IEB Working Paper 2026/19

ASTRATEGIC PRODUCT DISCLOSURE AND OBFUSCATION IN THE COFFEE MARKET UNDER ATIONAL INATTENTION

Stefano Castriota, Luigi Palumbo, Fernando Payró, Amedeo Piolatto

Version October 2026

IEB Institut d'Economia de Barcelona

IEB Working Paper 2026/19

STRATEGIC PRODUCT DISCLOSURE AND OBFUSCATION IN THE COFFEE MARKET UNDER RATIONAL INATTENTION

Stefano Castriota, Luigi Palumbo, Fernando Payró, Amedeo Piolatto

The **Barcelona Institute of Economics (IEB)** is a research centre at the University of Barcelona (UB) which specializes in the field of applied economics. The IEB is a foundation funded by the following institutions: "La Caixa" Foundation, Saba, the Barcelona City Hall, the Barcelona Metropolitan Area, the University of Barcelona, the Autonomous University of Barcelona, the Barcelona Provincial Council, Agbar, Cuatrecasas, Consorci Zona Franca Barcelona, Fundación Miquel y Costas, AENA and Fira de Barcelona.

Postal Address:
Institut d'Economia de Barcelona
Facultat d'Economia i Empresa
Universitat de Barcelona
C/ John M. Keynes, 1-11
(08034) Barcelona, Spain
Tel.: + 34 93 403 46 46
ieb@ub.edu
http://www.ieb.ub.edu

The IEB working papers represent ongoing research that is circulated to encourage discussion and has not undergone a peer review process. Any opinions expressed here are those of the author(s) and not those of IEB.

IEB Working Paper 2026/19

STRATEGIC PRODUCT DISCLOSURE AND OBFUSCATION IN THE COFFEE MARKET UNDER RATIONAL INATTENTION

Stefano Castriota, Luigi Palumbo, Fernando Payró, Amedeo Piolatto[i]

ABSTRACT: Measuring strategic obfuscation is challenging because firms deliberately conceal quality from researchers and consumers. We investigate how informational frictions generate pricing power using data on 145 ground coffee products from Italian supermarkets, where Arabica content is selectively disclosed or obfuscated. To overcome unobservability, we utilise a theoretical model of consumer rational inattention as empirical scaffolding. The framework predicts a threshold equilibrium: high-quality firms disclose attributes to capture a price premium, whilst lower-quality firms obfuscate to benefit from market pooling. Empirical results confirm selective disclosure, demonstrating how firms actively exploit consumer frictions to protect markups.

JEL Codes: D83, L15, L13
Keywords: Rational Inattention, Shannon Entropy in Consumer Choice, Product Obfuscation, Arabica vs. Robusta Pricing, Information Acquisition Costs, Selective Disclosure Thresholds.

Stefano Castriota
Università di Pisa
E-mail: stefano.castriota@unipi.it

Luigi Palumbo
Banca d'Italia
E-mail: luigi.palumbo@bancaditalia.it

Fernando Payró
Universitat Autònoma de Barcelona (UAB)
E-mail: Fernando.Payro@uab.es

Amedeo Piolatto
Universitat Autònoma de Barcelona (UAB)
BSE & IEB
E-mail: amedeo.piolatto@uab.es

[i]We gratefully acknowledge comments received from participants at the ENTER Jamboree, as well as seminar participants at Free University of Bolzano. We thank Philipp Weinhold for research assistance. Piolatto and Payró gratefully acknowledge support from the Spanish Ministry of Science and Innovation (Piolatto: PID2022-137707NB-I00 and PID2025-169447NB-I00; Payr´o: PID2023-147183NB-I00) and from the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics, CEX2024-001476-S), funded by MCIN/AEI/10.13039/501100011033. The views expressed in this paper are those of the authors only and do not necessarily represent the views of the Bank of Italy or the Eurosystem.

1 Introduction

Consumers routinely make decisions in environments where product information is incomplete, complex, or costly to process. While standard disclosure theory predicts that competition should lead firms to reveal all relevant information, anecdotal evidence across many markets suggests otherwise: firms frequently withhold or obscure salient product attributes. This paper contributes to the understanding of why such obfuscation arises, and how it may shape prices and the relationship between price and quality.

We develop and test a theory of strategic product disclosure under rational inattention. In our framework, firms decide whether to reveal or conceal a key quality attribute, while consumers face costs of processing information. When information is disclosed, consumers observe quality and make standard utility-maximising choices. When information is obfuscated, consumers optimally allocate limited attention, leading to probabilistic demand that responds only imperfectly to underlying quality differences. This departure from full information fundamentally alters firms' incentives to disclose.

The model delivers a simple and empirically relevant prediction: in equilibrium, there exists a disclosing-rule threshold. Firms with sufficiently high quality disclose to differentiate themselves, while lower-quality firms optimally choose to pool by withholding information. Obfuscation weakens the sensitivity of demand to quality, allowing low- and intermediate-quality firms to sustain higher prices than they could under full transparency. As a result, the price–quality relationship differs sharply across information regimes: it is steep under disclosure and flattened under obfuscation.

We test these predictions in the Italian ground coffee market, which provides a particularly clean setting to study strategic disclosure. Coffee quality is multi-dimensional, but the share of Arabica beans is a widely understood and salient proxy for quality.[1] Coffee consumption is widespread and culturally salient, and consumers are familiar with key quality attributes such as the Arabica share. At the same time, disclosure of this attribute is voluntary: some products prominently report the Arabica share on the label, while others do not, despite the absence of regulatory constraints. This variation allows us to observe firms' disclosure choices directly. Moreover, the market is well suited for empirical analysis because products are relatively homogeneous (we focus on standardised 250-gram packages sold in supermarkets) yet their true composition is not always directly observable to consumers.

A key empirical challenge in studying obfuscation is that the researcher, like the consumer, typically cannot observe what firms hide. We address this by manually recovering the Arabica share for non-disclosing products using technical documentation and direct enquiries with

[1] Cheng et al. (2016) show that Arabica is of superior quality compared to Robusta.

producers. This allows us to observe both true quality and disclosure decisions, and to estimate how disclosure affects the pricing of quality.

Our empirical findings strongly support the theoretical predictions. First, disclosure is highly selective: the probability that a firm reveals its Arabica share increases sharply with actual quality, consistent with a separating equilibrium. Second, products that disclose quality command higher prices on average, reflecting vertical differentiation. Third, and most importantly, the price–quality gradient is substantially steeper under disclosure. While Arabica content is positively correlated with prices in both regimes, its effect is more than twice as large when quality is explicitly revealed. This implies that obfuscation attenuates the transmission of quality into prices, consistent with the idea that costly attention reduces consumers' responsiveness to hidden attributes. The two price–quality schedules intersect at an intermediate level of quality, providing a reduced-form counterpart to the model's disclosure threshold.

Our paper contributes to several strands of the industrial organisation literature. First, we provide a micro-founded explanation for partial disclosure equilibria. Classical unravelling arguments (e.g. Milgrom, 1981) predict full disclosure when information is costless to reveal, yet many markets exhibit persistent opacity. We show that rational inattention breaks the unravelling logic by weakening consumers' response to hidden information, thereby sustaining equilibria in which some firms disclose while others obfuscate. Second, we contribute to the literature on obfuscation by shifting the focus from prices to product characteristics. While much of the existing work studies price complexity and search costs, we show how firms strategically manipulate the salience of quality attributes. Third, we provide empirical evidence linking information frictions to firms' pricing strategies, demonstrating that cognitive constraints not only affect consumer behaviour but also shape equilibrium market outcomes.

More broadly, our results highlight the role of informational frictions in determining how markets aggregate and transmit information. When consumers face costs of attention, transparency becomes an endogenous outcome rather than a default prediction of competition. This has implications for policy: voluntary disclosure may be insufficient to ensure efficient outcomes, as low-quality firms have strong incentives to remain opaque. Interventions that reduce information-processing costs, such as standardised labels or third-party certification, may therefore strengthen the link between prices and underlying quality.

Our findings align with broader evidence that informational and cognitive frictions distort market clearing. As shown across insurance and retail markets, these frictions generate consumer inertia and firm market power, dampening price-quality signals even in competitive settings (e.g. Coen et al., 2026). While structural research largely focuses on big-ticket durables like automobiles (Moraga-González et al., 2023), we demonstrate that cognitive con-

straints distort routine, low-stakes non-durables as well. By showing how attention costs alter voluntary disclosure incentives in fast-moving consumer goods, our analysis extends the friction paradigm to everyday retail goods.

The paper proceeds as follows: we first review the related literature in Section 1.1. Subsequently, Section 2 outlines the model and illustrates the theoretical results. Section 3 describes the data and methodology used in the empirical analysis. Section 3.1 establishes the econometric results. Section 4 concludes. We provide all mathematical proofs in Appendix A and additional empirical tests in Appendix B.

1.1 Related literature

At a fundamental level, our framework aligns with the industrial organisation literature on product market frictions and asymmetric information synthesised by Gavazza and Lizzeri (2021), who demonstrate how informational frictions alter market liquidity, buyer sorting, and price dispersion. Within this tradition, information environment manipulation is categorised into opacity and obfuscation. While opacity denotes structural non-transparency allowing full concealment, obfuscation represents a deliberate strategy to heighten price or attribute complexity, making information costly to acquire rather than impossible to find.

A prominent body of work examines opaque selling through intermediaries, which is particularly salient on online platforms where characteristics are observed post-purchase. Fay (2008) shows firms may intentionally disguise products using opaque channels hiding identity until transaction completion. Similarly, Shapiro and Shi (2008) analyse opaque travel agencies as market segmentation tools, where consumers trade lower prices for reduced attribute clarity. Although obfuscation here stems from the selling environment rather than direct disclosure decisions, the underlying economic mechanism is aligned: by hindering direct comparability, firms soften price competition and segment consumers by uncertainty tolerance. Extending this, Piolatto and Schuett (2025) show an intermediary operating dual platforms (one fully revealing, one opaque) to maximise rent extraction from both buyers and sellers.

Our setting focuses specifically on attribute obfuscation. Following Varian (1980), one strand of literature assumes hidden prices necessitate consumer search (Pennerstorfer et al., 2020; Moraga-González and Sun, 2025). This research explains persistent price dispersion in digital markets, showing firms obfuscate pricing structures via complex add-ons or hidden fees to increase search costs and protect markups (Ellison, 2005; Ellison and Ellison, 2009; Ellison and Wolitzky, 2012; Petrikaitė, 2018a,b). A second stream, including our study, focuses on attribute complexity, showing firms deliberately obscure characteristics to hinder comparison, especially for 'shrouded attributes' hidden from myopic or cognitively constrained consumers. Here, Gabaix and Laibson (2006) show that even in competitive markets, firms suppress post-

purchase attribute information (e.g. maintenance), while Spiegler (2006) considers rationally bounded consumers.

We contribute directly to characteristic-based obfuscation in oligopolies. Ivanov (2013) and Hwang et al. (2023) study competitive disclosure where firms set prices and information release, finding full disclosure as the unique symmetric equilibrium with sufficient firms (as disclosure maximises high-valuation probability). We provide a rationale for empirical evidence where only some firms disclose. Unlike Armstrong and Zhou (2022), where disclosure is set centrally by a platform, firms in our model choose independently. Furthermore, while Cheong and Kim (2004) argue oligopolists conceal information only when costly and yield pooling, our model shows firms conceal even at zero cost, with incentives depending on firm count so that revealing and non-revealing firms coexist. Crucially, obfuscation attenuates demand elasticity regarding underlying quality differences, allowing firms to pool heterogeneous goods and sustain higher prices. Our model formalises how obfuscation reduces demand quality-sensitivity, shaping disclosure incentives.

While classical voluntary disclosure theory predicts market unravelling to full information (Milgrom, 1981), our empirical findings support a rational-inattention equilibrium. In standard unravelling, above-average quality firms disclose, forcing all firms to reveal to avoid being inferred as the worst type. This departure aligns with work on strategic inference management: Miklos-Thal and Zhang (2013) show firms intentionally degrade or withhold information ('de-marketing') across product lines. We build on this logic, low-quality firms obfuscate attributes to avoid revealing their type, but whereas Miklos-Thal and Zhang (2013) rely on bounded expectations, cognitive attention costs in our framework prevent Milgromian unravelling from commencing.

Because processing detailed coffee attributes (e.g. the Arabica-Robusta ratio) incurs mental costs, consumers do not treat non-disclosure as a definitive signal of minimum quality. Instead, they pool non-disclosing firms into a coarse 'shrouded' category. This rational inattention creates a strategic 'dark pool' for firms below a quality threshold (v^*), making obfuscation more profitable than revealing a mediocre position. Persistent non-disclosure in the Italian coffee market thus reflects processing costs rather than verifiability failure.

This mechanism contrasts with traditional search cost models driven by consumer heterogeneity. While we assume limited attention induces coarse inferences, other settings feature consumers with search costs low enough to bypass frictions. For example, in the showrooming literature, Bar-Isaac and Shelegia (2023) show attentive consumers incur physical search costs to acquire verifiable pre-purchase information, altering firm disclosure choices. In our setting, the barrier is strictly cognitive: lacking a low-cost physical search option, unravelling fails, shielding lower-quality firms in a pooled market.

Methodologically, we model cognitive constraints via rational inattention (Sims, 2003a), where consumers optimise information processing against attention costs. Janssen and Kasinger (2024a) apply this to price obfuscation, deriving symmetric equilibria where all firms either reveal or obfuscate (yielding higher prices). We extend this framework to product characteristics, establishing a threshold rule: high-quality firms reveal, while lower-quality firms obfuscate and pool. This links disclosure theory with rational inattention, microfounding partial disclosure equilibria where obfuscated goods sell at lower prices.

Empirically, attribute-based obfuscation remains understudied compared to price obfuscation (Ellison and Ellison, 2009) or lab experiments (Fehr and Wu, 2025). In supermarkets, Richards et al. (2020) show varying package sizes hinder price comparisons to soften competition; we are the first to focus on product characteristics in this setting. Our work also connects to agricultural economics on coffee quality and information asymmetry. Hilger et al. (2011) show experimentally that consumers respond strongly to quality disclosures that reduce cognitive search costs. Because coffee is a credence good (McCluskey, 2000), firms face persistent incentives to shroud composition, making the Arabica share a strategic signal separating manufacturer from retailer brands. While McCluskey (2000) frames coffee as a credence good with retail-level asymmetry, upstream quality depends on supply chain architecture: Macchiavello and Miquel-Florensa (2017) show vertical integration and relational contracts secure high-quality bean production. We complement this supply-side focus by analysing demand-side disclosure choices. Finally, whereas Arnot et al. (2006) study how ethical labels (e.g. Fair Trade) affect price sensitivity, data limitations preclude Fair Trade attributes here. Together, these literatures form a unified framework for analysing disclosure strategies and pricing under partially observable quality.

2 The model

We consider a unit mass of atomless firms, each selling one differentiated product. A firm's product has quality $v \in [0,1]$, interpreted as the share of Arabica beans in the blend. Qualities are distributed according to a prior λ, which admits a strictly positive density on $[0,1]$. Each firm observes the quality of its own product.

Each firm faces a local residual-demand problem in which market size is normalised to one. Consumers choose between the firm's product and an outside option summarising the remaining products available in the market. The utility of the outside option is normalised to zero. We do not explicitly model substitution among all products; instead, the outside option captures the competitive environment faced by each firm.

The consumer's gross utility from consuming a product of quality v is $u(v)$, where $u : [0,1] \to \mathbb{R}_+$ is continuously differentiable and satisfies $u'(v) > 0$. If the product is sold at price p, the consumer's net utility from purchasing is $u(v) - p$. Production costs are normalised to zero.

After observing v, the firm chooses a disclosure decision $d \in \{0,1\}$ and a price $p \in \mathbb{R}_+$. If $d = 1$, quality is verifiably disclosed. If $d = 0$, quality is not directly revealed.

When quality is disclosed, consumers observe v without incurring an information-processing cost. They purchase if and only if $u(v) \geq p$. We assume that consumers purchase when indifferent. A firm that discloses therefore charges $p^D(v) = u(v)$ and obtains profit $\pi^D(v) = u(v)$.

2.1 Obfuscation and coarse price inference

When quality is not disclosed, consumers may acquire imperfect information about it through search, inspection, brand cues, reviews, or previous consumption experience. Processing this information is costly.

The price of a non-disclosing product may itself be informative about quality. Fully Bayesian consumers would therefore use both the disclosure decision and the observed price to update their beliefs. We instead impose a coarse-inference restriction: consumers condition on whether quality is disclosed, but do not use variation in prices within the non-disclosing category to infer quality.

Let $O \subseteq [0,1]$ denote the set of qualities for which firms choose not to disclose. For a fixed $\varepsilon \in (0,1)$, define

$$\mu_O^{\varepsilon}(A) = (1-\varepsilon)\frac{\lambda(A \cap O)}{\lambda(O)} + \varepsilon\lambda(A)$$

for every measurable set $A \subseteq [0,1]$, provided $\lambda(O) > 0$.

Assumption 1 (Coarse price inference)**.** *After observing any pair* $(p, d = 0)$*, the consumer uses the belief*

$$\mu(\cdot \mid p, d = 0) = \mu_O^{\varepsilon}.$$

In particular, the consumer uses the same belief at every price charged by a non-disclosing firm.

The coarse-inference restriction is related to cursed equilibrium, under which agents fail to fully account for the dependence of other players' actions on their private information (Eyster and Rabin, 2005). It is particularly close to cursed-expectations equilibrium, where agents underuse the information contained in equilibrium prices (Eyster et al., 2019), and to models of product pricing in which consumers underinfer hidden costs from observed prices (Eyster et al., 2021). It is also related to analogy-based expectation equilibria, where agents group distinct contingencies into coarse analogy classes and respond to average behaviour within each class (Jehiel, 2005; Jehiel and Koessler, 2008). In our model, all prices charged by non-disclosing firms belong to a single inference class. We do not derive Assumption 1 from any of these equilibrium concepts. Instead, we impose the coarse-inference rule directly and combine it with rational inattention in the consumer's subsequent information-acquisition problem.

The first component of μ_O^{ε} is anchored to the distribution of qualities among firms that do not disclose. The contamination component represents subjective uncertainty about the composition of the hidden category. Since λ has a strictly positive density and $\varepsilon > 0$, the belief μ_O^{ε} assigns positive probability to every open neighbourhood of every quality.

The parameter ε is fixed throughout the analysis. The belief μ_O^{ε} is not the Bayesian posterior conditional on non-disclosure. Moreover, under Assumption 1, the consumer ignores the information about quality contained in the price of a non-disclosing product.

Following the rational-inattention literature, the consumer chooses how much information to acquire by trading off expected utility against a Shannon mutual-information cost (Sims, 2003b; Matějka and McKay, 2015). Rational inattention has also been incorporated into strategic environments in which an informed seller chooses an offer and the buyer must allocate attention to the product's quality and the seller's offer (Ravid, 2020). Our use of rational inattention in an environment with product obfuscation is particularly related to Janssen and Kasinger (2024b). In the present model, rational inattention governs the direct acquisition of information about quality conditional on the coarse belief.

Fix a price p and a belief μ. The consumer's information strategy can be represented by conditional purchase probabilities $\beta(v, p; \mu) \in [0, 1]$, where $\beta(v, p; \mu)$ is the probability that the product is purchased when its true quality is v. Let

$$\bar{\beta}(p;\mu) = \int_0^1 \beta(v,p;\mu)\,\mu(dv)$$

denote the unconditional purchase probability.

The consumer solves

$$\max_{\beta:[0,1]\to[0,1]} \int_0^1 (u(v)-p)\beta(v,p;\mu)\,\mu(dv) - \kappa\mathcal{I}(\beta,\mu),$$

where $\kappa > 0$ measures the cost of processing information and

$$\mathcal{I}(\beta,\mu) = H\big(\bar{\beta}(p;\mu)\big) - \int_0^1 H\big(\beta(v,p;\mu)\big)\,\mu(dv)$$

is the Shannon mutual-information cost associated with the binary purchase decision. Here,

$$H(t) = -\left[t\log t + (1-t)\log(1-t)\right],$$

with the convention that $0\log 0 = 0$.

Conditional choice rules that agree μ-almost everywhere generate the same consumer payoff and information cost. We therefore identify such rules and define product-level demand using the continuous representative of the optimal conditional choice rule. Because the beliefs considered below have full support, this continuous representative is unique whenever it exists: two continuous functions that agree μ-almost everywhere must agree everywhere.

We impose the following regularity conditions.

Assumption 2 (Regularity)**.** *For every measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$, and under the belief μ_O^ε, the following conditions hold:*

1. *for every $p \geq 0$, the consumer's problem has a unique optimal unconditional purchase probability, and the optimal conditional choice rule admits a continuous representative;*
2. *in a neighbourhood of every price solving a firm's obfuscation problem, the consumer's solution is interior;*
3. *for every $v \in [0,1]$, the function $p \mapsto p\beta(v,p;\mu_O^\varepsilon)$ has a unique positive interior maximiser. The obfuscation objective is continuous on $[0,1]\times\mathbb{R}_+$. In a neighbourhood of its maximiser, it is continuously differentiable in (v,p), its derivative with respect to p is continuously differentiable in both v and p, and the strict second-order condition holds. The maximiser is locally bounded as a function of v.*

Lemma 1 (Consumer's best response)**.** *Fix a price p and a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. At an interior solution, the unique continuous representative of the consumer's optimal conditional choice rule satisfies, for every $v \in [0,1]$,*

$$\log\left(\frac{\beta(v,p;\mu)}{1-\beta(v,p;\mu)}\right)=\log\left(\frac{\bar{\beta}(p;\mu)}{1-\bar{\beta}(p;\mu)}\right)+\frac{u(v)-p}{\kappa}.$$

Equivalently,

$$\beta(v,p;\mu)=\frac{1}{1+\frac{1-\bar{\beta}(p;\mu)}{\bar{\beta}(p;\mu)}\exp\left(-\frac{u(v)-p}{\kappa}\right)}.$$

The unconditional purchase probability satisfies

$$\bar{\beta}(p;\mu)=\int_0^1\frac{1}{1+\frac{1-\bar{\beta}(p;\mu)}{\bar{\beta}(p;\mu)}\exp\left(-\frac{u(v)-p}{\kappa}\right)}\,\mu(dv).$$

Proof. See Appendix A. □

At an interior solution, for a fixed price and belief, the conditional purchase probability is continuous and strictly increasing in quality. Importantly, $\beta(v,p;\mu)$ is the actual demand faced by a firm of quality v. It is not replaced by the average purchase probability $\bar{\beta}(p;\mu)$.

For later use, define

$$V(p;\mu)=\int_0^1\left[\beta(v,p;\mu)-\bar{\beta}(p;\mu)\right]^2\mu(dv).$$

Since u is strictly increasing, μ has full support, and the solution is interior, $V(p;\mu)>0$.

Lemma 2 (Demand response)**.** *Fix a belief $\mu=\mu_O^\varepsilon$ for some measurable $O\subseteq[0,1]$ satisfying $0<\lambda(O)<1$. At an interior solution,*

$$\frac{\partial\bar{\beta}(p;\mu)}{\partial p}=-\frac{\bar{\beta}(p;\mu)\left[1-\bar{\beta}(p;\mu)\right]}{\kappa V(p;\mu)}\int_0^1\beta(v,p;\mu)\left[1-\beta(v,p;\mu)\right]\mu(dv)<0.$$

Moreover, for every $v\in[0,1]$,

$$\frac{\partial\beta(v,p;\mu)}{\partial p}=-\frac{\beta(v,p;\mu)\left[1-\beta(v,p;\mu)\right]\bar{\beta}(p;\mu)\left[1-\bar{\beta}(p;\mu)\right]}{\kappa V(p;\mu)}<0.$$

Proof. See Appendix A. □

2.2 The firm's problem under obfuscation

Given a coarse belief μ, a firm of quality v that does not disclose chooses its price to solve

$$\pi^O(v;\mu)=\max_{p\geq0}p\beta(v,p;\mu).$$

Let

$$p^O(v;\mu) = \arg\max_{p\geq 0} p\beta(v,p;\mu)$$

denote its unique optimal price.

Lemma 3 (Optimal obfuscation price)**.** *Fix a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. The optimal obfuscation price satisfies*

$$p^O(v;\mu) = \frac{\kappa V\big(p^O(v;\mu);\mu\big)}{\big[1-\beta\big(v,p^O(v;\mu);\mu\big)\big]\,\bar\beta\big(p^O(v;\mu);\mu\big)\,\big[1-\bar\beta\big(p^O(v;\mu);\mu\big)\big]}.$$

Proof. See Appendix A. □

The optimal price under obfuscation generally depends on the firm's true quality. Although consumers use the same coarse belief for every non-disclosing product, their direct information acquisition generates type-dependent demand.

Proposition 1 (Monotone pricing under obfuscation)**.** *Fix a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. The optimal obfuscation price $p^O(v;\mu)$ is strictly increasing in v.*

Proof. See Appendix A. □

To see the source of the result, consider the firm's objective in a neighbourhood of its positive interior optimum and define

$$R(v,p;\mu) = p\beta(v,p;\mu)$$

and

$$L(v,p;\mu) = \log R(v,p;\mu).$$

The logarithmic transformation does not change the optimal price. From Lemma 2,

$$L_{pv}(v,p;\mu) = -\frac{u'(v)}{\kappa}\frac{\partial\beta(v,p;\mu)}{\partial p} > 0.$$

At the optimal price, $R_p = 0$, $R > 0$, and the strict second-order condition gives $R_{pp} < 0$. Therefore,

$$L_{pp} = \frac{R_{pp}}{R} - \left(\frac{R_p}{R}\right)^2 = \frac{R_{pp}}{R} < 0.$$

The implicit-function theorem therefore gives

$$\frac{\partial p^O(v;\mu)}{\partial v} = -\frac{L_{pv}\big(v, p^O(v;\mu);\mu\big)}{L_{pp}\big(v, p^O(v;\mu);\mu\big)} > 0.$$

Uniqueness of the maximiser implies that the local solutions obtained from the implicit-function theorem coincide on overlapping neighbourhoods. They therefore define a continuously differentiable optimal-price function on $(0,1)$. Continuity of the objective and local boundedness of the maximiser imply that this function is continuous on $[0,1]$. Since $\partial p^O(v;\mu)/\partial v > 0$ at every $v \in (0,1)$, $p^O(v;\mu)$ is strictly increasing on $[0,1]$.

2.3 Disclosure incentives

For a fixed coarse belief μ, define the advantage of disclosure by

$$\Delta(v;\mu) = u(v) - \pi^O(v;\mu).$$

A positive value of $\Delta(v;\mu)$ means that type v prefers disclosure.

Proposition 2 (Single crossing)**.** *Fix a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. The disclosure advantage $\Delta(v;\mu)$ is strictly increasing in v.*

Proof. See Appendix A. □

By the envelope theorem,

$$\frac{\partial \pi^O(v;\mu)}{\partial v} = p^O(v;\mu)\beta\big(v, p^O(v;\mu);\mu\big)\left[1 - \beta\big(v, p^O(v;\mu);\mu\big)\right]\frac{u'(v)}{\kappa}.$$

Using the pricing condition in Lemma 3 gives

$$\frac{\partial \pi^O(v;\mu)}{\partial v} = u'(v)\beta\big(v, p^O(v;\mu);\mu\big)\frac{V\big(p^O(v;\mu);\mu\big)}{\bar\beta\big(p^O(v;\mu);\mu\big)\left[1 - \bar\beta\big(p^O(v;\mu);\mu\big)\right]}.$$

For every interior solution,

$$V(p;\mu) < \bar\beta(p;\mu)\left[1 - \bar\beta(p;\mu)\right],$$

because

$$\bar\beta(p;\mu)\left[1 - \bar\beta(p;\mu)\right] - V(p;\mu) = \int_0^1 \beta(v,p;\mu)\left[1 - \beta(v,p;\mu)\right]\mu(dv) > 0.$$

Since $\beta(v,p;\mu) < 1$, it follows that

$$\frac{\partial \pi^O(v;\mu)}{\partial v} < u'(v).$$

Therefore,

$$\frac{\partial \Delta(v;\mu)}{\partial v} > 0.$$

Holding the coarse belief fixed, the optimised payoff from obfuscation increases with quality, but it increases more slowly than the full-information payoff. Higher-quality firms consequently have a stronger incentive to disclose.

2.4 Equilibrium

An interior ε-contaminated coarse-inference equilibrium consists of measurable firm strategies $d^* : [0,1] \to \{0,1\}$ and $p^* : [0,1] \to \mathbb{R}_+$.

Let

$$O^* = \{v \in [0,1] : d^*(v) = 0\}$$

denote the set of qualities such that firms do not disclose. The strategies constitute an interior equilibrium if the following conditions hold.

First, $0 < \lambda(O^*) < 1$, and the consumer's coarse belief is

$$\mu^* = (1-\varepsilon)\lambda(\cdot \mid O^*) + \varepsilon\lambda.$$

Thus, the endogenous anchor of the consumer's belief is consistent with the equilibrium disclosure strategy. The final belief is nevertheless contaminated and is not fully Bayesian.

Second, after observing any pair $(p, d = 0)$, the consumer uses μ^*, independently of p, and uses the continuous representative of the rational-inattention purchase rule characterised in Lemma 1. After observing $(p, d = 1)$, the consumer observes v and purchases if and only if $u(v) \geq p$.

Third, firms choose disclosure and price optimally. In particular, $p^*(v) = u(v)$ whenever $d^*(v) = 1$, while $p^*(v) = p^O(v;\mu^*)$ whenever $d^*(v) = 0$. Moreover, $d^*(v) = 1$ whenever $u(v) > \pi^O(v;\mu^*)$, and $d^*(v) = 0$ whenever $u(v) < \pi^O(v;\mu^*)$. Either decision is optimal when the two payoffs are equal.

Since firms are atomless, the deviation of an individual firm does not alter O^* or the belief μ^*. Full support, together with the use of the continuous representative, makes the consumer's response well defined at the quality of every equilibrium or deviating firm.

For every $x \in (0,1)$, define

$$\mu_x^{\varepsilon}(A) = (1-\varepsilon)\frac{\lambda(A \cap [0,x])}{\lambda([0,x])} + \varepsilon\lambda(A).$$

Proposition 3 (Threshold equilibrium)**.** *Every interior ε-contaminated coarse-inference equilibrium has a threshold disclosure rule. In particular, there exists $v^* \in (0,1)$ such that, up to the behaviour of the indifferent type,*

$$d^*(v) = \begin{cases} 0, & v < v^*, \\ 1, & v > v^*. \end{cases}$$

The equilibrium coarse belief is

$$\mu^* = \mu^{\varepsilon}_{v^*},$$

and the marginal type satisfies

$$u(v^*) = \pi^O\left(v^*; \mu^{\varepsilon}_{v^*}\right).$$

Conversely, any $v^* \in (0,1)$ satisfying this equality generates an interior threshold equilibrium when firms use their corresponding optimal prices and the regularity conditions in Assumption 2 hold.

Proof. See Appendix A. □

For the fixed belief $\mu^{\varepsilon}_{v^*}$, Proposition 2 implies that

$$v \longmapsto u(v) - \pi^O\left(v; \mu^{\varepsilon}_{v^*}\right)$$

is strictly increasing. If the marginal type is indifferent, every lower-quality type prefers obfuscation and every higher-quality type prefers disclosure. Whether the marginal type discloses has no effect on the belief because λ is atomless.

Proposition 3 characterises an interior threshold equilibrium but does not establish that one exists. Existence requires that

$$v \longmapsto u(v) - \pi^O\left(v; \mu^{\varepsilon}_{v}\right)$$

cross zero. Because the coarse belief changes with the candidate cut-off, the fixed-belief single-crossing result does not guarantee such a crossing. Similarly, uniqueness of the equilibrium cut-off requires additional restrictions on how the composition of the hidden category affects obfuscation profits.

2.5 Implications

Any interior equilibrium of the model delivers three main implications.

First, disclosure is ordered by quality. In any interior equilibrium, firms with sufficiently low quality obfuscate, while firms with sufficiently high quality disclose. This follows because, holding the equilibrium coarse belief fixed, the full-information payoff increases more rapidly with quality than the optimised payoff from obfuscation.

Second, in any interior equilibrium, prices increase with quality in both regimes. Under disclosure, $p^D(v) = u(v)$, which is strictly increasing. Under obfuscation, $p^O(v; \mu^*)$ is also strictly increasing in quality.

The model does not establish that the obfuscation-price schedule is flatter than the disclosure-price schedule. A flatter empirical price–quality relationship among non-disclosing products is compatible with the model, but it is not a general theoretical prediction without additional functional-form or curvature restrictions.

Third, the theoretical disclosure cut-off is a profit-indifference point, not the intersection of the disclosed and obfuscated price schedules. The marginal type satisfies

$$u(v^*) = p^O(v^*; \mu^*)\beta\left(v^*, p^O(v^*; \mu^*); \mu^*\right).$$

Since the purchase probability is strictly below one at an interior solution,

$$p^O(v^*; \mu^*) > u(v^*).$$

The price associated with obfuscation therefore does not coincide with the disclosed price at the theoretical cut-off. An empirical intersection between fitted price–quality schedules should be interpreted as a descriptive reduced-form object rather than as a structural estimate of v^*.

Finally, the log-odds representation

$$\log\left(\frac{\beta(v, p; \mu)}{1 - \beta(v, p; \mu)}\right) = \log\left(\frac{\bar{\beta}(p; \mu)}{1 - \bar{\beta}(p; \mu)}\right) + \frac{u(v) - p}{\kappa}$$

shows how attention costs affect the consumer's responsiveness to hidden quality. Holding the unconditional purchase probability fixed, the sensitivity of the log odds of purchase to quality is

$$\frac{\partial}{\partial v}\log\left(\frac{\beta(v, p; \mu)}{1 - \beta(v, p; \mu)}\right) = \frac{u'(v)}{\kappa}.$$

Thus, a larger cost κ of processing the information reduces the responsiveness of the log odds of purchase to quality. A corresponding global comparative static for the level of

demand or for the equilibrium cut-off does not follow without additional assumptions because attention costs also affect conditional purchase probabilities, prices, obfuscation profits, and the endogenous composition of the hidden category.

3 The empirical analysis

Data and Variable Construction

Our empirical analysis is based on a novel dataset comprising 145 homogeneous 250-gram packages of ground coffee (stock-keeping units, SKUs) sold in Italy. The data spans 28 months, from November 2020 to February 2023, covering 555 stores across 18 major supermarket chains.[2] To ensure a clean identification of price variation, the sample is restricted to single-product packages comparable in size and form, thereby mitigating confounding factors related to packaging heterogeneity.

Price data were collected through daily web-scraping of major supermarket retailers. Given that retail prices typically exhibit short-run stickiness, we aggregate these observations at the monthly level for each product–supermarket pair. This aggregation reduces high-frequency noise while preserving meaningful longitudinal variation for the econometric analysis.

For each SKU, we observe a comprehensive set of product characteristics: whether the coffee is organic, its origin (single-country versus blend), the brand type (manufacturer's brand versus private label), and the intended preparation method (espresso versus moka). To control for firm-level heterogeneity, we incorporate the annual revenues of the producing firms to capture potential economies of scale and reputation-based price premia. Furthermore, we include indicators for whether a product or firm is featured in the definitive Italian coffee guide by Godina and Illiano (2024). As this guide utilises blind tasting sessions and selects participants based on market availability, it provides a quality proxy that is largely independent of brand-related price effects and selection bias.

The primary variables for identifying obfuscation are the Arabica bean share and its disclosure status. While coffee quality is multi-dimensional (encompassing altitude, roasting techniques, and agronomic practices), the Arabica percentage is the most salient and widely understood quality attribute among consumers (Yeager et al., 2023; Freitas et al., 2024). Despite its importance, the Arabica share is subject to selective disclosure (Mendes et al., 2024), providing a natural channel for strategic obfuscation.

In our initial sample of 171 products, 77 explicitly disclose the Arabica share on the pack-

[2]The sample includes stores from Coop, Conad, Crai, Deco, Despar, Dok, Emi, Emisfero, Famila (Nord, Adriatica, and Sud), Galassia, Gospesa, Mercatò, Pan, Sole 365 and Todis.

aging, whereas 94 do not. A central feature of our empirical design lies in an active data recovery strategy to uncover the composition of this 'hidden pool' of obfuscated products. By consulting technical documentation on producers' websites and conducting direct enquiries via telephone and email, we successfully retrieved the exact Arabica share for 68 of the 94 non-disclosing products—representing a high recovery rate of approximately 72%. Composition data could not be verified for only the 26 remaining non-disclosing products, which constitute a minor fraction of our initial observations. This data recovery enables us to observe true underlying quality across both information regimes and construct a binary disclosure indicator, $d_i \in \{0, 1\}$, representing the firm's strategic decision to reveal ($d_i = 1$) or withhold ($d_i = 0$) attribute information.

We emphasise that these unrecovered products are highly unlikely to introduce meaningful selection bias or alter our core findings for two primary reasons. First, from an empirical standpoint, these products correspond almost exclusively to small, localised brands with highly restricted distribution, accounting for a negligible share of the supermarket panel (ranging from a mere 2 to 32 observations per product). Consequently, their exclusion carries virtually no statistical weight in our regressions. Second, from an econometric standpoint, any remaining systematic selection would likely operate in our favour: firms that refuse to disclose composition even to academic researchers are plausibly concentrated in the lowest-quality, most opaque tier of the market. Excluding these lowest-quality obfuscators artificially inflates the average quality of the observed non-disclosing sample. As a result, any selection bias would operate in a downward direction, meaning our estimates conservatively understate both the true extent of strategic obfuscation and the steepness of the price–quality gradient under full disclosure.

We report the summary price statistics for the three groups—disclosing, non-disclosing but recovered, and non-disclosing unrecovered—in Table B.1. Additionally, Fig. C.1 illustrates the price distributions across these same three groups. All subsequent figures and tables exclude the 26 products for which composition could not be recovered.

Table B.2 in Appendix B reports the summary statistics for the main variables used in the empirical analysis. The final estimation dataset consists of 55,249 product–supermarket–month observations, covering the 145 homogeneous 250-gram packages of ground coffee sold across Italy for which composition data are available (77 disclosing, representing 49% of the total observations, and 68 recovered products). The average price is €3.58, exhibiting substantial dispersion from €0.89 to €11.53.

The sample's average Arabica content is 65%, slightly exceeding the global average production share in 2023 (International Coffee Organization, 2023), and spans the full range from 0% to 100%. Organic and single-origin coffees represent niche segments, accounting for

6% and 2% of products, respectively, whereas private-label products are relatively common (24%). Finally, most products are ground for moka pots, reflecting traditional Italian home consumption habits, though espresso-ground coffee also maintains a sizeable market share (30%).

Firm size exhibits considerable heterogeneity, ranging from small local roasteries with annual revenues of approximately €300,000 to multinational firms and supermarket chains with turnovers in the billions of euros. While a majority of producers are mentioned in the coffee guide (70%), only a small fraction of individual products receives such recognition (12%).

Table B.3 in Appendix B presents the correlation matrix for the variables included in the empirical analysis. The only correlation exceeding the conventional threshold of 0.8 is that between the private-label dummy and the indicator for whether the producing firm is mentioned in the coffee guide. Given this high collinearity, and since the private-label dummy is also more strongly correlated with firm size (measured by the logarithm of revenues), we exclude the private-label variable from the baseline econometric specifications.

Table B.4 in Appendix B compares average Arabica content and prices between products that disclose the blend composition on the label and those that do not. Products with disclosed composition exhibit a markedly higher Arabica share—89% on average, compared to 42% for products that obfuscate this information. Prices are also substantially higher in the disclosure group (€4.33 versus €2.88), corresponding to an increase of approximately 50%. These unconditional differences are consistent with the hypothesis that disclosure is associated with higher-quality positioning and provide preliminary evidence that selective disclosure may play a meaningful role in price formation.

Fig. 1 shows that the entire price distribution for disclosing products is shifted to the right relative to non-disclosing products, suggesting that the higher prices observed in the summary statistics reflect a systematic difference in pricing rather than the presence of a few outliers.

Figure 1: Kernel density of log-price by disclosure status.

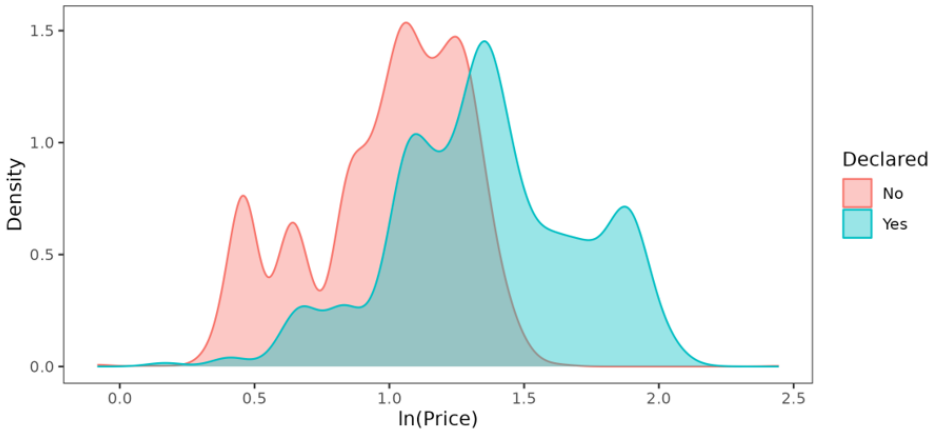

Fig. 2 shows the distribution of the share of Arabica for packages with declared and undeclared composition in the label. The distribution is already providing suggestive evidence that the choice to disclose information on quality is strongly driven by quality itself.

Figure 2: Histogram of Arabica composition by disclosure status

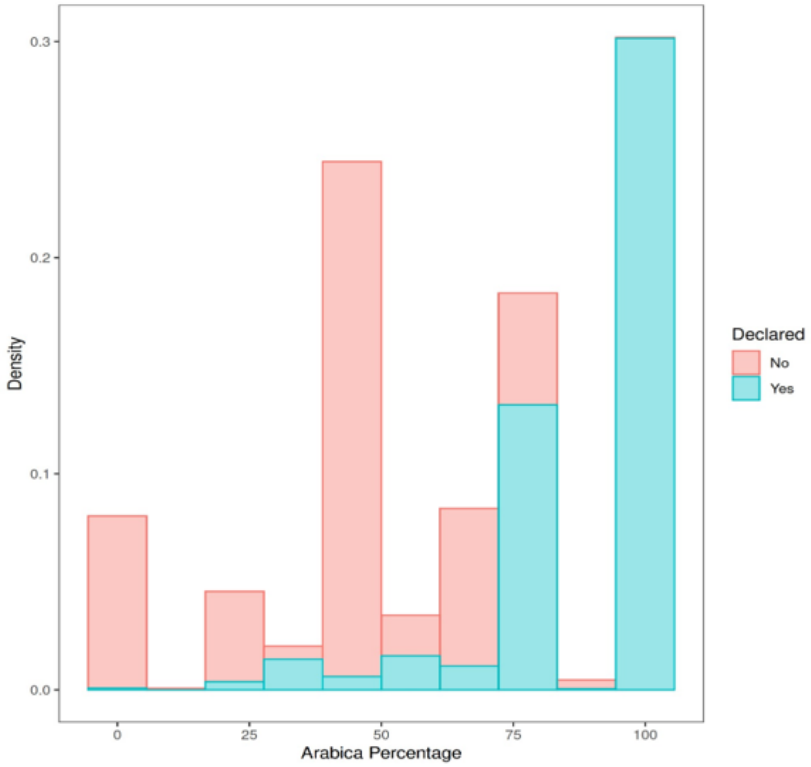

3.1 Econometric Analysis

To provide direct evidence on the strategic nature of disclosure, we begin the empirical analysis by focusing on the cross-section of 145 distinct coffee products and estimating a Probit model with robust standard errors of the decision to report the Arabica share on the label (Table 1). In this specification, the dependent variable equals one if the product explicitly discloses the percentage of Arabica. This exercise is closely aligned with the theoretical model, which predicts that disclosure is not random but instead depends systematically on product quality. The results strongly support this prediction: the Arabica share is a large and highly significant predictor of disclosure, indicating that higher-quality products are substantially more likely to reveal their composition. This finding provides direct evidence of selective disclosure and is consistent with a separating equilibrium in which firms with sufficiently high quality choose to disclose, while lower-quality firms pool by withholding information.

Table 1: Probit: decision to declare the % of Arabica (marginal effects)

	(1)	(2)
VARIABLES	Declared	Declared
Arabica	0.732***	0.763***
	(0.0570)	(0.0531)
Organic	0.264**	0.230**
	(0.119)	(0.101)
Single_country	0.0508	-0.0224
	(0.133)	(0.120)
Guide		-0.106*
		(0.0605)
Guide_score		0.0876
		(0.0837)
Espresso		-0.238***
		(0.0579)
Pseudo R2	0.34	0.42
Observations	145	145

Note: *The table reports Robust Standard Errors and shows Marginal Effects for the Probit specification. *** p<0.01, ** p<0.05, * p<0.1.*

Additional covariates further support this interpretation. Most notably, organic products exhibit a strong, positive, and statistically significant propensity to disclose their blend composition. In our sample, 99.7% of observations of organic products disclose their Arabica share.

Conversely, products intended for espresso preparation are significantly less likely to report Arabica content. This is consistent with the sensory requirements of Italian espresso: extraction relies on crema formation, which is typically enhanced by blends containing a non-

negligible share of Robusta (Illy and Viani, 2005), thereby reducing the strategic incentive to signal Arabica content. As a result, the Arabica share may be less salient as a quality signal in this segment, and firms can rely on alternative descriptors (such as body, intensity, or crema) rather than explicitly disclosing composition.

Importantly, this variable also plays a key role in our identification strategy. Since espresso orientation affects firms' incentives to disclose but is less likely to directly influence prices once observable quality and product characteristics are controlled for (see the next section), it provides a natural source of variation for addressing the potential endogeneity of disclosure in the subsequent hedonic price regressions via a control-function approach. To ensure the validity of our exclusion restriction, we must establish that the Espresso indicator does not directly affect the price of coffee through either supply-side costs or demand-side price discrimination. On the supply side, the industrial cost differences between grinding beans coarser (for Moka) or finer (for Espresso) are negligible; consequently, designating a blend for espresso does not inherently alter its marginal cost of production. On the demand side, the boundaries of our sample are explicitly restricted to traditional ground coffee, deliberately excluding single-serve pods, capsules, and whole beans. This restriction is vital because those formats cater to fundamentally different consumer segments with distinct pricing mechanics: capsules command a steep price premium for convenience and packaging, while whole beans attract a connoisseur segment willing to pay a premium for peak freshness. By stripping away both the extreme convenience of capsules and the connoisseurship of whole beans, we isolate a highly homogeneous baseline market. Within this tightly defined segment, consumers of Moka- and Espresso-ground coffee belong to broadly similar demographics shopping in the same supermarket aisles, precluding systematic price discrimination based purely on the grind size. Therefore, the Espresso indicator influences price almost exclusively through its functional requirement for Robusta beans to generate crema, which directly alters the Arabica share and the firm's subsequent incentive to obfuscate. Overall, this first-stage analysis establishes that disclosure is an endogenous and strategic choice, thereby motivating the subsequent analysis of how it interacts with the pricing of quality.

We estimate hedonic price regressions in which the dependent variable is the natural logarithm of the price of a 250-gram package. To test the effect of obfuscation, we adopt two complementary empirical strategies.

In the first one, the empirical relationship is estimated separately for the disclosing and obfuscating regimes to allow for parameter heterogeneity across different information envir-

onments.

$$
\begin{aligned}
d = 1 \text{ (Disclosure):} \quad & \ln(p_{it}) = \alpha^D + \beta_A^D \text{Arabica}_{it} + \mathbf{B}^D \mathbf{X}_{it} + \mathbf{\Gamma}^D \mathbf{Z}_{it} + \varepsilon_{it}^D \\
d = 0 \text{ (Obfuscation):} \quad & \ln(p_{it}) = \alpha^O + \beta_A^O \text{Arabica}_{it} + \mathbf{B}^O \mathbf{X}_{it} + \mathbf{\Gamma}^O \mathbf{Z}_{it} + \varepsilon_{it}^O
\end{aligned} \tag{1}
$$

In these specifications, $\ln(p_{it})$ represents the log-price of product i at time t, while Arabica_{it} denotes the observed Arabica shares. The vector $\mathbf{X}_{it}$ compactly represents the primary set of product and firm attributes[3] with $\mathbf{B}$ as the corresponding vector of coefficients. Finally, $\mathbf{Z}_{it}$ is a vector capturing time trends and regional and store-level fixed effects, allowing $\mathbf{\Gamma}$ to account for unobserved geographic and retail-specific characteristics that may influence pricing independently of product quality.

Table 2 shows the estimates under disclosure ($d = 1$), while Table 3 presents those for obfuscation ($d = 0$). We then formally compare the estimated coefficients across subsamples using chi-squared tests of equality in Table 4.

Table 2: Baseline Regression: Only Firms that Declare

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price
Constant	0.0645	0.0907	0.0804	0.117	0.362*	0.373
	(0.224)	(0.237)	(0.187)	(0.209)	(0.212)	(0.233)
Organic	0.299***	0.306***	0.256***	0.262***	0.242***	0.243***
	(0.0986)	(0.0941)	(0.0780)	(0.0757)	(0.0743)	(0.0729)
Single_country	-0.0320	-0.0414	-0.0240	-0.0452	-0.0550	-0.0852
	(0.0926)	(0.103)	(0.0835)	(0.0942)	(0.0781)	(0.0836)
Arabica	1.128***	1.049***	1.027***	0.955***	1.038***	0.957***
	(0.243)	(0.190)	(0.220)	(0.177)	(0.210)	(0.173)
Guide	0.399***	0.455***	0.427***	0.465***	0.413***	0.449***
	(0.110)	(0.124)	(0.0881)	(0.0976)	(0.0829)	(0.0892)
Ln_Revenues_mln		0.00521		0.00636		0.00798
		(0.0145)		(0.0141)		(0.0142)
Guide_score		-0.185*		-0.159		-0.170
		(0.107)		(0.120)		(0.122)
Observations	26,822	26,822	26,822	26,822	26,822	26,822
R-squared	0.532	0.569	0.614	0.639	0.659	0.685
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *Standard errors clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1.*

In the alternative specification, presented in Table 5, we estimate a pooled regression across the full sample to directly test for differences in the price–quality relationship between information environments.

[3]Specifically organic status, single-country origin, guide scores, and log-revenues.

Table 3: Baseline Regression: Only Firms that Obfuscate

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price
Constant	0.432***	0.378**	0.577***	0.507***	0.417***	0.290
	(0.0771)	(0.185)	(0.115)	(0.187)	(0.0855)	(0.205)
Organic	-0.276***	-0.264***	-0.246***	-0.215*	-0.258***	-0.226***
	(0.0553)	(0.0696)	(0.0901)	(0.108)	(0.0622)	(0.0780)
Single_country	0.301***	0.330***	0.235***	0.278***	0.188***	0.232**
	(0.0821)	(0.109)	(0.0772)	(0.0943)	(0.0684)	(0.0934)
Arabica	0.513***	0.524***	0.522***	0.539***	0.488***	0.511***
	(0.106)	(0.118)	(0.0998)	(0.109)	(0.0988)	(0.108)
Guide	0.345***	0.348***	0.335***	0.344***	0.337***	0.349***
	(0.0558)	(0.0552)	(0.0515)	(0.0511)	(0.0486)	(0.0498)
Ln_Revenues_mln		0.00567		0.00977		0.0122
		(0.0153)		(0.0149)		(0.0160)
Guide_score		0.0917		0.120		0.141
		(0.0933)		(0.0991)		(0.108)
Observations	28,427	28,427	28,427	28,427	28,427	28,427
R-squared	0.652	0.657	0.686	0.695	0.720	0.731
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *Standard errors clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1.*

Table 4: χ^2 test for difference in the coefficients of Arabica

Regressions	(1)	(2)	(3)	(4)	(5)	(6)
H0: $\beta(Const.\|d=0) = \beta(Const.\|d=1)$						
χ^2 (1)	2.41	0.92	5.17	1.95	0.06	0.07
$Prob > \chi^2$	0.1204	0.3381	0.0230	0.1629	0.8063	0.7869
H0: $\beta(Arabica\|d=0) = \beta(Arabica\|d=1)$						
χ^2 (1)	5.42	5.58	4.42	4.01	5.75	4.90
$Prob > \chi^2$	0.0199	0.0182	0.0356	0.0451	0.0165	0.0268

Note: χ^2 *tests are computed using covariance matrices clustered at the product level. The coefficients compared refer to Tables 2 and 3.*

$$\ln(p_{it}) = \alpha + \beta_1 \text{Arabica}_{it} + \beta_2(\text{Arabica}_{it} \times d_{it}) + \theta d_{it} + \mathbf{BX}_{it} + \mathbf{\Gamma Z}_{it} + \varepsilon_{it} \tag{2}$$

The coefficient β_1 captures the baseline effect of Arabica share on log-prices for the obfuscating group, while the interaction term β_2 identifies the differential premium associated with quality when it is transparently disclosed, and θ, instead, captures the effect of declaring the composition. This approach allows for a formal test of whether strategic obfuscation attenuates the sensitivity of prices to underlying product characteristics. As in the previous

specification, vector $\mathbf{X}_{it}$ controls for product and firm attributes (including organic status, origin, guide scores, and log-revenues), while $\mathbf{Z}_{it}$ accounts for regional and store-level fixed effects.

Table 5: Baseline Regression: All Firms

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price
Constant	0.464***	0.466***	0.476***	0.512***	0.495***	0.536***
	(0.0759)	(0.131)	(0.0835)	(0.136)	(0.0967)	(0.154)
Organic	0.266***	0.256***	0.233***	0.221***	0.215***	0.202***
	(0.0899)	(0.0870)	(0.0792)	(0.0766)	(0.0759)	(0.0744)
Single_country	-0.00341	-0.00773	-0.0120	-0.0134	-0.0318	-0.0337
	(0.0791)	(0.0813)	(0.0735)	(0.0777)	(0.0678)	(0.0712)
Declared	-0.411*	-0.373*	-0.371*	-0.354*	-0.377*	-0.354*
	(0.217)	(0.201)	(0.200)	(0.189)	(0.197)	(0.190)
Arabica	0.470***	0.466***	0.462***	0.457***	0.467***	0.464***
	(0.113)	(0.115)	(0.106)	(0.108)	(0.104)	(0.106)
Arabica × Declared	0.654**	0.628**	0.623**	0.612**	0.626**	0.607***
	(0.284)	(0.252)	(0.264)	(0.236)	(0.255)	(0.232)
Guide	0.367***	0.381***	0.369***	0.371***	0.355***	0.356***
	(0.0598)	(0.0645)	(0.0547)	(0.0592)	(0.0520)	(0.0560)
Ln_Revenues_mln		-0.00131		-0.00428		-0.00459
		(0.0101)		(0.0104)		(0.0109)
Guide_score		-0.101		-0.0861		-0.0936
		(0.0938)		(0.101)		(0.103)
Observations	55,249	55,249	55,249	55,249	55,249	55,249
R-squared	0.691	0.698	0.717	0.721	0.735	0.740
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *Standard errors clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1.*

In our baseline specifications, we cluster standard errors at the product level to account for the repeated observation of the same SKU across stores and over time, allowing for arbitrary within-product correlation while maintaining independence across products, as recommended by MacKinnon et al. (2023). For completeness, Tables B.5, B.6 and B.8 in Appendix B replicate, respectively, Tables 2, 3 and 5, but using Robust Standard Errors; results are unchanged.

Tables 2, 3 and 5 implement complementary strategies that differ in their identifying assumptions and statistical properties. The subsample regressions allow for a flexible comparison of price–quality relationships across disclosure regimes without imposing a common functional form, enhancing transparency and ease of interpretation at the cost of reduced statistical efficiency. By contrast, the pooled specification exploits the full sample and provides

a formal test of heterogeneous effects through interaction terms, improving efficiency while relying on stronger parametric assumptions. Comparing results across these specifications follows standard practice in applied econometrics to assess robustness (Lu and White, 2014), and the consistency of the findings strengthens their credibility.

Tables 2 and 3 show that Arabica content is statistically significantly and positively associated with prices in both subsamples. However, the magnitude of the Arabica coefficient is substantially larger for products that disclose this information, indicating that quality is more strongly capitalised into prices when it is observable to consumers. Even under obfuscation, consumers are able to extract some information about product quality from noisy signals, such as past experience, brand reputation, or costly search. As a result, quality differences are still partially reflected in prices, albeit to a much smaller extent than under full disclosure. At the same time, the intercept is higher for products that obfuscate the Arabica share, implying higher baseline prices at low levels of Arabica content. In other words, when quality is low, obfuscation allows firms to sustain higher prices by preventing consumers from accurately assessing product quality. Table 5 formally confirms these differences by rejecting the null hypothesis of equal Arabica coefficients across subsamples and allowing for differences in intercepts.

Together, these estimates imply two upward-sloping price–quality schedules: a disclosure regime characterised by a lower intercept and a steeper slope, and an obfuscation regime characterised by a higher intercept and a flatter slope. The pooled regressions in Table 5 confirm this pattern and imply that the two schedules intersect slightly below a 60% share of Arabica (58% using the coefficients from the linear regression (6) in Table 5, with statistical significance $p < 0.001$ assessed using the delta method, see Fig. 3). Below this intersection, the fitted price is higher under non-disclosure; above it, the fitted price is higher under disclosure.[4]

[4]These price comparisons do not establish which disclosure decision is more profitable, so the intersection does not identify the model's disclosure threshold.

Figure 3: Estimated price-quality ratio by disclosure status

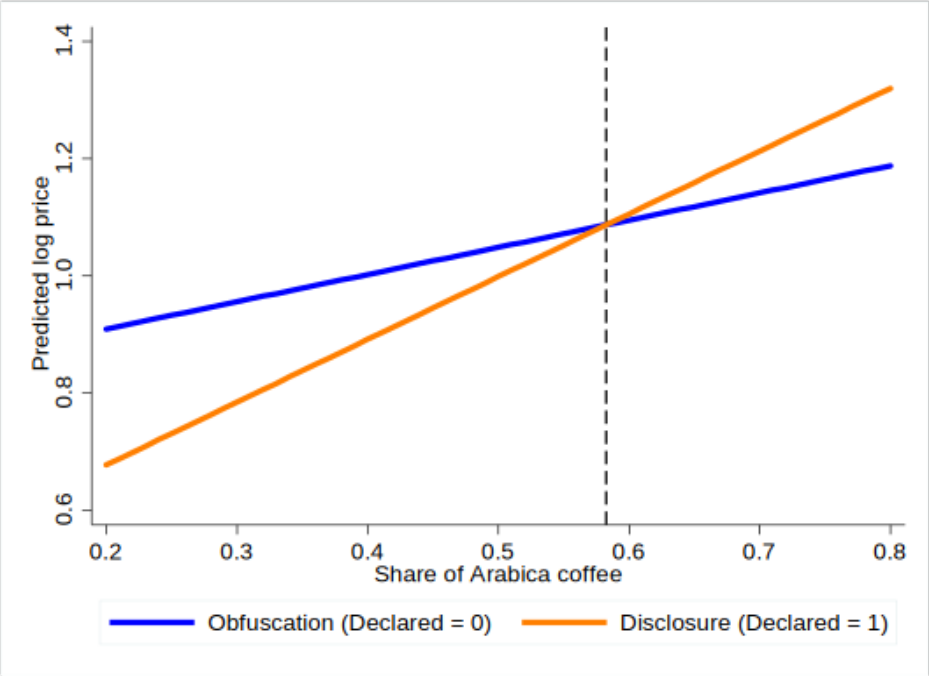

Note: *Plot based on Regression 6, Table 5*

In Tables 6 and 7, we repeat the analysis presented in Table 5 by splitting the sample according to the product's brewing destination ($s \in \{\text{Espresso, Moka}\}$).

$$\ln(p_{it}) = \alpha^s + \beta_1^s \text{Arabica}_{it} + \beta_2^s (\text{Arabica}_{it} \times d_{it}) + \theta^s d_{it} + \mathbf{B}^s \mathbf{X}_{it} + \mathbf{\Gamma}^s \mathbf{Z}_{it} + \varepsilon_{it}^s \qquad (3)$$

Specifically, we estimate Eq. (3) separately for espresso-specific and moka-specific coffee packages to verify whether the strategic pricing of quality and disclosure patterns are consistent across different consumption methods. The results, which remain largely similar to the full-sample estimates, confirm that the observed interaction between disclosure and quality remains robust even when accounting for the technical characteristics of the grind.

To address concerns related to limited overlap in the quality distribution across the two groups, we employ the overlap weighting estimator proposed by Li et al. (2018). In the first step, in line with the Probit model in Regression 2 of Table 1, we estimate each product's propensity score, namely the predicted probability that a product discloses its Arabica composition, conditional on the observed covariates (Arabica share, Organic, Single-country, Guide variables and Espresso orientation). Let $\hat{\pi}$ denote the estimated propensity score. Products that disclose receive weight $1 - \hat{\pi}$, whereas products that do not disclose receive weight $\hat{\pi}$. Intuitively, products whose disclosure decision is almost deterministic – for instance, products with very high Arabica content that are almost certain to disclose,

Table 6: Split Sample: Only Espresso

VARIABLES	(1) Ln_Price	(2) Ln_Price	(3) Ln_Price	(4) Ln_Price	(5) Ln_Price	(6) Ln_Price
Constant	0.559***	0.818***	0.655***	0.951***	0.368	0.857***
	(0.144)	(0.155)	(0.137)	(0.158)	(0.266)	(0.256)
Organic	0.393***	0.302***	0.317***	0.206***	0.275***	0.0867
	(0.0641)	(0.0669)	(0.0644)	(0.0595)	(0.0807)	(0.0896)
Single_country	0.0869	0.127**	0.0783	0.137**	0.00355	0.0734
	(0.0560)	(0.0622)	(0.0541)	(0.0607)	(0.0793)	(0.0839)
Declared	-0.515	-0.614**	-0.463	-0.588*	-0.475	-0.573*
	(0.314)	(0.287)	(0.303)	(0.305)	(0.307)	(0.296)
Arabica	0.406**	0.352*	0.380**	0.296*	0.378**	0.284*
	(0.200)	(0.198)	(0.175)	(0.168)	(0.151)	(0.144)
Arabica × Declared	0.623	0.730**	0.585	0.726*	0.614*	0.739**
	(0.391)	(0.363)	(0.374)	(0.370)	(0.357)	(0.344)
Guide	0.263***	0.123	0.262***	0.108	0.209**	-0.0277
	(0.0960)	(0.147)	(0.0867)	(0.104)	(0.0969)	(0.128)
Ln_Revenues_mln		-0.0254*		-0.0294***		-0.0459***
		(0.0129)		(0.0110)		(0.0126)
Guide_score		0.0904		0.122		0.118
		(0.115)		(0.128)		(0.140)
Observations	16,665	16,665	16,665	16,665	16,665	16,665
R-squared	0.453	0.475	0.505	0.535	0.580	0.615
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *Standard errors clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1.*

Table 7: Split Sample: Only Moka

VARIABLES	(1) Ln_Price	(2) Ln_Price	(3) Ln_Price	(4) Ln_Price	(5) Ln_Price	(6) Ln_Price
Constant	0.446***	0.340***	0.419***	0.344***	0.507***	0.426***
	(0.0572)	(0.108)	(0.0627)	(0.0953)	(0.0659)	(0.0968)
Organic	0.195	0.194	0.176*	0.177*	0.164*	0.166*
	(0.126)	(0.125)	(0.106)	(0.105)	(0.0972)	(0.0968)
Single_country	-0.0647	-0.0571	-0.0803	-0.0771	-0.0385	-0.0320
	(0.255)	(0.270)	(0.196)	(0.210)	(0.147)	(0.158)
Declared	-0.324	-0.251	-0.269	-0.221	-0.237	-0.192
	(0.221)	(0.194)	(0.208)	(0.183)	(0.207)	(0.190)
Arabica	0.519***	0.531***	0.493***	0.506***	0.501***	0.516***
	(0.142)	(0.146)	(0.120)	(0.123)	(0.0969)	(0.0988)
Arabica × Declared	0.605**	0.555**	0.574**	0.538**	0.538**	0.503**
	(0.267)	(0.227)	(0.250)	(0.217)	(0.237)	(0.211)
Guide	0.365***	0.384***	0.366***	0.379***	0.352***	0.366***
	(0.0888)	(0.0938)	(0.0726)	(0.0773)	(0.0587)	(0.0623)
Ln_Revenues_mln		0.0104		0.00754		0.00803
		(0.0113)		(0.0110)		(0.0105)
Guide_score		-0.120		-0.0804		-0.0639
		(0.0984)		(0.112)		(0.116)
Observations	38,584	38,584	38,584	38,584	38,584	38,584
R-squared	0.775	0.781	0.802	0.805	0.825	0.827
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *Standard errors clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1.*

or products with very low Arabica content that are very unlikely to disclose – receive very little weight because they have very limited comparable counterpart in the opposite disclosure regime. Conversely, products with intermediate propensity scores receive the largest weights, since they could plausibly belong to either group. In the second step, these overlap weights are used in the hedonic price regression. The resulting estimates therefore place most of the weight on products lying in the region of common empirical support, substantially reducing reliance on extrapolation from the extremes of the Arabica distribution.

Figure 4: Propensity-score distributions for declared and non-declared products

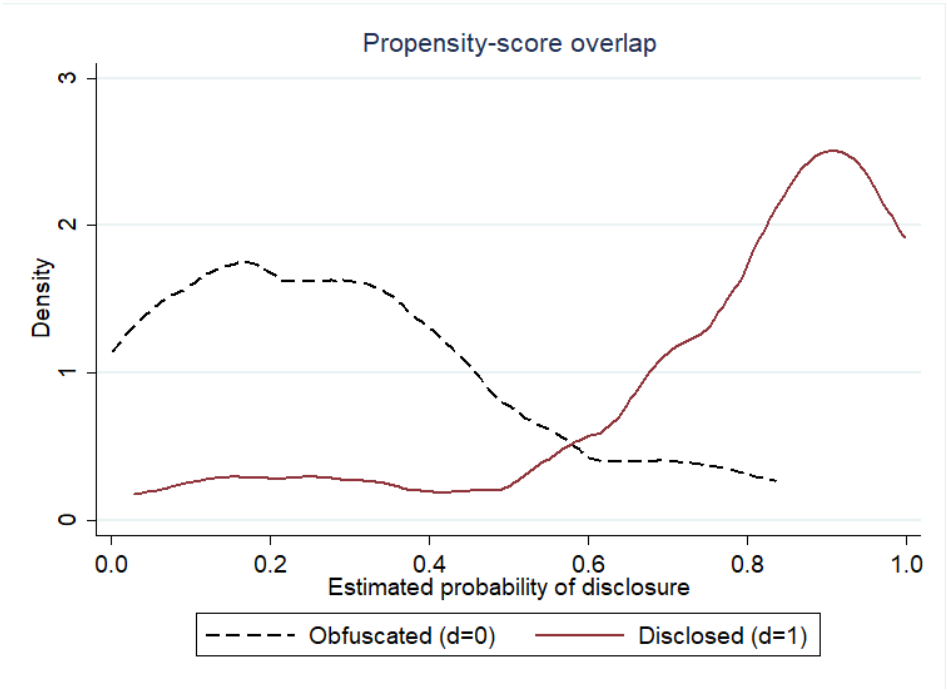

Fig. 4 reports the estimated propensity-score distributions for declared and non-declared products. Although the two groups remain clearly separated, a non-negligible overlap emerges for products with intermediate disclosure probabilities, precisely the region on which the overlap-weighting estimator concentrates estimation. The overlap-weighted estimates confirm the main findings. Table 8 shows that the overlap-weighted estimates are broadly consistent with those of Table 5. More specifically, comparing Regressions 6, the Arabica premium among non-disclosing products is virtually unchanged (0.460 in Table 8 versus 0.464 in Table 5), while the additional premium associated with disclosure declines from 0.607 in Table 5 to 0.372 in Table 8, suggesting that part of the larger interaction estimated in the baseline specification reflects comparisons involving products located in the tails of the Arabica distribution. Nevertheless, the price-quality relationship remains substantially steeper for products that disclose their composition: the implied Arabica slope for declaring products

equals 0.832, approximately 82% larger than that for non-disclosing products (0.459), and the equality of the two slopes is still rejected (F = 5.56, p = 0.019). Fig. 5 provides a graphical representation of the estimates of Regression 6 in Table 8: the fitted price-quality relationship remains visibly steeper for disclosing products, with the two functions intersecting at approximately 59% Arabica. Overall, the overlap-weighted estimates suggest that the steeper price-quality gradient for disclosing products is not an artefact of limited common support. Rather, it persists within the subset of products for which disclosure and non-disclosure are empirically comparable.

Figure 5: Estimated price-quality ratio by disclosure status

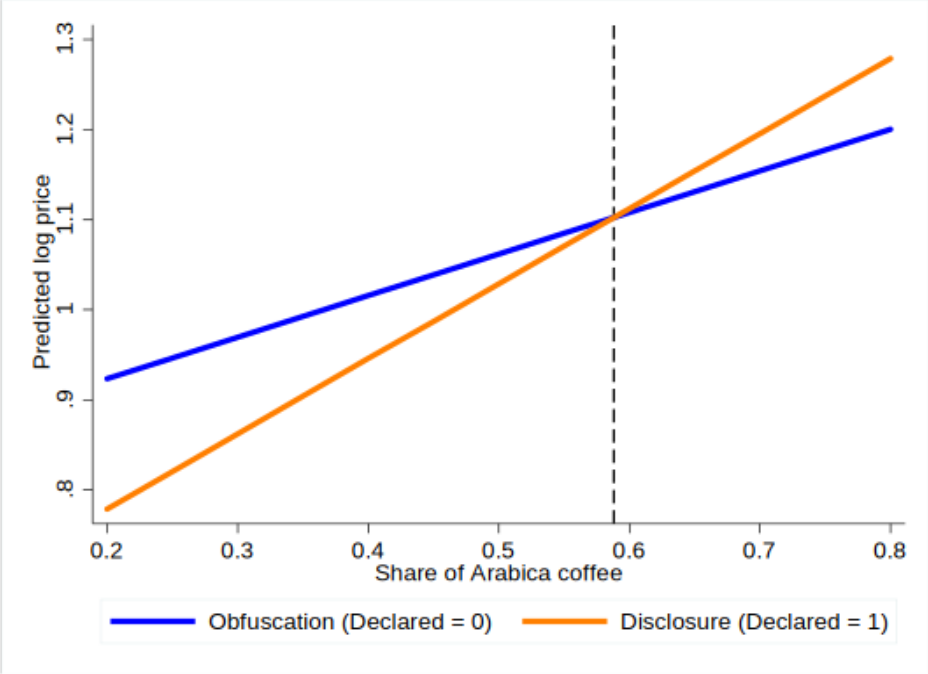

Note: Plot based on Regression 6, Table 8. Overlap-weighted estimator (Li et al., 2018).

As a further robustness check, Fig. 6 complements this analysis using a Local Linear Regression (LLR) (degree 1) that focuses on the region of common support (Arabica shares between 20% and 80%). We use a bandwidth of 0.2, which represents a compromise between achieving local identification and maintaining sufficient statistical power across the common support. This approach provides a flexible, nonparametric estimate of the price–quality relationship without imposing a global functional form, thereby allowing the price–quality relationship to be estimated locally using only comparable observations across disclosure regimes. Unlike OLS, which assumes a constant effect across the entire distribution and is therefore highly sensitive to extreme observations (outliers) and functional form specifications, LLR

Table 8: Overlap-weighted Regression

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price
Constant	0.492***	0.522***	0.474***	0.591***	0.454***	0.559***
	-0.0703	-0.134	-0.0591	-0.121	-0.0916	-0.134
Organic	0.438***	0.414***	0.370***	0.331***	0.345***	0.312***
	-0.106	-0.114	-0.0995	-0.108	-0.097	-0.102
Single_country	0.0771	0.0632	0.0582	0.0532	0.0487	0.0436
	-0.0633	-0.0665	-0.0616	-0.0676	-0.0639	-0.0699
Declared	-0.235*	-0.254**	-0.182	-0.228**	-0.187*	-0.219**
	-0.122	-0.12	-0.111	-0.103	-0.102	-0.0957
Arabica	0.466***	0.451***	0.485***	0.458***	0.484***	0.461***
	-0.105	-0.108	-0.0941	-0.0996	-0.0885	-0.0934
Arabica × Declared	0.366*	0.419**	0.311*	0.383**	0.313*	0.372**
	-0.188	-0.192	-0.172	-0.171	-0.163	-0.158
Guide	0.333***	0.344***	0.328***	0.325***	0.331***	0.328***
	-0.0472	-0.0526	-0.043	-0.0472	-0.0394	-0.044
Ln_Revenues_mln		-0.00373		-0.0115		-0.0106
		-0.00983		-0.00934		-0.00918
Guide_score		-0.0883		-0.102		-0.111
		-0.0762		-0.08		-0.0818
Observations	55249	55249	55249	55249	55249	55249
R-squared	0.602	0.612	0.643	0.656	0.677	0.69
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *Standard errors clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1.*

estimates the slope locally.

Indeed, Fig. 6 shows that, within the common support between 20% and 80% Arabica, local linear estimates continue to display a steeper price-quality gradient for disclosed products, with the two schedules crossing in the mid-to-high quality range. This suggests that the main result is not mechanically driven by the mass of 100% Arabica products, but reflects a broader difference in how quality is capitalised into prices across disclosure regimes.

Figure 6: Estimated price-quality ratio by disclosure status

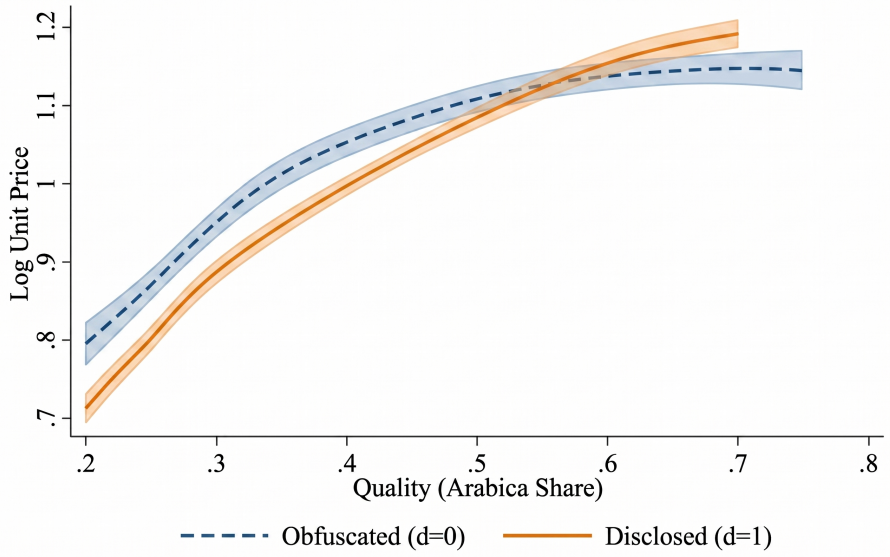

Note: Local linear regression (degree 1, bandwidth 0.2); shaded areas denote 95% confidence intervals

To verify that the estimated interaction between Arabica content and disclosure is not driven by a small number of influential products, we perform a leave-one-product-out robustness exercise. Specifically, we re-estimate the baseline specification (Table 5, Column 6) 145 times, excluding one product from the sample in each iteration. Figure 5 reports the resulting estimates of the Arabica × Declared coefficient, alongside their 95% confidence intervals. The estimated interaction effect remains consistently positive and economically similar across all specifications, retaining statistical significance in 144 out of 145 iterations. Even in the single specification where the 95% confidence interval spans zero, the point estimate remains positive and closely aligned with the baseline. This confirms that the steeper price–quality relationship observed under the disclosure regime is a systemic market feature, rather than an artefact driven by any single product or firm-specific outlier.

Figure 7: Leave-One-Product-Out Coefficient Estimates

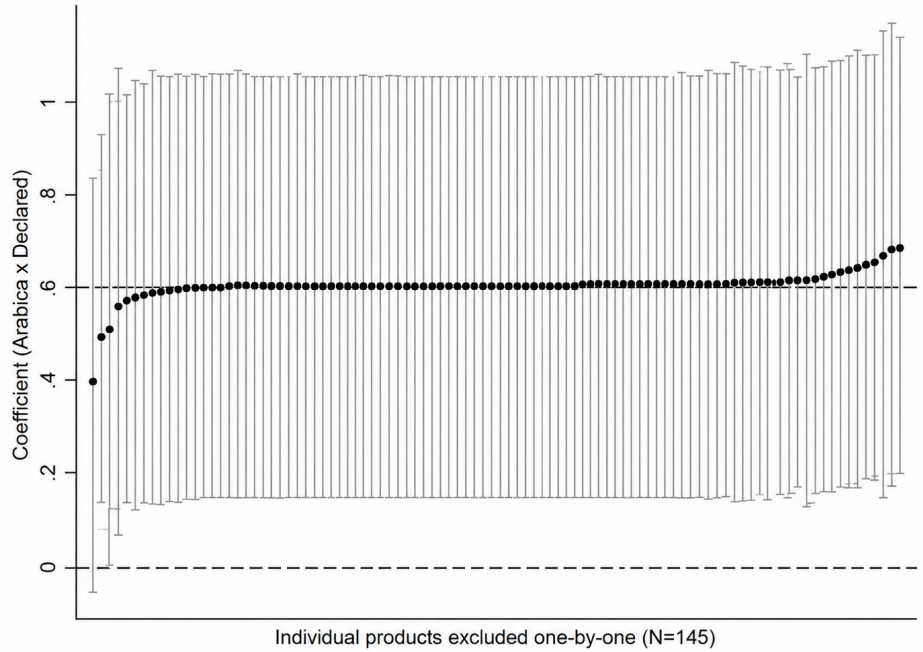

Note: The figure reports leave-one-product-out estimates of the coefficient on Arabica × Declared from the baseline hedonic price regression (Table 5, Column 6), ordered by coefficient size. Each point corresponds to a separate regression in which all observations for a single product are excluded. Vertical bars denote 95 % confidence intervals. The solid horizontal line represents the baseline coefficient from the full sample, while the dashed line marks zero. Regressions are sorted along the x-axis by the magnitude of the estimated coefficient.

Another potential concern is that our empirical specification does not directly control for other product characteristics, such as roast intensity or cultivation altitude, which may independently affect prices and correlate with disclosure decisions. We believe this concern is limited. Unlike Arabica share, which serves as a vertical quality signal, roast intensity primarily reflects horizontal taste heterogeneity and preparation style rather than a monotonic relationship with perceived quality. In the Italian market, darker roasts are typically associated with traditional espresso blends, while lighter roasts are more common among specialty coffees, with neither being unambiguously perceived as inherently higher quality by supermarket consumers. Moreover, roast intensity is rarely reported in a consistent, comparable metric on supermarket packaging, suggesting that firms do not treat it as a primary quality signal. Finally, omitted vertical attributes such as cultivation altitude are likely positively correlated with both Arabica content and the propensity to disclose. This implies that disclosing products may also be superior along other unobserved quality dimensions, reinforcing our core interpretation that disclosure is associated with genuinely higher-quality products.

Beyond these specific attributes, a broader econometric concern in our setting is that the

disclosure decision itself is strategic and may therefore be endogenous to aggregate unobserved product characteristics or brand positioning that also affect prices. In particular, firms with higher unobserved quality may be both more likely to disclose and able to charge higher prices, potentially biasing standard hedonic price coefficients. To systematically address this latent endogeneity, we implement a control-function approach. Specifically, we first estimate a probit model of the decision to disclose the Arabica share (as in Regression (2) of Table 1) and then include the corresponding generalised residual in the price equation. This procedure allows us to control for the unobserved component of the disclosure decision that may be correlated with pricing.

Identification relies on the product's brewing destination (Espresso vs Moka). While its role in shaping disclosure decisions is documented in Table 1, its exclusion from the price equation is motivated by the fact that, within the ground coffee segment, Espresso and Moka products are unlikely to command systematically different willingness to pay once observable characteristics are controlled for. Consistent with this interpretation, including the Espresso indicator in the price regressions does not yield statistically significant effects (results available upon request), and it is therefore excluded from the baseline specifications.

Table 9 reports the second-stage estimates. Arabica content is positively associated with prices for non-disclosing products, and the interaction between Arabica and disclosure is positive and statistically significant, implying a substantially steeper price–quality relationship under disclosure. The coefficient on the control-function term is small and statistically insignificant across specifications, suggesting that, conditional on observable characteristics, there is no strong evidence of residual selection on unobservables affecting prices. We also explored an instrumental variables approach; however, instrumenting both the disclosure indicator and its interaction with Arabica content leads to imprecise estimates. We therefore rely on the control-function specification and interpret the results as reflecting equilibrium relationships rather than causal treatment effects.

Table 9: Regressions with control function, all firms

VARIABLES	(1) Ln_Price	(2) Ln_Price	(3) Ln_Price	(4) Ln_Price	(5) Ln_Price	(6) Ln_Price
Constant	0.465***	0.493***	0.481***	0.537***	0.493***	0.558***
	(0.0758)	(0.120)	(0.0831)	(0.125)	(0.0971)	(0.142)
Organic	0.243***	0.235***	0.217***	0.204**	0.197**	0.184**
	(0.0916)	(0.0897)	(0.0823)	(0.0795)	(0.0803)	(0.0796)
Single_country	0.00963	0.00599	-0.00225	-0.00128	-0.0224	-0.0210
	(0.0767)	(0.0803)	(0.0721)	(0.0771)	(0.0666)	(0.0703)
Declared	-0.224	-0.243	-0.236	-0.250	-0.238	-0.247
	(0.331)	(0.335)	(0.306)	(0.314)	(0.307)	(0.320)
Arabica	0.334*	0.352*	0.364*	0.365**	0.367*	0.371**
	(0.198)	(0.187)	(0.192)	(0.178)	(0.190)	(0.176)
Arabica × Declared	0.625**	0.623**	0.603**	0.609**	0.604**	0.603**
	(0.278)	(0.251)	(0.260)	(0.236)	(0.252)	(0.233)
Guide	0.370***	0.377***	0.370***	0.368***	0.356***	0.352***
	(0.0585)	(0.0628)	(0.0540)	(0.0573)	(0.0511)	(0.0537)
Ln_Revenues_mln		-0.00393		-0.00645		-0.00704
		(0.00995)		(0.00994)		(0.0103)
Guide_score		-0.0928		-0.0793		-0.0871
		(0.0890)		(0.0962)		(0.0990)
Control Function	-0.111	-0.0874	-0.0809	-0.0705	-0.0826	-0.0713
	(0.135)	(0.125)	(0.129)	(0.119)	(0.130)	(0.121)
Observations	55,249	55,249	55,249	55,249	55,249	55,249
R-squared	0.695	0.700	0.719	0.723	0.737	0.741
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *Standard errors clustered at the product level. *** p<0.01, ** p<0.05, * p<0.1.*

4 Final remarks

This paper studies firms' decisions to disclose the Arabica share of ground coffee products. We develop a model combining rational inattention with coarse inference about undisclosed quality. Any interior equilibrium has a threshold disclosure rule: higher-quality firms disclose, while lower-quality firms withhold information. Using data on 145 products sold in Italian supermarkets, we find that products with higher Arabica content are more likely to disclose their composition. We also find that the relationship between Arabica content and price is steeper among disclosing products. This difference persists across several specifications and robustness checks, although it is not a general prediction of the model. The estimated price schedules intersect at approximately 58% Arabica content; this intersection describes the observed pricing relationship and does not identify the model's profit-based disclosure threshold.

While the theoretical model predicts a sharp threshold, implying that all firms above the limit declare their quality while those below do not, our empirical data reveals a more nuanced reality where a minority of firms deviate from this equilibrium. These deviations suggest that the disclosure point functions more as a probability gradient than a hard wall. Several factors likely rationalise these departures from the predicted optimum. If consumers possess varying cognitive costs, different brands may target specific sub-groups, shifting the optimal disclosure point based on the target audience's level of attentiveness. Furthermore, unlike the single-product firms in our baseline model, real-world firms must manage portfolios; they may obfuscate a high-quality product to avoid cannibalising a top-tier premium offering that sits just above it in quality. Obfuscation may also serve as insurance against agricultural shocks, allowing firms to maintain blending flexibility in response to volatile Arabica prices without violating disclosure promises. Finally, the market may be further segmented by product type or niche—such as Moka versus Espresso or Organic varieties—where different thresholds apply but remain difficult to estimate due to statistical power constraints. Similarly, established brand reputations may alter these incentives. Drawing a parallel to the upstream supply chain dynamics highlighted by Macchiavello and Miquel-Florensa (2017), where relational capital substitutes for formal enforcement, downstream firms with entrenched market reputations can leverage brand equity as a substitute for hard technical disclosure. For these prestigious roasters, the accumulated trust removes the marginal benefit of disclosing exact Arabica percentages, allowing them to remain in the hidden pool without incurring a low-quality penalty.

The presence of strategic obfuscation suggests that market competition alone is insufficient to achieve full transparency. This informational failure is likely to affect welfare by leading

consumers to select sub-optimal products relative to their idiosyncratic preferences and by allowing sellers to extract higher rents, thereby depressing trade volumes and total market surplus. To enhance market efficiency, regulatory bodies may consider mandating the disclosure of key quality proxies to ensure prices better reflect underlying characteristics market-wide. Policymakers could also promote standardised front-of-pack labelling, akin to the Nutri-Score or 'Traffic Light' systems, to lower the cognitive burden of quality assessment. Additionally, promoting QR-code-based transparency systems would lower search and processing costs for interested consumers, diminishing the strategic value of obfuscation and pushing more firms toward a disclosure equilibrium. Our results also indicate that external validators, such as the Godina and Illiano coffee guide, provide quality proxies independent of brand-related price effects. Supporting such independent certification bodies can provide consumers with objective benchmarks that counteract firm-led obfuscation. Ultimately, when consumers face limited attention budgets, firms have a structural incentive to hide; policy should therefore focus on reducing the cognitive burden of quality assessment to foster a market where price more accurately signals value.

References

Armstrong, M., Zhou, J., 2022. Consumer information and the limits to competition. American Economic Review 112, 534–77.

Arnot, C.M., Boxall, P.C., Cash, S.B., 2006. Do ethical consumers care about price? A revealed preference analysis of Fair Trade coffee purchases. Canadian Journal of Agricultural Economics 54, 555–565.

Bar-Isaac, H., Shelegia, S., 2023. Search, showrooming, and retailer variety. Marketing Science 42, 251–270.

Cheng, B., Furtado, A., Smyth, H.E., Henry, R.J., 2016. Influence of genotype and environment on coffee quality. Trends in Food Science & Technology 57, 20–30.

Cheong, I., Kim, J.Y., 2004. Costly information disclosure in oligopoly. The Journal of Industrial Economics 52, 121–132.

Coen, J., Gavazza, A., Gottlieb, D., Ramadorai, T., 2026. Regulating inaction: The case of price walking. Mimeo.

Ellison, G., 2005. A model of add-on pricing. The Quarterly Journal of Economics 120, 585–637.

Ellison, G., Ellison, S.F., 2009. Search, obfuscation, and price elasticities on the internet. Econometrica 77, 427–452.

Ellison, G.M., Wolitzky, A., 2012. A search cost model of obfuscation. The RAND Journal of Economics 43, 417–441.

Eyster, E., Madarász, K., Michaillat, P., 2021. Pricing under fairness concerns. Journal of the European Economic Association 19, 1853–1898. doi:10.1093/jeea/jvaa041.

Eyster, E., Rabin, M., 2005. Cursed equilibrium. Econometrica 73, 1623–1672. doi:10.1111/j.1468-0262.2005.00631.x.

Eyster, E., Rabin, M., Vayanos, D., 2019. Financial markets where traders neglect the informational content of prices. The Journal of Finance 74, 371–399. doi:10.1111/jofi.12729.

Fay, S., 2008. Selling an opaque product through an intermediary: The case of disguising one's product. Journal of Retailing 84, 59–75.

Fehr, E., Wu, K., 2025. Obfuscation in Competitive Markets. Discussion Paper 59/25. ROCKWOOL Foundation Berlin (RFBerlin).

Freitas, V.V., Borges, L.L.R., Castro, G.A.D., Almeida, L.F., Crepalde, L.T., Kobi, H.B., Vidigal, M.C.T.R., dos Santos, M.H., Fernandes, S.A., Maitan-Alfenas, G.P., Stringheta, P.C., 2024. Influence of roasting levels on chemical composition and sensory quality of arabica and robusta coffee: A comparative study. Food Bioscience 59, 104171.

Gabaix, X., Laibson, D., 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. The Quarterly Journal of Economics 121, 505–540.

Gavazza, A., Lizzeri, A., 2021. Frictions in product markets, in: Ho, K., Hortaçsu, A., Lizzeri, A. (Eds.), Handbook of Industrial Organization. Elsevier. volume 4, pp. 433–484.

Godina, A., Illiano, M., 2024. Guida dei caffè e delle torrefazioni d'Italia 2023-2024. Mondadori Electa, Milan, Italy.

Hilger, J., Nordyke, E., Lusk, J.L., 2011. Can experts inform consumer choice? The case of wine and coffee. Journal of Agricultural and Resource Economics 36, 153–170.

Hwang, I., Kim, K., Boleslavsky, R., 2023. Competitive advertising and pricing. Unpublished manuscript.

Illy, A., Viani, R., 2005. Espresso Coffee: The Science of Quality. 2nd ed., Academic Press, London, UK.

International Coffee Organization, 2023. Coffee Development report 2022/23. Technical Report. International Coffee Organization.

Ivanov, M., 2013. Information revelation in competitive markets. Economic Theory 52, 337–365.

Janssen, A., Kasinger, J., 2024a. Obfuscation and rational inattention. The Journal of Industrial Economics 72, 390–428.

Janssen, A., Kasinger, J., 2024b. Obfuscation and rational inattention. The Journal of Industrial Economics 72, 390–428. doi:10.1111/joie.12362.

Jehiel, P., 2005. Analogy-based expectation equilibrium. Journal of Economic Theory 123, 81–104. doi:10.1016/j.jet.2003.12.003.

Jehiel, P., Koessler, F., 2008. Revisiting games of incomplete information with analogy-based expectations. Games and Economic Behavior 62, 533–557. doi:10.1016/j.geb.2007.06.006.

Li, F., Morgan, K.L., Zaslavsky, A.M., 2018. Balancing covariates via propensity score weighting. Journal of the American Statistical Association 113, 390–400.

Lu, X., White, H., 2014. Robustness checks and robustness tests in applied economics. Journal of Econometrics 178, 194–206.

Macchiavello, R., Miquel-Florensa, J., 2017. Vertical integration and relational contracts: Evidence from the costa rica coffee chain. American Economic Journal: Microeconomics 9, 142–170.

MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2023. Cluster-robust inference: A guide to empirical practice. Journal of Econometrics 232, 272–299.

Matějka, F., McKay, A., 2015. Rational inattention to discrete choices: A new foundation for the multinomial logit model. American Economic Review 105, 272–298. doi:10.1257/aer.20130047.

McCluskey, J.J., 2000. A game theoretic approach to organic foods: An analysis of asymmetric information and policy. Agricultural and Resource Economics Review 29, 1–9.

Mendes, G.A., de Oliveira, M.A.L., Rodarte, M.P., dos Anjos, V.C., Bell, M.J.V., 2024. Determination of arabica and robusta species in blends of roasted coffee by mid-infrared spectroscopy in association with mixture design. Food Chemistry Advances 4, 100592.

Miklos-Thal, J., Zhang, J., 2013. (De)marketing to manage consumer quality inferences. Journal of Marketing Research 50, 55–69.

Milgrom, P.R., 1981. Good news and bad news: Representation theorems and applications. The Bell Journal of Economics 12, 380–391.

Moraga-González, J.L., Sun, Y., 2025. Consumer search costs and the provision of service quality. International Journal of Industrial Organization 103, 103208.

Moraga-González, J.L., Sándor, Z., Wildenbeest, M.R., 2023. Consumer search and prices in the automobile market. The Review of Economic Studies 90, 1394–1440.

Pennerstorfer, D., Schmidt-Dengler, P., Schutz, N., Weiss, C., Yontcheva, B., 2020. Information and price dispersion: theory and evidence. International Economic Review 61, 871–899.

Petrikaitė, V., 2018a. Consumer obfuscation by a multiproduct firm. The RAND Journal of Economics 49, 206–223.

Petrikaitė, V., 2018b. Vertical bargaining and obfuscation. International Journal of Industrial Organization 57, 106–130.

Piolatto, A., Schuett, F., 2025. Platform design and rent extraction. Journal of Industrial Economics 73, 124–166.

Ravid, D., 2020. Ultimatum bargaining with rational inattention. American Economic Review 110, 2948–2963. doi:10.1257/aer.20170620.

Richards, T.J., Klein, G.J., Bonnet, C., Bouamra-Mechemache, Z., 2020. Strategic obfuscation and retail pricing. Review of Industrial Organization 57, 859–889.

Shapiro, D., Shi, X., 2008. Market segmentation: The role of opaque travel agencies. Journal of Economics & Management Strategy 17, 803–837.

Sims, C.A., 2003a. Implications of rational inattention. Journal of Monetary Economics 50, 665–690.

Sims, C.A., 2003b. Implications of rational inattention. Journal of Monetary Economics 50, 665–690. doi:10.1016/S0304-3932(03)00029-1.

Spiegler, R., 2006. Competition over agents with boundedly rational expectations. Theoretical Economics , 207–231.

Varian, H.R., 1980. A model of sales. American Economic Review 70, 651–59.

Yeager, S.E., Batalia, M.E., Guinard, J.X., Ristenpart, W.D., 2023. Consumer perception of coffee quality: The role of species, roast level, and sensory information. Critical Reviews in Food Science and Nutrition 63, 1010–1036.

Appendix A Proofs

A.1 Proof of Lemma 1

Proof. Fix a price p and a belief $\mu = \mu_O^{\varepsilon}$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. To simplify notation, write $\beta(v) = \beta(v, p; \mu)$ and $\bar{\beta} = \int_0^1 \beta(v)\, \mu(dv)$.

For a fixed prior, mutual information is convex in the conditional choice rule. Hence, the consumer's objective is concave in β, and the first-order condition characterises an interior optimum.

At an interior solution, the first-order condition with respect to $\beta(v)$ is

$$u(v) - p - \kappa \left[H'(\bar{\beta}) - H'(\beta(v)) \right] = 0$$

for μ-almost every v. Since

$$H'(t) = \log\left(\frac{1-t}{t}\right) = -\log\left(\frac{t}{1-t}\right),$$

the first-order condition can be rewritten as

$$u(v) - p = \kappa \left[\log\left(\frac{\beta(v)}{1-\beta(v)}\right) - \log\left(\frac{\bar{\beta}}{1-\bar{\beta}}\right) \right].$$

Therefore,

$$\log\left(\frac{\beta(v)}{1-\beta(v)}\right) = \log\left(\frac{\bar{\beta}}{1-\bar{\beta}}\right) + \frac{u(v)-p}{\kappa}.$$

Solving for $\beta(v)$ gives

$$\beta(v, p; \mu) = \frac{1}{1 + \frac{1-\bar{\beta}(p;\mu)}{\bar{\beta}(p;\mu)} \exp\left(-\frac{u(v)-p}{\kappa}\right)}.$$

Integrating this expression with respect to μ yields

$$\bar{\beta}(p; \mu) = \int_0^1 \frac{1}{1 + \frac{1-\bar{\beta}(p;\mu)}{\bar{\beta}(p;\mu)} \exp\left(-\frac{u(v)-p}{\kappa}\right)} \mu(dv).$$

By Assumption 2, the optimal unconditional purchase probability $\bar{\beta}(p; \mu)$ is unique. Hence, every interior optimal conditional choice rule satisfies the same logit expression and all such rules agree μ-almost everywhere.

The expression above is continuous in v because u is continuous. It therefore defines a continuous representative of the optimal conditional choice rule. This continuous representative is unique. Suppose that two continuous representatives coincide μ-almost everywhere but differ at some v_0. By continuity, they must differ on an open neighbourhood of v_0. Since μ

has full support, this neighbourhood has strictly positive μ-probability, contradicting equality μ-almost everywhere. Hence, the continuous representative is uniquely determined at every $v \in [0,1]$. □

A.2 Proof of Lemma 2

Proof. Fix a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$, and consider an interior solution. Write $\beta(v) = \beta(v, p; \mu)$ and $\bar{\beta} = \bar{\beta}(p; \mu)$, and define

$$J(p;\mu) = \int_0^1 \beta(v)[1-\beta(v)]\,\mu(dv).$$

We first establish that $\bar{\beta}(p;\mu)$ is differentiable. Let

$$\Lambda(z) = \frac{1}{1+\exp(-z)}$$

and define

$$F(q,p) = q - \int_0^1 \Lambda\left(\log\left(\frac{q}{1-q}\right) + \frac{u(v)-p}{\kappa}\right)\mu(dv).$$

At an interior solution,

$$F(\bar{\beta}(p;\mu), p) = 0.$$

Differentiating F with respect to q and evaluating at the fixed point gives

$$F_q(\bar{\beta}, p) = 1 - \frac{J(p;\mu)}{\bar{\beta}(1-\bar{\beta})}.$$

Recall that

$$V(p;\mu) = \int_0^1 [\beta(v) - \bar{\beta}]^2 \mu(dv).$$

Moreover,

$$\bar{\beta}(1-\bar{\beta}) - J(p;\mu) = \int_0^1 \beta(v)^2 \mu(dv) - \bar{\beta}^2 = V(p;\mu).$$

Consequently,

$$F_q(\bar{\beta}, p) = \frac{V(p;\mu)}{\bar{\beta}(1-\bar{\beta})} > 0.$$

The implicit-function theorem therefore implies that $\bar{\beta}(p;\mu)$ is continuously differentiable in a neighbourhood of the price under consideration. Since the solution is interior, q remains bounded away from zero and one in a sufficiently small neighbourhood. The integrand and

its derivatives are therefore locally bounded, which justifies differentiation under the integral sign.

Differentiating the log-odds representation in Lemma 1 with respect to p gives

$$\frac{\beta_p(v)}{\beta(v)[1-\beta(v)]} = \frac{\bar{\beta}_p}{\bar{\beta}(1-\bar{\beta})} - \frac{1}{\kappa}.$$

Consequently,

$$\beta_p(v) = \beta(v)[1-\beta(v)]\left[\frac{\bar{\beta}_p}{\bar{\beta}(1-\bar{\beta})} - \frac{1}{\kappa}\right].$$

Integrating over v yields

$$\bar{\beta}_p = J(p;\mu)\left[\frac{\bar{\beta}_p}{\bar{\beta}(1-\bar{\beta})} - \frac{1}{\kappa}\right].$$

Rearranging,

$$\bar{\beta}_p\left[1 - \frac{J(p;\mu)}{\bar{\beta}(1-\bar{\beta})}\right] = -\frac{J(p;\mu)}{\kappa}.$$

Using

$$1 - \frac{J(p;\mu)}{\bar{\beta}(1-\bar{\beta})} = \frac{V(p;\mu)}{\bar{\beta}(1-\bar{\beta})},$$

we obtain

$$\frac{\partial\bar{\beta}(p;\mu)}{\partial p} = -\frac{\bar{\beta}(p;\mu)[1-\bar{\beta}(p;\mu)]}{\kappa V(p;\mu)}\int_0^1 \beta(v,p;\mu)[1-\beta(v,p;\mu)]\mu(dv).$$

The expression is strictly negative because the solution is interior and $V(p;\mu) > 0$.

Substituting this expression into the derivative of the conditional purchase probability gives

$$\beta_p(v) = -\frac{\beta(v)[1-\beta(v)]}{\kappa V(p;\mu)}[J(p;\mu) + V(p;\mu)].$$

Since

$$J(p;\mu) + V(p;\mu) = \bar{\beta}(1-\bar{\beta}),$$

we conclude that

$$\frac{\partial\beta(v,p;\mu)}{\partial p} = -\frac{\beta(v,p;\mu)[1-\beta(v,p;\mu)]\bar{\beta}(p;\mu)[1-\bar{\beta}(p;\mu)]}{\kappa V(p;\mu)} < 0.$$

□

A.3 Proof of Lemma 3

Proof. Fix a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. A firm of quality v that obfuscates chooses p to maximise

$$R(v, p; \mu) = p\beta(v, p; \mu).$$

By Assumption 2, the maximiser is positive and interior. Its first-order condition is

$$\beta(v, p; \mu) + p\frac{\partial \beta(v, p; \mu)}{\partial p} = 0.$$

Using Lemma 2,

$$\frac{\partial \beta(v, p; \mu)}{\partial p} = -\frac{\beta(v, p; \mu)[1 - \beta(v, p; \mu)]\bar{\beta}(p; \mu)[1 - \bar{\beta}(p; \mu)]}{\kappa V(p; \mu)}.$$

Substitution into the first-order condition gives

$$\beta(v, p; \mu) - p\frac{\beta(v, p; \mu)[1 - \beta(v, p; \mu)]\bar{\beta}(p; \mu)[1 - \bar{\beta}(p; \mu)]}{\kappa V(p; \mu)} = 0.$$

Since the solution is interior, $\beta(v, p; \mu) > 0$. Dividing through by $\beta(v, p; \mu)$ and solving for p yields

$$p = \frac{\kappa V(p; \mu)}{[1 - \beta(v, p; \mu)]\bar{\beta}(p; \mu)[1 - \bar{\beta}(p; \mu)]}.$$

Evaluating this expression at $p^O(v; \mu)$ gives

$$p^O(v; \mu) = \frac{\kappa V\big(p^O(v; \mu); \mu\big)}{\big[1 - \beta\big(v, p^O(v; \mu); \mu\big)\big]\, \bar{\beta}\big(p^O(v; \mu); \mu\big)\, \big[1 - \bar{\beta}\big(p^O(v; \mu); \mu\big)\big]}.$$

□

A.4 Proof of Proposition 1

Proof. Fix a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. Define obfuscation revenue by

$$R(v, p; \mu) = p\beta(v, p; \mu).$$

By Assumption 2, for every v, the function $R(v, \cdot; \mu)$ has a unique positive interior maximiser $p^O(v; \mu)$.

In a neighbourhood of this maximiser, define

$$L(v, p; \mu) = \log R(v, p; \mu) = \log p + \log \beta(v, p; \mu).$$

Since the logarithm is strictly increasing, L and R have the same maximiser.
From Lemma 1,

$$\frac{\partial \beta(v,p;\mu)}{\partial v} = \beta(v,p;\mu)[1-\beta(v,p;\mu)]\frac{u'(v)}{\kappa}.$$

Therefore,

$$L_v(v,p;\mu) = [1-\beta(v,p;\mu)]\frac{u'(v)}{\kappa}.$$

Differentiating with respect to p gives

$$L_{pv}(v,p;\mu) = -\frac{u'(v)}{\kappa}\frac{\partial \beta(v,p;\mu)}{\partial p}.$$

By Lemma 2, $\partial\beta(v,p;\mu)/\partial p < 0$. Since $u'(v) > 0$, it follows that

$$L_{pv}(v,p;\mu) > 0.$$

At the optimal price,

$$R_p\big(v, p^O(v;\mu);\mu\big) = 0.$$

Moreover, $R\big(v, p^O(v;\mu);\mu\big) > 0$, and the strict second-order condition in Assumption 2 implies

$$R_{pp}\big(v, p^O(v;\mu);\mu\big) < 0.$$

Hence,

$$L_{pp}\big(v, p^O(v;\mu);\mu\big) = \frac{R_{pp}\big(v, p^O(v;\mu);\mu\big)}{R\big(v, p^O(v;\mu);\mu\big)} - \left[\frac{R_p\big(v, p^O(v;\mu);\mu\big)}{R\big(v, p^O(v;\mu);\mu\big)}\right]^2 < 0.$$

The first-order condition for the log objective is

$$L_p\big(v, p^O(v;\mu);\mu\big) = 0.$$

The implicit-function theorem therefore gives

$$\frac{\partial p^O(v;\mu)}{\partial v} = -\frac{L_{pv}\big(v, p^O(v;\mu);\mu\big)}{L_{pp}\big(v, p^O(v;\mu);\mu\big)} > 0$$

for every $v \in (0,1)$.

The implicit-function theorem gives a locally continuously differentiable optimal-price function around every $v \in (0,1)$. Uniqueness of the maximiser implies that these local solutions coincide on overlapping neighbourhoods. Hence, they define a continuously differentiable function $p^O(\cdot;\mu)$ on $(0,1)$.

It remains to establish continuity at all qualities. Fix $v_0 \in [0,1]$. Local boundedness of the maximiser implies that there exist a relative neighbourhood U of v_0 and a compact set $K \subseteq \mathbb{R}_+$ such that $p^O(v;\mu) \in K$ for every $v \in U$. Since R is continuous on $U \times K$ and its maximiser is unique, the maximum theorem implies that $p^O(\cdot;\mu)$ is continuous at v_0. Hence, $p^O(\cdot;\mu)$ is continuous on $[0,1]$.

Since its derivative is strictly positive on $(0,1)$, $p^O(v;\mu)$ is strictly increasing on $[0,1]$. □

A.5 Proof of Proposition 2

Proof. Fix a belief $\mu = \mu_O^\varepsilon$ for some measurable $O \subseteq [0,1]$ satisfying $0 < \lambda(O) < 1$. Recall that

$$\Delta(v;\mu) = u(v) - \pi^O(v;\mu),$$

where

$$\pi^O(v;\mu) = p^O(v;\mu)\beta\big(v, p^O(v;\mu);\mu\big).$$

By Assumption 2 and the argument in the proof of Proposition 1, $p^O(v;\mu)$ is locally continuously differentiable for $v \in (0,1)$. Differentiating the optimised profit and using the firm's first-order condition gives

$$\frac{\partial \pi^O(v;\mu)}{\partial v} = p^O(v;\mu)\frac{\partial \beta\big(v, p^O(v;\mu);\mu\big)}{\partial v}.$$

Lemma 1 implies

$$\frac{\partial \beta(v,p;\mu)}{\partial v} = \beta(v,p;\mu)[1-\beta(v,p;\mu)]\frac{u'(v)}{\kappa}.$$

Therefore,

$$\frac{\partial \pi^O(v;\mu)}{\partial v} = p^O(v;\mu)\beta\big(v, p^O(v;\mu);\mu\big)\left[1-\beta\big(v, p^O(v;\mu);\mu\big)\right]\frac{u'(v)}{\kappa}.$$

Using the optimal-pricing condition in Lemma 3 gives

$$\frac{\partial \pi^O(v;\mu)}{\partial v} = u'(v)\beta\big(v, p^O(v;\mu);\mu\big)\frac{V\big(p^O(v;\mu);\mu\big)}{\bar{\beta}\big(p^O(v;\mu);\mu\big)\left[1-\bar{\beta}\big(p^O(v;\mu);\mu\big)\right]}.$$

For every interior solution,

$$\bar{\beta}(p;\mu)[1-\bar{\beta}(p;\mu)] - V(p;\mu) = \int_0^1 \beta(v,p;\mu)[1-\beta(v,p;\mu)]\mu(dv) > 0.$$

Hence,

$$V(p;\mu) < \bar{\beta}(p;\mu)[1-\bar{\beta}(p;\mu)].$$

Moreover, $0 < \beta(v,p;\mu) < 1$. It follows that

$$\beta(v,p;\mu)\frac{V(p;\mu)}{\bar{\beta}(p;\mu)[1-\bar{\beta}(p;\mu)]} < 1.$$

Consequently,

$$\frac{\partial \pi^O(v;\mu)}{\partial v} < u'(v).$$

Therefore,

$$\frac{\partial \Delta(v;\mu)}{\partial v} = u'(v) - \frac{\partial \pi^O(v;\mu)}{\partial v} > 0$$

for every $v \in (0,1)$.

Continuity of the objective and local boundedness of the unique maximiser imply that $\pi^O(\cdot;\mu)$ is continuous on $[0,1]$. Hence, $\Delta(\cdot;\mu)$ is also continuous on $[0,1]$. Since its derivative is strictly positive on $(0,1)$, $\Delta(v;\mu)$ is strictly increasing on $[0,1]$. □

A.6 Proof of Proposition 3

Proof. Consider an interior ε-contaminated coarse-inference equilibrium, and let μ^* denote its equilibrium coarse belief. For this fixed belief, define

$$\Delta(v;\mu^*) = u(v) - \pi^O(v;\mu^*).$$

By Proposition 2, $v \mapsto \Delta(v;\mu^*)$ is continuous and strictly increasing.

Firm optimality implies that a type v obfuscates whenever $\Delta(v;\mu^*) < 0$, discloses whenever $\Delta(v;\mu^*) > 0$, and may choose either action whenever $\Delta(v;\mu^*) = 0$.

Because the equilibrium is interior, both the set of obfuscating types and the set of disclosing types have strictly positive λ-measure. Since $\Delta(\cdot;\mu^*)$ is strictly increasing, its zero set contains at most one quality. It follows that there exist qualities v_O and v_D such that

$$\Delta(v_O;\mu^*) < 0 \qquad \text{and} \qquad \Delta(v_D;\mu^*) > 0.$$

Strict monotonicity implies that $v_O < v_D$. By continuity, there exists a unique $v^* \in (v_O, v_D)$ satisfying

$$\Delta(v^*; \mu^*) = 0.$$

Equivalently,

$$u(v^*) = \pi^O(v^*; \mu^*).$$

Strict monotonicity implies that $\Delta(v; \mu^*) < 0$ for every $v < v^*$ and $\Delta(v; \mu^*) > 0$ for every $v > v^*$. Thus, up to the behaviour of the indifferent type,

$$d^*(v) = \begin{cases} 0, & v < v^*, \\ 1, & v > v^*. \end{cases}$$

Because λ admits a density, the singleton $\{v^*\}$ has zero probability. It is therefore immaterial for the equilibrium belief whether the marginal type discloses or obfuscates. Up to a λ-null set, the set of non-disclosing qualities is $[0, v^*]$. Hence,

$$\mu^* = (1 - \varepsilon)\lambda(\cdot \mid [0, v^*]) + \varepsilon\lambda = \mu^{\varepsilon}_{v^*}.$$

The marginal indifference condition can therefore be written as

$$u(v^*) = \pi^O\left(v^*; \mu^{\varepsilon}_{v^*}\right).$$

For the converse, suppose that $v^* \in (0, 1)$ satisfies

$$u(v^*) = \pi^O\left(v^*; \mu^{\varepsilon}_{v^*}\right).$$

Fix the belief $\mu^* = \mu^{\varepsilon}_{v^*}$. By Proposition 2, the function

$$v \longmapsto u(v) - \pi^O(v; \mu^*)$$

is strictly increasing. Since it equals zero at v^*, it is negative for every $v < v^*$ and positive for every $v > v^*$. Therefore, every type below the cutoff strictly prefers obfuscation, while every type above the cutoff strictly prefers disclosure.

Let firms below the cutoff obfuscate and charge $p^O(v; \mu^*)$, and let firms above the cutoff disclose and charge $u(v)$. The marginal type may choose either action. The resulting set of non-disclosing firms agrees with $[0, v^*]$ up to a λ-null set and therefore generates the belief $\mu^{\varepsilon}_{v^*}$. Consumers respond optimally according to Lemma 1, and firms choose optimal disclosure and pricing decisions. Hence, these strategies constitute an interior threshold equilibrium whenever the regularity conditions in Assumption 2 hold. □

Appendix B Tables

Table B.1: Full sample summary

Disclosure	Unique Products	Obs	Mean Price	Std. Dev.
Yes	78	26,822	4.33	1.470
No - Recovered	67	28,247	2.88	0.735
No - Unrecovered	26	494	2.65	0.799

Table B.2: Summary statistics

Variable	Description	Obs	Mean	Std. Dev.	Min	Max
Price	Price in €	55,249	3.58	1.36	0.89	11.53
Declared	DV=1 if the share of Arabica is declared in the label	55,249	0.49	0.50	0	1
Arabica	Share of Arabica	55,249	0.65	0.31	0	1
Organic	DV=1 if the coffee is organic	55,249	0.06	0.24	0	1
Single_country	DV=1 if the coffee comes from a single country	55,249	0.02	0.13	0	1
Time trend	Month-Year time trend	55,249	15.80	7.08	1	28
Private_label	DV=1 if the coffee is packaged under the name of the retailer rather than that of the manufacturer	55,249	0.24	0.43	0	1
Espresso	DV=1 if the coffee is minced for the espresso machine	55,249	0.30	0.46	0	1
Revenues_mln	Revenues of the coffee company in mln €	55,249	3,765	5,952	0.3	18,500
Guide	DV=1 if the company is in the coffee guide	55,249	0.70	0.46	0	1
Guide_score	DV=1 if the coffee is rated or mentioned in the coffee guide	55,249	0.12	0.32	0	1

Table B.3: Correlation matrix

Variable	Ln_Price	Declared	Arabica	Organic	Single_c.	Time	Priv_lab	Espresso	Ln_Rev.	Guide	G_score
Ln_Price	1										
Declared	0.5270	1									
Arabica	0.7099	0.7472	1								
Organic	0.1911	0.2636	0.2464	1							
Single_country	0.0117	0.1054	0.1089	-0.0296	1						
Time	-0.0410	-0.1293	-0.2484	0.0574	-0.0571	1					
Private_label	-0.4638	-0.0148	-0.1688	0.0631	0.1637	0.0779	1				
Espresso	-0.1497	0.0429	0.1153	-0.0657	0.1428	-0.1531	0.2719	1			
Ln_Revenues	-0.3511	-0.3461	-0.3411	-0.1020	0.1242	0.0786	0.5776	-0.1238	1		
Guide	0.3754	-0.0891	0.0625	-0.3895	-0.1460	-0.0513	-0.8617	-0.2656	-0.3812	1	
Guide_score	0.0692	0.1835	0.0992	-0.0946	-0.0487	0.0273	-0.2083	0.3365	-0.3355	0.2417	1

Table B.4: Difference in mean between the two subgroups

	Declared	Undeclared	Difference	$Pr(\|T\| > \|t\|)$
Arabica	0.89	0.42	0.47	0.000
Price	4.33	2.88	1.45	0.000
H_0 : Difference = 0				

Table B.5: Regressions, only firms who declare

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price	Ln_Price
Constant	0.0645***	0.0907***	0.0804***	0.117***	0.362***	0.373***
	(0.0108)	(0.0120)	(0.0101)	(0.0114)	(0.0222)	(0.0224)
Organic	0.299***	0.306***	0.256***	0.262***	0.242***	0.243***
	(0.00515)	(0.00547)	(0.00483)	(0.00522)	(0.00520)	(0.00539)
Single_country	-0.0320***	-0.0414***	-0.0240***	-0.0452***	-0.0550***	-0.0852***
	(0.00930)	(0.00989)	(0.00897)	(0.00935)	(0.00911)	(0.00918)
Arabica	1.128***	1.049***	1.027***	0.955***	1.038***	0.957***
	(0.0102)	(0.00967)	(0.00983)	(0.00963)	(0.00968)	(0.00952)
Guide	0.399***	0.455***	0.427***	0.465***	0.413***	0.449***
	(0.00454)	(0.00504)	(0.00428)	(0.00476)	(0.00431)	(0.00485)
Ln_Revenues_mln		0.00521***		0.00636***		0.00798***
		(0.000915)		(0.00100)		(0.00108)
Guide_score		-0.185***		-0.159***		-0.170***
		(0.00320)		(0.00378)		(0.00394)
Observations	26,822	26,822	26,822	26,822	26,822	26,822
R-squared	0.532	0.569	0.614	0.639	0.659	0.685
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *The Table reports Robust Standard errors. *** p<0.01, ** p<0.05, * p<0.1.*

Table B.6: Regressions, only firms who obfuscate

VARIABLES	(1) Ln_Price	(2) Ln_Price	(3) Ln_Price	(4) Ln_Price	(5) Ln_Price	(6) Ln_Price
Constant	0.432***	0.378***	0.577***	0.507***	0.417***	0.290***
	(0.00404)	(0.0112)	(0.00852)	(0.0122)	(0.0146)	(0.0177)
Organic	-0.276***	-0.264***	-0.246***	-0.215***	-0.258***	-0.226***
	(0.00603)	(0.00663)	(0.00738)	(0.00814)	(0.0154)	(0.0151)
Single_country	0.301***	0.330***	0.235***	0.278***	0.188***	0.232***
	(0.0119)	(0.0120)	(0.0120)	(0.0115)	(0.00950)	(0.0106)
Arabica	0.513***	0.524***	0.522***	0.539***	0.488***	0.511***
	(0.00502)	(0.00556)	(0.00483)	(0.00557)	(0.00463)	(0.00558)
Guide	0.345***	0.348***	0.335***	0.344***	0.337***	0.349***
	(0.00225)	(0.00269)	(0.00215)	(0.00261)	(0.00207)	(0.00249)
Ln_Revenues_mln		0.00567***		0.00977***		0.0122***
		(0.000961)		(0.00101)		(0.00104)
Guide_score		0.0917***		0.120***		0.141***
		(0.00555)		(0.00587)		(0.00614)
Observations	28,427	28,427	28,427	28,427	28,427	28,427
R-squared	0.652	0.657	0.686	0.695	0.720	0.731
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *The Table reports Robust Standard errors. *** p<0.01, ** p<0.05, * p<0.1.*

Table B.7: χ^2 test for difference in the coefficients of Arabica

Regressions	(1)	(2)	(3)	(4)	(5)	(6)
H0: $\beta(Const.\|d=0) = \beta(Const.\|d=1)$						
χ^2 (1)	1022.54	305.99	1406.22	545.74	4.4	8.57
$Prob > \chi^2$	0.0000	0.0000	0.0000	0.0000	0.0359	0.0034
H0: $\beta(Arabica\|d=0) = \beta(Arabica\|d=1)$						
χ^2 (1)	2941.94	2222.02	2134.66	1399.23	2682.99	1668.72
$Prob > \chi^2$	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Note: χ^2 *tests are computed using robust covariance matrices. The coefficients compared refer to Tables B.5 and B.6.*

Table B.8: Baseline Regression (Robust Standard Errors), all firms

VARIABLES	(1) Ln_Price	(2) Ln_Price	(3) Ln_Price	(4) Ln_Price	(5) Ln_Price	(6) Ln_Price
Constant	0.464***	0.466***	0.476***	0.512***	0.495***	0.536***
	(0.00372)	(0.00752)	(0.00605)	(0.00852)	(0.0201)	(0.0210)
Organic	0.266***	0.256***	0.233***	0.221***	0.215***	0.202***
	(0.00403)	(0.00414)	(0.00374)	(0.00384)	(0.00379)	(0.00380)
Single_country	-0.00341	-0.00773	-0.0120	-0.0134*	-0.0318***	-0.0337***
	(0.00861)	(0.00870)	(0.00768)	(0.00770)	(0.00702)	(0.00697)
Declared	-0.411***	-0.373***	-0.371***	-0.354***	-0.377***	-0.354***
	(0.00964)	(0.00972)	(0.00911)	(0.00911)	(0.00895)	(0.00909)
Arabica	0.470***	0.466***	0.462***	0.457***	0.467***	0.464***
	(0.00513)	(0.00541)	(0.00491)	(0.00527)	(0.00488)	(0.00518)
Arabica × Declared	0.654***	0.628***	0.623***	0.612***	0.626***	0.607***
	(0.0115)	(0.0113)	(0.0108)	(0.0106)	(0.0106)	(0.0105)
Guide	0.367***	0.381***	0.369***	0.371***	0.355***	0.356***
	(0.00239)	(0.00260)	(0.00222)	(0.00239)	(0.00212)	(0.00227)
Ln_Revenues_mln		-0.00131**		-0.00428***		-0.00459***
		(0.000616)		(0.000624)		(0.000652)
Guide_score		-0.101***		-0.0861***		-0.0936***
		(0.00298)		(0.00319)		(0.00327)
Observations	55,249	55,249	55,249	55,249	55,249	55,249
R-squared	0.691	0.698	0.717	0.721	0.735	0.740
Time trend	Yes	Yes	Yes	Yes	Yes	Yes
Region F.E.	No	No	Yes	Yes	No	No
Store F.E.	No	No	No	No	Yes	Yes

Note: *The Table reports Robust Standard errors. *** p<0.01, ** p<0.05, * p<0.1.*

Appendix C Figures

Figure C.1: Kernel density of log-price by disclosure status and availability of composition.

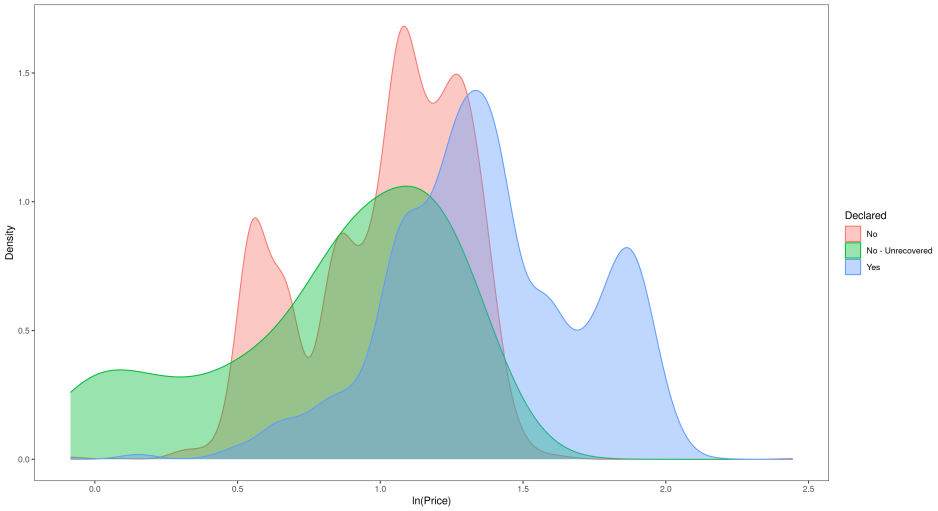

IEB Institut d'Economia de Barcelona