Melons as Lemons:

Asymmetric Information, Consumer Learning and Quality Provision*

Jie Bai†

Harvard Kennedy School

March 2017

Abstract

There is often a lack of reliable quality provision in many markets in developing countries. I designed an experiment to understand this phenomenon in a setting that features typical market conditions in developing countries: the retail watermelon markets in a major Chinese city. I first demonstrate empirically that there is substantial asymmetric information between sellers and buyers on sweetness, the key indicator of quality, yet sellers do not differentiate and price watermelons by quality. I then randomly introduce one of two branding technologies into 40 out of 60 markets—one sticker label that is widely used and counterfeited and one novel laser-cut label. I track sellers' quality, pricing and sales over an entire season and collect household panel purchasing data to examine the demand side's response. I find that laser-branding induced sellers to provide higher quality and led to higher sales profits. However, after the intervention was withdrawn, all markets reverted back to baseline. To rationalize the experimental findings, I build an empirical model of consumer learning and seller quality provision. The results suggest that consumers are hesitant to upgrade their perception under stickers, which makes quality provision a low-return investment. While the new technology enhances learning, the resulting increase in profits is not sufficient to cover the fixed cost of the technology for small individual sellers. Counterfactual analysis shows that information frictions and fragmented markets lead to significant under-provision of quality. Third-party interventions that subsidizing the initial demand and learning process or entry of large firms could improve welfare.

JEL Classifications: D22, D61, D82, D83, L11, L14, L15, O12, O14, O31
Total Word Count: About 12,500

*The paper was previously circulated under the title "Melons as Lemons: Asymmetric Information, Consumer Learning and Seller Reputation." I am deeply grateful to Benjamin Olken, Abhijit Banerjee, Robert Townsend and Nikhil Agarwal for their continuous encouragement and guidance over the course of this project. I thank David Atkin, Esther Duflo, Marcel Fafchamps, Hanming Fang, Glenn Ellison, Robert Gibbons, German Sergio Gieczewski, Seema Jayachandran, Yan Ji, Siyuan Liu, Rocco Macchiavello, Yuhei Miyauchi, Xiaosheng Mu, Ariel Pakes, Frank Schilbach, Peng Shi, Maheshwor Shrestha, Paulo Somaini, Tavneet Suri, Jean Tirole, Juuso Toikka, Richard Zeckhauser, Juanjuan Zhang, Hongkai Zhang, and seminar participants at NEUDC, AFE, Berkeley, Chicago (Economics/Booth), CICER, Columbia, Cornell, GWU, Harvard Economics, Harvard Kennedy School, Microsoft Research, MIT Sloan Applied Economics, MIT Sloan Marketing, Northwestern, NYU Stern, PKU, PSU, SITE, Stanford GSB, UMich, UPenn, Yale SOM Marketing, and the development and IO lunches at MIT for helpful discussions and feedback. Funding for this project was generously provided by the Weiss Family Fund, the Abdul Latif Jameel Poverty Action Lab, and the George and Obie Shultz Fund. All errors are my own.

†Address: 79 JFK Street, Cambridge, MA 02142. Phone: 617-6157771. E-mail: jie_bai@hks.harvard.edu.

1 Introduction

A key problem in developing countries is the lack of reliable provision of high quality goods and services. The problem is exacerbated in markets for experience goods, such as food products and pharmaceuticals. It is well recognized in economics that when contracting on quality is difficult information frictions can lead to quality deterioration and firms need a good reputation to succeed. However, a reputation for quality is precisely what many firms in developing countries are lacking. The question is, then, what are the barriers that hinder firms' ability and incentive to provide quality? Answering this question can have important policy implications: first, on a broad level, it helps governments to structure policies to facilitate industrial quality upgrading; second, it may offer new solutions to address the information problem, alternative to direct government regulations and quality controls, which can be very costly to enforce in countries with weak legal institutions.

In this paper, I designed an experiment to understand what hinders quality provision in a typical developing country market setting: the retail watermelon markets in a major Chinese city. Several features of this market make it particularly suited for studying this topic. First, there are a large number of small independent local markets, which allows randomization at the market level. Second, the quality of a watermelon can be very well captured by its sweetness, which can be objectively measured (ex-post) using a sweet meter. I document substantial asymmetric information between sellers and buyers on this key dimension of quality and a stark absence of quality premium at baseline. The goal of this research is to understand this phenomenon and provide a framework for thinking about quality provision in other similar settings.

I first propose a model of a long-run seller choosing quality to maximize the expected discounted sum of profits, subject to a dynamic demand system rooted in consumer learning. The model highlights two broad explanations for the lack of quality differentiation at baseline: first, it could be that the cost of reliably providing high quality is too high relative to consumers' valuation for quality. Consequently, higher quality is neither demanded nor supplied. Second, due to the information problem, a seller's claim of offering high quality cannot be immediately verified, and therefore consumers' initial perception and speed of learning matter for the seller's incentive to provide quality. In particular, pessimistic prior can make quality provision a low-return investment, and markets can stuck in an equilibrium with no quality differentiation.

The welfare and policy implications under these two explanations are very different. In order to tease the stories apart, I conducted a field experiment with 60 sellers in 60 different markets in Shijiazhuang, China. I randomly introduced one of two branding technologies into 40 out of the 60 markets—one sticker label that is widely used and often counterfeited, and one novel laser-cut label. Pilot surveys suggest that consumers regard laser-branding as being more effective at deterring counterfeits because laser machines are very expensive. Hence, the new technology could potentially dispel negative historical stereotypes associated with stickers,[1] thereby allowing sellers to establish trust faster. The model suggests that sellers in the laser group may have stronger incentives to provide quality. For a cross-randomized subset of sellers, I further provided them with a temporary monetary incentive to invest in high quality. By facilitating the initial demand and learning process, the model suggests that higher quality may sustain even in the post-incentive period.

The intervention lasted over eight weeks, spanning the entire peak season for watermelons. Each of the 60 sellers was asked to sell two piles of watermelons at the retail site: a premium pile and a normal pile. Sellers were *free* to set the quality, price, and quantity for each pile. For watermelons in the premium pile, sellers either received a laser engraving of the words "premium watermelon", or a sticker label with the same words, or no labeling at all (in which price serves as the main signal for quality). Quality differentiation was mandatory for the first two weeks but sellers were free to decide afterwards. This allows me to examine the differential incentives across the treatment groups. The incentive treatment was enforced through biweekly quality checks, and was lifted at the end of the sixth week. I kept track of sellers' quality, pricing and sales over the entire season, and collected household panel purchasing data to examine the demand side's response to quality differentiation.

There are three main experimental findings: first, laser branding induced sellers to provide a genuine quality-price premium, establishing that reputational incentives can potentially motivate quality. On the other hand, sellers in the label-less group sharply reverted back to no differentiation after the first two weeks; evidence for the sticker group is quite mixed: on average, quality of the premium pile was not significantly higher than the market average. Second, the incentive treatment successfully induced sellers to provide higher quality than their non-incentivized counterparts, but higher quality was only sustained for the laser incentive group. Overall, these findings are consistent

[1] This relates to the theory of collective reputation in Tirole (1996) and informational free-riding in Fang (2001).

with the model's predictions. Finally, in terms of sales outcomes, quality differentiation under sticker did not outperform no differentiation. In contrast, sellers in the laser group earned 30-40% higher sales profits on average as a result of both higher prices and higher total sales as a result of attracting more high-end consumers over time. This result demonstrates that there is a high demand for quality. Having said that, one year after the intervention when the laser technology was no longer provided for free, all markets reverted back to baseline. This suggests that individual sellers would not have the incentive to invest in the new technology themselves.

The experimental findings provide a qualitative explanation for the lack of quality differentiation at baseline. In the second part of the paper, I estimate an empirical model of consumer learning and seller quality provision to rationalize these findings and perform welfare and policy analysis.

The structural estimation proceeds in two steps. In the first step, I estimate a discrete choice demand system that explicitly models consumers' learning process and prior beliefs, which vary across different branding technologies. The demand model incorporates rich heterogeneity in household preferences, and controls for a full set of market and time fixed effects to correct for price endogeneity. The model is estimated using simulated maximum likelihood and I exploit purchasing patterns and experience realizations in the household panel data for identification. In the second step, I solve numerically for the seller's optimal pricing and quality, taking the demand side as estimated, and apply a minimum distance estimator to recover the seller's discount factor and unobserved effort costs from observed empirical policies. Overall, the structural estimates describe the data well: purchasing patterns generated by the Bayesian learning process fit the actual purchasing behavior, and the simulated sales mimic the actual sales dynamics in the data.

The structural estimates indicate that consumers' prior perception is more "stubborn" under sticker than under laser. As a result, trust can take a long time to establish, which explains why sellers do not have the incentive to provide quality under the existing "contaminated" signaling technology. While the new technology enhances consumer learning and thereby strengthens sellers' incentives, the increase in the discounted return, taking into account effort costs, is still not large enough to justify the fixed cost of the technology for individual sellers. There are two reasons: (a) each seller's size is very small; and (b) it may be difficult for sellers to extract all the consumer surplus due to market competition.

These structural results point to the importance of understanding the role of market structure in the presence of information problems. To further highlight these interactions and trade-offs faced by policy makers, I conduct several counterfactual exercises to examine the role of firm size and market competition. The results indicate that information frictions and fragmented markets lead to significant under-provision of quality in this setting. The gain in consumer surplus from quality differentiation is large as a result of both enlarged choice set and allocative efficiency. While an individual seller would not undertake such costly investment, a third-party could invest in the new technology and subsidize it for sellers to improve society's welfare. Alternatively, since sellers' net profits scale up with market size, the results suggest that there could be a profitable entry opportunity for a large upstream firm.

This paper contributes to the empirical literature on consumer learning, firm reputation and quality provision in markets with information problems.[2] While many studies examine online trading environments,[3] empirical work in the offline world is relatively sparse (Banerjee and Duflo, 2000; Jin and Leslie, 2009; Macchiavello, 2010; List, 2006; Björkman-Nyqvist, Svensson, and Yanagizawa-Drott, 2013; Macchiavello and Morjaria, 2015). As discussed in Bar-Isaac and Tadelis (2008), the empirical challenge is that researchers typically do not observe all information available to buyers, and sellers' behavior beyond what the buyers observe. This study takes advantage of a field experiment and collects data that directly keep track of the both sides. The results demonstrate that the way consumers gather information and learn shapes seller's incentive. This recalls the finding in Björkman-Nyqvist, Svensson, and Yanagizawa-Drott (2013) that quality provision of anti-malaria drugs in Uganda is hampered by consumers' misconceptions. Although the contexts differ, the policy conclusions are remarkably alike: to motivate high quality provision, policies that enhance consumer learning or entry of large firms may be needed. Ultimately, the external validity of the results is an empirical question as the exact learning dynamics and quality production technology vary across industries. The general framework and the experimental design proposed here can be applied to other settings.

[2]This study is motivated by the extensive body of economics and marketing literature on the role of advertising as signals for product quality (Bagwell, 2007). Most theoretical work focuses on equilibrium predictions between advertising and quality, where quality is exogenous. This study examines sellers' endogenous quality choice. The experimental design and counterfactual exercises are in the same spirit with the theoretical proposal in Henze, Schuett, and Sluijs (2015). Findings also speak to the role of rebranding as disrupting the negative link between consumers and the origin brand (Prasad and Dev, 2000).

[3]For example, see Jin and Kato (2006); Cabral and Hortacsu (2010); Klein, Lambertz, and Stahl (2016)

The study also relates to the broad literature on firm performance and quality upgrading in development and trade.[4] Previous studies have addressed: (1) supply side constraints, including credit access, lack of quality inputs, and managerial constraints;[5] and (2) demand side factors, including access to high-income markets (e.g., Verhoogen (2008); Atkin, Khandelwal, and Osman (forthcoming)).[6] This study highlights another potential barrier to quality upgrading, which is due to the information problem and low collective reputation.[7] Rising concerns among the public regarding product quality and safety in developing countries can lead to general distrusts at the bigger group level (either industry or country), which generates an important externality: it not only hampers individual firm's incentive to move up the quality ladder and its ability to penetrate higher-end markets (as we see for sellers in the sticker group), but also hurts new firms which are "endowed" with the damaged reputation of the ancestors (Macchiavello, 2010).

Findings of the study also speak to the role of firm size and market structure on product quality (e.g., see Kugler and Verhoogen (2012)).[8] Most of the empirical work focuses on settings where quality is observable. This study examines a setting with information asymmetry and highlights an important tradeoff: while market competition among small firms helps to expand sales, it can also discourage quality improvements especially when the returns take time to accrue.[9]

The paper also adds to a growing body of work on field experiments across firms and markets

[4]See De Loecker and Goldberg (2014) for a comprehensive review of the empirical literature. Most of the literature has relied on price and input costs to infer quality (e.g., Schott (2004); Hallak (2006)), or used structural approach to infer quality from market shares (Hummels and Klenow, 2005; Khandelwal, 2010; Hallak and Schott, 2011; Feenstra and Romalis, 2014; De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016). Previous studies have offered various critiques of the different methods. More recently, there is a growing set of work studying this topic by focusing on particular products (e.g., Chen and Juvenal (2016); Crozet, Head, and Mayer (2012); David Atkin and Osman (forthcoming). Watermelons map very nicely into a one-dimensional vertical taste model and I collect direct measures on quality using objective measures of sweetness.

[5]For example, see De Mel, McKenzie, and Woodruff (2008); Harrison and Rodríguez-Clare (2009); Kugler and Verhoogen (2012); Banerjee (2013); Bloom, Eifert, Mahajan, McKenzie, and Roberts (2013). There is also a rich development literature on the under-adoption of technology and profitable business strategies. Most studies have focused on credit access, risk aversion and under-experimentation (e.g., Hausmann and Rodrik (2003); Munshi (2004); Foster and Rosenzweig (2010); Duflo, Kremer, and Robinson (2011); Bryan, Chowdhury, and Mobarak (2014)). The costs of experimentation is conceivably low in this setting. However, small firm size and market fragmentation may prevent the take-up of a new technology that involves high fixed costs.

[6]Also see Manova and Zhang (2012) and Park, Yang, Shi, and Jiang (2010) in the context of international trade.

[7]Studies have found that Information frictions play an important role in trade–e.g., see Allen (2014); Startz (2016)

[8]Relatedly, prior literature has also studied how competition affects innovation (Aghion, Bloom, Blundell, Griffith, and Howitt, 2005) and how vertical contracts affect technology adoption (Acemoglu, Antràs, and Helpman, 2007).

[9]For theoretical insights, see Kranton (2003) and Villas-Boas (2004). Macchiavello, Morjaria, et al. (2016) also considers an asymmetric information setting and shows competition can backfire by weakening the value of relationships.

in developing countries (e.g., Atkin, Khandelwal, and Osman (forthcoming); Casaburi and Reed (2014); Hardy and McCasland (2016); Bergquist (2016)). While intermediary traders and other supply chain actors have received significant attention, the final retail markets are far less studied. This study focuses on the downstream retailers, but in fact the lack of quality differentiation is seen at every stage of the watermelon value chain, from the farmers to the middlemen, and in the wholesale markets. One bottom-up solution is that if quality can be priced in the downstream, such incentive may trickle up and generate pressure to improve quality for the upstream producers, much as the spillovers via backward linkages documented in the FDI literature (Javorcik, 2004).

To estimate the demand model, this paper builds on the literature of estimating consumer learning models for experience goods (Ching, Erdem, and Keane, 2013). The framework can be used for analyzing the introduction of new goods in other settings where researchers could combine market-level price data with individual-level purchasing data. With the exception of Huffman, Rousu, Shogren, and Tegene (2007), most empirical papers have tended to focus on learning along the horizontal dimension of taste (e.g., Dickstein (2014)). This paper instead examines consumer learning on the vertical dimension. I further integrate the demand model with a supply-side model to study firms' incentive and endogenize quality.

The remainder of this paper is organized as follows. Section 2 describes the setting. Section 3 outlines a conceptual framework. Section 4 describes the experimental design and the data. Section 5 presents the experimental results. Section 6 estimates an empirical model of learning and quality provision. Section 7 uses the structural estimates to examine the welfare implications of information frictions and fragmented markets. Section 8 concludes. Alternative theory, sampling and data collection, additional reduced form results, and technical details for the structural analysis are provided in online appendix.[10]

2 The Retail Watermelon Markets in China

Most consumer goods transactions in developing countries take place in semi-formal, open air, local retail markets, which are typically located near clusters of gated residential communities and operate throughout the year with little turnover (Appendix Figure 1). Each market houses a fairly large

[10]The online appendix can be found here: https://sites.google.com/site/jiebaiecon/research

number of fragmented small-scale retailers operating side by side and selling many different types of daily food products. Most of the retailers procure their products from the same big wholesale market in the city–they are the final link in the long supply chain for agricultural products. This study focuses on watermelon sales of the downstream fruit sellers across local markets in Shijiazhuang, China. I begin by describing four stylized facts:

Fact 1. The markets are highly localized with repeated interactions between sellers and consumers.

Most watermelon transactions take place in local markets. On average, households in summer purchase 1 to 2 watermelons per week, and around 80% of the purchases are made from the local markets (see Table 1). Given this long-term repeated interactions among local sellers and consumers, market-based reputational incentives could potentially motivate quality provision in this setting.

Fact #2. Quality varies considerably across watermelons (of the same breed) within the same batch. While consumers find it difficult to detect the underlying quality at the point of transaction, sellers can assess quality based on less obvious observables.

This fact is supported by ample anecdotal evidence. To formally establish the presence of information asymmetry in this setting, I conducted a sorting test with 30 fruit sellers in 30 different local markets in the city. Each of them was asked to sort 10 watermelons into two piles: one for high quality and one for low quality.[11] The watermelons were randomly picked by surveyors from the sellers' stores with no obvious distinguishable differences in outlook. The same test was repeated with 5 randomly chosen local consumers in each market. Finally, quality was measured using a sweetness meter. A baseline blind tasting test shows that sweetness strongly correlates with consumer's taste: among 210 consumers who were asked to compare two watermelons of high and low sweetness measures, 97% preferred the sweeter one.

The lightest gray line of Figure 1 plots the cumulative sweetness distribution of all 300 watermelons. To give a sense of the scale, a sweetness difference of 0.5 matters significantly for taste–sweetness above 10.5 is considered to be very good and that below 9 is very bad. A one-way analysis of variance shows that 70% of the variation is explained within sellers; in other words, quality varies within single batches of watermelons at each given store. The darker grey lines compare the sweetness dis-

[11] Specifying a fixed number of watermelons for each pile may wash out differences between skilled and unskilled subjects, while not doing so can lead to trivial sorting. In practice, the maximum and minimum for each pile are set to be 7 and 3 respectively. On average, sellers sorted 4.4 watermelons to the premium pile and consumers sorted 3.5.

tribution of the premium piles sorted by sellers and consumers. A two-sample Kolmogorov-Smirnov test rejects the equality of distributions at 1% level.

There are two main takeaways from the CDF plots: first, sellers are much better than consumers at assessing quality, demonstrating asymmetric information between the two sides of the markets. Second, however, sellers' ability to differentiate is also far from being perfect. This inherent noise in quality control can significantly impede sellers' ability to signal quality–unless consumers are willing to experiment and upgrade their perception, trust can take a long time to establish.

Fact #3. Consumers are heterogeneous in their willingness to pay for quality.

To elicit willingness to pay for quality, households were asked in the baseline survey to consider a hypothetical situation wherein two piles of watermelons are sold in the local markets: one pile of ordinary quality sells at 1.5 RMB/Jin; the other of premium quality sells at a higher price. Surveyors announced the premium price from high to low and recorded the highest number that led to the choice of the premium pile.[12] Figure 2 plots the cumulative distributions of the reported willingness to pay for households in different income and age groups. Willingness to pay is higher for households with higher income (left figure) and for non-elderly households (right figure).

Fact #4. In contrast to many other fruits sold in these markets, there is a stark absence of quality differentiation for watermelons at baseline.

Despite the underlying variation in quality within each batch of watermelons, sellers sell an undifferentiated pile and do not price watermelons by quality. Within each local market, there is also little price variation across sellers. This contrasts sharply with other fruits, including peaches, cherries, bananas, and grapes, for which we observe substantial quality differentiation: for example, sellers usually sell multiple bins of peaches (of the same breed) at different prices. One feature that distinguishes watermelons from many other fruits is that the quality of the latter can be relatively easy to assess by consumers at the point of transaction—for example, a nice peach or banana looks different from a rotten one; for cherries and grapes, consumers can just pick one and taste it.[13]

[12] 1 Jin $\approx$ 1.1 pounds. The rest of the paper uses Jin as the unit for price. Prices (in RMB/Jin) for the premium pile were announced in the following order: 2.5, 2.2, 2, 1.9, 1.8, 1.7, 1.6, and 1.5.

[13] In many settings, giving out free samples can be an easy way to signal quality. However, since quality varies within single batches of watermelons, the quality of one is not indicative of the quality of others; it is too costly for sellers to cut open every single watermelon because once open, it is hard to preserve under high temperature.

This leads us to think that asymmetric information may be playing a role for the absence of quality differentiation for watermelons. However, this does not immediately imply a market failure: first, it depends on people's valuation for quality and effort costs of providing quality; second, relational contracts, which commonly exist in these markets (Fafchamps, 2002), may have fully mitigated the information problem. For example, sellers can pick higher quality watermelons and give to repeat customers, and may use watermelons as a tool to build relationships and maintain businesses for other fruits they sell. To the extent that we observe substantial switchings across sellers among households (only 1 out of 675 households in the study sample had all of its reported fruit purchases from a single designated local seller) and such preferential treatment may not perfectly align with people's willingness to pay, there would still be welfare losses due to allocative inefficiency.[14]

The next section sets up a conceptual framework for understanding the lack of quality premium in this setting and motivates an experimental design to tease apart the competing hypotheses.

3 Model: Quality Provision with Asymmetric Information

The framework is adapted from Shapiro (1982). I first set up the model and specify the assumptions.

3.1 Basic Setup

Supply side: A single long-run seller faces a fixed pool of consumers. Time horizon is discrete and infinite. The seller maximizes expected discounted sum of profits with discount factor $\delta \in (0, 1)$.[15]

In each period, the seller could choose to sell just one "normal" product, or she could choose to introduce a new "premium" product and sell both. I call the former "no differentiation" and the latter "differentiation."

The per-unit cost (P_W) and price (P_N) of the normal product are assumed to be fixed. Let $\underline{\gamma}$ denote the quality of the normal product, where quality is operationalized as the probability that a

[14]The study collects data on both watermelons and peach sales, the second popular summer fruit. If relational contracting has already *perfectly* allocated high quality watermelons to high valuation customers (i.e., there is no market failure), we would not expect to see an effect on sales outcomes when sellers were induced to differentiate quality under the experiment (see Section 5.3). Similarly, if there are significant "relational spillovers" to other fruits, we would expect to see an impact on their peach sales (see results in Appendix D.1).

[15]The model here abstract away from market competition. The assumption is made to match the experimental setting: in each market, only one seller was incentivized to differentiate quality, and there was little strategic response from the others (see Section 6). Therefore, I work with a monopoly model here and defer the counterfactual analysis of oligopolistic competition to Section 7.

consumer finds the product satisfactory. $\underline{\gamma}$ is exogenously fixed and known by consumers.

If the seller chooses to introduce a premium product, she chooses the quality γ_H, which is initially unobserved by consumers.[16] The extra marginal cost of the premium product is $C(\gamma_H; \underline{\gamma})$, where $C_{\gamma_H} = \partial C/\partial\gamma_H > 0, C_{\gamma_H\gamma_H} = \partial^2 C/\partial\gamma_H^2 > 0$. In this setting, C can be thought of as inspection effort costs made by downstream sellers when sourcing watermelons from the upstream. The seller also sets the price of the premium product, denoted as P_H^t. To focus on the seller's optimal policies of the premium option, I assume that the price and quality for the normal product are held the same as that under no differentiation. The main qualitative takeaways from the model do not hinge on this assumption. In the empirical analysis, I shall take a closer look at sellers' actual quality and pricing behavior when they start to differentiate under the experiment.

Demand side: There are many ways that one could model consumers' behavior and beliefs when the seller introduces a premium option. The model here focuses on the aspect of consumer learning, which may play an important role for newly introduced experience goods. In this setting, consumers are not informed about the experiment, therefore it is plausible from their perspective to regard the new product as coming from some alternative *upstream source* with some underlying quality that is initially unknown but can be learned over time via actual consumption experiences. We can think of the behavior learning dynamics below as a reduced form way of capturing learning in a larger Bayesian game in which consumers are trying to infer the *supplier's type*.

To model the learning process, I adopt a similar framework to that in Dickstein (2014). Suppose that prior beliefs about γ_H follow a beta distribution with parameters (a_0, b_0), where a_0 can be interpreted as the number of prior good experiences and b_0 as the number of prior bad experience. The prior mean is given by $\mu^0 = \frac{a_0}{a_0+b_0}$. Let e_t denote period t's experience realization, which is a Bernoulli random variable with success (satisfactory) probability γ_H. For analytical tractability, I assume that all consumers receive the same experience shock e_t in each period when they purchase the premium product and that information is shared to those who do not purchase by word of mouth.[17] Since beta distribution is the conjugate prior for Bernoulli likelihood, beliefs in period t,

[16]The model analyzes the case of a once-for-all quality choice. In principle, it is possible for sellers to adjust quality and price in every period, however that period is defined. Section 3 of Shapiro (1982) considers such a case and the qualitative conclusions are similar: (1) asymmetric information could lead to quality deterioration and (2) prior beliefs matter for seller's incentive to provide quality.

[17]In reality, consumers receive different experience shocks in each period and it is more natural to think of γ_H as

after a sequence of experience realizations $\mathbb{e}^{t-1} = (e_1, \dots, e_{t-1})$, simply follow a beta distribution with parameters $(a_0 + s_{t-1}, b_0 + f_{t-1})$, where s_{t-1} and f_{t-1} are the number of satisfactory and non-satisfactory experiences up to time $t-1$.

In each period, consumers either buy one unit of the product or do not buy any product at all. The utility of not buying is normalized to 0. Consumers' valuation is uniformly distributed between $[\underline{\theta}, \overline{\theta}]$ with mass M. For a consumer with valuation θ who buys a product at price P, the utility is $\theta - P$ if the product is satisfactory and $-P$ if it is not. In each period, consumers make their purchase decisions to maximize the expected current period utility.

The seller's problem: The seller chooses whether to differentiate by quality or not.[18] Let $Q^t_{N,\text{nodiff}}$ denote the demand under no differentiation, $Q^t_{H,\text{diff}}$ and $Q^t_{N,\text{diff}}$ denote the demand for the premium and normal products under differentiation. Under no differentiation, the seller's discounted sum of profits are fixed, given by the parameters of the model:

$$\Pi_{\text{nodiff}} = \sum_{t=1}^{\infty} \delta^{t-1}(P_N - P_W) Q^t_{N,\text{nodiff}} \qquad \text{where} \quad Q^t_{N,\text{nodiff}} = (\overline{\theta} - \frac{P_N}{\underline{\gamma}})\frac{M}{\overline{\theta} - \underline{\theta}} \tag{1}$$

Under differentiation, the seller faces a dynamic demand system. In particular, $Q^t_{H,\text{diff}}$ and $Q^t_{N,\text{diff}}$ are functions of $\mu^{t-1}(\mathbb{e}^{t-1}(\gamma_H); a_0, b_0)$, which evolves over time as consumers learn.[19] For a given γ_H, the optimal P^t_H is imposed by static profit maximization. (Because the stylized model assumes complete information diffusion, there is no dynamic implication of current sales.) The expected discounted sum of profits under γ_H is

$$\Pi_{\text{diff}}(\gamma_H) \equiv \mathbb{E}\left[\sum_{t=1}^{\infty} \delta^{t-1} \max_{P^t_H} \Big(\big(P^t_H - P_W - C(\gamma_H; \underline{\gamma})\big) Q^t_{H,\text{diff}} + (P_N - P_W) Q^t_{N,\text{diff}} \Big)\right] \tag{2}$$

---

the mix of good watermelons at a given point in time. In the structural estimation, I enrich the model by allowing individuals' beliefs to diverge over time with observed experience realizations in the data.

[18] It is possible that the profit maximization decision is to only sell the premium product. This happens when costs of providing quality is very low. However, such behavior is not observed and I exclude the case here for convenience.

[19] Demands are determined by cutoff types as in standard vertical taste models. For interior solutions, we have:

$$\begin{aligned}
Q^t_{H,\text{diff}}(P^t_H, P_N, \mu^{t-1}(\mathbb{e}^{t-1}(\gamma_H); a_0, b_0); \underline{\gamma}) &= \left(\overline{\theta} - \frac{P^t_H - P_N}{\mu^{t-1}(\mathbb{e}^{t-1}(\gamma_H); a_0, b_0) - \underline{\gamma}}\right)\frac{M}{\overline{\theta} - \underline{\theta}} \\
Q^t_{N,\text{diff}}(P^t_H, P_N, \mu^{t-1}(\mathbb{e}^{t-1}(\gamma_H); a_0, b_0); \underline{\gamma}) &= \left(\frac{P^t_H - P_N}{\mu^{t-1}(\mathbb{e}^{t-1}(\gamma_H); a_0, b_0) - \underline{\gamma}} - \frac{P_N}{\underline{\gamma}}\right)\frac{M}{\overline{\theta} - \underline{\theta}}
\end{aligned}$$

where the expectation is taken over sequences of experience shocks $\{e_t\}_{t=1}^{\infty}$ generated by γ_H. Let γ_H^* denote the argmax of $\Pi_{\text{diff}}(\gamma_H)$ and $\Pi_{\text{diff}}(\gamma_H^*)$ the maximized expected value under differentiation.

Suppose there is an initial fixed cost F of introducing a premium option, and the seller chooses to differentiate if and only if $\Pi_{\text{diff}}(\gamma_H^*) - F > \Pi_{\text{nodiff}}$.[20]

This completes the setup of the model. In Section 6, I provide some descriptive evidence on the model's key assumptions and enriches the basic setup in the empirical model by incorporating greater dimensions of consumer heterogeneity, private experience shocks, and market competition. For the remainder of this section, I work with the basic framework to derive some testable implications, which motivate the experimental design and inform the reduced form analysis.

3.2 The Effects of Prior Beliefs

In light of the model, there are two broad explanations for the lack of quality differentiation at baseline: high costs and asymmetric information. First, if cost is high relative to consumers' valuation for quality, for instance $\overline{\theta} < C(\underline{\gamma};\underline{\gamma}) + P_W$, then higher quality may not be demanded and supplied even under symmetric information. Second, since a seller's claim of offering high quality cannot be immediately verified, consumers' prior beliefs matter. Sellers who rationally discounts future profits may lack the incentive to provide quality if trust takes a long time to establish. Hence, markets can be stuck in an equilibrium with no quality differentiation.

In reality, these two aspects act jointly. However, the welfare implications are very different: under the former, the distortion on quality provision caused by the information problem is small, whereas under the latter it could be large.

In practice, it is hard to directly infer costs. To understand the main barrier for quality provision, the experiment aims to create exogenously variations in prior beliefs. These variations should have minimal effects if the key barrier for quality provision is high costs. On the other hand, if the information problem is the key barrier, enhancing prior beliefs could significantly strengthen sellers' incentives to provide quality. The effects are stated in the following two propositions:

Proposition 1: *(Incentive to provide quality)* $\Pi_{\text{diff}}(\gamma_H)$ increases with a_0 and decreases with b_0.

[20] F is not needed for deriving the comparative statics. Without F, if non-differentiation is the optimal strategy under asymmetric information, it is also the optimal strategy under symmetric information as well as under the first best because only the highest valuation ($\overline{\theta}$) matters for the decision on the extensive margin. However, this is a knife-edge scenario–any positive initial costs of introducing the premium product could break it.

Enhancing prior beliefs, either by increasing a_0 or decreasing b_0, raises the seller's return under differentiation. The intuition is straightforward. In particular, holding a_0 fixed, a lower b_0 implies a higher prior mean and a larger prior variance, and hence a faster speed to establish trust and larger discounted returns. The next proposition examines how the seller's optimal quality choice responds to prior beliefs if she differentiates.

Proposition 2: *(Optimal quality choice)* If $\frac{\partial^2 \Pi_{\text{diff}}(\gamma_H^*; a_0, b_0)}{\partial \gamma_H \partial a_0} > 0$, γ_H^* increases with a_0. Similarly, if $\frac{\partial^2 \Pi_{\text{diff}}(\gamma_H^*; a_0, b_0)}{\partial \gamma_H \partial b_0} < 0$, γ_H^* decreases with b_0.

Proposition 2 states a simple monotone comparative statics: if prior beliefs and quality are complementary, the seller will be induced to provide higher quality when prior beliefs improve.[21] In reality, optimistic beliefs encourage sales, which enables information to spread faster, and thus rewards good behavior and punishes bad behavior faster. This channel is absent in the stylized model (with perfect information diffusion) but will be captured in the empirical model.

The next section describes the experiment and relate the treatments to this framework for thinking about their effects.

4 Experimental Design and Data Collection

4.1 Experimental Design and Timeline

Overview. The experiment was conducted in Shijiazhuang, the capital city of Hebei province, China.[22] The city has over 800 gated communities and more than 200 local markets. Randomization happened at the market level. 60 sellers located in 60 different markets were recruited to participate in the study following an initial screening and a sequential selection procedure to minimize heterogeneity in the study sample for power concerns and logistical purposes. Details for the screening process and selection criteria are described in Appendix C.1.

There are typically 3 to 5 sellers in each local market; only 1 was selected. In what follows, I call the 60 sellers the *sample sellers*, as opposed to the *other sellers*, who were not directly involved in the experiment but nonetheless operate businesses in these markets. All sample sellers were asked

[21] In general, γ_H^* is non-monotonic in a_0 and b_0. When $a_0 + b_0$ is sufficiently large, as one of the two parameters tends to 0 (i.e. very pessimistic or very optimistic beliefs), the incentive to provide quality vanishes.

[22] Urban area: 399.3 sq km (154.2 sq mi); urban population: 2,861,784; urban density: 7,200/sq km (19,000/sq mi)

to experiment with selling two piles of watermelons: a premium pile and a normal pile. Sellers were free to choose quality, price, and quantity for each pile. They were randomized into 6 groups:

Branding treatments. Sellers were randomized into one of the three branding groups: laser, sticker and label-less. Every morning, surveyors visited the sellers' stores and performed a free branding service. For the laser group, surveyors used a laser-engraving machine to print a laser-cut label of the words "premium watermelon" ("Jing Pin Xi Gua" in Chinese Pinyin) on the watermelons in the premium pile. For the sticker group, surveyors pasted a sticker with the same words. For the label-less group, surveyors did nothing. Note that the branding treatment was only for watermelons in the premium pile, picked by the sellers themselves; those in the normal pile were left as they were. Figure 3 shows pictures of the branding treatments. Most sellers sold two piles of watermelons at the beginning of the experiment, but some reverted back to non-differentiation after some time. For those sellers, branding was withdrawn because there was no longer a premium pile.

A cross-randomized incentive treatment. Within each branding group, half of the sellers were randomly given an incentive to maintain quality for the premium pile. The incentive treatment was enforced via unannounced quality checks twice per week. At every check, surveyors randomly picked one watermelon from the premium pile and one from the normal pile. The quality of both was measured using a sweetness meter (Appendix Figure 2). For sellers in the incentive group, if the sweetness of the former attained 10.5 at both checks, sellers received a monetary reward of 100 RMB at the end of the week, on par with daily sales profits. Sellers in the non-incentive group received the same quality checks, but were not given any reward. The incentive was removed in the later part of the intervention, and that was unanticipated by the sellers.

In total, there were 6 distinct treatment units. Randomization was stratified on housing prices, i.e. a dummy variable indicating whether the baseline average housing price in the surrounding gated communities is below or above the median. Appendix Figure 3 shows a map of the 60 sellers, marked by groups. Note that these markets are geographically segregated and the average distance between any two markets is 3 kilometers. Since watermelon transactions are highly localized, spillover effects across markets should be minimal.

Figure 4 describes the timeline of the study. The market intervention was rolled in from July 13 to July 19, 2014. Two weeks into the intervention, a universal announcement was made to all sellers that they were free to decide whether they wanted to continue with quality differentiation or not. This allows me to examine differential incentives across groups. Six weeks into the intervention, the incentive was removed. The intervention was phased out from September 6 to September 12. An endline survey was conducted upon the surveyors' final visit to sellers' stores, and two follow-up surveys were conducted to examine longer-term outcomes.

4.2 Testable Implications

To predict the effects of the treatments on sellers' behavior and market outcomes, I relate the experimental design to the framework in Section 3 and discuss two ways in which laser branding could potentially affect consumers' prior perception in the context of China.

First, the cost of laser machines is very high ($\approx$ 8k USD). Evidence from consumer pilot survey suggests that consumers regard laser branding as more effective at deterring counterfeits than stickers, which can be cheaply fabricated and highly "contaminated" by rampant counterfeiting activities in the past. This is also true in many other developing countries where brand protection is weak.[23] This discussion relate to the theory of collective reputation studied in Tirole (1996). A key takeaway from the model is that equilibrium could be history dependent. The new "uncontaminated" laser branding could potentially wipe out bad historical stereotypes associated with stickers, thereby allowing trust to establish faster. Proposition 1 and 2 suggest that sellers in the laser group would have a stronger incentive to differentiate and provide higher quality.

Second, laser branding represents a large conspicuous sunk investment, which could signal the presence of a price premium that is high enough to motivate high quality. The formal argument is known as forward induction (discussed in Appendix B). This argument also suggests that consumers' prior perception may be more optimistic under laser.

Yet a third potential effect of laser branding is that it simply represents something "cool" and directly affects utility other than signaling quality. However, sales dynamics between laser incentive and laser non-incentive groups rule out a pure "coolness" story. The empirical model explicitly include a laser-specific constant in consumers' indirect utility function to account for this possibility.

Finally, we can think of the incentive treatment as shifting the posterior beliefs. The idea is

[23] See for example studies by Björkman-Nyqvist, Svensson, and Yanagizawa-Drott (2013) and Qian (2008).

based on the following: if the incentive could motivate sellers to provide higher quality, then upon its removal after some period T, sellers who have had the incentive are essentially endowed with higher beliefs than those sellers who have not had the incentive. Proposition 2 suggests that higher quality may be sustained even in the post-incentive period. This, therefore, provides a further test for the forces in the model, and also helps to identify effort costs in the structural estimation.

4.3 Data: Overview

Baseline surveys. Table 1 summarizes the sample characteristics. On average, a local market houses 3 to 5 fruit sellers. Most sellers sell fruits all year long and do not expect to relocate. The median household consumes 1 watermelon per week in the summer, and 75.6% of the households list the local market as the main source for watermelon purchases.

Supply side: quality, prices, and sales. Quality data (measured in sweetness) were collected from the biweekly random quality checks as described in Section 4.1. Surveyors' collected daily retail prices for both the sample sellers and the other sellers in these markets, as well as the daily wholesale price. Sellers were asked to record down their daily sales for watermelons and peaches by quality category. In total, there were 60,806 transaction records over the course of the intervention. 81% of transactions were for watermelons and 19% were for peaches. On average, sellers sell 257 Jin of watermelons per day, and the average daily sales profit is 103 RMB.[24] For the empirical analysis, I aggregate the transaction-level sales to seller-day-quality category level.

Demand side: household panel purchasing. 675 households in 27 communities, evenly distributed across the treatment groups, were recruited to record the family's summer fruit consumption experiences. For each fruit purchase, households were asked to record the date of the purchase, the place of the purchase, the quantity bought, the amount paid, whether the purchase was made from the sample seller or from other places (including other sellers in the local market, nearby supermarkets and other places), and whether the purchased fruit had any branding on it or not. Important, households were asked to report a satisfaction rating ranging from 1 to 5, where a higher number indicates higher level of satisfaction. In total, there were 15,292 purchase records, of which 30.8%

[24] Here and in all subsequent analysis, sales profit is computed using sales quantity and prices: Sales profit = premium pile price $\times$ premium pile sales quantity + normal pile price $\times$ normal pile sales quantity - total sales quantity $\times$ wholesale price. Alternatively, I can use the recorded sales values to calculate profits. Results are qualitatively robust.

were for watermelons. The median for the number of watermelons consumed per week is 1 and the mean is 1.15 with standard deviation 1.06. These numbers match well with Table 1.

Endline and follow-up surveys. The seller endline survey was conducted during the surveyors' final visit to the sellers' stores and elicited sellers' willingness to pay for different branding technologies. The household endline survey was distributed together with the last week's recording sheet and elicited households' willingness to pay for quality under different branding technologies. Two follow-up surveys were conducted one week and one year after the intervention to collect longer-term quality differentiation and pricing behavior. Attrition rate is small: 1 seller dropped out during the intervention because the market was closed for road construction. For the second follow-up, surveyors were able to locate 57 of the original 60 sellers.

Details of the recruiting procedure, data collection, and issues with omissions and errors in sellers' and households' recording were discussed in detail in Appendix C.1-C.4. Balance checks on market, seller and household baseline characteristics are shown in Appendix Tables 1 to 3.

5 Experimental Evidence on the Effects of Branding and Incentive

This section presents experimental evidence on the effects of the branding and the incentive treatments, and tests the predictions of the model in Section 3.

Figure 5 plots the number of sellers who differentiated quality at sale in each treatment group over time. We see that sellers in the label-less groups reverted back to baseline right after the announcement. On the other hand, most sellers in the sticker and laser groups continued to differentiate till the end of the intervention.

The rest of the section examines demand side's responses, sellers' quality, pricing and sales in order to understand the differential incentives, in particular, why sellers who were induced to differentiate during the intervention had not already done so at baseline.

5.1 How Do Different Branding Technologies Affect Consumers' Prior Beliefs?

This question is difficult to examine in a regression framework as a household's purchase decision in a period is affected by the entire history of past purchasing and consumption experiences, which are in turn affected by prior beliefs. Table 2 provides some suggestive evidence.

The data is collapsed to the household-week level. The dependent variable is a dummy for whether a household purchased any premium watermelons in a given week, regressed on two measures of past experiences: (1) the average lagged satisfaction rating of all past premium purchases, and (2) the percentage of past purchases that attain the highest rating. Note that if a household has never purchased any premium watermelons in the past, these measures are not defined. Therefore, the coefficients are only estimated from household-week observations conditioning on a positive number of premium watermelons being purchased prior to that given week.

Column 1 and 2 of Panel A show that lagged experiences strongly predict repurchasing decisions for households in the laser markets. To interpret the magnitudes, take the estimate in column 2, which shows that for two similar households at a given point in time, the household that has had only very good past experiences is 45% more likely to repurchase a premium watermelon than the household that has not had any very good experiences (but that has experienced the product). On the other hand, the coefficients are much smaller and noisier for households in the sticker groups, as shown in columns 3 and 4. These patterns are consistent with the discussions in Section 4.2: prior beliefs may be more "stubborn" under stickers, which implies that purchasing decisions would be less responsive to past consumption experiences.

As a sanity check, Panel B repeats the same exercise for purchase decisions of the normal pile. Since consumers have experienced unlabeled watermelons for a long time, each additional consumption experience should weigh less. Indeed, we see that the coefficients are small and insignificant.

5.2 How Do Different Branding Technologies Affect Sellers' Quality Choice?

Panel A of Table 3 compares the premium pile quality, measured in sweetness, for sellers in the sticker and laser groups by pooling together the quality checks data. Standard errors are clustered at the seller (market) level, which is the unit of randomization. This applies to all the regression analysis below unless otherwise stated. We see that on average, sellers in the laser group provide significantly higher quality than sellers in the sticker group, both with and without the incentive.

To further understand sellers' quality differentiation behavior, I look at how the quality of the premium pile compares with that of the normal pile, and at how the two compare with the market average. Specifically, I run the following regression:

$$y_{ipt} = \alpha + \beta \text{Premium}_p + \gamma_i + \lambda_t + \epsilon_{it} \tag{3}$$

The outcome variable y_{ipt} is sweetness measured for pile p of seller i at week t. The key explanatory variable is a dummy for the premium pile. Thus, α represents the mean of the normal pile and β measures the average difference between the two piles. To focus on the effect of the branding treatment, I restrict the sample to sellers in the non-incentive groups and estimate Equation 3 separately for the laser and sticker groups, controlling for seller (γ_i) and time (λ_t) fixed effects.

Panel A of Table 4 shows that the average quality of the premium pile is significantly higher than that of the normal pile. However, the difference could be either due to a genuine quality improvement of the premium pile or a deterioration of the normal pile. To examine the latter possibility, Panel B runs the same regression but with quality difference from the market average as the outcome variable. I use the average sweetness of randomly picked watermelons from sellers in the label-less group after they had reverted back to non-differentiation as a proxy for the average quality.

Column 3 shows that sellers in the laser group maintained a higher quality for the premium pile and kept the normal pile quality on par with the market average. This suggests that sellers may have spent more efforts on sourcing good watermelons in the upstream. This result alone demonstrates that reputational incentives are present and can potentially motivate quality provision. As long as providing higher quality involves positive efforts, in a one-shot game, sellers would not exert such additional efforts and would randomly label some watermelons as "premium" and sell them at a higher price.

The evidence for the sticker group is quite mixed. On average, the quality of the normal pile appears to be lower than the market average and the quality supremacy of the premium pile (sum of α and β) is not significantly different from 0 (p-value = 0.584). The large standard errors indicate that there could be considerable heterogeneity across sellers in the sticker group. Anecdotally, some sellers in the sticker group simply labeled all watermelons except for a few observationally bad ones, which they then marked down and sold as a low-end product. While the sample size is too small to formally examine heterogeneity within a treatment group, I note the difference between the sticker group and the genuine quality-price premium observed for the laser group.

5.3 How Do Different Branding Technologies Affect Sellers' Return?

To focus on the effects of the branding treatments, I restrict the sample to the non-incentive groups and run the following regression:

$$y_{it} = \alpha + \beta_1 \text{Sticker}_i + \beta_2 \text{Laser}_i + \lambda_t + \epsilon_{it} \tag{4}$$

The outcome variables are log sales profits (in RMB), markup from market average price (in RMB/Jin)[25] and sales quantity (in Jin) for each pile, and the total sales quantity for seller i on day t. If a seller stops differentiating quality, the unit price for the premium pile is defined to be the same as that for the normal pile and sales quantity for the premium pile is coded as 0. Results are shown in Table 5. Sticker and laser are group dummies and the omitted group is the label-less group. All regressions include day fixed effects (λ_t) to control for time-specific aggregate shocks, such as weather. The even columns control in addition for community and seller baseline characteristics.

Column 1 and 2 show that on average, the laser group earns 30-40% higher sales profits than the label-less group. This is due to both a higher unit price (columns 3 and 4) and higher sales quantity for the premium pile (columns 5 and 6). Sales of the normal pile are not significantly different from those of the label-less group. The results suggest that sellers in the laser group attract more high-end customers without losing sales on the normal product. On the other hand, for the sticker group, sales of the premium pile appear to be lower on average than the laser group (the p-value of a one-sided test is 0.238) despite a lower markup. Furthermore, sales of the premium pile (columns 5 and 6) are offset by a reduction in the sales of the normal pile (columns 9 and 10). As a result, total sales and profits are not significantly different from those of the label-less group.

These results explain why sellers did not differentiate quality at baseline despite the fact that stickers have long been cheaply available. While sellers in the laser group earned higher sales profits, the relevant consideration is whether the increase in profits, netting out the effort costs of providing higher quality, justifies the fixed cost of the laser machine. One year after the intervention, when surveyors revisited these markets, none of the 57 sellers that could be tracked continued with quality

[25] Here and in all subsequent analysis with prices, I use the listed prices as observed by surveyors during the morning visits to the markets. Alternatively, for the sample sellers, I can use the effective prices, which are calculated as total daily sales values divided by total daily sales quantity for each quality category. Results look very similar and all the qualitative conclusions are robust.

differentiation. This suggests that individual sellers would not have the incentive to take up this new technology themselves. The structural analysis will help to shed light on why that is the case, which is related to the structure of these markets.

5.4 How Is Seller's Quality Choice Affected by the Incentive Treatment?

Panel B of Table 3 shows that the incentive did lead sellers to provide higher quality for both sicker and laser groups. To examine whether higher quality was sustained in the post-incentive period, Table 6 runs a difference-in-difference regression. The coefficient for the interaction term between the incentive treatment and the post-incentive dummy is close to zero and non-significant for the laser group. On the contrary, sellers in the sticker incentive group seemed to revert to a lower quality level after the incentive was removed. These results are consistent with the fact that it may take longer to establish trust under stickers. Therefore, it is not clear how much the incentive has facilitated initial learning within this short intervention.

Appendix D.2 and D.3 present richer analysis exploring the time dynamics of sales and changes in household perception elicited at endline to corroborate the above findings. Overall, the reduced form results are consistent with the model's predictions and provide a qualitative explanation for the lack of quality differentiation at baseline. The next section structurally estimate the model to rationalize the experimental findings and perform counterfactual analysis.

6 An Empirical Model of Learning and Quality Provision

The empirical model follows the same setup as the model outlined in Section 3. I first enrich the basic setup in Section 6.1. Estimation proceeds in two steps (Section 6.2). First, the dynamic demand model is estimated using the household panel data. Second, the supply-side parameters are calibrated by solving for the sellers' optimal policies, taking the demand estimates as given. Section 6.3 discusses the results and examines model fit. Section 6.4 uses the structural estimates to simulate consumers' beliefs and sellers' net returns evolution under each treatment.

6.1 Setup and Assumptions

6.1.1 Demand Side: A Model of Consumer Learning and Purchasing

I start by restating the key assumptions and providing some qualitative justifications.

Assumption 1 (*Demand side*): (1) Consumers share a common prior about the unobserved quality, which is believed to be fixed over time (for a given type of watermelon); (2) Consumers update only on the premium option; (3) Consumers make purchasing decisions based on current expected utility.

Assumption 1.1 is discussed in Section 3. Quality is operationalized as the probability of being good. Assumption 1.2 is consistent with the reduced form results in Panel B of Table 2 (discussed in Section 5.1). To model forward-looking behavior, one needs to solve a dynamic discrete choice problem. Besides the usual computational difficulties (discussed in detail in Ching, Erdem, and Keane (2013)), the current setting poses an additional challenge, which is that it may be hard to model the value of experimentation in the context of a new good as consumers' perceptions about future product availability, price and quality would matter. The goal of the empirical exercise is to estimate a parsimonious model that describes consumers' actual purchasing behavior, and that is also tractable enough to be integrated with the supply side.[26] As a first pass, given the seasonal nature of the fruit, if the option value of experimentation plays an important role, we would expect that the number of first-time buyers for the less-known premium option to be higher in the initial period. However, there does not appear to be such a pattern in the data (Appendix Table 7).

Priors, distribution of outcomes, and updating

The prior distribution and the updating process are described in Section 3. Here, I enrich the setup by incorporating private experience shocks and an enlarged choice set that includes buying from other sellers. Let $e_{imjt} \in \{0, 1\}$ indicate whether a type j watermelon is satisfactory or not for individual i in market m at time t. There are three types of watermelons: $j \in \{1, 2, 3\}$, where

[26] In practice, it is difficult to combine a complex dynamic demand system with a supply-side model (e.g., see Ching (2010) and Hendel and Nevo (2011)), and estimating a full dynamic game between forward-looking heterogeneous consumers and sellers under asymmetric information remains as an empirical challenge. Fershtman and Pakes (2012) propose an equilibrium concept, called the Experience Based Equilibrium, where players choose their optimal strategies based on their own observable experiences. The authors provide an estimation framework that is based on a reinforcement learning algorithm. In a similar spirit, one could view the beliefs evolution in this empirical model as consumers learning to converge to a steady state.

$j=1$ indicates the premium pile from the sample seller, $j=2$ indicates the normal pile from the sample seller, and $j=3$ indicates those from all other sources. Prior beliefs about the quality of the premium option are assumed to follow beta distribution with parameters (a_0, b_0). The posterior at time t is given by $(a_{im1t}, b_{im1t}) = (a_0 + s_{im1t}, b_0 + f_{im1t})$, where s_{im1t} and f_{im1t} are the numbers of satisfactory and non-satisfactory experiences individual i has had after time t. By Assumption 1.2, consumers do not update on the other options. Let q and $q + \Delta q$ denote the (degenerate) beliefs about the quality of other sources and the normal option. Δq captures any spillover effect.

Decision rule

Consumer's expected utility of purchasing option $j \in \{1,2,3\}$ at time t is

$$u_{imjt} = (\theta_0 + \theta_1 \text{WTP}_i)\mu_{imj,t-1} - (\alpha_0 + \alpha_1 \text{Highinc}_i) P_{mjt} + \beta \text{Num}_i + \eta_i \mathbb{1}_{(j=1)} + \xi_i \mathbb{1}_{(j\in\{1,2\})} + \lambda_m + \lambda_t + \epsilon_{imjt}$$

where $\mu_{imj,t-1}$ denotes consumer i's posterior for option j at the end of time $t-1$. P_{mjt} is j's price in market m at time t. θ captures vertical taste differentiation, and is allowed to vary across consumers with different baseline self-reported willingness to pay for quality. The price coefficient α is allowed to be different for high- and low-income groups. Num_i is the number of watermelons consumed per week reported at baseline, which seeks to capture heterogeneous love for watermelons in general. η_i and ξ_i are unobserved preferences for the premium option and for the sample seller. For example, some consumers may have a predilection for expensive products, and some may be more likely to buy from a particular seller than from the others (i.e. horizontal taste differentiation). λ_m are market fixed effects, capturing time-invariant differences across markets. λ_t are time fixed effects, capturing aggregate time shocks that affect all markets, such as weather shocks.[27] ϵ_{imjt} are idiosyncratic random utility shocks, which are realized in each period before the purchasing decision is made. Let V_{imjt} denote the mean utility, excluding the random shock component.

There is an outside option with mean utility 0 for not purchasing any watermelon in a given period (denoted as $j=0$). Consumer chooses j with the highest expected utility:

$$\max_{J=\{0,1,2,3\}} V_{imjt} + \epsilon_{imjt}$$

[27] In estimation, I exclude the time fixed effect for the first period, thus the estimated time effects are relative to the first week. I estimate the full set of market fixed effects (as the utility specification does not contain a constant term).

Further assuming that the idiosyncratic shocks ϵ_{imjt} follow i.i.d. Type 1 extreme value distribution, the choice probability takes a logit form:

$$\text{Prob}_{imjt} = \frac{\exp(V_{imjt})}{\sum_{k=0}^{3}\exp(V_{imkt})}$$

6.1.2 The Supply Side: A Model of Quality Provision

For the supply side, I focus on the laser groups, for which we have seen clear evidence for providing quality. Sellers choose prices and quality to maximize the net present value of profits:

$$\max_{\{p_H^t, p_N^t, \gamma_H^t\}} \sum_{t=1}^{\infty} \delta^{t-1}\mathbb{E}\Big(\left(p_H^t - p_W^t - C(\gamma_H^t)\right) Q_H^t + (p_N^t - p_W^t)Q_N^t + \mathbb{1}_{\text{Inc}} \times \phi(\gamma_H^t)B\Big) \tag{5}$$

$$\text{s.t.} \quad \{p_H^t, p_N^t\}, g(\mu^t, X) \rightarrow Q_H^t, Q_N^t \rightarrow_{\gamma_H^t} g(\mu^{t+1}, X), \quad g(\mu^0, X) \text{ given}$$

where $g(\mu^t, X)$ is the joint distribution of household beliefs (μ^t) and characteristics (X) included in the demand model, and it constitutes the seller's state variable. The evolution depends on the prior, the learning dynamics and seller's policies.

The per-unit cost of the normal product is the wholesale price p_W, and the additional marginal cost of providing the premium product is parameterized as:

$$C(\gamma_H) = c\log\left(\frac{1-\underline{\gamma}}{1-\gamma_H}\right)$$

where $\underline{\gamma}$ denotes the average quality of the undifferentiated pool. $C(\gamma_H)$ captures the effort costs of sourcing better watermelons in the upstream. In the extreme case, if $\gamma_H = \underline{\gamma}$, the cost simply reduces to 0. Finally, the objective function for the incentive group contains an additional term $\phi(\gamma_H^t)B$ capturing the expected incentive payment: $B = 100$ RMB; $\phi(\gamma_H) = \gamma_H^2$ to match the empirical setting (since quality checks were conducted twice per week).

The main estimation challenge for solving the dynamic optimization problem is that the state space is of infinite dimension. To make progress, I make an important simplifying assumption:

Assumption 2 (*Supply side*): Seller pegs the normal pile price at the market average in each period and chooses a once-for-all quality (γ_H) and markup (m_H) for the premium pile: $p_H^t = p_N^t + m_H$.

Appendix Figure 8 and 10 plot the price and quality trajectories for the laser groups. We do not observe a clear time pattern.[28] Appendix Table 11 further examines the time dynamics in a regression framework, and the coefficients for the time variables are very close to zero.

The empirical patterns above provide some qualitative justification for the assumption. One explanation could be that frequent price adjustments may send some negative signals to consumers, and although quality differentiation happens daily, to actually fine-tune that to actual demand conditions may be hard and mentally costly. Having said that, a seller may well increase markup in longer-time horizons as beliefs evolve. Unfortunately, the data, which only lasts for 8 weeks, is limited in addressing these important long-term dynamics. Given this limitation, the approach undertaken here searches for the optimal policies within the restricted class of policies.

6.2 Estimation Strategies and Identification

6.2.1 Demand Side: Simulated Maximum Likelihood Estimation

The demand model is estimated using simulated maximum likelihood (Train (2009)). I collapse the household panel purchasing data to household-week level and merge that with the market-week level average prices. For each purchase experience, the household reports a satisfaction rating from 1 to 5. I recode 5 to be satisfactory and $\{1, 2, 3, 4\}$ as well as missing values to be non-satisfactory.[29] To allow prior beliefs to be different under different branding technologies, I estimate separate a_0 and b_0 for laser and sticker. We can think of households living in different markets as facing different choice sets: households in the laser and sticker markets face a premium option with either a laser or a sticker label. For households in the label-less markets, they face a restricted choice set without the premium option (from week 3 onwards). Finally, to allow for different spillover effects across a seller's multiple products, I estimate separate belief shifters, $\Delta q(s)$ and $\Delta q(l)$, for the normal option for sticker and laser groups.

Details for the estimation procedure and standard error calculation are provided in Appendix E.1. I now briefly discuss the identifying assumptions: first, market and time fixed effects fully capture unobserved time-varying shocks that directly affect both prices and demands for a market. Second, η and ξ fully capture unobserved persistent individual heterogeneity. Under these two assumption,

[28] There appears to be discrete jump in quality for the incentive group after the first week. This could be due to sellers having initial doubts about receiving the monetary rewards at the beginning of the intervention.

[29] Appendix E.2 discusses alternative prior specifications, including a Dirichlet's prior and different rating thresholds.

with one period data on market shares, we can identify the market specific constants, the mean of the prior beliefs multiplied by the vertical taste parameters, the price coefficients, the coefficient for *Num*, and the distributions of η and ξ (following standard arguments in the discrete choice literature). Parameters θ, a_0 and b_0 are identified from the dynamic purchasing patterns. Intuitively speaking, if repurchasing decisions are very responsive to past experiences, it could because households either care a lot about quality (large θ) or the variance of the prior is large (small a_0 and b_0). However, the *difference* in the *change* in the repurchasing probabilities between going from zero to one good (or bad) experience and that going from one to two separately identifies these parameters. In particular, the difference should be bigger under the large variance story than it is under the large willingness to pay story because belief updating is more salient for the former case.[30]

6.2.2 Supply Side: Minimum Distance Estimator

Taking the demand estimates as given, the supply side parameters are estimated using a minimum distance estimator. Ideally, one would like to solve for the optimal policies market by market and apply the minimum distance estimator to the full vector of policies for all sellers. Unfortunately, γ_H is not observed for each individual seller and cannot be reliably approximated using the empirical satisfaction rate due to the small sample size for each market. Given this data limitation, I first construct a hypothetical *average market* by pooling together all households in the laser markets and averaging the market fixed effect estimates. I then solve for the optimal policies, m_H^* and γ_H^*, for a seller facing this hypothetical *average market*. The structural parameters are calibrated by minimizing the distance between the optimal policies and the empirical average policies:

$$v(\delta, c) = \sum_{g \in \{\textbf{laser inc,laser non-inc}\}} \left(\gamma_{Hg}^*(\delta, c) - \overline{\gamma}_{Hg}\right)^2 + \left(m_{Hg}^*(\delta, c) - \overline{m}_{Hg}\right)^2 \tag{6}$$

where $\overline{\gamma}_{Hg}$ and $\overline{p}_{Hg}$ are the empirical average quality and markup. Details of constructing the hypothetical market and measuring the empirical counterparts are discussed in Appendix E.3.

For each given set of δ and c, the optimal policies are found using grid search. The objective function is minimized by searching over grids of δ and c. Intuitively speaking, low quality provision could be either due to high costs or low discount factors, but the former implies a larger quality gap

[30] Appendix Table 8 summarizes the repurchasing probabilities conditioning on the number of satisfactory and non-satisfactory experiences. The patterns are largely consistent with the reduced form results in Table 2.

between the incentive and non-incentive groups: the more convex the cost function (larger c), the steeper the increase in the costs of improving quality, which dampens the effect of the incentive.

6.3 Results and Model Fit

The simulated ML estimates are presented in Table 7. Market and time fixed effects are reported in Appendix Table 9. To match the small market shares of the premium option in the first week, I constrain the prior mean (a_0) to be zero in actual estimation (see further discussion in Appendix E.1). Estimates of other key parameters are qualitatively similar to the unconstrained case.

Looking at column 1, the estimate for b_0 is 0.938 for laser and 2.578 for sticker. The point estimates are consistent with the discussion in Section 4.2 and suggest that the prior beliefs are more *stubborn* under sticker than under laser. In particular, one satisfactory experience updates the posterior mean to 0.52 under laser, but only to 0.28 under sticker.

Beliefs about the quality of the undifferentiated option from the other sellers is estimated to be 0.307. This number matches well with the 30% empirical satisfaction rate in the household data. The negative $\Delta q(s)$ suggests that consumers in the sticker markets seem to perceive the normal pile as having lower quality if sellers sell it beside another pile that is labeled with a sticker and that is purported to be of a higher quality, which is in fact consistent with sellers' actual behavior shown in column 4 of Table 4. The signs of the other estimates are consistent with expectations.

Appendix Figure 5 and 6 examine model fit by looking at dynamics of market share and repurchasing probabilities conditioning on experiences. Overall, the purchasing patterns generated by the prior estimates and the Bayesian learning process mimic the actual purchasing patterns well.

Columns 2 to 4 of Table 7 consider three extensions to the baseline model by considering direct utility of laser, correlated learning, and information diffusion. While the measures and the approaches are not perfect, the results are reassuring. Overall the ML estimates stay quite robust across various specifications and the likelihood ratio test does not reject the baseline model.

Taking the demand estimates in column 1 of Table 7, δ and c are estimated to be 0.98 and 0.64.[31] We see that the model is able to generate a quality gap between the incentive and non-incentive groups (0.48 versus 0.41), which is fairly close to the empirical gap (0.53 versus 0.40). The optimal

[31] Appendix Figure 11 plots the value of the objective function as δ and c vary and Appendix Table 12 reports the optimal policies under various δ and c in comparison with the empirical policies.

markup for the incentive group is also higher than that for the non-incentive group, though the magnitudes are larger than the empirical values. Table 8 simulate aggregate sales outcomes using the parameter estimates and the average empirical policies. Overall, the simulated weekly average sales quantity and profits are in line with the actual sales outcomes in the data.

6.4 Beliefs and Net Returns Evolution

I now use the structural estimates to examine how beliefs endogenously evolve over time, and how prior beliefs affect seller's incentive to provide quality. The goal is to rationalize the experimental findings and provide a quantitative explanation to the lack of quality differentiation.

Panel A of Figure 6 plots the group average beliefs evolution for the quality of the premium pile. We see that the average beliefs are the highest for the laser incentive group by the end of the intervention. Conditioning on the incentive treatment, average beliefs rise faster under laser than under sticker, in line with the reduced form patterns in Table 2.

The average beliefs evolution is a result of two underlying effects: first, laser branding induces faster belief updating; second, laser branding induces sellers to provide higher quality, resulting in more satisfactory experiences. To decompose the two effects, Panel B simulates counterfactual beliefs evolution under three scenarios: (1) sticker prior and sticker group's empirical policy (dashed line); (2) laser prior and sticker group's empirical policy (dotted line with square markers); (3) laser prior and laser group's empirical policy (dotted line with diamond markers. Comparing (1) and (2), we see that holding the supply-side behavior as fixed, laser branding alone has a significant impact on beliefs evolution. This difference shape sellers' incentives to provide quality, which further drive markets to different outcomes over time. The gap (1) and (3) represents the total effect.

Figure 7 plots seller's net profits evolution. An extrapolation to 5 seasons suggests that there might be large gains under laser (Table 8): the five-season discounted sum of net profits is ≈ 13 kRMB higher than baseline (≈ 11 kRMB). However, this increase is still not large enough to justify the initial investment cost of the laser machine ($\approx$ 50-60k RMB). There are two reasons: first, each individual seller's market size is small; and second, it may be difficult for sellers to extract all the consumer surplus. The former indicates a collective action failure since one laser machine can serve multiple markets and the total gain in producer surplus can exceed the costs. The latter points to the importance of understanding the role of fragmented markets in the presence of information problems, which I turn to in the last section.

7 Welfare Effects of Information Frictions & Fragmented Markets

In a second-best world with multiple frictions, the welfare implication of each friction is theoretically ambiguous as the different frictions could counteract. In particular, while market power generally distorts quality provision from the first best (i.e., Spence distortion), it also internalizes the return of investing in quality by allowing sellers to capture a larger portion of the gain in consumer surplus. To examine the interaction, I conduct counterfactual exercises that remove one imperfection at a time in order to isolate the effect of the other. These exercises involve extrapolation beyond the sample period and assumptions on demand supply side conducts in various counterfactual scenarios. The primary goal is to highlight some general economic forces and tradeoffs faced by policy makers in regulating markets with both information frictions and imperfect competition.

Table 9 presents the results. The numbers reflect five-season discounted sum of surpluses for the same *average market* described in Section 6.2. Details of the calculation are in Appendix E.5.

The baseline benchmark. Column 1 calculates the welfare for the baseline scenario with no quality differentiation. Using column 1 as the benchmark, I next examine the counterfactual outcomes without information frictions. That is, for any quality that a seller chooses, she could immediately convey that information to consumers.

Symmetric information: one seller deviation. Column 2 considers a single seller deviation. I first solve for the seller's optimal quality and markup for the premium pile, holding the other sellers' strategies the same as in column 1. The optimal quality of the premium pile is 0.769, much higher than that of the normal pile. The seller's net profit is almost 7 times higher than the baseline case. This result demonstrates that without information frictions, baseline cannot be an equilibrium as there is a large single-seller profitable deviation.

Symmetric information: separating equilibrium. Column 3 computes the equilibrium outcome under symmetric information. For each γ_H and m_H chosen by the other sellers, I first solve for the optimal γ_H^* and m_H^* of the sample seller. A symmetric Nash equilibrium is found by searching for the fixed point. Here and in subsequent analyses, I focus on the best equilibrium for sellers in case of multiple equilibria. We see that competition puts a downward pressure on price and increases quality. Consumer surplus is significantly higher than that in column 2 because of the lowered price and enlarged choice set. A comparison of the total surplus in columns 1 and 3 shows that information frictions result in a welfare loss of about 66.4% in this setting.

Symmetric information: first best. Column 4 solves for the first-best outcome. The key takeaway is that in this setting the welfare loss caused by market power (column 3 versus column 4) is small relative to that caused by the information problem (column 1 versus column 3), suggesting that these markets are already quite competitive. Next, I turn to welfare under asymmetric information.

Asymmetric information: one seller. The bottom panels of columns 5 and 6 compute the discounted sum of surpluses, taking into account the learning process. Compared to column 1, the increases in total surplus are 49k and 65k RMB for the non-incentive and incentive cases respectively. Balk of the gains comes from gain in consumer surplus as a result of both enlarged choice set and allocative efficiency (i.e., allowing high-valuation consumers self-sort into buying higher quality, albeit more expensive, product). In fact, the total gains are on par with the cost of a laser machine. While an individual seller would not undertake such an investment, a third-party could invest in the technology and subsidize/rent it to the sellers. The result also implies a profitable entry opportunity of a large upstream firm to invest in the technology and build a reputation for quality over time (although this would depend on whether there will be profitable entries of large counterfeits in the longer term).

Asymmetric information: competition. To examine the effect of market competition in the presence of information asymmetry, I compute the symmetric Nash equilibrium outcome when all sellers in a market are given access to the new technology and simultaneously choose once-for-all quality and markup. The result is shown in column 7.[32] We see that competition induces sellers to provide higher quality (compare to the monopoly case in column 5); however, quality is still quite low compared to the first-best (column 4) as competition on price may in fact discourage quality.

To further highlight this tradeoff, imagine a counterfactual policy in which government could regulate the price for the premium product and still let sellers compete on the quality dimension.

[32]There is another low-quality equilibrium with $\gamma_H^* = 0.4$ and $m^* = 0.12$. See Appendix Table 15.

What would happen? We can think of this as analogous to the first-best but under asymmetric information when it is hard to directly enforce quality. The result is shown in the last column of Table 9. In line with the discussion above, the social planner would want to set a higher markup to ease competition, which leads to higher quality provision compared to the free market outcome in column 7. That being said, the additional welfare gain is small because higher markup also directly discourages sales, and thus beliefs take an even longer time to take off.

One important caveat with this counterfactual exercise is that consumers' learning dynamics are held the same under the case when it is provided by all sellers in the market and just by a single seller. In reality, the prior beliefs and learning dynamics may be quite different under these two scenarios. Next, I conclude with some potential avenues for exploring this issue in future research.

8 Conclusion

This study empirically examines the dynamic interactions between sellers and consumers in an experience good market setting, the local watermelon markets in China. I find that information frictions and fragmented markets lead to significant under-provision of quality in this setting. Though there is a high demand for quality, trust could take a long time to establish under the existing "contaminated" signaling technology. While there is a new technology that could enhance consumer learning, small individual sellers do not have the incentive to invest in this technology due to their small market size and market competition. The results suggest that third-party interventions that subsidize the initial demand and learning process could enhance welfare. Alternatively, the results suggest that there may be a profitable entry opportunity for a large upstream firm. Indeed, one of the largest watermelon seed companies in China, Hebei Shuangxing Seed Co., Ltd., is starting a new business venture to contract with farmers, invest in high quality production and establish a premium brand using the laser technology.

Though the exact learning processes and quality production technology are different for different goods, the study aims to highlight two broad takeaways:

First, good reputations may take a long time to establish, as is the case with the Wholefoods brand in the United States. In developing countries that lack such reputable entities, consumers' beliefs and learning dynamics matter for firms' incentive to provide quality. Rampant counterfeiting activities under the old sticker technology created distrusts among consumers. In this environment, it can be hard for a single firm to signal its quality and establish trust; in equilibrium, firms' incentive to provide quality is also low, which breeds more mistrusts tomorrow. The findings highlight an important externality due to collective reputation. These forces may be particularly relevant in the context of international trade, which I leave for future work.

Second, not much unlike these local retail markets, many industries in developing countries are characterized by fragmented markets with a large number of small- and medium-sized players. While market competition helps to expand sales, it can also discourage quality improvements since many innovations require large fixed costs and small firms would not be able to fully internalize the surpluses. The results highlight the importance to understand the effects of competition on quality provision especially in the presence of information problems. A possible extension is to consider a similar intervention but vary the number of firms treated in a market. I leave that as a potential avenue for exploration in the future.

(Author affiliation: Harvard University)

References

Acemoglu, D., P. Antràs, and E. Helpman (2007): "Contracts and technology adoption," The American economic review, 97(3), 916–943.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005): "Competition and Innovation: An Inverted-U Relationship," The Quarterly Journal of Economics, 120(2), 701–728.

Allen, T. (2014): "Information frictions in trade," Econometrica, 82(6), 2041–2083.

Atkin, D., A. K. Khandelwal, and A. Osman (forthcoming): "Exporting and firm performance: Evidence from a randomized trial," Quarterly Journal of Economics.

Bagwell, K. (2007): "The economic analysis of advertising," Handbook of industrial organization, 3, 1701–1844.

Banerjee, A. V. (2013): "Microcredit under the microscope: what have we learned in the past two decades, and what do we need to know?," Annu. Rev. Econ., 5(1), 487–519.

Banerjee, A. V., and E. Duflo (2000): "Reputation Effects And The Limits Of Contracting: A Study Of The Indian Software Industry," The Quarterly Journal of Economics, 115(3), 989–1017.

Bar-Isaac, H., and S. Tadelis (2008): Seller reputation. Now Publishers Inc.

Bergquist, L. F. (2016): "Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya," Discussion paper.

Björkman-Nyqvist, M., J. Svensson, and D. Yanagizawa-Drott (2013): "The market for (fake) antimalarial medicine: Evidence from uganda," .

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2013): "Does Management Matter? Evidence from India*.," Quarterly Journal of Economics, 128(1).

Bryan, G., S. Chowdhury, and A. M. Mobarak (2014): "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh," Econometrica: Journal of the Econometric Society, 82(5), 1671–1748.

Cabral, L., and A. Hortacsu (2010): "The dynamics of seller reputation: Evidence from ebay," The Journal of Industrial Economics, 58(1), 54–78.

Casaburi, L., and T. Reed (2014): "Interlinked Transactions and Pass-Through: Experimental Evidence from Sierra Leone," .

Chen, N., and L. Juvenal (2016): "Quality, trade, and exchange rate pass-through," Journal of International Economics, 100, 61–80.

Ching, A. T. (2010): "A dynamic oligopoly structural model for the prescription drug market after patent expiration," International Economic Review, 51(4), 1175–1207.

Ching, A. T., T. Erdem, and M. P. Keane (2013): "Invited paper-learning models: An assessment of progress, challenges, and new developments," Marketing Science, 32(6), 913–938.

Crozet, M., K. Head, and T. Mayer (2012): "Quality sorting and trade: Firm-level evidence for French wine," The Review of Economic Studies, 79(2), 609–644.

David Atkin, A. K., and A. Osman (forthcoming): "Exporting and Firm Performance: Evidence from a Randomized Experiment," Quarterly Journal of Economics.

De Loecker, J., and P. K. Goldberg (2014): "Firm performance in a global market," Annu. Rev. Econ., 6(1), 201–227.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016): "Prices, markups, and trade reform," Econometrica, 84(2), 445–510.

De Mel, S., D. McKenzie, and C. Woodruff (2008): "Returns to capital in microenterprises: evidence from a field experiment," The Quarterly Journal of Economics, 123(4), 1329–1372.

Dickstein, M. J. (2014): "Efficient provision of experience goods: Evidence from antidepressant choice," .

Duflo, E., M. Kremer, and J. Robinson (2011): "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya," American Economic Review, 101, 2350–2390.

Fafchamps, M. (2002): "Spontaneous market emergence," Topics in Theoretical Economics, 2(1).

Fang, H. (2001): "Social culture and economic performance," American Economic Review, pp. 924–937.

Feenstra, R. C., and J. Romalis (2014): "International Prices and Endogenous Quality*.," Quarterly Journal of Economics, 129(2).

Fershtman, C., and A. Pakes (2012): "Dynamic games with asymmetric information: A framework for empirical work," The Quarterly Journal of Economics, 127(4), 1611–1661.

Foster, A. D., and M. R. Rosenzweig (2010): "Microeconomics of technology adoption," Annual Review of Economics, 2(1), 395–424.

Hallak, J. C. (2006): "Product quality and the direction of trade," Journal of international Economics, 68(1), 238–265.

Hallak, J. C., and P. K. Schott (2011): "Estimating cross-country differences in product quality," The Quarterly Journal of Economics, 126(1), 417–474.

Hardy, M., and J. McCasland (2016): "It Takes Two: Experimental Evidence on the Determinants of Technology Diffusion," Discussion paper.

Harrison, A., and A. Rodríguez-Clare (2009): "Trade, foreign investment, and industrial policy for developing countries," Handbook of Development Economics, 5, 4039–4214.

Hausmann, R., and D. Rodrik (2003): "Economic development as self-discovery," Journal of development Economics, 72(2), 603–633.

Hendel, I., and A. Nevo (2011): "Intertemporal price discrimination in storable goods markets," Discussion paper, National Bureau of Economic Research.

Henze, B., F. Schuett, and J. P. Sluijs (2015): "Transparency in markets for experience goods: experimental evidence," Economic Inquiry, 53(1), 640–659.

HUFFMAN, W. E., M. ROUSU, J. F. SHOGREN, AND A. TEGENE (2007): "The effects of prior beliefs and learning on consumers' acceptance of genetically modified foods," Journal of Economic Behavior & Organization, 63(1), 193–206.

HUMMELS, D., AND P. J. KLENOW (2005): "The variety and quality of a nation's exports," The American Economic Review, 95(3), 704–723.

JAVORCIK, S. B. (2004): "Does foreign direct investment increase the productivity of domestic firms? In search of spillovers through backward linkages," The American Economic Review, 94(3), 605–627.

JIN, G. Z., AND A. KATO (2006): "Price, quality, and reputation: Evidence from an online field experiment," The RAND Journal of Economics, 37(4), 983–1005.

JIN, G. Z., AND P. LESLIE (2009): "Reputational incentives for restaurant hygiene," American Economic Journal: Microeconomics, 1(1), 237–267.

KHANDELWAL, A. (2010): "The long and short (of) quality ladders," The Review of Economic Studies, 77(4), 1450–1476.

KLEIN, T. J., C. LAMBERTZ, AND K. O. STAHL (2016): "Market transparency, adverse selection, and moral hazard," Journal of Political Economy, 124(6), 1677–1713.

KRANTON, R. E. (2003): "Competition and the incentive to produce high quality," Economica, 70(279), 385–404.

KUGLER, M., AND E. VERHOOGEN (2012): "Prices, plant size, and product quality," The Review of Economic Studies, 79(1), 307–339.

LEGGETT, C. G. (2002): "Environmental valuation with imperfect information the case of the random utility model," Environmental and Resource Economics, 23(3), 343–355.

LIST, J. A. (2006): "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions," Journal of Political Economy, 114(1), 1–37.

MACCHIAVELLO, R. (2010): "Development uncorked: Reputation acquisition in the new market for Chilean wines in the UK," .

MACCHIAVELLO, R., AND A. MORJARIA (2015): "The value of relationships: evidence from a supply shock to Kenyan rose exports," The American Economic Review, 105(9), 2911–2945.

MACCHIAVELLO, R., A. MORJARIA, ET AL. (2016): "Competition and relational contracts: Evidence from rwanda's coffee mills," .

MANOVA, K., AND Z. ZHANG (2012): "Export prices across firms and destinations," The Quarterly Journal of Economics, p. qjr051.

MUNSHI, K. (2004): "Social learning in a heterogeneous population: technology diffusion in the Indian Green Revolution," Journal of development Economics, 73(1), 185–213.

PARK, A., D. YANG, X. SHI, AND Y. JIANG (2010): "Exporting and firm performance: Chinese exporters and the Asian financial crisis," The Review of Economics and Statistics, 92(4), 822–842.

Prasad, K., and C. S. Dev (2000): "Managing hotel brand equity: A customer-centric framework for assessing performance," The Cornell Hotel and Restaurant Administration Quarterly, 41(3), 22–4.

Qian, Y. (2008): "Impacts of entry by counterfeiters," Quarterly Journal of Economics, 123(4), 1577–1609.

Schott, P. K. (2004): "Across-product versus within-product specialization in international trade," The Quarterly Journal of Economics, 119(2), 647–678.

Shapiro, C. (1982): "Consumer information, product quality, and seller reputation," The Bell Journal of Economics, 13(1), 20–35.

Startz, M. (2016): "The value of face-to-face: Search and contracting problems in Nigerian trade," .

Tirole, J. (1996): "A theory of collective reputations (with applications to the persistence of corruption and to firm quality)," The Review of Economic Studies, 63(1), 1–22.

Train, K. E. (2009): Discrete choice methods with simulation. Cambridge University Press.

Verhoogen, E. A. (2008): "Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector," The Quarterly Journal of Economics, 123(2), 489–530.

Villas-Boas, J. M. (2004): "Consumer learning, brand loyalty, and competition," Marketing Science, 23(1), 134–145.

Figure 1: Asymmetric Information Between Sellers and Consumers
in the Watermelon Market

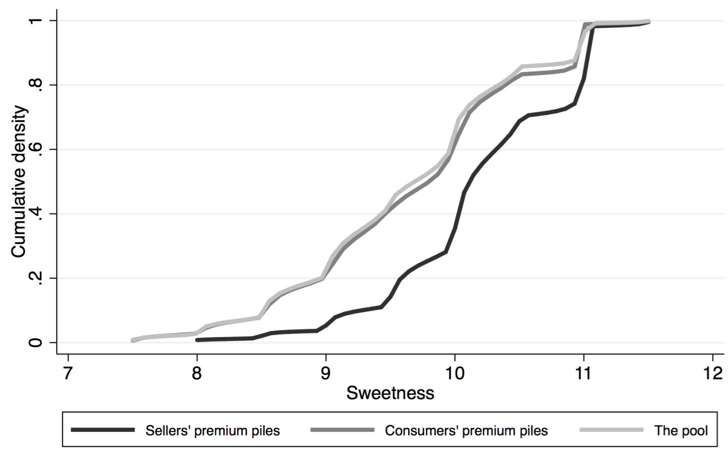

Note: This figure shows the empirical cumulative quality distribution for: (1) all 300 randomly picked watermelons used in the sorting tests; (2) the premium piles sorted by sellers; (3) the premium pile sorted by consumers. Quality is measured using a sweetness meter. For each watermelon, two measures are taken, one at the center and the other at the side, and the measures are then averaged.

Figure 2: Heterogeneity in Consumers' Willingness to Pay for Quality

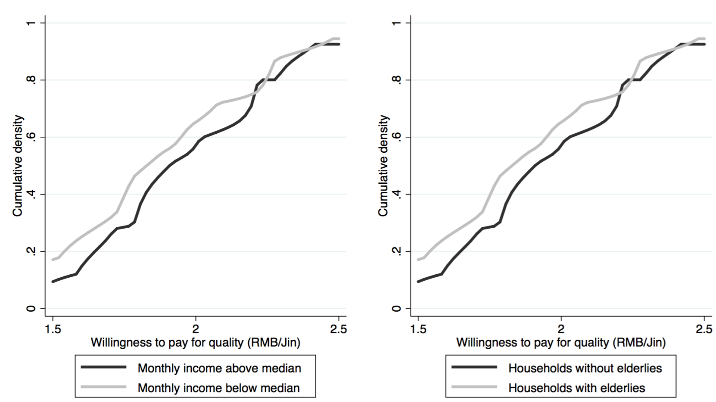

Note: This figure shows the heterogeneity of households' self-reported willingness to pay for quality elicited in the baseline survey. Households were asked to consider a hypothetical situation where they see two piles of watermelons sold in the local market, one pile of ordinary quality at 1.5 RMB/Jin and the other pile of premium quality but at a higher price. Surveyors announced the price for the premium pile from high to low and recorded the highest number that led to the choice of the premium pile. The sequence of prices (in RMB/Jin) were announced in the following order: 2.5, 2.2, 2, 1.9, 1.8, 1.7, 1.6, 1.5. The left figure plots the empirical cumulative distributions for households with monthly income above and below the median. The right figure shows the distributions for households with and without elderly members.

Figure 3: Pictures of the Branding Treatments

Panel A. The Label-less Group

Panel B. The Sticker Group

Panel C. The Laser Group

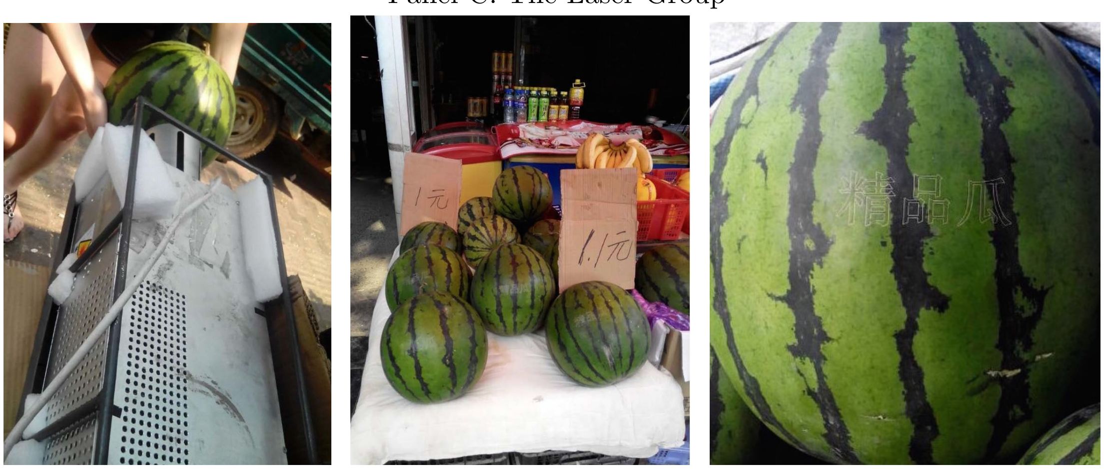

Note: This figure depicts the actual implementation of the branding treatments. Sellers sold two piles of watermelons, a premium pile and a normal pile, and put up two price boards. Surveyors visited the markets every morning and branded the watermelons in the premium pile. Nothing was done for the label-less group (Panel A). For the sticker group, a sticker label reading "premium watermelons" was pasted on the watermelons (Panel B). For the laser group, the same words were printed on the watermelons using a laser-engraving machine (Panel C).

Figure 4: Timeline of the Study

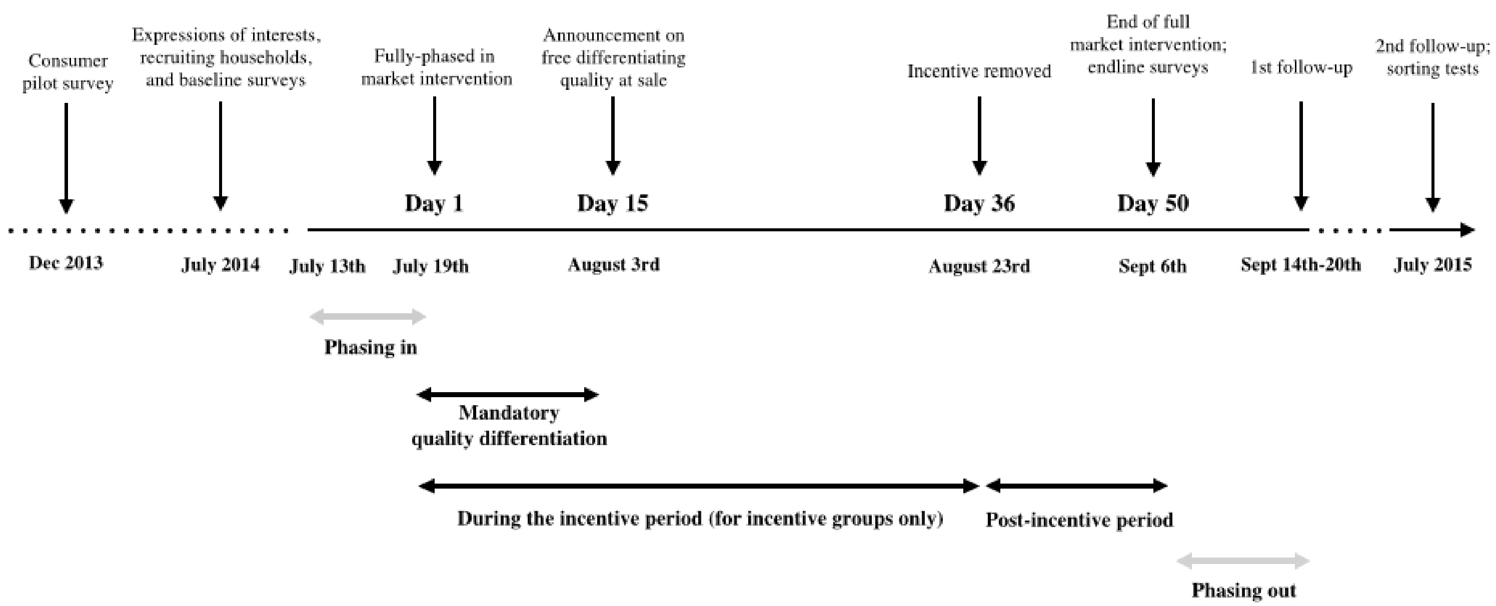

Note: This figure gives an overview of the time of the study.

1. A consumer pilot survey was conducted in December 2013 to elicit consumers' perceptions of different branding technologies.
2. Expressions of interests and baseline surveys were conducted in July 2014.
3. The market intervention was rolled in from July 13 to 19, 2014. The intervention was kicked off with the label-less group on July 13 and 14, followed by the sticker group on July 16 and 17, and finally the laser group on July 18 and 19th. July 19 is defined to be day 1 of the full-market intervention.
4. Quality differentiation was mandatory for the first 2 weeks, from July 19 to August 3. An announcement was made to all sellers on August 3 that they were free to differentiate or not afterwards.
5. On August 23, 35 days (6 weeks) into the intervention, the incentive (for the incentive groups) was lifted.
6. September 6 is the last day of the full-market intervention. An endline survey was conducted at surveyors' final visits to sellers' stores. Most of data analysis focuses on the period from July 19 (day 1) to September 6 (day 50).
7. The market intervention was gradually phased out from September 6 to September 12, 2014.
8. A short follow-up survey was conducted from September 14 to 20, 2014, and another one was conducted a year later, in July 2015.

Figure 5: Quality Differentiation at Sale

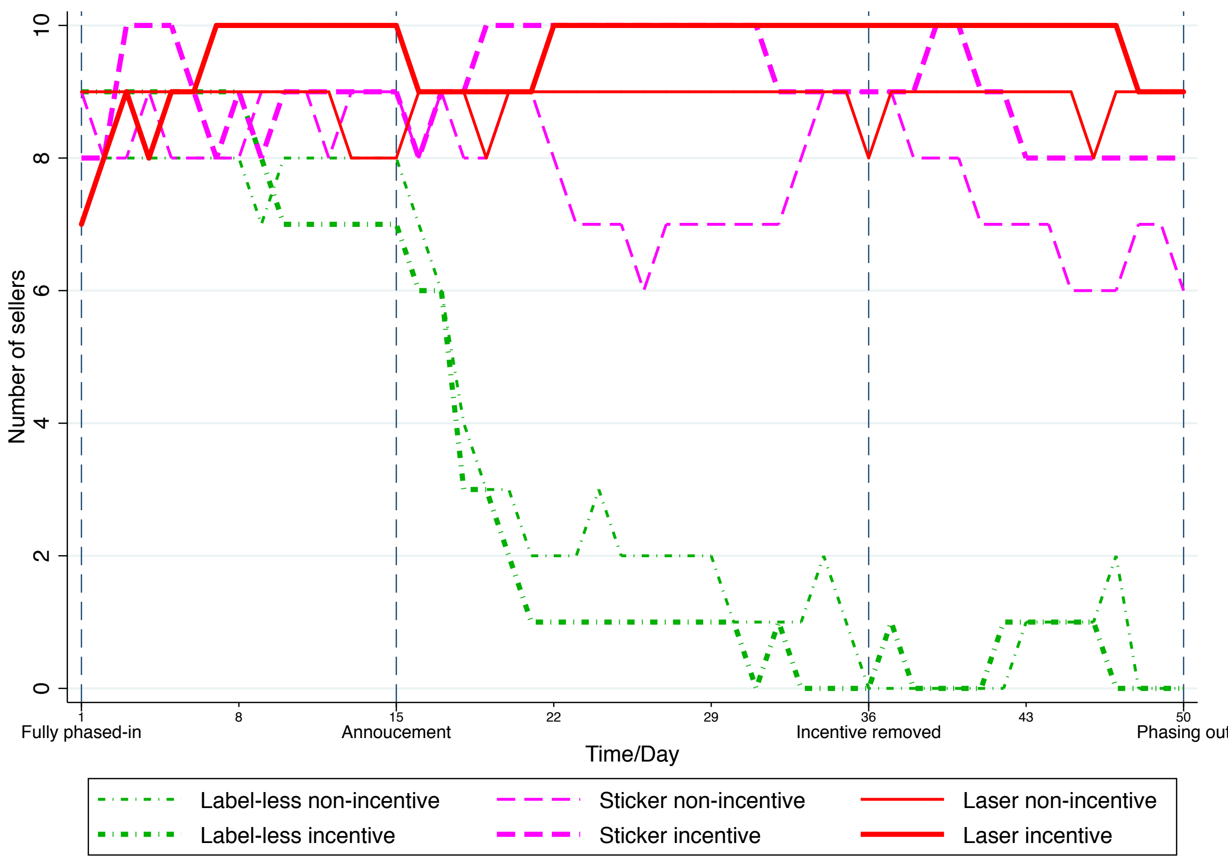

Note: This figure plots the number of sellers who differentiated quality at sale in each treatment group over time. The time axis runs from July 19 (day 1) to September 6 (day 50), 2014, corresponding to the period of the fully phased-in market intervention. The panel is not fully balanced because not all sellers operated their businesses on all days.

Figure 6: Beliefs Evolution

Panel A. Average Beliefs Evolution by Treatment Group

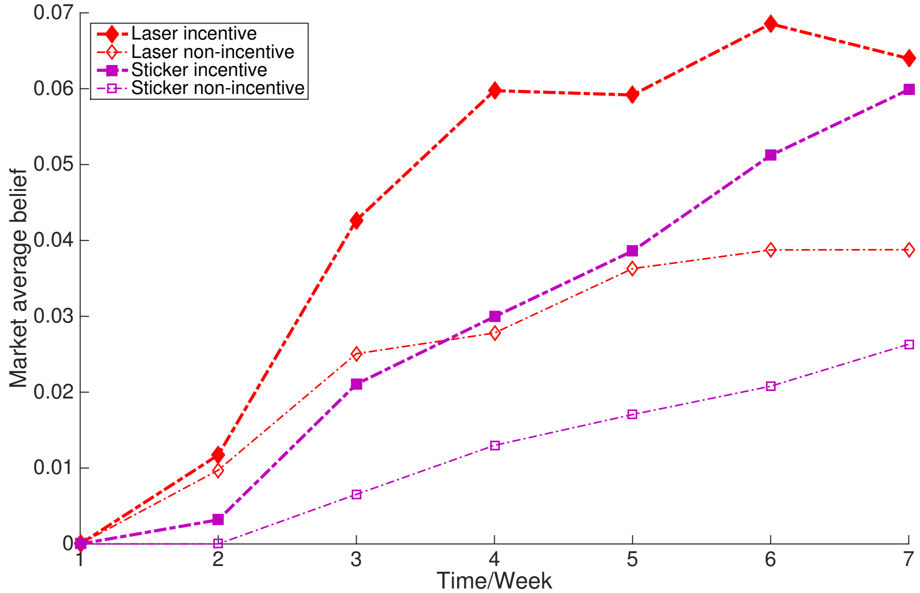

Panel B. Counterfactual Beliefs Evolution

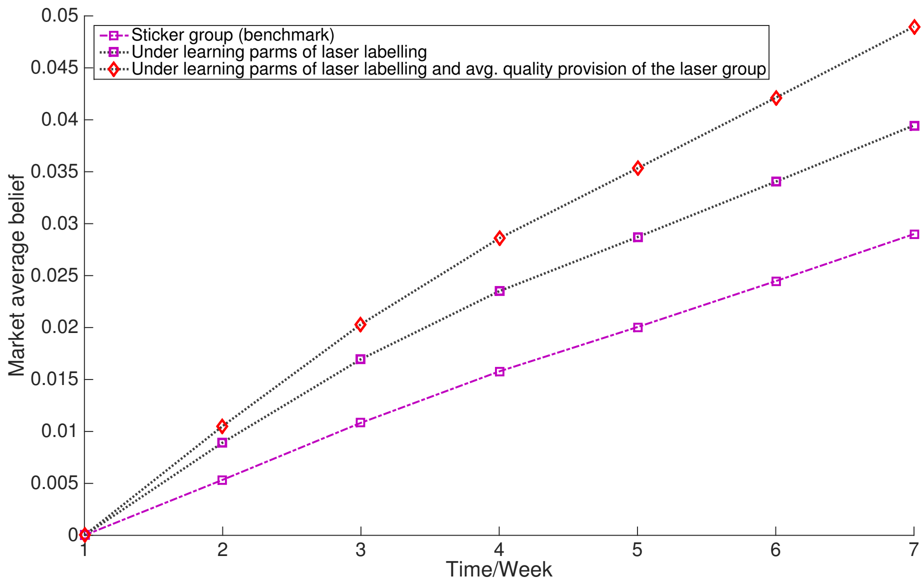

Note: This figure plots the average beliefs evolution about the quality of the premium pile. Panel A plots the market average beliefs calculated using the estimated prior beliefs (see Table 7) and the actual experience realizations for households in each treatment group. In particular, I take the demand estimates in column 1 of Table 7 and feed them through the actual purchasing and experience realizations to compute the posterior for each household in each period. I then average that across all households in a given treatment group to get the group average beliefs. Panel B simulates the counterfactual beliefs evolution for the sample of households in the sticker group under three different scenarios: (1) under sticker group's average empirical quality (measured in terms of the empirical satisfaction rate for sticker-labeled watermelons); (2) the same quality as in (1) but replacing the prior beliefs with that under laser; (3) replacing both the prior beliefs and the average empirical quality with that for the laser group. The simulation procedure is discussed in Appendix E.4.

Figure 7: Net Profits Evolution

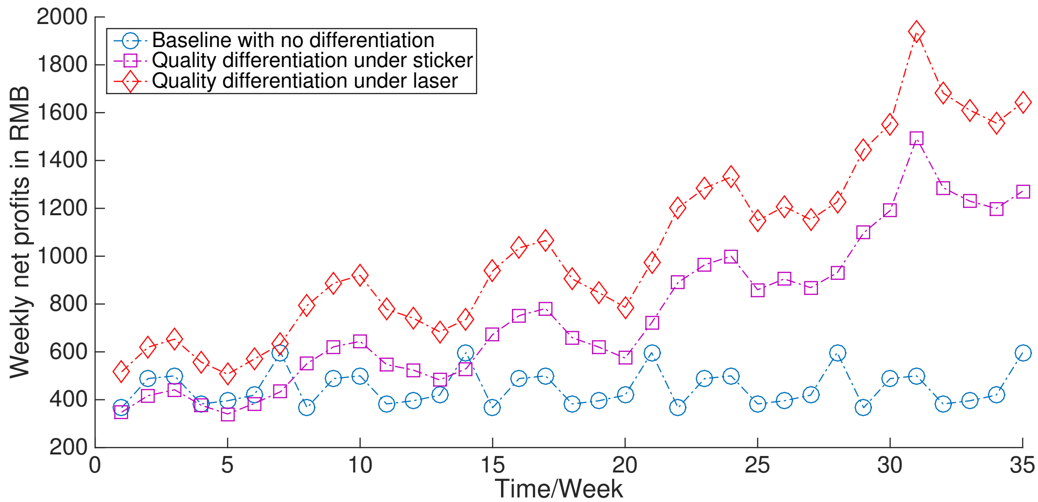

Note: This figure plots the simulated net profits evolution (sales profits minus effort costs) for a seller facing the hypothetical *average market* under the following three scenarios: (1) baseline with no differentiation; (2) quality differentiation under laser branding and the average empirical policies (markup and quality) of the laser non-incentive group; (3) quality differentiation under sticker branding but following the same policies as (2). Details for constructing the hypothetical market is explained in Section 6.2. The simulation procedure is discussed in Appendix E.4.

Table 1: Baseline Summary Statistics

	Observations	Median	Mean	Std. Dev
Panel A. Community and market characteristics				
Size measured in the number of housing units	60	1350	1915	1930
Housing price (in thousand RMB/meter^2)	60	8.95	8.291	1.594
Fraction of elderly	60	0.25	0.28	0.123
Distance to the nearest supermarket (in kilometer)	60	1.5	1.567	1.046
Years since establishment	60	15.5	17.633	11.242
Number of competitors in the local market	60	3	3.533	2.273
Panel B. Seller characteristics				
Gender (female=1 and male=0)	60	0	0.483	0.504
Age	60	42	41.067	9.189
Years of schooling	59	9	10.254	2.509
Selling fruits as primary income source (dummy)	60	1	0.95	0.22
Selling fruits only in the summer (dummy)	60	0	0.033	0.181
Planning to stop selling fruits (dummy)	60	0	0.017	0.129
Number of years selling fruits	60	8	9.017	6.035
Number of years selling fruits at this location	60	6.5	7.867	6.239
Planning to relocate (dummy)	60	0	0	0
Purchasing from fixed wholesaler(s) (dummy)	60	0	0.217	0.415
Panel C. Household characteristics				
Household size	658	3.5	3.76	1.366
Fraction of elderly	657	0	0.169	0.272
Fraction of female	657	0.5	0.498	0.154
Household monthly income (in thousand RMB)	647	4	5.250	3.235
Fruit as % of total food consumption	602	30	32.01	17.906
Watermelon as % of total fruit consumption	626	30	35.627	25.292
Number of watermelons consumed per week	654	1	1.308	.695
Local markets as main purchase source (dummy)	675	1	0.756	0.43
Supermarkets as main purchase source (dummy)	675	0	0.227	0.419
Willingness to pay for quality (RMB/Jin)	633	2	1.926	0.312

Note: This table provides the summary statistics for sample characteristics of communities, sellers and households measured in the baseline surveys. In total, 60 sellers in 60 communities (markets) and 675 households were recruited for this study. Variation in the number of observations are due to missing responses in the baseline surveys. The measure for household's willingness to pay for quality is explained under the footnote of Figure 2.

Table 2: Purchasing Dynamics under Different Branding Technologies

	Households in the Laser Markets		Households in the Sticker Markets	
	(1)	(2)	(3)	(4)
Panel A. Purchasing decision of the premium pile				
Lagged avg. satisfaction rating	0.280**		0.049	
	(0.090)		(0.044)	
Lagged % of very good experiences		0.454**		0.110
		(0.129)		(0.075)
Observations	165	167	183	183
Panel B. Purchasing decision of the normal pile				
Lagged avg. satisfaction rating	0.035		-0.014	
	(0.029)		(0.039)	
Lagged % of very good experiences		0.010		-0.016
		(0.032)		(0.086)
Observations	520	576	497	530
Household Baseline Controls	✓	✓	✓	✓
Week Fixed Effects	✓	✓	✓	✓

Note: This table examines the purchasing dynamics under different branding technologies. Each observation is at the household-week level. The dependent variable for Panel A is whether the household has purchased any watermelon from the premium pile for a given week. The dependent variable for Panel B is the corresponding purchasing dummy for the normal pile. The purchasing dummies are regressed on two measures of lagged experiences: (1) the average lagged satisfaction rating (ranging from 1 to 5) of all premium watermelons purchased prior to the period; (2) the percentage of past consumption experiences that attained the highest satisfaction rating of 5. Note that if a household has never purchased any premium watermelons, these lagged experience measures are not defined. Therefore, the coefficients are only estimated from household-week observations for which a positive number of premium watermelons have been consumed by the household prior to the given week. All regressions control for week fixed effects and the following set of household baseline characteristics: household size, percentage of elderly, monthly income, average number of watermelons consumed per week reported in the baseline survey, and the baseline self-reported willingness to pay for quality (measured in RMB/Jin). Standard errors are clustered at the seller level. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 3: Quality Provision by Treatment Group

Dep var: Quality of the premium pile (measured in sweetness)

	A. By branding treatments (sticker and laser)				B. By incentive treatment (during incentive)			
	Non-incentive		Incentive		Laser		Sticker	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Laser	0.711***	0.619**	0.282*	0.309**				
	(0.222)	(0.266)	(0.136)	(0.128)				
Incentive					0.496*	0.563**	1.033***	1.006***
					(0.246)	(0.266)	(0.176)	(0.176)
Observations	238	238	230	230	197	197	194	194
Baseline Controls		✓		✓		✓		✓
Time Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Omitted group mean	9.738		10.654		10.451		9.738	
Std. dev	(1.104)		(0.886)		(1.04)		(1.104)	

Note: This table examines quality provision by treatment group. Quality is measured in sweetness. Each observation is at the seller-check level. The key explanatory variables are the group dummies. The mean and standard deviation for the omitted group are shown in the bottom two rows. Panel A examines the heterogeneity across different branding groups. Columns 1 and 2 restrict the sample to the non-incentive groups only. Columns 3 and 4 restrict to the incentive groups. Panel B examines the heterogeneity for sellers with and without the incentive. Since sellers in the label-less group reverted back to non-differentiation after the mandatory period, the sample for this analysis includes only sellers in the sticker and laser groups. The time period is from week 1 to week 6, before the incentive was lifted. Columns 5 and 6 look within the laser group. Columns 7 and 8 look within the sticker group. All regressions control for check fixed effects. The even columns control for the following set of seller and community baseline characteristics: number of competitors in the local market, average housing price, and distance to the nearest supermarket. Standard errors are clustered at the seller level. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 4: Quality Differentiation Behavior

Sample: sticker and laser non-incentive groups

	Dep var: Quality measured in sweetness			
	A. Level		B. Diff. from the avg. pool	
	Laser	Sticker	Laser	Sticker
	(1)	(2)	(3)	(4)
Premium pile	0.735***	0.378**	0.786***	0.453**
	(0.157)	(0.163)	(0.129)	(0.172)
Observations	212	184	142	116
Seller Fixed Effects	✓	✓	✓	✓
Time Fixed Effects	✓	✓	✓	✓
Normal pile mean	9.787	9.366	0.102	-0.285
Std. dev.	(0.99)	(0.923)	(0.774)	(0.965)

Note: This table examines the quality differentiation behavior of sellers in the sticker and laser non-incentive groups. Quality is measured in sweetness. Each observation is at the seller-pile-check level. The key explanatory variable is a dummy for the premium pile. The mean and standard deviation for the normal pile are shown in the bottom two rows. The dependent variable for Panel A is in the level of the measured sweetness and that for Panel B is the difference from the market average quality. The average is computed as the average sweetness of randomly picked watermelons from the undifferentiated piles of the label-less group at each check (from week 3 and onwards). All regressions control for seller and time fixed effects. Standard errors are clustered at the seller level. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 5: Effects of the Branding Treatments on Price, Quantity and Profits

Sample: non-incentive groups

	Ln(Sales Profits)		Premium Markup		Premium Quantity		Normal Markup		Normal Quantity		Total Quantity	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Sticker	0.031	-0.038	0.039**	0.045***	49.852*	49.454*	0.001	-0.001	-40.374	-55.550**	9.478	-6.096
	(0.199)	(0.196)	(0.015)	(0.015)	(28.758)	(28.506)	(0.010)	(0.009)	(24.860)	(23.831)	(39.378)	(41.676)
Laser	0.297*	0.396**	0.069***	0.065***	62.041***	70.450***	-0.006	-0.001	-12.445	-4.449	49.596	66.002**
	(0.154)	(0.156)	(0.020)	(0.019)	(22.073)	(23.296)	(0.010)	(0.010)	(26.705)	(18.699)	(36.728)	(31.906)
Observations	1452	1452	1456	1456	1462	1462	1456	1456	1462	1462	1462	1462
Baseline Controls		✓		✓		✓		✓		✓		✓
Time Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Label-less Mean	4.284		0.055		56.313		0.011		180.475		236.788	
Std. dev.	(0.687)		(0.091)		(136.508)		(0.084)		(124.07)		(156.597)	

Note: This table examines sales profits, price and quantity for sellers in the non-incentive groups. Each observation is at the seller-day level. Sticker and laser are group dummies, and the omitted group is the label-less group, the mean and standard deviation for which are shown in the last two rows. Markup is defined to be the difference between the unit price (RMB/Jin) charged by the seller and the market average. Quantity is measured in Jin and profits are measured in RMB. If a seller stops to differentiate quality at sale, the unit price of the premium pile is defined to be the same as that of the normal pile, and the sales quantity of the premium pile is coded as 0. The even columns control for the following set of seller and community baseline characteristics: number of competitors in the local market, average housing price, and distance to the nearest supermarket. All regressions control for day fixed effects. Standard errors are clustered at the seller level. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 6: Effects of Removing the Incentive on Quality Provision

Dep var: Quality of the premium pile (measured in sweetness)

	Laser		Sticker	
	(1)	(2)	(3)	(4)
Incentive	0.502**	0.550**	1.026***	1.034***
	(0.239)	(0.256)	(0.171)	(0.169)
Post	0.013	0.014	0.224	0.226
	(0.299)	(0.301)	(0.255)	(0.256)
Post X Incentive	-0.008	-0.008	-0.683*	-0.674*
	(0.401)	(0.405)	(0.376)	(0.380)
Observations	236	236	232	232
Seller (Market) Baseline Controls		✓		✓

Note: This table runs a difference-in-difference regression to examine the effect of removing the incentive. The dependent variable is the measured sweetness of watermelons in the premium pile. Incentive is a dummy for the incentive group. Post is a dummy for the period after the incentive was lifted (i.e. week 7 and 8). The key explanatory variable is the interaction term. Each observation is at the seller-check level. Columns 1 and 2 look within the laser groups; columns 3 and 4 look within the sticker groups. In addition, the even columns control for a set of baseline characteristics, including the number of competitors in the local market, the average housing price, and distance to the nearest supermarket. Standard errors are clustered at the seller level. ***, **, and * indicate significance at the 1, 5, and 10 percent level, respectively.

Table 7: Simulated Maximum Likelihood Estimation Results of Consumer Learning Models

Parameters	Baseline Model (1)		Direct Utility of Laser (2)		Correlated Learning (3)		Information Diffusion (4)	
$a_0(s)$	0.000	(-)	0.000	(-)	0.000	(-)	0.000	(-)
$b_0(s)$	2.578	(0.733)	2.383	(0.683)	2.639	(0.818)	2.453	(0.757)
$a_0(l)$	0.000	(-)	0.000	(-)	0.000	(-)	0.000	(-)
$b_0(l)$	0.938	(0.471)	1.037	(0.510)	0.995	(0.554)	0.850	(0.498)
q	0.307	(0.088)	0.313	(0.089)	0.283	(0.089)	0.309	(0.098)
θ_0	8.549	(1.197)	8.500	(1.185)	9.149	(1.577)	8.518	(1.533)
θ_1	0.346	(0.285)	0.309	(0.277)	0.373	(0.312)	0.330	(0.286)
α_0	0.169	(0.046)	0.170	(0.045)	0.166	(0.046)	0.168	(0.046)
α_1	-0.007	(0.006)	-0.007	(0.006)	-0.007	(0.006)	-0.007	(0.006)
β	0.061	(0.035)	0.062	(0.035)	0.057	(0.035)	0.057	(0.035)
$m(\eta)$	0.479	(0.195)	0.406	(0.236)	0.451	(0.108)	0.442	(0.216)
$\sigma(\eta)$	0.426	(0.182)	0.436	(0.196)	0.433	(0.188)	0.433	(0.191)
$m(\xi)$	-1.583	(0.046)	-1.585	(0.046)	-1.583	(0.046)	-1.584	(0.046)
$\sigma(\xi)$	0.784	(0.056)	0.786	(0.056)	0.784	(0.056)	0.784	(0.056)
$\Delta q(s)$	-0.081	(0.022)	-0.082	(0.023)	-0.064	(0.025)	-0.081	(0.029)
$\Delta q(l)$	-0.001	(0.012)	-0.003	(0.013)	-0.003	(0.011)	-0.003	(0.012)
$\nu(l)$	n.a.	-	0.399	(0.278)	n.a.	-	n.a.	-
$\phi_{\text{spillover}}$	n.a.	-	n.a.	-	1.218	(0.839)	n.a.	-
ϕ_{info}	n.a.	-	n.a.	-	n.a.	-	2.176	(3.597)
Market FE (abbreviated)	✓		✓		✓		✓	
Time FE (abbreviated)	✓		✓		✓		✓	
Log likelihood	-3709.749		-3708.752		-3708.578		-3708.383	
D (-2×Log(likelihood ratio))			1.993		2.341		2.732	

Note: This table shows the simulated maximum likelihood estimation results of the consumer learning models. a_0 and b_0 are constrained to be non-negative. Details for the estimation procedures are explained in Appendix E.1. Column 1 shows the estimates for the baseline model. Column 2 includes a product-specific constant ν for the premium option under laser label to account for any direct utility of laser. Column 3 incorporates correlated learning by allowing the posterior for the premium pile to enter linearly into the mean utility of the normal pile (i.e. good experiences from the premium pile may lead consumers to favor the sample seller in general). Column 4 includes a linear function of the market average beliefs (computed as the average beliefs of households in a given market at a given time) in the mean utility of the premium option to account for information diffusion. The log-likelihood ratio statistics for testing the baseline model against these alternative models are presented in the last row. Estimates for the market and time fixed effects are abbreviated from this table and are reported in Appendix Table 9. Standard errors shown in parentheses are calculated using the outer product of gradients (OPG) estimate for the asymptotic covariance matrix (see Appendix E.1 for details).

Table 8: Simulated Market Outcomes

	Laser non-incentive		Laser incentive		Counterfactual I Prior beliefs under sticker	Counterfactual II No differentiation
Structural parameters						
Market size : 4.5 × 194 households (to match initial sales quantity)						
$\delta = 0.98$, $c = 0.64$						
	(1)		(2)		(3)	(4)
Empirical average policies						
Average quality of the undifferentiated pile ($\underline{\gamma}$)	0.300		0.300		0.300	0.300
Average quality of the premium pile ($\overline{\gamma}_H$)	0.400		0.530		0.400	0.300
Average markup of the premium pile in RMB/Jin ($\overline{m}_H$)	0.142		0.178		0.142	0.000
Average weekly outcomes for the first season						
	Simulated	Actual	Simulated	Actual	Simulated	Simulated
Sales quantity of the premium pile (number)	53	50	58	62	41	-
Sales quantity of the normal pile (number)	81	76	80	74	48	85
Total sales quantity (number)	133	126	138	136	89	85
Total sales quantity of other sellers (number)	311	-	303	-	331	321
Sales profits (in RMB)	657	748	760	875	461	450
Net profits (sales profits minus effort costs) (in RMB)	579	-	550	-	392	450
Sales profits of other sellers (in RMB)	1,345	-	1,390	-	1,428	1,754
Simulated longer term outcomes						
Disc. Σ of net profits for two seasons (in RMB)	8,361		7,554		5,777	5,524
Disc. Σ of net profits for five seasons (in RMB)	24,408		23,165		13,281	11,367

Note: This table simulates market outcomes for the hypothetical *average market* using the estimated dynamic demand system and the estimated supply-side parameters. Details for constructing the hypothetical market are explained in Section 6.2. Column 1 simulates the market outcomes under the average empirical policies of the laser non-incentive group and column 2 does that for the laser incentive group. Column 3 performs a counterfactual exercise by replacing the learning parameters (including $a_0, b_0, \Delta q$) under laser with those under sticker (see Table 7). Column 4 simulates the outcomes for the baseline case with no quality differentiation. Details for the simulation procedures are explained in Appendix E.4.

Table 9: Welfare Effects of Information Frictions and Fragmented Markets

	Baseline	Symmetric information			Asymmetric information			
		One seller deviation	Oligopolistic competition	First-best	One seller w/o incentive	One seller w incentive	Oligopolistic competition	Price regulation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quality and markup								
Average quality of the undifferentiated pile ($\underline{\gamma}$)	0.300	0.300	0.300	0.300	0.300	0.300	0.300	0.300
Quality of the premium pile (γ_H)	-	0.769	0.787	0.825	0.400	0.530	0.440	0.530
Markup of the premium pile (m_H)	-	1.156	1.080	0.577	0.142	0.178	0.170	0.340
No adjustment (disc. Σ of 5 seasons)								
Sales profits	11,367	237,102	83,515	59,315	52,963	91,515	28,292	38,859
Effort costs	0	147,736	46,178	56,895	14,801	58,009	7,863	14,764
Net profits (PS_{own})	11,367	89,365	37,337	2,420	38,162	33,505	20,429	24,095
Sales profits of other sellers	44,330	23,568	335,177	241,773	31,793	21,199	102,983	149,404
Effort costs of other sellers	0	0	188,691	233,224	0	0	31,973	60,181
Net profits of other sellers (PS_{other})	44,330	23,568	146,486	8,550	31,793	21,199	71,010	89,222
Expected maximum utility in RMB (CS)	207,419	370,370	598,265	804,228	305,196	394,443	484,279	531,841
Total surplus ($= \text{PS}_{\text{own}} + \text{PS}_{\text{other}} + \text{CS}$)	263,116	483,303	782,088	815,198	375,151	449,147	575,718	645,158
Ratio relative to baseline	1.000	1.837	2.972	3.098	1.426	1.707	2.188	2.452
With adjustment (disc. Σ of 5 seasons)								
Net profits (PS_{own})	-	-	-	-	24,408	23,165	14,695	15,400
Net profits of other sellers (PS_{other})	-	-	-	-	39,357	39,134	68,011	71,448
Expected maximum utility in RMB (CS)	-	-	-	-	248,408	266,130	361,737	363,430
Total surplus ($= \text{PS}_{\text{own}} + \text{PS}_{\text{other}} + \text{CS}$)	-	-	-	-	312,173	328,429	444,443	450,278

Note: This table examines the welfare effects of information frictions and market competition. The top panel solves for the optimal policies under each counterfactual scenario. Quality is the probability of being good and markup is the difference between the prices of the premium and the normal pile, measured in RMB/Jin. The middle and bottom panel calculate the 5-season discounted sum of surpluses (in RMB) under the corresponding policies for the same hypothetical *average market* as that for Table 8 (see in Section 6.2 for details on constructing the hypothetical market). Details for calculating the consumer and producer surpluses are discussed in Section 7.