \frac{\delta^2}{v} \left(\frac{v_1}{v} - \frac{1}{2} \right)$.

If firm 1's cost advantage is sufficiently large, namely $v_1 < v/2$, then firm 1 strictly increases quality provision. The equilibrium prices are then $p_2 = c_2$, $p_1 = c_2 \frac{q_1}{q_2}$.

If instead firm 1's cost advantage is small, $v_1 > v/2$, then for low enough levels of the unit cost F , it is optimal for firm 1 to keep quality provision at the equilibrium level prior to the shock, $q_1^* = \delta/v$, and translate its cost advantage into profits by setting price $p_1 = c(\delta/v)$. Finally, firm 2 has no incentive to deviate because decreasing quality (thereby diminishing average costs) also decreases perceived surplus. The sharing rule ensures that at prices $p_1 = p_2 = c(\delta/v)$, firm 1 captures the entire market, as only firm 1 can reduce price and make a profit.

Finally, note that the proof above extends to the case of $N > 2$ identical firms that start out in the symmetric equilibrium of Proposition 2, and where *one* firm then receives an idiosyncratic cost shock. As in the proof of the symmetric equilibrium (Proposition 2), the best response q_j of the innovating firm j , when the remaining $N - 1$ firms play the same quality q_{-j} , does not depend on $N \geq 2$. Therefore, when these firms are at the longterm symmetric equilibrium $q^S(F)$, firm j 's best response is as described above. Moreover, none of the remaining $N - 1$ firms then has a positive incentive to deviate from $q^S(F)$. To see this, recall that when $F > \frac{\delta^2}{v} \left(\frac{v_1}{v} - \frac{1}{2} \right)$, firm j optimally increases quality. Then a firm $-j$ cannot

profitably engender a salient advantage relative to firm j (either a salient quality increase relative to firm j 's or a lower quality that makes j 's higher price salient). This is because firm $-j$ is already minimising average costs. When instead $F \leq \frac{\delta^2}{v} \left(\frac{v_1}{v} - \frac{1}{2} \right)$, firm j does not change quality provision and the reasoning for the 2 firm case applies. ■

Lemma 4 (returns competition under rationality). This setting is similar to the price competition game of Lemma 1. While the costs facing investors are fixed at v (the security's risk), intermediaries compete in terms of the return they provide investors. Since intermediaries provide identical securities, this competition game admits symmetric equilibria. As in the main text, we focus on symmetric equilibria: both firms offer the maximum return to investors, $R_i - F = \bar{R} - F$, and share the market. No intermediary has an incentive to deviate from this configuration: increasing the returns offered to investors would lead to negative profits, while decreasing the returns would lead to the loss of the market share and would not increase profits. ■

Proposition 5 (financial innovation under salient thinking). Suppose firm 2 creates a security of fixed total return and cost, $(\bar{R} - F, \rho)$. Firm 1 develops a financial innovation and can create a family of securities $(\bar{R} + \alpha - F, \rho + \frac{v}{2} \cdot \alpha^2)$, indexed by α , the increase in returns relative to the competition. The firms play a two stage game: in the first stage firm 1 chooses α , and in the second stage both firms choose how big a return to pledge to investors. Firm 1 pledges return $R_\alpha - F$ where $R_\alpha \in [\bar{R}, \bar{R} + \alpha]$ so that in the return competition stage it sells security $(R_\alpha - F, \rho + \frac{v}{2} \cdot \alpha^2)$ and maximizes profits $\bar{R} + \alpha - R_\alpha$.

To determine the optimal choice of α , we begin by noticing that, for α sufficiently small, the marginal cost of quality for firm 1 is lower than its average cost. This is because returns increase linearly in α , while risk increases quadratically. As a result, firm 1 finds it optimal to provide a salient increase in returns. The pledged returns R_α must satisfy both the constraint that returns are salient, and the valuation constraint. The salience constraint reads $R_\alpha - F > (\bar{R} - F) \cdot \frac{\rho + \frac{v}{2}\alpha^2}{\rho}$ (recall that firm 1 provides higher returns at a higher risk), while the valuation constraint reads $R_\alpha > \bar{R} + \delta \frac{v}{2} \alpha^2$. The valuation constraint is binding when $\bar{R} > F + \delta \rho$. In this case, firm 1 must provide at least $R_\alpha = \bar{R} + (\bar{R} - F) \frac{\delta v}{2} \alpha^2$. To maximize profits $\bar{R} + \alpha - R_\alpha$, firm 1 sets $\alpha = \frac{1}{\delta v}$.

The salience constraint is binding when $\bar{R} \geq F + \delta\rho$. In this case, firm 1 must provide at least $R_\alpha = F + (\bar{R} - F) \left(1 + \frac{v}{2\rho}\alpha^2\right)$. To maximize profits $\bar{R} + \alpha - R_\alpha$, firm 1 sets $\alpha = \frac{1}{\bar{R} - F} \cdot \frac{1}{v}$.

■