Retail Price Discrimination and Food Waste

Timothy J. Richards and Stephen F. Hamilton*

January 14, 2019

Abstract

When consumers are willing to pay a premium for high-quality fruit and vegetables, retailers may be motivated to sell fresh produce from only the upper portion of the quality-distribution, even when such behavior results in unsold products in the lower tail. In this paper, we consider the economic incentive for retailers to truncate the quality distribution of fresh produce at the farm level as part of a second-degree price-discrimination strategy to achieve higher prices in the consumer market. We estimate a structural model of retail price discrimination and conduct a series of counter-factual experiments using data from a major US food retailer and demonstrate that observed behavior is consistent with quality-based price discrimination in the consumer market. Our evidence indicates that quality standards on fresh produce can explain a substantial proportion of food waste by retailers in the US food system.

*Richards is Professor and Morrison Chair of Agribusiness, Morrison School of Agribusiness, W. P. Carey School of Business, Arizona State University; Hamilton is Professor and Department Head, Economics Department, Orfalea College of Business, California Polytechnic State University, San Luis Obispo, San Luis Obispo, CA. Contact author: Richards, Address: 7231 E Sonoran Arroyo Mall, Mesa, AZ, 85212. Ph. 480-727-1488, email: trichards@asu.edu. Funding from Agricultural and Food Research Initiative (National Institute for Food and Agriculture, USDA) grant no. 2016-09900 is gratefully acknowledged.

1 Introduction

Each year, US retailers reject, discard, or donate some 19.5 million metric tons of perfectly edible food products, resulting in the loss of considerable economic, social, and ecological value (Buzby and Hyman 2012). While there are many possible explanations for why loss occurs at the retail level, the rise of "sharing economy" firms that sell produce rejected by traditional retailers suggests fundamental failure in the retail market for fresh produce (Richards and Hamilton 2018). In this paper, we study how market failure in US grocery markets contributes to food waste at the retail level through the use of quality standards as a price-discrimination strategy for retailers to extract rents from consumers in the fresh produce category.

We frame our analysis around a model of second-degree price-discrimination in which retailers select quality standards to raise retail prices for fresh produce. Food waste emerges in the model through the resulting mismatch between the distribution of consumers' valuations for food quality and the distribution of fresh produce quality supplied by farms.

We test our theory using a unique data set on fresh apple sales by a major US retailer. These data allow us to exploit variation in quality standards across apple varieties to compare the empirical distribution of fresh apple quality purchased by consumers to the distribution of fresh apple quality supplied by farmers. Our findings demonstrate that retail quality standards effectively truncate the distribution of consumers' valuations above the minimum quality level supplied by farmers, resulting in food waste as a consequence of profit-maximizing behavior by retailers.[1]

[1] We define food waste as the quantity of food that meets retail quality standards but is subsequently not sold in the retail market due to retail pricing behavior. While food products not sold by retailers may be donated or diverted to animal feed or compost markets, as opposed to ending up in landfills (Bellemare et al. 2017), food used for secondary purposes nevertheless represents an economically wasteful misallocation of resources, which we refer to generically as "food waste".

While it is well-known that grocery retailers operate in relatively competitive environments, retail margins for fresh fruits and vegetables remain among the largest of all categories in the supermarket (Ailawadi and Harlam 2004). Retailers enjoy a certain measure of market power due to fixed shopping costs among consumers that lead to preferences for geographically proximate retailers (Bell, Ho, and Tang 1998; Bell and Lattin 1998), and this creates a potentially profitable opportunity for retailers to engage in quality-based price discrimination strategies to raise prices in fresh produce markets. Yet, although there is abundant empirical evidence that retailers price discriminate in consumer product markets,[2] there is a surprising lack of empirical research to date on the use of food quality standards as a price-discrimination tool in fresh fruit and vegetable markets.

In markets with quality-differentiated products, economists have long recognized the incentive for imperfectly competitive firms to distort product quality from the socially-efficient level (Mussa and Rosen 1978; Shaken and Sutton 1982, 1987). Because market demand for high-quality products is driven by consumers with the highest quality valuations, firms have an incentive to distort product quality downward to increase sales on high-end products, while raising prices to a level that potentially excludes consumers with lower valuations of product quality from the market.[3]

In our context, identifying quality-differentiation is challenging, because fresh foods are differentiated both vertically (by objective quality standards) and horizontally (by subjectively preferred attributes). For example, all consumers may prefer crisp apples, while some prefer green apples and others prefer red. We resolve this challenge by building on recent

[2] See, e.g., Cohen (2008, paper towels), Besanko, Dube, and Gupta (2003, ketchup), Verboven (2002, automobiles), Clerides (2002, books), Busse and Rysman (2005, yellow pages), Leslie (2004, broadway shows), Crawford and Shum (2007, cable television), McManus (2007, coffee shops), Courty and Pagliero (2012, concert tickets), Borenstein (1991, airlines) and Borenstein and Rose (1994) and Shephard (1991, gas stations).

[3] Crawford and Shum (2007) demonstrate this to be the case for cable consumers in the US, and McManus (2007) finds the same for coffee shops in Virginia.

advances in the emerging empirical literature that examines price discrimination strategies in settings with both vertical and horizontal product differentiation, following Leslie's (2004) analysis of seats for Broadway shows, and McManus's (2007) study of pricing by college-town coffee shops.[4]

Our empirical test of retail price discrimination in the fresh produce market is further complicated by the nature of the data. Unlike the case of manufactured consumer products, where product quality is relatively uniform across products sold by retailers, the quality of fresh produce grown on the farm is variable (Gallardo, et al. 2017), and we observe retail prices and sales only for the portion of the fresh apple quality distribution that retailers choose to buy from the farm market.

We turn this feature of the data to our advantage two ways. First, we exploit natural variation in quality among different fresh apple varieties to test for retail price discrimination in the quality standard set for each variety in the fresh apple market. Second, we use the estimated coefficients from our model to derive inferences on food waste by comparing the empirical distribution of consumers' valuations for fresh apple quality to the physical distribution of fresh apple quality supplied by the farm sector. This latter approach is possible, because under the null hypothesis that retailers price discriminate according to variation in fresh produce quality, the empirical distribution of consumers' valuations for quality reveals the underlying distribution of the products sold by retailers.[5] Under circumstances where the distribution of consumers' valuations for quality is truncated relative to what is supplied

[4] To our knowledge, Crawford and Shum (2007) is the only other study that accounts for both horizontal and vertical differentiation in an empirical model of price discrimination. However, they maintain the null hypothesis of quality degradation by price-discriminating monopolists, and test for departures using a non-parametric approach, while we test for evidence of price discrimination using a parametric model of demand.

[5] Our notion of price discrimination is second-degree price discrimination, or varying price by an attribute of the product, rather than third-degree price discrimination, or charging different prices to different market segments with no difference in cost (Stole 2007).

by the farm sector, this allows us to infer the quantity of fresh produce that is excluded by retailers under a given quality standard. Our novel identification strategy thus relies on the distribution of consumer willingness-to-pay for quality to estimate the extent of food loss at the retail level of the food system that arises through the use of retail quality standards.[6]

Our analysis contributes both to the empirical literature on quality-based price discrimination and to the growing literature on the economic causes of food waste. While others estimate empirical models of quality-based price discrimination (Verboven 2002; Leslie 2004; Crawford and Shum 2007; McManus 2007; Cohen 2008), we are the first to rely on the distribution of consumers' preference for quality to study the types of market failure that derive from this distribution. Thus, our analysis provides novel insights that reveal an important mechanism that contributes to food loss in the retail distribution channel.

We find evidence that the distribution of fresh apple quality offered by retailers differs significantly from the distribution of fresh apple quality sold in the farm product market. Specifically, we demonstrate that the distribution of harvested-apple quality is approximately log-normal, whereas the distribution of fresh apple quality sold in the retail market exhibits substantial right-skewedness, with a large probability mass evident in the left tail. Based on our empirical estimates, retailers effectively exclude approximately 10% of harvested fresh apple products from the market, an estimate remarkably similar to the overall level of retail food loss found by Buzby and Hyman (2012) using highly-aggregated data. Further, numerical simulations reveal that, as the willingness to pay for quality rises by a small amount (1%), the amount of food loss rises markedly (to over 25%). Price-discriminating behavior on the part of retailers therefore appears to be an essential factor determining retail

[6] We define "food loss" as food that is not sold for its intended use. That is, below-grade apples may be sold for juicing purposes, but often at a loss. Farmers purchase inputs with the intent that they will produce output that is sold for primary-market prices, and anything less is an economic loss.

food loss.

In the next section, we provide a brief background on the price-discrimination literature, and the implications for food loss. We also develop an analytical model of quality-based price discrimination to demonstrate how truncating the distribution of consumer-types is profitable for retailers, and at the same time, contributes to retail-level food waste. In the third section, we present our empirical model, which controls for both vertical and horizontal product-differentiation in examining the distribution of retail quality. We describe our data and identification conditions in a fourth section, while section five presents our results and provides a detailed discussion on the implications of our findings for food waste.

2 Price Discrimination in Retail Food Markets

2.1 Background on Price Discrimination

In most retail markets, sellers have a well-defined mechanism for price discrimination. Manufacturers are able to produce a range of quality over which retailers can match products to consumers according to their willingness-to-pay for higher quality items.[7] In this setting, the conditions for profitable price-discrimination are well understood (Mussa and Rosen 1978; Anderson and Dana 2009; Aguirre, Cowan, and Vickers 2010).

While there is a relatively large literature on the theoretical conditions for profit-enhancing

[7] Much of the literature on price discrimination concerns durable goods, based on the notion that sellers of items that degrade in value over time face a time-consistency problem. While sellers would like to be able to commit to low production levels in the present in order to create higher resale values in the future, when the future arrives they desire to produce more in order to take advantage of the higher prices. Consumers know that sellers have the incentive to over-produce in future periods, so wait to purchase (Coase 1972; Anderson and Ginsburg 1994). When the future arrives, sellers validate consumers' expectations, and discount older goods. The presence of secondary markets facilitate this mechanism as low-valuation consumers can simply buy used goods, which provide viable competition for new goods. Economic models of the role of secondary markets in durable good price-formation are legion (Corts 1998; Takeyama 2002; Estaban and Shum 2007), but rely fundamentally on the fact that durable goods can be expected to provide a stream of services into the future, where the content of the service-stream degrades very little with the passage of time.

price discrimination, there are relatively few empirical studies of price discrimination in consumer non-durable good markets. Moreover, most empirical studies do not have access to cost-data, confounding the analysis of whether price-cost margins reflect price discrimination or not. Verboven (2002) addresses this issue using aggregate data on gasoline and diesel car sales and prices from three European markets. Assuming low-mileage drivers will prefer gasoline engines and high-mileage drivers prefer more efficient diesel engines, he shows that auto manufacturers do indeed price discriminate among drivers through quality-based standards. Similarly, Leslie (2004) uses data on ticket sales for a single Broadway play in which seat location forms the basis for quality differentiation. Controlling for many other factors that may explain variation in price-cost margins in a structural model of demand, he concludes that the show producer used second-degree price discrimination to successfully increase profits, but that the net effect on consumer welfare was negligible.

Consumer products that are both highly differentiated and frequently purchased would appear to offer less latitude for price discrimination. Repeat-purchases allow consumers to become well-attuned to the horizontal differences among products, providing little latitude for vertical separation. Among the few empirical studies that examine consumer-good markets, Cohen (2008) considers second-degree price-discrimination among paper towel manufacturers by package-size differentiation, and uses a series of counter-factual simulations to show that consumer surplus is higher as a result. Examining a similar quality-distortion question to Crawford and Shum (2007), McManus (2007) also uses package-size as the key principle of vertical differentiation among coffee shops, and finds that there is virtually no distortion from the optimal attribute set among the largest sizes of coffee, but the distortion increases as sizes become smaller, as predicted by theoretical models of endogenous quality choice (Rochet and Stole 2002).

In a retail food context similar to our own, Nevo and Wolfram (2002) examine couponing activity in the breakfast cereal market, and use a price-discrimination argument to explain why manufacturers make extensive use of coupons. Although price-discrimination is the orthodox explanation for why coupons seem to make sense, their empirical results find the opposite, namely that coupon use is positively-related to shelf prices, resulting in non-coupon prices that fall at the same time coupons are used. Besanko, Dube, and Gupta (2003), on the other hand, use a structural model of the vertical relationship between ketchup manufacturers and retailers to show that third-degree price discrimination is not a prisoner's dilemma as commonly thought (Shaffer and Zhang 1995), but can be profit-increasing. However, these prior studies are constrained by the nature of their data. In the packaged-good industry, retailers are given products of fixed-quality, so the notion that retailers price discriminate over the distribution of consumers' willingness-to-pay for quality cannot be directly assessed. Fresh produce markets therefore provide a unique empirical opportunity, because consumer willingness-to-pay for quality is critical both in setting retail prices and in determining the quality standard the retailer imposes on suppliers in the farm product market.

Our research is related to the literature on minimum quality standards (MQS), both in agriculture (Bockstael 1984; Lapan and Moschini 2007; Saitone and Sexton 2010) and in other markets (Leland 1979). Bockstael (1984) shows that MQS generate welfare losses in an otherwise competitive farm product market simply because products that consumers would have willingly purchased are precluded from the market by the standard. Similarly, Saitone and Sexton (2010) find that MQS lead to welfare-losses due to the inefficient "upgrading" of products that would otherwise be sold at a lower quality. Lapan and Moschini (2007) apply a similar Mussa-Rosen (1978) model of vertical quality differentiation that we examine here; however, they consider a market with asymmetric information to show how MQS for genetically-modified foods can improve social welfare. Our context differs in the sense that we do not consider retail market failure driven by asymmetric information, but by retail grading standards that are typically set far above the fresh produce standards set by USDA. Fresh food products sold by retailers are typically upgraded, making the quality of products available to consumers in the retail market a function of buying practices by retailers that further screen fresh produce quality from within the broader sample of farm products that meet the MQS. Our point of departure from the literature on MQS is that we consider a retailer's objective in setting a quality standard on fresh produce to attract consumers with high valuations for fresh produce quality, rather than to support the underlying farm markets for fresh produce.

We examine fresh apple sales and quality standards by a major US retailer. Retailers selling fresh apples cannot perfectly control the quality of the apples they sell, and instead impose a minimum quality selection criteria on their purchases from fresh apple wholesalers and distributors. While visual inspection and random testing are helpful in maintaining fresh apple quality, apple quality is nonetheless subject to the vagaries of a biological production process, particularly when retailers purchase from different growers, in different regions, that rely on different cultural practices. Climatic and cultural differences at the farm level result in substantial variation in fresh apple quality above the retailer's purchase standard. Horizontal differentiation within the category of fresh apples is also important, as retailers typically stock different varieties of apples in side-by-side, pyramid displays. Because retailers are able to market each apple variety at a different price point according to the distribution of consumers' valuations for that variety, and this profit motive is independent of the (uncensored) quality distribution of the individual varieties, retail prices for fresh apples consequently exhibit a distribution of prices that reflects the distribution of consumers' willingness-to-pay for different varieties within the fresh apple category. This feature of our retail data allows us to exploit natural variation in product quality in the fresh apple retail market to test whether price discriminating behavior by a retailer facing consumers with heterogeneous preferences for quality can generate retail food loss.

2.2 Economic Model of Food Waste

In this section, we derive a model of retail price discrimination to demonstrate how truncating the distribution of quality offered to consumers is profit-maximizing and contributes to food waste in the form of excess supply of retail goods. For clarity we consider the case of a monopoly retailer who seeks to match consumer preferences for product quality with a distribution of farm products that varies in terms of product quality.

Suppose farm supply is distributed uniformly in terms of quality on a unit interval with density of one.[8] Accordingly, for any product quality standard set by the retailer, $q \in (0, 1)$, the average quality of products sold by the retailer is

$$\widehat{q} = \frac{1+q}{2}. \tag{1}$$

Following Moorthy (1988), the cost of implementing a retail quality standard is $c(q) = \alpha q^2/2$. This cost function reflects the fact that retailers faces fixed costs of implementing a higher standard that are rising in the quality level the retailer seeks to maintain, for instance when providing higher quality fresh apples requires more costly displays.

Consumers vary in their preference for quality, with types θ distributed uniformly over $[0, 1]$ with a density of one. Consumers cannot perceive the quality of a given product prior to consumption and either consume the low-quality, non-graded product as an outside option, or consume a randomly selected product from the retailer that meets the retailer's

[8] The qualitative results of our analysis go though for general distributions of food quality.

quality standard. A type θ consumer derives utility $v - p_L$ from the purchase of a non-quality differentiated good at price p_L and utility $v + \theta\widehat{q} - p$ from the purchase of a randomly-selected product from the retailer that meets the minimum quality standard q at price p.

Consumer type $\widehat{\theta} = 2(p - p_L)/(1 + q)$ is indifferent between the two types of products. As a result, consumer demand for the retailer's quality-graded product is $Q_d = 1 - \widehat{\theta} = 1 - 2(p - p_L)/(1 + q)$ and demand for the outside good is $Q_L = \widehat{\theta} - 0 = 2(p - p_L)/(1 + q)$, provided all consumers make a purchase in equilibrium. To focus on on food waste at the retail level arising from the equilibrium choice of the quality standard q^e, we consider the case in which $p_L = 0$.[9]

The retail quality standard creates a premium market in which the supply of the graded product, $Q_s(q) = 1 - q$, is matched with retail demand, $Q_d(q, p) = 1 - 2p/(1 + q)$. Because demand for graded products that meet the quality standard cannot exceed the supply of farm products that meet the standard, it follows that $p \geq q(1 + q)/2$. For retail prices such that $p = q(1 + q)/2$, the retail market clears the quantity of the farm product that meets or exceeds the retail quality standard. Conversely, food waste occurs in the model whenever the supply constraint is slack, $p > q(1 + q)/2$, as the quantity of farm products that meet the retailer's grading standard in this case exceeds the quantity demanded by consumers.

Retail profit is given by

$$\pi(p, q; \alpha) = p\left(1 - \frac{2p}{1 + q}\right) - \frac{\alpha q^2}{2}. \tag{2}$$

The retailer's problem is to maximize profits in (2) subject to the constraint that sufficient supply of graded farm products exists to meet retail demand,

$$p \geq q(1 + q)/2. \tag{3}$$

[9] In general, the model is capable of distinguishing between food waste at the farm level under circumstances in which a portion of consumers do not purchase the farm product at a positive price for the ungraded product, $p_L > v$.

Consumer surplus in the retail market for graded products is

$$CS(p,q) = \int_{\widehat{\theta}}^{1} (\theta \widehat{q} - p)\, d\theta. \tag{4}$$

Making use of expressions (1), (2) and (4), welfare can be expressed as

$$W(q,p;\alpha) = CS(p,q) + \pi(p,q;\alpha) = \frac{1+q}{4} - \frac{p^2}{1+q} - \frac{\alpha q^2}{2}. \tag{5}$$

Notice that welfare is monotonically decreasing in p. This implies that a social planner seeking to maximize welfare would wish to set $p = 0$ and then select a retail quality standard to maximize welfare expression (5); however, $Q_d(q,p) = 1$ at $p = 0$, which violates the constraint that sufficient supply of graded product exists to meet retail demand, $Q_s(q) \geq Q_d(q,p)$ in the case of a positive quality standard, $q > 0$. It follows immediately that supply constraint (3) always binds for the social planner.

The social planner's problem is to maximize welfare expression (5) subject to supply constraint (3). Substituting the pricing constraint into expression (5), the optimal grading standard satisfies

$$\frac{dW(q;\alpha)}{dq} = \frac{1}{4}\left(1 - 2q(1+2\alpha) - 3q^2\right) = 0.$$

This condition results in an optimal retail quality standard of

$$q^* = \frac{2\sqrt{1+\alpha+\alpha^2}}{3} - \frac{(1+2\alpha)}{3}.$$

Now consider the optimal grading standard of a price-setting retailer. The retailer maximizes $\pi(p,q;\alpha)$ in (2) with respect to p and q, subject to the supply constraint (3). In the event of a non-binding supply constraint in the retail market, maximizing (2) with respect to p and q, respectively, results in the first-order necessary conditions

$$\frac{d\pi(p,q;\alpha)}{dp} = 1 - \frac{4p}{1+q} = 0. \tag{6}$$

and

$$\frac{d\pi(p,q;\alpha)}{dq} = 2\left(\frac{p}{1+q}\right)^2 - \alpha q = 0. \tag{7}$$

Simultaneously solving equations (6) and (7) yields

$$\begin{aligned} q^m &= \frac{1}{8\alpha}, \\ p^m &= \frac{1+8\alpha}{32\alpha}. \end{aligned}$$

Substituting the retail equilibrium values into demand yields $Q_d^m \equiv Q_d(q^m, p^m) = \frac{1}{2}$. It follows immediately that the supply constraint binds when $\alpha \leq \frac{1}{4}$.

In the case of a binding supply constraint, $Q_s(q) = Q_d(q,p)$, the retailer maximizes

$$\pi^c(q;\alpha) = \frac{q(1+q)(1-q)}{2} - \frac{\alpha q}{2},$$

which yields the solution

$$q^c = \frac{\sqrt{3+\alpha^2}}{3} - \frac{\alpha}{3}.$$

It is straightforward to verify that $q^c > q^*$ for all $\alpha \geq 0$.

Combining these results, the equilibrium retail grading standard is

$$q^e = \begin{cases} \left(\sqrt{3+\alpha^2} - \alpha\right)/3 & \text{for } \alpha \leq \frac{1}{4} \\ 1/8\alpha & \text{for } \alpha > \frac{1}{4} \end{cases}.$$

A retail minimum quality standard results in surplus food at the retail level whenever the cost of imposing a minimum quality standard is sufficiently high ($\alpha \geq \frac{1}{4}$). The reason is that the cost of raising product quality with a higher retail standard exceed the revenue that can be earned from the resulting increment in average food quality, providing retailers with an incentive to lower stocking costs by reducing the retail quality standard while maintaining a relatively high retail price level. The quantity of farm products that meets the lower grading standard rises, while the quantity demanded by consumers at the equilibrium quality-price pair (q^m, p^m) remains fixed through pairwise adjustments in product quality and price. Food waste emerges as a positive share of farm products meeting the retailer's quality standard is left unsold in the retail marketplace at the equilibrium price.

To fix these ideas, consider the case in which $\alpha = 1$. For $\alpha = 1$, the optimal retail quality standard is $q^* = \frac{2\sqrt{3}-3}{3}$, which implies roughly 85% of the farm product is sold in the retail market at a retail price that clears the retail market without food waste, $p^* = \frac{2-\sqrt{3}}{3}$. In contrast, the equilibrium retail price in the private market is $p^m = \frac{9}{32} > p^*$, and the equilibrium retail quality standard is $q^m = \frac{1}{8}$. This implies a quantity of food supplied by the farm sector of $Q_s(\frac{1}{8}) = \frac{7}{8}$ at the retail quality standard, while consumer demand is $Q_d(\frac{1}{8}, \frac{9}{32}) = \frac{1}{2}$, resulting in an equilibrium level of food waste of 37.5% in the retail market.

Notice that in this case that the retailer reduces the grading standard below the socially-optimal level to economize on grading costs: $q^m = \frac{1}{8} < q^*$. The implication is that the retailer artificially reduces the quality standard from the socially optimal level to lower retailing costs, redistributing food from the farm level to the retail level, where graded food subsequently goes unsold. Whether this mechanism prevails in retail markets more generally, however, is an empirical question.

3 Empirical Model of Price Discrimination

In this section, we use the theory of price discrimination to develop an empirical test of the hypothesis developed above, namely whether retailers price discriminate based on the quality of their fresh produce. We then examine the implications for the amount of surplus food left on the market.

To examine the implications for retail price discrimination on food waste in the US apple market, we are interested in statistically identifying differences in the distribution of quality produced on US apple orchards and sold by US retailers. To do so, we define quality in terms of the measurable indicators of eating quality produced on-farm, but not actually sold to retailers. Eating quality is generally measured using three parameters: (1) crispness, (2) sweetness, and (3) acidity. While there is considerable horizontal differentiation among apple varieties, there is empirical evidence that some apple varieties are generally preferred, due to their favorable attribute-profiles (Harker, et al. 2003, 2008; Hampson, et al. 2000; Hampson and Kemp 2003; Gallardo, et al. 2017). Further, eating quality is randomly distributed within each variety (Miller, et al. 2007; Henroid, et al. 2008) due to the fact that apples are subject to the normal variation inherent in any biological production process, and the geographic dispersion of production activity. Therefore, unless retail prices are set to perfectly match the distribution of fresh apple quality produced on farms, some consumers will be excluded from the market by retailers' price discriminating behavior.

We develop an empirical model that we use to examine the potential for price-discriminating behavior by food retailers. Our model is based on a discrete-choice process in which apple items, defined as varieties that are differentiated both horizontally, and vertically across hundreds of geographic markets within the same retail chain. We focus on the sale of bagged apples for the reason cited above, and include six different varieties that differ substantially in terms of their measurable quality metrics (Miller, et al. 2007; Henroid, et al. 2008). We assume consumers purchase one of 14 different combinations of bag-size and apple variety from a store that is chosen in a previous, unmodeled decision stage.

Consumers decide to purchase one alternative from the 14 items in our data according to a random utility framework. That is, consumers choose the item that provides the most utility from all alternatives available, subject only to our assumption regarding the unobserved distribution for variety (horizontal differentiation), and the unobserved preference for quality (vertical differentiation) as in McManus (2007). Let $i = 1, 2, ..., I$ index consumers, $j = 1, 2, ..., J$ index the items (UPCs, or variety-bag combinations) offered by stores, $r = 1, 2, ..., R$, in week $t = 1, 2, ..., T$, so the indirect utility function is written:

$$U_{ijrt} = \theta_{ik} q_j^{\gamma_k} + \alpha_i p_{jrt} + \sum_{l=1}^{L} \beta_l X_{ljrt} + \delta_j + \xi_r + \tau CF_{jrt} + \varepsilon_{ijrt}, \tag{8}$$

where: $\theta_{ik} = \theta_{ok} + \theta_{1k}\mu_k,\ \mu_k \sim \log N(0, \sigma_{\mu k})$ describes the distribution of quality-preference for each variety, k, q_j is the observed quality index for each item, γ_k is the curvature of utility in quality, which is allowed by vary by variety, $\alpha_i = \alpha_o + \alpha_1 \nu,\ \nu \sim N(0, \sigma_\nu)$ is the marginal utility of income, assumed to be normally distributed over consumers (p_{jrt} is the price (per lb.) of item j in retailer r week t, X_{jlrt} is a set of covariates that describe store-specific marketing-mix activity specific to item j and seasonal indicator variables, δ_j is a set of item-fixed effects, ξ_r is a set of store-fixed effects, CF_{jrt} is the control-function value that varies by item, store, and week, and ε_{ijrt} is a random variable, assumed to be unobserved, and uncorrelated with the random elements in quality preference, and the marginal utility of income.[10]

Our set of marketing-mix variables includes a binary indicator of whether a particular item was on temporary price promotion in a specific store and week (defined as at least a 10% reduction from the previous week), and an interaction term between the promotion indicator and the shelf price, intended to capture both the expected shift and rotation of the demand curve during promotional periods. We assume the distribution of ε_{ijrt} is Type I Extreme Value, so the underlying demand model is a non-linear variant of a logit discrete choice model. Our outside option consists of all other fresh fruit purchases at each store, during each week of the sample period. That is, we implicitly assume that the consumer

[10] Leslie (2004) and McManus (2007) each assume a log-normal distribution for the preference for quality as it is not tenable that any consumer would place a negative value on a higher quality product, *ceteris paribus*.

enters the store seeking fruit, so if they are not satisfied with either the quality or price of the apples on offer, they will purchase a different type of fresh fruit. If they do so, we assume they earn an indirect utility that is also Extreme Value distributed such that: $U_{iort} = \varepsilon_{iort}$.

Although the underlying utility model is a relatively standard multinomial logit model (MNL), the fact that we allow for heterogeneity in horizontal item-preferences (ε_{ijrt}), and unobserved heterogeneity in both the willingness to pay for quality (θ_{ik}) and the marginal utility of income (α_i) means that we have to integrate over all distributions in order to estimate the market share for item j in store r. We assume that the distributions for each of these parameters are independent, for the sake of tractability. In this sense, our model is a non-linear variant of a mixed-logit model, which we estimate using simulated maximum likelihood in the absence of closed-form expressions for the aggregate market share variable. Therefore, we write the market share for each item as:

$$s_{jrt}(\Omega) = \int_{A_{jrt}} dF\varepsilon(\varepsilon) dF\mu_k(\mu_k) dF_\nu(\nu), \tag{9}$$

for parameters Ω, where:

$$A_{jrt} = [(\varepsilon, \mu_k, \nu) \mid U_{ijrt}(p_{jrt}, q_j, , X_{ljrt}, CF_{jrt}, \delta_j, \xi_r) \geqslant U_{imrt}(p_{mrt}, q_m, , X_{lmrt}, CF_{mrt}, \delta_m, \xi_r)]$$

represents the set of items that are chosen by consumers. With our assumption that ε is EV distributed, we simplify this expression by including the known functional form for the EV density to arrive at the estimated share function:

$$s_{jrt}(\Omega) = \int_{A_{jrt}} \left(\frac{\exp(\eta_{jrt})}{\sum_{m=1}^{M} \exp(\eta_{mrt})} \right) dF\mu_k(\mu_k) dF_\nu(\nu), \tag{10}$$

where η_{jrt} is the mean utility for item j at store r in week t. But, because μ_k and ν have no similar closed-form, we simulate the likelihood function over the assumed distributions for each. Given that the market consists of N_{jrt} market-share observations, the log-likelihood function is written:

$$\pounds(p_{jrt}, q_j, , X_{ljrt}, CF_{jrt}, \delta_j, \xi_r \mid \Omega) = \sum_{t=1}^{T}\sum_{r=1}^{R}\sum_{j=1}^{J} N_{jrt} \log s_{jrt}(\cdot, \Omega),$$

which we simulate using Halton draws ($H = 50$) in order to improve the efficiency of the simulation process (Bhat 2003).

As is clear from the description above, we estimate the model using a control function approach (Petrin and Train 2010) in order to control for the clear endogeneity of retail prices. We discuss our identification strategy in greater detail below, but summarize the quality of our instruments here.

In general, instruments should be correlated with the endogenous variable of interest, yet independent of the unobservable in the demand equation. While there is, logically, no test for the latter condition, instruments are generally considered appropriate if economic theory suggests there should be a correlation, and the relationship is borne out in the data. In short, we chose our instruments in the expectation that they would be correlated with the cost of producing and / or retailing fresh apples. To that end, our instruments consist of input price indices at the grower (fuel, fertilizer, chemicals, labor, and business services) as well as the packing (wholesale labor) and retail levels (retailing labor and utilities) in addition to variety-specific wholesale apple prices. We include a set of item-specific fixed-effects, as well as a set of seasonal indicator variables in order to control for either seasonal production, or seasonal releases from storage.

Our instruments also include shipments, and wholesale prices. Variation in supply is largely due to planting decisions made many years prior to our data sample and, to a lesser extent, weather considerations. Therefore, supply is pre-determined, and a valid instrument. Conditional on shipment levels, FOB prices are determined in a competitive upstream market, so are also exogenous to the retail pricing decision.

In the first-stage instrumental-variables regression on the retail price, our preferred set of instruments produces an R^2 value of over 91% and an F-statistic of 37.772.8 (table 3). Because the estimated test-statistic value is far greater than the threshold value of 10, our instruments are not weak in the sense of Staiger and Stock (1997).

4 Data

4.1 Data Summary

We derive our data from two sources. First, we rely on highly granular Nielson ScanTrack data for every store of a major US retail chain. These data describe prices and sales volumes for bagged items for each of 6 apple varieties over 52 weeks from October 31, 2014 through October 31, 2015.[11] There are fully 2,800 stores represented in the data, providing a total of 93,700 store-level observations.

We choose to examine the fresh apple retail market for a number of reasons. First, apples are second only to bananas in terms of their sales volume across all US retail food stores. Second, unlike bananas, apples are sold in a range of varieties that capture both horizontal and vertical elements of product differentiation. While individual apple consumers may have a "favorite" variety of apple, giving rise to horizontal-differentiation in the retail market, all consumers may have objective measures of apple quality that differ genetically across apple varieties, allowing retailers the latitude to control retail prices through the use of minimum quality standards on all varieties. Third, among the fresh-produce items offered by retailers, apples tend to be relatively durable, so there is very little promotional activity intended to

[11] Although the majority of apples sold through food retailers are displayed in bulk, retailers do not track such "random-weight" sales in a manner sufficiently consistent for detailed analysis. For bagged apples, we know the size is relatively consistent in each package, so we remove size-variation as a potential, unobserved source of price-variation that may confound our identification of quality-preference.

move apples that would otherwise become unsalable. Fourth, despite their durability, apples, like other fresh produce items, are subject to a range of biological factors that cause quality to vary from season to season, and from shipment to shipment, even within each horizontal tier of product variety. For that reason, our price data contains a considerable amount of price-variation that is useful in identifying variation in demand.

Our unit of observation is the universal product code (UPC), which is unique to each bag-variety combination. That is, our data describes 6 apple varieties (Ambrosia, Fuji, Gala, Honeycrisp, Jazz, and Pink Lady), and 6 different bag sizes (2 lb., 4 lb., 5 lb., 6 lb., 7 lb., and 8 lb.). Because not all varieties are offered in all 6 bag-types, we observe a total of 14 different UPCs, or bag-variety combinations in the data. We refer to each of the unique UPCs as an "item" below to avoid confusion with varieties and bags.

We use wholesale (FOB) price data on fresh apples from the Washington Tree Fruit Association (WTFA). The WTFA maintains a database of farm-gate apple prices on a weekly basis for all Washington-produced varieties. Although these prices are averaged over all size categories, the averages are likely to reflect variation that is more specific to each variety than it is to apples in general. Moreover, there is no guarantee that our retailer purchased all of their apples from Washington state. As the dominant apple-producing state, however, price trends in Washington are likely to capture similar price-variation in each variety in other regions. Because we use these prices only as an instrument for retail-market prices, there is no error in our estimation process induced by errors in approximating the true wholesale price paid by the retailer.

We include other instruments to capture variation in cost that is not reflected in wholesale prices, and the cost of distributing and retailing fresh apples. Weekly wages for workers in the food retailing and fresh-fruit packing industries are from the Bureau of Labor Statistics

Current Employment Statistics Survey (BLS 2018a). Price indices for electricity are also from the Bureau of Labor Statistics Consumer Price and Inflation data tool (BLS 2018b). Other farm-level input prices are from the United States Department of Agriculture (National Agricultural Statistics Service, NASS) Index of Prices Paid data base (USDA), and include monthly measures of chemical, fuel, labor, services, and a composite index of interest, taxes, and wages. We also include shipment quantities from the WTFA to proxy the total amount of production.

We construct indices of eating quality for each variety using data from the horticulture and post-harvest literature (Miller et al. 2007; Henroid, et al. 2008). While observed quality is likely to vary by region, crop-year, and even by grower, these studies are conducted under controlled conditions, designed to produce attribute-measures that are at least broadly representative of the inherent differences in eating quality among the different varieties. To the extent that they do not capture the quality differences between the apple varieties in our data, we estimate the response to quality using a random-parameters approach, which is intended to control for any unobserved heterogeneity that is otherwise impossible to account for in a more systematic way. The indices are constructed by summing the mean-centered values for crispness (pounds pressure), sweetness (soluble solids content, %), and acidity (%). The resulting quality indices produce rather sharp differences between the measured quality of each variety, so they do not appear to be "too similar" to distinguish econometrically. We summarize the price, wholesale movement, retail sales, and quality data for each variety in Table 1, and the retail data by item (UPC) in Table 2.

[Table 1 in here]

[Table 2 in here]

Based on these data, it is clear that there is sharp variation in quality, and substantial variation in market shares, retail prices, and wholesale prices. Pink Lady apples appear to be the highest quality among our 6 varieties, and they also have the highest implied markup over wholesale price (124%). Honeycrisp apples, on the other hand, have the highest retail price, by far, and the second-largest market share; however, they have the lowest markup (71%) because their wholesale price is also highest among all varieties. Gala apples tend to have the highest market share, are relatively inexpensive, and are among the top-quality apples, but they are widely adopted by farmers, so production quantity may constrain their ability to generate large margins. Overall, the data in this table suggests that there is sufficient variation in prices, market shares, and quality indices across apple varieties to identify the parameters of the demand model.

Retail prices are clearly driven by something other than just quality considerations. In the first-stage, instrumental-variables regression described above, we control for many of the factors that are likely to influence retail prices, but are independent from any unobservable factors that retailers are likely to take into account in forming retail prices. The parameters of this regression are shown in Table 3 below. Based on these estimates, it is apparent that different variety / pack combinations are indeed priced different from each other, and marketing costs do have an important effect on retail pricing decisions. Any remaining variation, therefore, is likely due to market-preferences for quality. We return to this issue in more detail next.

[Table 3 in here]

4.2 Price Variation and Demand Identification

There are many reasons why observed prices in Table 1 vary across the items in our data set, some of which are useful in identifying variation in demand, and others not. Price variation that is due to endogenous factors is likely to be unobserved by the econometrician, left to the error term, and thus not useful for demand estimation. In the model presented below, we attempt to remove as many observable factors as possible from the error term, so the remaining price variation identifies variation in demand due to price and quality as cleanly as possible.

First, apples of the same variety and package vary in price from store to store for reasons of geography. Prices for the same item may vary from region to region for many reasons. Because the spatial distribution of apple production differs from the spatial distribution of stores (most apples are produced in Washington, New York, and Michigan, for example), prices will differ to the extent that transportation costs cause the delivered-price to vary from store to store. Further, the retailer practices "zone pricing" which means that prices for the same item are priced in order to reflect differences in willingness-to-pay associated with observed market attributes.[12] We account for geographic variation in demand by using a panel-data estimator that accounts for variation in demand over stores, and time. While there may be several stores in each geographic market, our data does not reveal where each store is situated. However, our panel-estimation strategy is preferred because it captures the fact that, even if two stores are in the same "market" for aggregation purposes, the demand conditions for sub-markets may be radically different. For example, stores in the Chicago area are typically described as being located in the same market, but stores on the south side of Chicago and the north side are in substantially different markets.

Second, apple prices vary over time. Depending on the size of the harvest, changes in FOB prices, seasonal trends in preferences, prices for competitive goods in the store, and projections of supply-availability (storage) much of the price-variation in our data occurs over time. We control for the inherent seasonality of fresh fruit prices by allowing for a set of

[12] Zone pricing is an example of price discrimination, but third-degree price discrimination and not the second-degree price discrimination that we seek to identify.

seasonal indicator variables in the demand model. We control for elements of the temporal variation in prices that reflect endogenous concerns, that is, pricing and stocking strategies by the retailer that are responses to its perception of demand-variation, through our control-function approach. Exogenous price-variation that remains helps identify changes in demand from week to week.

Third, our apple prices vary from package to package. While most apples are sold in bulk form, data on bulk apple sales would be inappropriate for this study as it would be impossible to control for intra-variety variation in size, quality, and appearance. Bagged apples tend to be very similar in size, and are more generally uniform in appearance than apples sold in bulk. That said, prices vary from bag to bag in a way that may suggest a different form of second-degree price discrimination, one intended to take advantage of volume pricing discounts. We control for this effect by including a set of item fixed-effects, removing the source of variation that may be due solely to volumetric concerns.

Fourth, and most importantly, apple prices vary from variety to variety. Variety-based price-variation may be due to consumer preference for particular varieties, higher acquisition costs, or a combination of the two. We control for variety-based preferences through the item fixed-effects described above, and for higher costs of production through the control function. Because varieties differ in their measurable quality attributes, and are assumed to affect utility in a non-linear way, controlling for other sources of observed price-variation among varieties allows us to cleanly identify consumers' willingness-to-pay for quality.

We first examine the data for any reduced-form, or model-free, evidence of price-discriminating behavior. Based on our understanding of second-degree price discrimination from first principles, we know that price differences across items, when the difference in price is not related to the cost of acquisition or selling, constitutes *prima facie* evidence of price discrimination.

5 Results and Discussion

We first present our estimates of a reduced-form model of demand in which we reveal some basic features of our data, and then present more detailed findings from the quality-based price discrimination model described above. We then conclude this section by presenting some findings from a series of counter-factual simulations in which we demonstrate the concrete implications of how the distribution of willingness-to-pay for quality differs from the true distribution of quality delivered off the farm.

5.1 Consumer Valuation of Quality

Estimates of a linear version of the consumer willingness-to-pay for quality model are presented in Table 4. As suggested by the full non-linear model in (8) above, we allow the quality-preference parameter to vary by store (as a proxy for market-level preferences). We then recover the individual-level preferences by calculating the implied individual-level parameters, conditional on both the data, and the population estimates (Train 2003). We then describe the distribution of the individual parameters non-parametrically, by fitting a kernel density estimator to the individual-level parameter vector. Our kernel density estimator allows us to make inferences regarding the distribution of the willingness-to-pay for quality without imposing assumptions on the underlying distribution of parameters.[13] Namely, we use a Bowman and Shenton (1975) Chi-square test of normality to compare the fitted kernel density to a normal distribution, and calculate coefficients of skewness and kurtosis to examine exactly how the resulting distribution departs from normality, if indeed it does.

[Table 4 in here]

[13] A kernel density function is composed of a weighting function, which is any mathematical function that integrates to 1.0. We assume an Epanechnikov (1969) weighting function in order to maintian maximum flexiblity in the shape of the kernel density.

In estimating a linear version of our maintained model in (8), we remove the potential for non-linear returns to quality. That is, consumers are assumed to prefer one increment in quality to any other.[14] While this is a simplification, it is more tractable than the full non-linear model, and provides a robustness check of the main hypothesis of the paper, namely that the distribution of willingness-to-pay for quality does not match the distribution of quality produced on the farm. Using the results reported in this table, we first examine whether a more parsimonious version of the maintained utility model, one with fixed parameters, provides a better fit to the data. Comparing the fit of Model 1 with Model 2 in Table 4 using a likelihood-ratio (LR) test shows that the random parameter model is preferred ($\chi^2 = 19,272.9$ compared to a critical value at a 5% level of 5.99). Further, we see that the standard deviation of the preference for quality is significantly different from zero, supporting the random-parameter model. In fact, comparing the point estimates of the two models, the fixed-coefficient model appears to over-state the preference for quality by nearly double. Clearly, unobserved heterogeneity represents a substantial component of the preference for variety, and the bias in the willingness-to-pay for individual items reflects this observation.

Figure 1 shows the non-parametric distribution of the willingness-to-pay for quality implied by these estimates, relative to a log-normal distribution of quality. In the context of apple production, Din et al. (2003) show that the actual distribution of quality on the farm is distributed log-normal, indicating that use of the log-normal distribution is appropriate. It is clear from Figure 1 that the empirical distribution of consumer willingness-to-pay for quality is not normal. Moreover, it is significantly different from log-normal as well ($\chi^2 = 6.55$,

[14] McManus (2007) rules out this case as his price schedules are concave in q_i (coffee size) so that any consumer with a positive preference for size will always purchase a large coffee. In our data, price schedules are not concave, so this problem does not arise.

compared to a critical value of 3.84). Indeed, there appears to be a much steeper gradient in consumer preferences for product quality than exists in the actual quality of apples produced on farms, an observation that is consistent with the use of minimum quality standards as an instrument for retail price discrimination.

[Figure 1 in here]

It is also possible that quality preferences are non-linear (McManus 2007). In this case, we assume consumers perceive declining returns to quality, so that the difference between two high-quality apples is smaller than the difference between high-quality and low-quality apples. We conduct the same analysis for the non-linear model, first comparing a fixed and random-coefficient version of the model, and then estimating the kernel density of the quality-preference parameter that emerges from the preferred model. These estimates are shown in Table 5 below. According to the LR specification test, we prefer the random-coefficient version of the model, as unobserved heterogeneity indeed appears to be important in determining fresh apple preferences ($\chi^2 = 7,230.5$). As in the linear model, the preference for quality in the random-coefficient version of the model is substantially greater than the fixed-coefficient estimate, suggesting that the bias in not accounting for unobserved heterogeneity understates the preference for quality.

[Table 5 in here]

Similar to the linear-model case, we compare the empirical distribution of quality-preference to a log-normal comparator (Figure 2) by estimating the kernel density of the willingness-to-pay for quality parameter. Again, we reject the null hypothesis that the log of quality-preference is normally distributed, and conclude that the empirical distribution of preferences differs from the log-normal distribution that governs quality of harvested apples ($\chi^2 = 52.73$). Comparing the linear and the non-linear densities, the departure from normality is more clear in the non-linear case, which suggests that linearity may mask some of the true distribution of quality-preferences.

[Figure 2 in here]

5.2 Implications for Food Waste

In this section we numerically characterize the implications of retailers' price-discrimination behavior on the amount of induced food waste, using our bagged-apples example as a case study. Recall that the underlying logic of our approach is to compare the distribution of what is purchased at retail, with what is produced on the farm. Once the distribution of willingness-to-pay for quality is established, the key empirical challenge is to identify the distribution of *actual* quality, or the quality that leaves the farm.

There are two alternatives in characterizing this distribution. First, we could use a "revealed preference" argument similar to Ellickson, Misra, and Nair (2012) to suggest that produce buyers for our supermarket are sufficiently sophisticated to pay wholesale prices that reflect true apple quality. If this assumption is valid, then we could estimate another kernel density on wholesale apple prices and compare the wholesale and retail distributions of revealed quality. However, the weakness in this argument is clear: Buyers for retail chains are not likely to purchase apples that they know will not meet minimum quality standards. Therefore, we rely on a second option, which involves referring to the horticultural literature to obtain field-trial estimates of the distribution of quality available to buyers for the retail chains.

Despite the importance in developing an understanding of how much fruit grown at the farm level actually makes it to the retail market, there is very little empirical evidence on this issue. In the horticulture and post-harvest literatures, Din et al. (2003) remains the only research that specifically addresses this issue.[15] Specifically, they estimate the distribution of harvested-apple quality using metrics for firmness and sweetness similar to the ones we use in our quality index above. Using a number of apple-varieties, and harvest-dates, they find that a log-normal distribution best characterizes the natural distributions of both firmness and sweetness. Because the particular parameterization for their log-normal distribution is unique to their field-trial, we compare our empirical distribution of quality to a generic log-normal distribution that is comparable as possible to the apples in our study. That is, we need to be able to compare the two distributions, both with substantial right skew, for the amount that is likely to remain in the left-tail of the farm-level distribution of quality particular to our sample of apples.

We normalize the distribution of quality produced on-farm by assuming that the means of the retail and farm-level quality distributions are equal. This assumption is reasonable as it provides a measure of central tendency that is likely to be of empirical importance to both apple packers, and retail buyers. That is, neither of these players is likely to have sufficient data to form expectations as to the location of the median or the mode, two other potential candidates, so their assessment of what the "average" apple looks like is the most sensible. With this assumption, the more skewed the distribution of farm-quality, the less mass there will be to the left of the minimum quality purchased for sale. Further, because the mean of a log-normal distribution is higher than either the median or mode, this assumption represents a very conservative estimate of the amount of food loss. We explore the sensitivity of our findings to this assumption in a set of counter-factual simulations below.

Numerical estimates of the both non-parametric kernel densities, their associated log-normal comparisons, and the associated levels of food loss are in Table 6. In constructing

[15] Although this evidence is from Israel, the production conditions are very similar to those found in the primary growing region of the US (Washington state).

this value, we simply find the area between the willingness-to-pay for quality density function, and the log-normal benchmark. By integrating the difference between the two density functions numerically, we arrive at an estimate of the percentage of apples that are grown that do not make retail standards. For the linear model, the estimates in Table 6 show that approximately 10% of the apple harvest does not make retail grade. In the non-linear case, the steeper gradient shown in Figure 2 implies a greater amount of food loss, nearly 12.1% for our parameterization. While food loss at the retail level forms a relatively small part of the total amount of food lost in the entire food system (Buzby and Hyman 2012), based on the value of the 2016 apple harvest (USDA ERS) the more conservative of these two estimates represents a loss of some $350 million per year.[16] Aggregated over all fresh fruit and vegetable categories, the amount of loss is substantial indeed.

[Table 6 in here]

We examine the sensitivity of our food loss estimates to our normalization assumption as it is clearly key to determining how much retailers sell, and how much they reject. Our key assumption is that the mean of the empirical and implied-harvest distributions are the same. Therefore, we shift the distribution of the mean willingness-to-pay for quality to the right, moving away from our conservative baseline above, toward cases that are likely more realistic. For illustrative purposes, we shift the mean of the distribution to the right by 1%, 2%, 5%, and 10% to show the potential effect on food loss. Another way of interpreting this sensitivity is to consider each distribution as representing a different type of store. While a 1% shift relative to the consumers in our data (a large discount retailer) may represent Kroger shoppers, a 10% may capture Whole Foods rejection policy. With a 1% shift, we find that the amount of implied food loss rises from 10% to over 26.3%. At 2%, the loss rises

[16] Clearly, the term "loss" in this case implies that the secondary market for apples generates much less value than the retail market.

to 31.4%, to 44.3% at 5%, and 49.2% at a 10% shift. In other words, if shoppers demand even a 10% increase in quality from their retailers, the amount of apples discarded on the packhouse floor would increase by nearly half.

Our findings are immediately relevant to the case of apples, but likely generalize to any perishable food product with an inherent distribution of quality that differs from the distribution of willingness to pay for quality. The fundamental economic principal that governs the amount of food loss in our analysis is the tendency of retailers to price discriminate on the basis of quality. With increasing consumer demands for local, GMO free, organic, and premium fresh produce, retailers' ability to extract surplus through quality-based, second-degree price discrimination strategies will only increase. An unintended consequence of the demand for quality, therefore, will be more loss in the system.

In our model, food loss results from the rational, optimizing behavior of imperfectly competitive retailers. Although the extent of market failure associated with the degree of power exercised by food retailers is likely small, the market opportunity is much larger, so the implications for policy are clear. Emerging platforms, such as Imperfect Produce or Food Cowboy (Richards and Hamilton 2018), suggest that entrepreneurs are making markets for the type of fresh produce that traditional retailers leave on the packing-house floor. Our conversations with executives from these firms provide only anecdotal evidence, but the common theme in these "how can we help you?" discussions is the need for fundamental policy reform: Tax breaks for gleaning operations, exemption from liability laws and marketing-order regulations, and rational guest-worker policies are the most common. Policy of this type need not supplant existing markets, but provide conditions for new markets to arise, and flourish.

6 Conclusions

In this paper, we study whether retailers' second-degree, quality-based price discrimination policies can explain food loss between the farm and retail levels. If retailers price-discriminate based on the quality of the produce their shoppers demand, then they will essentially truncate the distribution of quality they sell through a system of minimum quality standards. Yet, fresh produce that arrives at the packinghouse door is more likely to be continuously distributed, ranging from low-quality products below saleable levels in the retail market, to high-quality products that appeal to only the most discerning consumers. The difference between the distribution of quality purchased by retailers, and that grown on the farm, is the amount of food loss between the farm and retail levels of the food supply chain.

We develop an empirical estimate of the extent of food loss through this mechanism using data on fresh apple sales from a major US retailer, and agronomic data on the actual distribution of harvested apple quality. We estimate a structural model of horizontal and vertical product differentiation in order to separate the element of retail pricing that is due to natural variation in consumer tastes for different types of apples from consumers' preferences for apples that are simply better, as measured by objective quality metrics. We then form non-parametric kernel density estimates of consumers' willingness-to-pay for quality in order to better understand how well the demand for quality matches its supply.

We find a substantial departure from the distribution of consumers' willingness to pay for quality, and the natural distribution of quality that arrives from the orchard. Using a very conservative method of comparing the distributions of the supply of and the demand for quality, we find that the amount of food loss in apples due to quality-based price discrimination is approximately 10%. However, if we move to a less conservative method of comparing the distributions, we find food-loss estimates of up to 50% if consumers' willingness-to-pay for quality is only 10% greater than our baseline estimates.

The policy implications that arise from our analysis are clear. Our findings suggest that much of the economic loss in the supply chain, whether through discarding sub-grade produce or through value degradation, is the result of rational, profit-maximizing decisions on the part of retailers. Public intervention that prevents retailers from price discriminating on the basis of quality would be one way to solve the problem; however, such an approach would involve a level of regulation that is unlikely to be acceptable. Rather, market solutions such as the platforms described above, with either independent suppliers or the retailers themselves selling blemished or low-quality produce in secondary markets as part of zero-waste marketing programs, may be a more viable long-term solution.

Future research in this area may consider data on a wider range of fresh-produce categories, sold by a different set of retailers. Our findings are specific to fresh apples sold by a single, low-priced, retailer, so may be substantially different for a different item, and different seller. There are also other ways of conceptualizing retailers' produce-selling strategies beyond simple quality-based price discrimination. Food banks and blemished-produce platforms may serve as secondary markets (Anderson and Ginsburgh 1994; Chen, Estaban, and Shum 2013) that effectively facilitate retailers' high-margin produce merchandising strategies. Further, it may also be the case that retailers with multiple platforms, serving different socioeconomic markets, may sell lower-quality produce as "damaged goods" in attempting to protect their premium markets (Deneckere and McAfee 1996; McAfee 2007). We leave these issues for future research.

References

[1] Aguirre, I., Cowan, S.G.B., Vickers, J. (2010). Monopoly price discrimination and demand curvature. *American Economic Review*, 100 (4), 1601–1615.

[2] Ailawadi, K. L., & Harlam, B. (2004). An empirical analysis of the determinants of retail margins: the role of store-brand share. *Journal of Marketing*, 68(1), 147-165.

[3] Anderson, E. T., & Dana Jr, J. D. (2009). When is price discrimination profitable? *Management Science*, 55(6), 980-989.

[4] Anderson, S. P., & Ginsburgh, V. A. (1994). Price discrimination via second-hand markets. *European Economic Review*, 38(1), 23-44.

[5] Bell, D. R., Ho, T. H., & Tang, C. S. (1998). Determining where to shop: Fixed and variable costs of shopping. *Journal of Marketing Research*, 35(3), 352-369.

[6] Bell, D. R., & Lattin, J. M. (1998). Shopping behavior and consumer preference for store price format: Why "large basket" shoppers prefer EDLP. *Marketing Science*, 17(1), 66-88.

[7] Bellemare, M. F., Çakir, M., Peterson, H. H., Novak, L., & Rudi, J. (2017). On the measurement of food waste. *American Journal of Agricultural Economics*, 99(5), 1148-1158.

[8] Besanko, D., Dubé, J. P., & Gupta, S. (2003). Competitive price discrimination strategies in a vertical channel using aggregate retail data. *Management Science*, 49(9), 1121-1138.

[9] Bockstael, N. E. (1984). The welfare implications of minimum quality standards. *American Journal of Agricultural Economics*, 66(4), 466-471.

[10] Borenstein, S. (1991). Selling costs and switching costs: Explaining retail gasoline margins. *RAND Journal of Economics*, 22, 354-369.

[11] Borenstein, S. and N. L. Rose. (1994). Competition and price dispersion in the U.S. airline industry. *Journal of Political Economy*, 102, 653-683.

[12] Bowman, K.O.; Shenton, L.R. (1975). Omnibus contours for departures from normality based on b1 and b2. *Biometrika*, 62(2), 243–250.

[13] Busse, M., & Rysman, M. (2005). Competition and price discrimination in yellow pages advertising. *RAND Journal of Economics*, 36(2), 378-390.

[14] Buzby, J. C., & Hyman, J. (2012). Total and per capita value of food loss in the United States. *Food Policy*, 37(5), 561-570.

[15] Chen, J., Esteban, S., & Shum, M. (2013). When do secondary markets harm firms? *American Economic Review*, 103(7), 2911-2934.

[16] Clerides, S. K. (2002). Book value: intertemporal pricing and quality discrimination in the US market for books. *International Journal of Industrial Organization*, 20(10), 1385-1408.

[17] Coase, R. H. (1972). Durability and monopoly. *Journal of Law and Economics*, 15(1), 143-149.

[18] Cohen, A. (2008). Package size and price discrimination in the paper towel market. *International Journal of Industrial Organization*, 26(2), 502-516.

[19] Corts, K. (1998). Third-degree price discrimination in oligopoly: all-out competition and strategic commitment. *RAND Journal of Economics*, 29(2), 306-323.

[20] Courty, P., & Pagliero, M. (2012). The impact of price discrimination on revenue: Evidence from the concert industry. *Review of Economics and Statistics*, 94(1), 359-369.

[21] Crawford, G. S., & Shum, M. (2007). Monopoly quality degradation and regulation in cable television. *Journal of Law and Economics*, 50(1), 181-219.

[22] Deneckere, R. J., & McAfee, R. P. (1996). Damaged goods. *Journal of Economics & Management Strategy*, 5(2), 149-174.

[23] Din, G. Y., Zugman, Z., Saks, Y., & Kim, M. (2003). Lognormal Distribution of Apples in Relation to Maturity Restriction at Harvest. In *International Conference on Quality in Chains. An Integrated View on Fruit and Vegetable Quality* 604, 605-610.

[24] Ellickson, P. B., Misra, S., & Nair, H. S. (2012). Repositioning dynamics and pricing strategy. *Journal of Marketing Research*, 49(6), 750-772.

[25] Epanechnikov, V. A. (1969). Non-parametric estimation of a multivariate probability density. *Theory of Probability and its Applications*, 14(1), 153–158.

[26] Gallardo, R.K., I. Hanrahan, J. Luby, C. Yue, V. McCracken, J. McFerson, C. Ross, and L. Carrillo-Rodriguez. (2017). Combining sensory evaluations and experimental auctions to elicit willingness-to-pay for fresh fruit quality characteristics. *Agribusiness: An International Journal*. http://onlinelibrary.wiley.com/doi/10.1002/agr.21534/full

[27] Hampson, C. R., & H. Kemp. (2003). Characteristics of important commercial apple varieties, in *Apples: botany, production and uses* (pp. 61–89).

[28] Hampson, C. R., Quamme, H., Hall, J., MacDonald, R., King, M., & Cliff, M. (2000). Sensory evaluation as a selection tool in apple breeding. *Euphytica*, 111, 79–90.

[29] Harker, F. R., Gunson, F. A., & Jaeger, S. R. (2003). The case for fruit quality : an interpretive review of consumer attitudes , and preferences for apples. Postharvest Biology and Technology, 28, 333–347.

[30] Harker, F.R., Kupferman, E.M., Marin, A.B., Gunson, F.A., & Triggs, C.M. (2008). Eating quality standards for apples based on consumer preferences. Postharvest Biology and Technology, 50, 70-78.

[31] Henroid, R. E., Tustin, D. S., Breen, K. C., Oliver, M., Dayatilake, G. A., Palmer, J. W., Diack, R., & Johnston, J. (2008). Thinning effects on'Scifresh'apple fruit quality at harvest and after storage. In *IX International Symposium on Integrating Canopy, Rootstock and Environmental Physiology in Orchard Systems* 903: 783-788.

[32] Lapan, H., & Moschini, G. (2007). Grading, minimum quality standards, and the labeling of genetically modified products. *American Journal of Agricultural Economics*, 89(3), 769-783.

[33] Leland, H. (1979). Quacks, lemons, and licensing: A theory of minimum quality standards. *Journal of Political Economy*, 87(6), 1328-1346.

[34] Leslie, P. (2004). Price discrimination in Broadway theatre. *RAND Journal of Economics* 35, 520–41.

[35] McAfee, P. (2007). Pricing damaged goods. *Economics: The Open-Access, Open-Assessment E-Journal,* 1 (2007-1): 1–19.

[36] McManus, B. (2007). Nonlinear pricing in an oligopoly market: The case of specialty coffee. *RAND Journal of Economics*, 38(2), 512-532.

[37] Miller, S., McNew, R., Belding, R., Berkett, L., Brown, S., Clements, J., Cline, J., Cowgill, R., Crassweller, E., Garcia, D., Greene, D., Greene, G, Hampson, C., Merwin, I., Moran, R., Roper, T., Schupp, J., & Stover, E. (2007). Performance of apple cultivars in the 1995 NE-183 Regional Project Planting. *Journal of the American Pomological Society*, 58(2), 65-77.

[38] Moorthy, K.S. (1988). Product and price competition in a duopoly. *Marketing Science*, 7(2), 141-168.

[39] Mussa, M., & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18(2), 301-317.

[40] Nevo, A., & Wolfram, C. (2002). Why do manufacturers issue coupons? An empirical analysis of breakfast cereals. *RAND Journal of Economics,* 319-339.

[41] Petrin, A., & Train, K. (2010). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research*, 47(1), 3-13.

[42] Richards, T. J., & Hamilton, S. F. (2018). Food waste in the sharing economy. *Food Policy*, 75, 109-123.

[43] Risk Management Association. *Annual Statement Studies*, various years. See: www.rmahq.org

[44] Rochet, J. C., & Stole, L. A. (2002). Nonlinear pricing with random participation. *Review of Economic Studies*, 69(1), 277-311.

[45] Saitone, T. L., & Sexton, R. J. (2010). Impacts of minimum quality standards imposed through marketing orders or related producer organizations. *American Journal of Agricultural Economics,* 92(1), 164-180.

[46] Shaffer, G., & Zhang, Z. J. (1995). Competitive coupon targeting. *Marketing Science*, 14(4), 395-416.

[47] Shaked, A., & Sutton, J. (1982). Relaxing price competition through product differentiation. *Review of Economic Studies*, 49(1), 3-13.

[48] Shaked, A., & Sutton, J. (1987). Product differentiation and industrial structure. *Journal of Industrial Economics*, 36(2), 131-146.

[49] Shephard, A. (1991). Price discrimination and retail configuration. *Journal of Political Economy*, 99, 30-53.

[50] Staiger, D., & Stock, J. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3), 557-586.

[51] Stole, L. A. (2007). Price discrimination and competition. *Handbook of Industrial Organization*, 3, 2221-2299.

[52] Takeyama, L. N. (2002). Strategic vertical differentiation and durable goods monopoly. *Journal of Industrial Economics*, 50(1), 43-56.

[53] Train, K. (2003). *Discrete choice methods with simulation.* Cambridge, UK: Cambridge University Press.

[54] United States Department of Agriculture, National Agricultural Statistics Survey. (2018). Prices Paid Survey and Indices.

(https://www.nass.usda.gov/Surveys/Guide_to_NASS_Surveys/Prices_Paid_and_Prices_Paid_I
Accessed April 25, 2018.

[55] United States Department of Agriculture, Economic Research Service. (2018). Fruit and tree nut yearbook tables. (https://www.ers.usda.gov/data-products/fruit-and-tree-nut-data/fruit-and-tree-nut-yearbook-tables/). Accessed May 14, 2018.

[56] United States Department of Labor, Bureau of Labor Statistics. (2018a). Employment, Hours, and Earnings. Current Employment Statistics Survey. (https://www.bls.gov/data/). Accessed April 25, 2018.

[57] United States Department of Labor, Bureau of Labor Statistics. (2018b). Inflation and Prices. (https://www.bls.gov/data/). Accessed April 25, 2018.

[58] Varian, H. (1985). Price discrimination and social welfare. *American Economic Review*, 75(4), 870-875.

[59] Verboven, F. (2002). Quality-based price discrimination and tax incidence: Evidence from gasoline and diesel cars. *RAND Journal of Economics*, 33(2), 275-297.

Table 1. Distribution of Retail Data by Variety

Variety	Measure	Units	Value	Std. Dev.
Ambrosia	Retail Price	$ / lb	1.7898	0.1660
Fuji	Retail Price	$ / lb	1.1745	0.2329
Gala	Retail Price	$ / lb	1.1435	0.2480
Honeycrisp	Retail Price	$ / lb	2.5800	1.4551
Jazz	Retail Price	$ / lb	1.5258	0.1851
Pink Lady	Retail Price	$ / lb	1.7138	1.0811
Ambrosia	Wholesale Price	$ / lb	0.3536	0.0469
Fuji	Wholesale Price	$ / lb	0.5835	0.0577
Gala	Wholesale Price	$ / lb	0.5553	0.0381
Honeycrisp	Wholesale Price	$ / lb	1.5121	0.3699
Jazz	Wholesale Price	$ / lb	0.3912	0.0950
Pink Lady	Wholesale Price	$ / lb	0.7638	0.0623
Ambrosia	Market Share	%	7.4413	0.6971
Fuji	Market Share	%	20.7214	2.2970
Gala	Market Share	%	31.8340	3.7501
Honeycrisp	Market Share	%	24.6542	3.2368
Jazz	Market Share	%	5.3947	0.6218
Pink Lady	Market Share	%	9.9545	1.1737
Ambrosia	Quality Index	Value	2.5420	0.0000
Fuji	Quality Index	Value	2.6694	0.0000
Gala	Quality Index	Value	3.1077	0.0000
Honeycrisp	Quality Index	Value	3.0082	0.0000
Jazz	Quality Index	Value	2.6022	0.0000
Pink Lady	Quality Index	Value	4.0706	0.0000
Ambrosia	Shipments	,000 Lbs. / wk.	3.1622	0.9063
Fuji	Shipments	,000 Lbs. / wk.	13.1204	4.8946
Gala	Shipments	,000 Lbs. / wk.	19.6129	8.9026
Honeycrisp	Shipments	,000 Lbs. / wk.	6.0061	2.6901
Jazz	Shipments	,000 Lbs. / wk.	0.4256	0.3477
Pink Lady	Shipments	,000 Lbs. / wk.	3.6735	1.0254

Source: Nielsen scanner data for Retailer X, and quality indices from Miller, et al. (2004, 2007) and Henroid, et al. (2008).

Table 2. Distribution of Retail Data by UPC

Item	Description	Measure	Units	Value	Std. Dev.
Item 1	Ambrosia, 4 lb.	Retail Price	\$ / lb	1.7894	0.1667
Item 2	Fuji, 5 lb.	Retail Price	\$ / lb	1.2133	0.2095
Item 3	Fuji, 6 lb.	Retail Price	\$ / lb	1.2366	0.0984
Item 4	Fuji, 7 lb.	Retail Price	\$ / lb	1.0241	0.0954
Item 5	Gala, 5 lb.	Retail Price	\$ / lb	1.2059	0.2415
Item 6	Gala, 6 lb.	Retail Price	\$ / lb	1.1973	0.1408
Item 7	Gala, 7 lb.	Retail Price	\$ / lb	0.9899	0.1032
Item 8	Gala, 8 lb.	Retail Price	\$ / lb	0.8614	0.1225
Item 9	Honeycrisp, 4 lb.	Retail Price	\$ / lb	2.3584	0.4606
Item 10	Jazz, 4 lb.	Retail Price	\$ / lb	1.6063	0.1942
Item 11	Jazz, 4 lb.	Retail Price	\$ / lb	1.3948	0.0383
Item 12	Pink Lady, 2 lb.	Retail Price	\$ / lb	3.4389	0.1603
Item 13	Pink Lady, 4 lb.	Retail Price	\$ / lb	1.4132	0.1810
Item 14	Pink Lady, 5 lb.	Retail Price	\$ / lb	1.3632	0.0950
Item 1	Ambrosia, 4 lb.	Market Share	%	2.9357	0.6971
Item 2	Fuji, 5 lb.	Market Share	%	7.4694	2.0379
Item 3	Fuji, 6 lb.	Market Share	%	8.0883	2.1941
Item 4	Fuji, 7 lb.	Market Share	%	9.7725	2.7565
Item 5	Gala, 5 lb.	Market Share	%	12.3489	3.2499
Item 6	Gala, 6 lb.	Market Share	%	14.3848	4.3224
Item 7	Gala, 7 lb.	Market Share	%	13.5813	3.6283
Item 8	Gala, 8 lb.	Market Share	%	8.1562	4.6149
Item 9	Honeycrisp, 4 lb.	Market Share	%	9.7265	3.2368
Item 10	Jazz, 4 lb.	Market Share	%	1.8121	0.5086
Item 11	Jazz, 4 lb.	Market Share	%	2.6321	0.7285
Item 12	Pink Lady, 2 lb.	Market Share	%	1.0448	0.3607
Item 13	Pink Lady, 4 lb.	Market Share	%	4.3635	1.0928
Item 14	Pink Lady, 5 lb.	Market Share	%	3.6838	1.2924

Source: Nielsen scanner data for Retailer X. Conditional market shares.

Table 3. Instrumental Variables Regression

Variable	Estimate	Std. Err.	t-ratio
Type 2	-0.6027*	0.0018	-336.7095
Type 3	-0.8678*	0.0026	-336.3605
Type 4	-0.9580*	0.0025	-378.6364
Type 5	-0.6320*	0.0017	-373.9763
Type 6	-0.8766*	0.0026	-342.4063
Type 7	-0.9465*	0.0027	-345.4489
Type 8	-1.1780*	0.0034	-347.4838
Type 10	-0.2269*	0.0022	-104.5438
Type 11	-0.4975*	0.0027	-182.2454
Type 12	1.3279*	0.0058	228.9431
Type 13	-0.5691*	0.0021	-271.0000
Type 14	-0.7476*	0.0028	-266.0320
Electricity	-0.0303*	0.0005	-63.0833
Retail Wage	0.0138*	0.0005	28.7083
Packing Wage	0.0017*	0.0001	12.0714
Fertilizer	-0.0226*	0.0007	-30.5811
Pest Chemicals	-0.2207*	0.0063	-35.0333
Services	0.0499*	0.0016	31.1938
Fuel	-0.0087*	0.0003	-28.9000
Farm Labor	-0.0659*	0.0021	-30.9296
Interest, Taxes	0.2643*	0.0073	36.4593
Plant Chemicals	0.0351*	0.0006	60.5862
Quarter 2	0.0719*	0.0030	23.9003
Quarter 3	0.1043*	0.0033	31.3183
Quarter 4	0.0983*	0.0032	31.1968
FOB Price	0.5879*	0.0014	411.1469
Shipments	-0.0010*	0.0000	-22.6190
R^2	0.9129		
F	37,772.8		

Note: A single asterisk indicates significance at a 5% level.

Table 4. Empirical Model of Price Discrimination: Linear Case

	Model 1: Fixed		Model 2: Random	
Variable	Estimate	Std. Err.	Estimate	Std. Err.
Constant	-11.9334*	0.1138	-20.3240*	4.2088
Item 2	-0.8141*	0.0275	0.1202	0.4707
Item 3	0.1431*	0.0379	1.0031*	0.4712
Item 4	-0.4658*	0.0425	0.3761	0.4717
Item 5	-1.9721*	0.0520	-7.8549*	1.6106
Item 6	-1.2963*	0.0616	-7.1599*	1.6104
Item 7	-2.0278*	0.0671	-7.9229*	1.6104
Item 8	-3.3808*	0.0749	-9.2741*	1.6109
Item 10	-1.0077*	0.0191	3.3035*	1.1722
Item 11	-0.7510*	0.0290	3.6070*	1.1720
Item 12	-5.1108*	0.0631	-8.4414*	1.7979
Item 13	-6.4321*	0.1027	-9.1338*	1.7987
Item 14	-6.5755*	0.1110	-9.4303*	1.7984
Promotion	0.4585*	0.0410	0.4534*	0.0333
Prom*Price	-0.1474*	0.0308	-0.1553*	0.0282
Quarter 2	-0.0032	0.0130	-0.0064	0.0144
Quarter 3	-0.4430*	0.0181	-0.4616*	0.0128
Quarter 4	-0.9061*	0.0198	-0.8966*	0.0155
Control Function	1.6334*	0.0484	1.5577*	0.0370
Random Parameter Means				
Quality	3.9696*	0.0612	1.9794*	0.2289
Price	-1.3700*	0.0369	-1.3519*	0.0307
Random Parameter Std. Devs.				
Quality			0.0211*	0.0046
Price			0.1984*	0.0019
Random Parameter Function				
Quality (Variety 2)			-0.0657*	0.0239
Quality (Variety 3)			0.1651*	0.0198
Quality (Variety 4)			-0.0746*	0.0204
Quality (Variety 5)			-0.2702*	0.0318
Quality (Variety 6)			-0.0771*	0.0237
Variance of Error Term			1.0190*	0.0011
LLF	-147,903.88		-138,267.44	
χ^2	114,292.24		133,565.12	

Note: A single asterisk indicates significance at a 5% level. Estimated with simulated maximum likelihood.

Table 5. Empirical Model of Price Discrimination: Non-Linear Case

	Model 1: Fixed		Model 2: Random	
Variable	Estimate	Std. Err.	Estimate	Std. Err.
Constant	-3.5306*	0.0264	-0.3106*	0.0343
Item 2	-0.1916*	0.0063	-0.5425*	0.0583
Item 3	0.0013	0.0087	-0.3619*	0.0586
Item 4	-0.1161*	0.0097	-0.4688*	0.0585
Item 5	-0.4707*	0.0117	-0.3639*	0.0496
Item 6	-0.3161*	0.0139	-0.2133*	0.0498
Item 7	-0.4847*	0.0151	-0.3803*	0.0499
Item 8	-0.7932*	0.0169	-0.6756*	0.0501
Item 10	-0.2328*	0.0044	-0.6622*	0.0906
Item 11	-0.1992*	0.0066	-0.6170*	0.0909
Item 12	-1.0028*	0.0125	-0.2346*	0.0835
Item 13	-1.4600*	0.0203	-0.5576*	0.0832
Item 14	-1.4847*	0.0222	-0.6025*	0.0830
Promotion	0.1044*	0.0093	0.1038*	0.0079
Prom*Price	-0.0311*	0.0070	-0.0327*	0.0069
Quarter 2	-0.0104*	0.0030	-0.0111*	0.0030
Quarter 3	-0.0964*	0.0041	-0.1013*	0.0027
Quarter 4	-0.1927*	0.0045	-0.1965*	0.0032
Control Function	0.3526*	0.0110	0.3304*	0.0077
Random Parameter Means				
Quality	0.0538*	0.0007	0.5832*	0.0636
Price	-0.3597*	0.0084	-0.3408*	0.0059
Random Parameter Std. Devs.				
Quality			0.0207*	0.0023
Price			0.0770*	0.0003
Random Parameter Function				
Qual (Variety 2)			0.0132	0.0197
Qual (Variety 3)			0.0513*	0.0199
Qual (Variety 4)			0.0119	0.0110
Qual (Variety 5)			0.0419	0.0483
Qual (Variety 6)			0.0495	0.0315
Variance of Error Term			0.2340*	0.0002
LLF	-3,851.23		-235.974	
χ^2	7,702.47		471.949	

Note: Quality values transformed from logs to represent positive-constraint. Asterisk indicates significance at a 5% level. Estimated with simulated maximum likelihood.

Table 6. Non-Parametric Kernel Density Estimates

	Linear Model		Non-Linear Model	
	Empirical	Log-Normal	Empirical	Log-Normal
Bandwidth	0.0354	0.0353	0.1853	0.1850
Mean	1.6011	1.6011	0.2379	0.2379
Standard	0.2089	0.2086	1.0947	1.0929
Skewness	0.2918	0.0000	1.5326	0.0000
Kurtosis-3	-1.3139	-0.0380	2.8910	-0.0380
χ^2	6.5485	0.0047	52.7325	0.0047
Minimum	1.2844	0.8705	0.0066	0.0005
Maximum	1.9373	2.3317	0.6137	1.0934
Points	1062		1062	
% Food Loss	10.0814		12.0732	

Note: Kernel densities estimated with Epanechnikov weighting function.
Critical Chi-square value at 5% is 3.84 for null of log-normality.

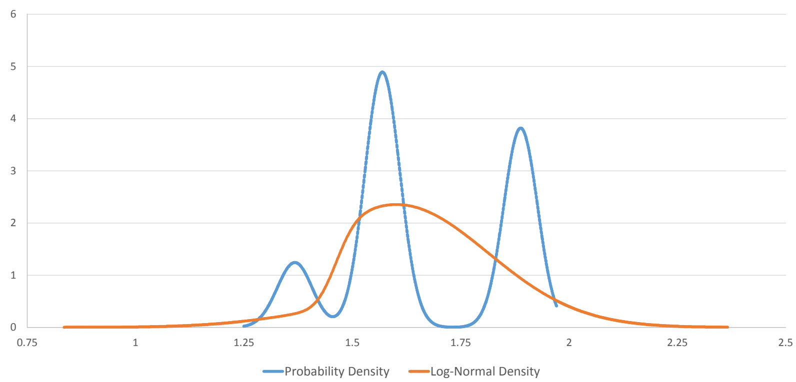

Figure 1: Linear Quality Density Function

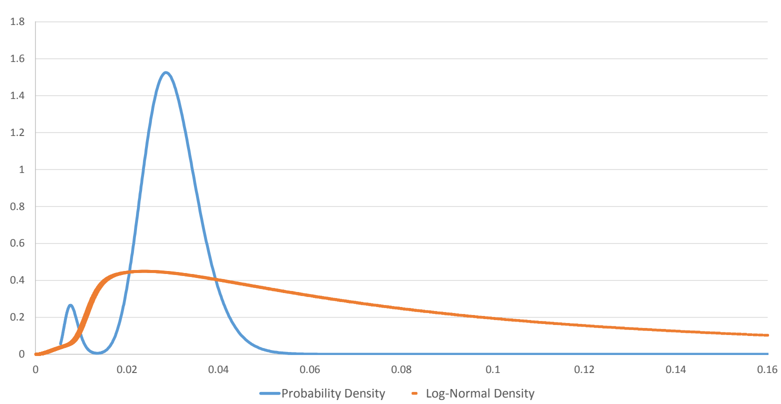

Figure 2: Non-Linear Quality Density Function