Poverty Alleviation through Social Labeling Programs? Information Valuation and Willingness to Pay for Fair Trade Coffee

Arnab K. Basu*
Robert L. Hicks†

This Version: April 2013

Abstract: What product attributes do consumers account for in their willingness to pay for labeled products? Based on the eligibility criteria of the Fair Trade program for coffee we develop a theoretical model and design a conjoint experiment to test whether consumers' willingness to pay a price premium varies with the levels of poverty eradication achieved by the "Fair Trade" label. In particular, we focus on two attributes of the Fair Trade label for coffee (participation of poor producers, and the associated income guarantee) to analyze how respondents value information regarding both the Fair Trade program and the label performance in their willingness to pay. Starting from either very low levels of participation by poor producers, or low income guarantees to them by the Fair Trade label, respondents' willingness to pay points to an interplay between their preferences for poverty eradication and worsening relative deprivation amongst poor producers that a Fair Trade label engenders.

JEL Classification: I1, O1, D1
Keywords: Poverty Aversion; Relative Deprivation; Fair Trade Label; Conjoint Analysis

*Dyson School of Applied Economics and Management, Cornell University, Ithaca, NY 14853; Center for Development Research (ZEF), University of Bonn, 53113 Germany and Institute for the Study of Labor (IZA), Bonn, Germany. E-mail: arnab.basu@cornell.edu

†Department of Economics, College of William and Mary, Williamsburg, VA 23187. USA. E-mail: rob .hicks.wm.edu

1 Introduction

The history of Social-labeling dates back to the pioneering "White Label" initiative in the United States in 1899. Sponsored and monitored by the National Consumers League, the "White Label" guaranteed manufacturers' compliance with the maintenance of decent working conditions and the absence of child labor in production. In recent years, relatively more recognizable social labeling initiatives include the RUGMARK label which promotes carpets made without child labor originating from India, Pakistan and Nepal, and the "Fair Trade" label, which provides income support to poor producers worldwide for a variety of products[1]. The popularity of products exhibiting socially responsible production and process methods, particularly "Fair Trade" labeled ones, has grown at a rapid pace in Western consumer markets over the past couple of decades. According to the Fair Trade Labelling Organizations International (FLO), global sales of "Fair Trade" labeled products that are ethically produced exceed $2.3 billion worldwide. Specifically, and in the context of this study, sales of Fair Trade Coffee has yielded average annual growth rates of 65% in Canada and 47% in the US since its introduction in North America. Interestingly, Fair Trade Coffee accounted for only 0.5% of the US coffee market in 2004 as compared to around 6% in Switzerland and 20% in the UK (Ronchi 2006; Richardson and Stahler 2007).

The popularity of Fair Trade products is primarily based on its market-driven approach to achieve social goals. These social goals encompass minimum wages, reasonable working hours and conditions, no child labor in production and a guaranteed price floor for poor participating producers in developing countries. In effect, the Fair Trade label aims to correct the information distortion created by a lack of product differentiation along "socially acceptable" standards by allowing consumers in the developed world to express their preference for these products through a higher willingness to pay. Nevertheless, empirical evidence on consumers' willingness to pay for labeled products in Western markets fails to provide conclusive evidence regarding the existence of a positive price-premium. A number of empirical studies have shown the existence of a positive, albeit small, premium for numerous products ranging from canned tuna to organic textiles (Teisl et al. (2002); Nimon and Beghin (1999); Bjorner et al. (2004)). However, in the context of Fair Trade labeled products specifically, recent studies present con-

[1]Concurrently with the social labeling initiatives there has been a widespread proliferation of Eco-labels. Starting with Germany's "Blue Angel" label launched in 1978 to designate "environmentally friendly" products, various eco-labels promoting environmentally friendly process and production methods (PPMs) like "dolphin-safe tuna", furniture made from wood harvested from sustainably managed forests, fish from sustainable fisheries and a variety of "organic" products has proliferated Western consumer markets recently.

flicting results. While a group of studies based on consumer surveys find evidence of consumers' willingness to pay a positive price premium, revealed preference results based on field studies paint a different picture[2].

Amongst the first group of studies involving consumer surveys, a NBER study done in 1999 through interviews found 80% of polled consumers in the United States willing to pay more for products made under good working conditions (Elliot and Freeman (2003)). A second survey undertaken by the Program on International Policy Attitudes (PIPA) at the University of Maryland in 1999 found that 76% of respondents were willing to pay an extra $5 for a $20 garment that was certified free of sweatshop labor (PIPA 2000). Becchetti and Rosati (2005) analyzed the behavior and motivation of one thousand consumers purchasing socially responsible Fair Trade labeled goods that were being sold at a price premium in European markets. The study found that Fair Trade products have less than unit income elasticity while demand for them is negatively correlated with geographical distance from the nearest Fair Trade shop, but positively correlated with the age of the consumer and awareness about the socially responsible criteria. The awareness of the socially responsible criteria depended in turn on conditioning factors such as consumption habits and membership of volunteer organizations.

Amongst the second group of studies involving field experiments, Milkman (2004) reports an experiment where t-shirts labeled Sweat-X, a prominent California Fair Trade Cooperative, was sold alongside a better known brand All-Star. Even though both labels were priced the same and had the same quality, more All-Star t-shirts were sold. The study concluded that consumers did not trust the Sweat-X label. Second, Kimeldorf et.al (2004) at the University of Michigan carried out a field experiment at a department store with plain white socks. One

[2]Some recent papers have also looked at the reasons behind the existence (or lack thereof) of a fair trade premium by focusing on the supply side of the issue. For instance, de Janvry, McIntosh and Sadoulet (2009) point out that unlike other certifications like "Dolphin-safe Tuna" the fair trade mechanism for coffee explicitly aims at delivering rents generated through ethical demand directly to producers. Using data from a producer cooperative in Guatemala they generate an inverse relationship between market price of *unlabeled* coffee and the Fair Trade coffee quality. The argument here is that with discrete differences in coffee quality, over-certification (increasing the number of producers) under the Fair Trade umbrella without improvements in product quality can dissipate the rent over time. The issue of Fair Trade product quality is also analyzed by Richardson and Stahler (2007) where increasing the size of the Fair Trade market creates a moral hazard problem. Increased membership into a cooperative without a corresponding increase in the cost of monitoring dilutes product quality as farmers "free ride". This in turn tends to reduce the price premium for Fair Trade products in a quality differentiated market. Reinstein and Song (2009) consider the linkages amongst three groups of agents — consumers, retailers and suppliers, and where a supplier's investment in production reduces retailers' costs. In this scenario, the price premium generated by consumers' altruism can translate into retailers paying more to the suppliers to secure increased investment when retailers and suppliers interact repeatedly.

batch of socks (treatment group) were labeled as made with good working conditions and no child labor while the other (control group) was unlabeled and priced at $1.00 a pair. By varying the price of the labeled socks from $1.00 - $1.40 in 5 cent increments, Kimeldorf et.al found that 30.3% of consumers bought the labeled socks when the price was higher for the treatment group while 43% of consumers bought the labeled pairs when the price was identical to the unlabeled pairs. Overall, only 26.5% of consumers actually paid more to buy labeled socks. Third, in a revealed preference study undertaken by Hiscox and Smyth (2005) in a major New York City retail store for towels, dolls and candles labeled as made under good working condition / no sweatshop labor found that sales rose for these labeled products, and that consumers exhibited inelastic demand (-0.5%) for price increases up to 20% above the baseline unlabeled product.

More specific to our paper, Koppel and Schulze (2008) conducted a field experiment with Fair Trade labeled coffee on the effectiveness of direct and indirect transfer mechanisms. Koppel and Schulze shows that consumers when faced with the choice between making an indirect transfer to alleviate poverty by paying a Fair Trade premium through coffee purchases and making a direct donation through a charitable organization for the poor coffee producers decided in favor of the indirect transfer mechanism. Hiscox et al. (2011) examine consumer willingness to pay for fair trade coffee using EBAY auctions. By posting otherwise identical products on EBAY, differences in winning auction prices were observed and from this consumers paid approximately 23% more for the fair trade coffee. In a similar study set in a grocery store chain, Hainmueller et al. (2011) find a price premium of nearly 10% and two distinct types of consumers- those who were price sensitive and unwilling to pay more for the fair trade label and those willing to pay more.

This apparent contradiction in the revealed preference studies based on field experiments regarding consumers' willingness to pay for labeled products can be attributed to a number of factors, notably, credibility of the label itself, the type of product in question, consumer heterogeneity (age, gender, education and income, as examples), and quality of the labeled product vis-à-vis its' unlabeled counterpart. However, one of the impediments to observing the existence of a higher price premium for labeled products might be attributed to label transparency, an issue that has received scant attention in the literature. Take the case of label performance as an example. First, in the United States, the "Dolphin-safe Tuna" is an ubiquitous label which guarantees that the Tuna was caught in ways not intentionally harmful to Dolphins. Yet the label does not specify the impact of this technology on current Dolphin stocks. Second,

while it is well known that the Fair Trade label for coffee guarantees a price floor to marginal growers in developing countries, the label does not specify either the fraction of poor producers served by the label or the income guarantee received by them. In both these examples it is only natural that consumers' willingness to pay is based on their perception about the extent of poverty alleviation achieved by the Fair Trade program rather than the actual performance of the program itself.

In two earlier empirical studies, Hicks (2006) and Basu and Hicks (2008) used stated preference methods to investigate the impact of advertised label performance on consumer willingness to pay for Fair Trade coffee. By varying the percentage of poor producers participating in the program these papers found that consumers are willing to pay a positive premium over unlabeled coffee only when participation reached a certain threshold. Our basic motivation in this paper follows the empirical observations generated by the above two papers, and lies in disentangling the impact of performance attributes associated with the Fair Trade program on consumers' willingness to pay a positive price premium. Specifically, this paper relates performance attributes to the existence of a price premium of a Fair Trade labeled product in the following ways:

- We develop a theoretical model that underscores how consumers' aversion or affinity to poverty and relative deprivation faced by producers, included as well as excluded, by a Fair Trade program translate into observed altruism towards socially responsible production objectives[3].

- The experimental design of this paper (i) incorporates both producer participation and revenue increases engendered by a Fair Trade program as performance attributes and (ii) analyzes how respondents' willingness to pay a Fair Trade premium varies with these performance attributes conditional on the information provided about the costs and benefits of the Fair Trade program. Noteworthy in the context of label transparency is the fact that aside from the stated objective that the Fair Trade Labeling Organization (FLO) targets only poor/marginal producers, and that once selected into the program poor pro-

[3]In this respect, our theoretical model can also be viewed as related to the recent work of Benebou and Tirole (2009) that develops the foundations of pro-social behavior by drawing on the recent developments in the psychology and economics field. Benebou and Tirole models pro-social behavior as a complex link between genuine altruism, social or self image concerns and material incentives. Our theoretical model and experimental design linking consumers' concern for poverty and relative deprivation for both participants within a Fair Trade program, as well as those outside, with the willingness to pay the Fair Trade premium can thus be viewed as disentangling further the components of genuine altruism.

ducers are guaranteed a price floor, the label does not provide updated information on the either the scale (share of poor producers in the program) or the scope (the size of income increase guaranteed to poor producers) on the label itself. Thus, while the existence of a positive price premium for Fair Trade products can be easily attributed to consumers' preference towards poverty alleviation, there is no way to predict how the willingness to pay itself might vary with the *degree* of poverty eradication achieved by the Fair Trade label.

- Our empirical strategy can also be viewed as one that identifies the value of added information on product labels. By comparing two groups of respondents — one with information about label performance and the other without — we are able to identify the gap between the actual willingness to pay (when label performance is known) and the potential willingness to pay (when label performance is unknown).

In order to understand how consumers account for the degree of poverty eradication in their willingness to pay, we proceed as follows: First, we incorporate the two basic tenets, the scale and scope of the program for Fair Trade labeled coffee, to build a theoretical model that exploit the following critical insight hitherto neglected in the literature on the extrapolation of revealed preferences for consumers in their choice to buy a Fair Trade product[4]. If indeed poverty alleviation is the main motive behind the purchase of a Fair Trade product (in this case coffee) then consumers should theoretically account for not only the share of poor producers whose income improves through their decision to purchase, but also the degree of income shortfall that is generated as a consequence of this purchase decision between those poor producers who are a part of the Fair Trade program and those who remain outside its purview within a developing country. Specifically, the Fair Trade label transforms a homogeneous product into a differentiated one by dividing producers within a country into three categories: (i) poor producers who are part of the Fair Trade program; (ii) poor producers who are currently excluded by the program but can potentially be a part of the Fair Trade label in the future; and (iii) relatively rich producers who do not meet the eligibility criteria. Therefore, unlabeled coffee is produced by both rich *and* poor producers, and while consumers may not be able to distinguish the share of poor producers within the unlabeled group, their decision to purchase unlabeled coffee may well be driven by the concern for deepening relative deprivation amongst poor producers rather than just their preference towards poverty alleviation for producers already participating in a Fair Trade program *per se*. In essence, we exploit the *sorting* effect the Fair Trade label engenders

[4]These two metrics are cited by FLO certifiers themselves, as evidence of the success of Fair Trade programs.

amongst poor producers — those included and those excluded by the program — to analyze to what extent aversion to poverty and relative deprivation affect consumers' willingness to pay for Fair Trade Coffee[5].

Our theoretical model predicts that consumers' willingness to pay for Fair Trade labeled coffee could be subject to two opposing effects as both the scale (participation) and scope (income) of the Fair Trade program increases. On one hand, aversion to poverty may lead to an increase in the willingness to pay as poor producers included by the Fair Trade program benefit while on the other hand, concern for relative deprivation of poor producers not yet included by the program may very well lead to an opposite effect. However, the central question as to over what range of the scale and scope of a Fair Trade program can one expect increasing consumer willingness to pay, should be in principle closely tied to the type of information consumers have about the *overall* costs and benefits associated with the Fair Trade program for poor producers both excluded and included by the program within a country. To account for this second layer of information asymmetry our choice experiments additionally separate both groups of respondents (choosing performance and traditionally labeled coffee respectively) into three distinct information regimes about the costs and benefits of the Fair Trade program as follows: (i) full-information — where respondents are given information about both the costs and benefits of the program; (ii) benefits only — where respondents are given information about just the benefits of the program but not the costs and (iii) no-information — where respondents do not receive any information on either the costs or benefits of the program. In effect, by varying the information about the costs and benefits of the Fair Trade program, we put our theory of poverty aversion and relative deprivation in consumers' willingness to pay for coffee to an empirical test by conducting experiments using stated preference discrete choice methods (Louviere et. al. (2000)) undertaken by University students in the United States.

Using the choice experiment data, we estimate a mixed logit model (Train (2003)) that allows for consumer heterogeneity over preferences for product attributes. Our results show that there is significant consumer heterogeneity with respect to label parameters such as certifying agency and program performance attributes. The results from the mixed logit model are

[5]Unlike green labeling programs where labeling itself encourages environmentally friendly production choices to be applied to the production of labeled products but not unlabeled products (Basu, Chau and Grote (2003, 2005, 2006)), Fair Trade addresses an existing issue (poverty among some but not other farmers). In other words, while the existing literature on social and eco-labeling treats unlabeled products as uniformly "bad" (produced with child labor, or with eco-unfriendly techniques), unlabeled products in the Fair Trade setting are produced by both poor and rich farmers. Accordingly, the informational role in terms of label performance is distinctly different for Fair Trade labeled products.

further used to simulate the distribution of parameters and investigate consumer willingness to pay for Fair Trade program attributes. While the mixed logit allows for tremendous richness in heterogenous preferences, our results show that the majority of consumers exhibit an inverted-U relationship with respect to their willingness to pay for increases in the scope and scale respectively for Fair Trade labeled coffee. On the other hand, consumers willingness to pay for Fair Trade labeled coffee without any performance criteria regarding the scale and scope exhibit a uniform mark-up over their unlabeled counterpart. We also provide some robustness checks for the label performance results. Thus, it bears emphasis that the primary objective of this paper lies in understanding how label transparency in the form of additional information about the scale and scope of a Fair Trade program affects consumers' willingness to pay a price premium *conditional* on the information related to the fulfillment of social objectives by the Fair Trade program itself.

Finally, it should be noted that our choice of coffee as the Fair Trade product is driven by the following considerations: (i) coffee has a relatively inelastic demand (lack of close substitutes), (ii) Fair Trade is a highly recognizable label in the United States particularly amongst University students who are the respondents in our choice experiments[6]. Additionally, our choice of University students as respondents allows for automatic controls on the age, education and income of consumers. Considerable heterogeneity amongst respondents along these variables could have potentially made it harder to identify consumer attitudes towards poverty and relative deprivation in our experiments.

The remainder of the paper is organized as follows. In section 2 we highlight the criteria for eligibility for producers to join the Fair Trade label for coffee and develop a theoretical model to pin down the interplay between poverty aversion, relative deprivation and willingness to pay. Section 3 explains our survey methodology and empirical results while section 4 concludes.

[6]Also, there is considerable awareness amongst University students that many developing countries are highly dependent on coffee exports. For instance, in 2008 Brazil exported 28,209,376 bags of coffee (each containing 60 kgs), Colombia 11,551,966 bags, Vietnam 15,478,622 bags, Ethiopia 2,871,191 bags and Kenya 607,022 bags of coffee. At the same time the composite price index of coffee in world markets have been US$1.24 per/lb in 2008 (International Coffee Organization (http://www.ico.org/trade_statistics.asp)).

2 Theoretical Model

In order to motivate our theory of how consumers might incorporate the scale and scope of Fair Trade program in their willingness to pay, we first summarize the certification criteria of Fair Trade organizations with regards to coffee. The main goal behind Fair Trade certification is to establish a price floor thereby guaranteeing small scale farmers a living income. For example, the Fair Trade organization "Equal Exchange" guarantees a price floor of $1.26/lb of washed Arabica to producers, plus an additional 15c/lb if the coffee is organically produced. Thus, the price guarantee to growers of organic coffee is $1.41/lb and $1.26/lb for non-organic coffee. If the world market price were to rise above $1.26, the program requires importers to pay the higher market price instead, plus an additional 5c/lb Fair Trade premium along with the 15c/lb organic premium. Again, as an example, if the world market price is $1.30/lb, organic coffee producers in Fair Trade cooperatives would receive $1.30 + 5c + 15c = $1.50/lb, while producers of non-organic coffee would receive $1.35/lb[7]. Second, a farmer has to be a small scale operator to be eligible for a Fair Trade program. Typically small scale farmers are defined as those who cultivate less than 3 hectares of land and harvest 1,000-3,000 pounds of unroasted coffee a year (usually with a cash income of $500-$1,000 a year) as well as principally rely on their own families' labor[8].

There are thus three potential issues affecting consumers' willingness to pay Fair Trade coffee. First, the guaranteed price floor directly targets absolute poverty of participating producers. Second, fairness concerns are addressed through eligibility criteria imposed on participating producers. However, this eligibility criteria has a sorting effect: it *sorts out* the group of relatively poor farmers from a pool of otherwise anonymous and indistinguishable producers. In addition, the eligibility criteria also impacts the *composition* of non-participating producers. With Fair Trade, these can be of two types: (i) non-eligible rich producers, and (ii) poor producers who cannot participate, due perhaps to insufficient (quantity) demand, and / or the transaction costs required to do so (for example, geographical distance to a cooperative). In fact, Milford (2004) notes that while Fair Trade cooperatives are technically open to new

[7]The guaranteed price floor acts as an insurance mechanism that shifts the risks of market volatility to the Fair Trade program. If the market price is low, growers receive, at least, the Fair Trade minimum price of $1.26/lb (or $1.41/lb for the organic variety). The insurance mechanism is especially valuable since more often than not, market prices are sometimes 25-50% below the Fair Trade minimums. For instance, during the 2000-2004 coffee crisis, market prices remained low reaching $0.42 in June 2002. Even when market prices are high, poor growers receive substantially less than the market price (between $.30-.50/lb) from middlemen.(http://www.globalexchange.org/campaigns/fairtrade/coffee/faq.html).

[8]This criteria alone makes Fair Trade potentially representative of an estimated 75% of all coffee farmers.

members, these co-operatives use membership regulations actively to obtain a specific number of producers and a corresponding volume of coffee production[9]. Third, Fair Trade products also encompass a bundle of attributes (organically grown, environmental sustainability, the provision of credit etc.) This introduces yet another layer for consideration in the willingness to pay for Fair Trade coffee.

Based on the preceding discussion we consider for our theoretical model three distinct stakeholders for a Fair Trade labeling program: (i) labeling organization, (ii) producers and (iii) consumers.

2.1 Labeling Organization

Throughout, we assume that a labeling organization determines the scale and scope of the Fair Trade program. The objective of the Fair Trade program is to include poor producers who meet the eligibility criteria with a price guarantee for their products. Thus, two sets of rules govern the execution of the labeling program where: (i) only poor producers are eligible and (ii) a pre-determined level of income increase is guaranteed for those producers participating in the program.

2.2 Producers

Prior to the Fair Trade program, let there be two groups of producers: N poor producers in the first group which includes some that are eligible to participate in the Fair Trade program, and N^* producers in the second group who represent high income producers who do not meet the eligibility criteria for joining the Fair Trade program.

The introduction of Fair Trade program further divides the group of poor producers into participating (l) and non-participating (u) producers. Together with the ineligible (u^*) rich producers, there are thus now three groups of producers, denoted with a subscript l, u, u^*, respectively. Let q denote the fraction of poor producers whose income is guaranteed by the Fair Trade program. Thus, Nq is the number of poor producers participating in the program while $N(1-q)+N^*$ is the number of producers who remain outside the program and sell unlabeled coffee from a given country-of-origin. Note that $N(1-q)$ is the number of poor producers excluded from the

[9]See Richardson and Stahler (2007) and de Janvry, McIntosh and Sadoulet (2009) for theoretical and empirical work that highlights the importance of membership size within the coffee cooperatives to keep monitoring costs tractable and maintain high quality standards.

Fair Trade program.

As population shares, we denote

$$s_l = \frac{Nq}{N+N^*} \equiv nq, \;\; s_u = \frac{N(1-q)}{N+N^*} \equiv n(1-q), \;\; s_{u*} = \frac{N^*}{N+N^*} \equiv 1-n.$$

Accordingly, we specify $W_i, i = l, u, u^*$ as the income levels of the three groups of producers, respectively, the participating poor, the non-participating poor, and the ineligible rich. Denoting $\delta > 0$ as the degree of output price increase guaranteed by the Fair Trade program, we shall assume throughout that the following income ranking holds:

$$W_{u*} > W_l(\delta) > W_u.$$

where $W_l(\cdot)$ is assumed to be strictly increasing in δ.

Poverty and Relative Deprivation

By virtue of the price guarantee, fair trade labeling impacts both the absolute degree of poverty and the relative income deprivation facing producers in the three groups. In particular, let z denote the poverty line, and P_i be the Sen's (1970) measure of absolute poverty for individuals in group i:

$$P_i(q, \delta) = \max\left(\frac{z - W_i}{z}, 0\right).$$

We shall assume that only rich producers live at an income level above the poverty line ($W_{u*} > z$), and thus $P_{u*} = 0$. But the measure of absolute poverty for the participating and non-participating poor are given as,

$$\begin{aligned} P_l &= \frac{z - W(\delta)}{z} \\ P_u &= \frac{z - W_u}{z} \end{aligned}$$

To gauge the extent to which fair trade labeling promotes fairness in income distribution across the three groups of producers, we employ the following well-known measure of relative deprivation (Deaton (2001) and Yitzhaki (1979)), R_i, where

$$R_i = \sum_{j \neq i} s_j \max\left(\frac{W_j - W_i}{W_i}, 0\right).$$

Thus, R_i measures income shortfall of group i producers relative to others. The weights s_j are simply the population share of producers in group j. Given the income ranking $W_{u*} > W_l(\delta) >$

W_u, we thus have:

$$\begin{aligned} R_{u*} &= 0 \\ R_l &= s_{u*}\frac{W_{u*}-W(\delta)}{W(\delta)} = (1-n)\frac{W_{u*}-W(\delta)}{W(\delta)} \\ R_u &= s_l\frac{W(\delta)-W_u}{W_u} + s_{u*}\frac{W_{u*}-W_u}{W_u} = nq\frac{W(\delta)-W_u}{W_u} + (1-n)\frac{W_{u*}-W_u}{W_u} \end{aligned}$$

The absolute poverty and relative deprivation impact of a Fair Trade program can now be seen clearly. First, as the scope (income guarantee) of the program improves, or as δ rises, absolute poverty and relative deprivation facing participating poor producers improve as both P_l and R_l are decreasing in δ. An increase in the scale (percentage of participating poor producers) of the program, or as q increases, naturally implies that more poor producers can participate, and does not by itself impact absolute poverty or relative deprivation of existing participating producers.

It is also of interest to examine the role of the scale and scope of the Fair Trade program on non-participating poor producers. In particular, whereas both P_{u*} and R_{u*} are equal to zero and thus independent of the program, an increase in either q or δ implies a worsening of the relative income status of non-participating poor producers. Specifically, as the scale (q) increases, the group of non-participating producers becomes even more marginalized as R_u increases. Similarly, as the scope (δ) rises, the income gap between participating and non-participating poor producers rises through an increase in R_u.

2.3 Consumer Preferences

The willingness to pay for Fair Trade relative to unlabeled coffee is potentially affected by: (i) product attributes — the bundle of added attributes associated with Fair Trade products, including both its social and environmental impacts; (ii) sorting — the extent to which a Fair Trade program sorts producers into (a) participating poor farmers and (b) a mix of the rich *and* the non-participating poor farmers and (iii) attitude towards poverty and relative deprivation — the degree to which the scale and scope of a Fair Trade program changes the absolute poverty and relative deprivation faced by the participating and non-participating producers.

Our model of consumer preferences thus accounts for consumers' concern for (i) poverty — whether consumers prefer, are indifferent to, or are averse to buying products made by poverty stricken producers, (ii) inequality — whether consumers prefer, are indifferent to, or are averse to buying products made by producers who are poor relative to other producers producing an otherwise similar product, (iii) scale of the program — measured by the share of poor producers

participating in the program, (iv) scope of the program — the size of the income increase guaranteed to participating producers and (v) other product attributes that a Fair Trade labeled product with performance metrics typically addresses. Our objective in the ensuing theoretical formulation and the experimental design is to trace out any systematic relationship between the price premium and the above concerns.

Taking the above as inputs, we are now in a position to write down consumer preferences. For sake of clarity, and in consonance with our experimental design, we consider four specifications of the indirect utility function for consumers contingent on the label characteristics: (i) the indirect utility of consumers associated with the purchase of Fair Trade labeled coffee that advertises the *performance* of the label (i.e., producer participation and revenue guarantee) and (ii) the indirect utility of consumers who choose unlabeled coffee when presented with the alternative of fair trade labeled coffee with the performance metrics. Needless to say, by setting the performance metrics to zero we can easily identify (iii) the indirect utility of consumers associated with the purchase of traditional Fair Trade labeled coffee (without any performance metric on scale or scope attached), and (iv) the indirect utility of consumers who choose unlabeled coffee when faced with the alternative of a traditional Fair Trade labeled one.

Let consumers' preference take the following form[10]:

$$U^p(V, \theta_i, P_i, R_i) = V + \theta_i + \theta_p P_i + \theta_r R_i + m - p^p$$

$U^p(V, \theta_i, P_i, R_i)$ represents the indirect utility of a consumer who purchases a product made by producers of type i. V denotes consumers inherent valuation of the product, in our case, a unit of coffee while θ_j, $j = i, p, r$ are scalars that may be of either positive or negative values, and represents respectively the added environmental / social attributes associated with the product (θ_i), attitude towards poverty (θ_p) and relative deprivation (θ_r). Finally, m is income and p^p is the price of the product.

[10]The indirect utility function is derived from the following utility maximization problem of a representative consumer. Preferences are quasi-linear and defined over three goods, x_1, x_2, x_3 where x_1 is coffee — one unit of which is consumed with attribute θ. Thus, a consumer maximizes $u(x_1, x_2, x_3) = \theta + \hat{u}(x_2, x_3)$ subject to the budget constraint $m = p_1 + p_2 x_2 + p_3 x_3$. The associated indirect utility function is therefore

$$\begin{aligned} v(p_1, p_2, p_3) &= \theta + \hat{v}(p_2, p_3, m) \\ &= \theta + \frac{(m - p_1)}{\phi(p_2, p_3)} \end{aligned}$$

given that $\hat{v}$ is homothetic, we assume that the price index $\phi(p_2, p_3) = 1$ which allows us to substitute in equation (1), $\theta = V + \theta_l + \theta_p P_i + \theta_r R_i$ and $p_1 = p^p$.

Note that the parameters θ_p and θ_r represent respectively the nature of consumer preferences in regards to the poverty and relative deprivation facing the producers. For instance, $\theta_p < 0$ pertains to a poverty averse consumer whose indirect utility decreases while $\theta_p > 0$ pertains to a poverty loving consumer whose indirect utility increases from purchasing a product made by producers with a higher associated poverty index. Similarly, $\theta_r < 0$ pertains to a consumer who is averse to relative deprivation implying that their indirect utility falls while $\theta_r > 0$ pertains to a consumer with affinity towards relative deprivation implying that their indirect utility increases from purchasing a product from a producer whose income is relatively higher as compared to others who produce the same product.

Thus for a consumer, the indirect utility from unit consumption of the fair trade product when performance of the label is advertised is thus given simply by:

$$U_l^p(\delta, p) = U^p(V, \theta_l, P_l, R_l) = V + \theta_l + \theta_p P_l + \theta_r R_l + m - p_l^p \tag{1}$$

Meanwhile, the indirect utility of consuming a unit of unlabeled product (when faced with the choice of a Fair Trade label that advertises performance), produced by a mix of non-participating producers (u and u^*), is

$$\begin{aligned} U_{nl}^p(\delta, p) &= \frac{s_u}{s_u + s_{u*}} U^p(V, 0, P_u, R_u) + \frac{s_u}{s_u + s_{u*}} U^p(V, 0, P_{u*}, R_{u*}) \\ &= V + \frac{n(1-q)}{1-nq}\theta_p P_u + \frac{n(1-q)}{1-nq}\theta_r R_u + m - p_{nl}^p \end{aligned} \tag{2}$$

Noting that $\theta_u = \theta_{u*} = 0$, consumers would be willing to choose the Fair Trade product as long as $U_l^p(\delta, p) \geq U_{nl}^p(\delta, p)$. Subsequently, the price premium that consumers are at most willing to pay for Fair Trade labeled products when performance is advertised is thus

$$p_l^p - p_{nl}^p = \theta_l + \theta_p \left(P_l - \frac{n(1-q)}{1-nq} P_u \right) + \theta_r \left(R_l - \frac{n(1-q)}{1-nq} R_u \right) \tag{3}$$

It is trivial to note from equation (3) that if consumers are indifferent to poverty and inequality ($\theta_p = \theta_r = 0$), the price premium, $p_l^p - p_{nl}^p = \theta_l$, is a constant and independent of the scope (δ) and scale (q) of the program. As long as the social and environmental attributes of the program is deemed valuable to the consumer ($\theta_l > 0$), there will be a positive price premium.

Consider now the case of a traditional label where performance metrics are not advertised. In this case the indirect utility of a consumer who chooses to purchase either a traditionally labeled product or an unlabeled one is given respectively as

$$U_l^t(V, \theta_i, p_l^t) = V + \theta_l + m - p_l^t$$

$$U^t_{nl}(V, p^t_u) \quad = \quad V + m - p^t_{nl}$$

and the price premium equals $p^t_l - p^t_{nl} = \theta_l$. Here again, a positive price premium would exist for the traditional Fair Trade labeled product as long as consumers value the environmental/social attribute of the label ($\theta_l > 0$). Thus, the above scenario where consumers are indifferent to poverty and relative deprivation when label performance is advertised is observationally equivalent to the case of a traditional label that does not advertise performance in terms of either producer participation and revenue.

Of particular interest to us in this paper, is the way in which the price premium changes as the scale (q) and the scope (δ) of the program varies.

2.4 Change in the Scale of a Fair Trade Program

We first consider the impact of a change in scale (percentage of participating poor producers) on the price premium for performance labeled coffee. Differentiating equation (3) with respect to q we have,

$$\frac{\partial(p^p_l - p^p_{nl})}{\partial q} = \frac{n(1-n)}{(1-nq)^2}\left[\theta_p P_u + \theta_r R_u\right] - \theta_r \left[\frac{n(1-q)}{(1-nq)}\frac{\partial R_u}{\partial q}\right] > (<)\ 0 \tag{4}$$

Given that absolute poverty of the participating *and* the non-participating poor as well as the relative deprivation of the participating poor are independent of a change in the percentage of participating producers, the effect on the price premium is driven by (i) consumers' attitude towards poverty and relative deprivation and (ii) the impact of a change in the scale on the relative deprivation of non-participating producers (rich and poor). Consider therefore the following cases.

Case I: Consumers are averse to poverty and relative deprivation ($\theta_p < 0$ and $\theta_r < 0$).
In this case the fair trade premium exhibits an "inverted U" shape with respect to an increase in the scale. This result is generated by the opposing *composition* and the *individual level* impact on relative deprivation triggered by a change in the scale of the program. To begin with, consider first the composition impact. An increase in q increases the share of participating poor producers and decreases the share of poor non-participating producers. All else equal, this implies that the average non-participating producer (which includes rich and poor producers) is no longer as relatively poor as before. For poverty averse consumers this implies an increase in utility from purchasing the unlabeled product. Further, an increase in q changes

the composition of non-participating producers as the share of rich producers amongst all non-participating producers increases. Accordingly, the average degree of relative deprivation of non-participating producers (rich and the non-participating poor combined) decreases as well. For consumers averse to poverty and relative deprivation this combined effect increases their incentive to purchase the unlabeled product, and is captured by the first term in equation (4) above. On the other hand, an increase in q marginalizes those poor producers who do not participate in the Fair Trade program even further as captured by an increase in R_u for the individual non-participating poor producer. This individual level impact induces consumers averse to relative deprivation to shift their preference towards the labeled Fair Trade product. This effect is captured by the second term in equation (4) above. While these two effects occur simultaneously, their relative impacts depend on the existing level of q. Starting from a situation where q is very low, the second effect on relative deprivation (the individual level effect which is relative deprivation deepening) dominates since the share of producers $1-q$ who suffer a worsening in relative deprivation is still high. As q approaches 1, however, the first effect on poverty and relative deprivation (the composition effect which is relative poverty and deprivation improving) dominates since the average non-participating producer is richer. Taken together, for a consumer that is averse to poverty and relative deprivation, the willingness to pay a Fair Trade price premium first rises and then falls as the scale of the program increases.

Case II: Consumers are poverty loving and show affinity towards relative deprivation ($\theta_p > 0$ and $\theta_r > 0$).

This is the exact opposite of Case I above, and the interplay between the composition and the individual effect would lead the price premium to exhibit a regular "U" shape as the scale of the program increases.

Case III: Consumers are poverty loving but averse to relative deprivation ($\theta_p > 0$ and $\theta_r < 0$). Since consumers are poverty loving, an increase in the scale of the program implies that the share of participating poor producers increase. Thus the average non-participating producer (amongst the rich and poor producers combined) is no longer as relatively poor as before. For poverty loving consumers this implies an increase in utility from purchasing the labeled product. However, since consumers are averse to relative deprivation the composition and individual effects starting from a small q suggests that the willingness to pay for the Fair Trade labeled product first rises and then falls (as in Case I above) as the scale of the program increases. The combination of poverty loving but aversion to relative deprivation in consumer preferences thus exhibit an inverted "U" relationship between the price premium and the scale of the Fair Trade program.

Case IV: Consumers are poverty averse but show affinity towards relative deprivation ($\theta_p < 0$ and $\theta_r > 0$).
This is the opposite of Case III above, and the price premium would exhibit a regular "U" pattern with increases in the scale of a Fair Trade program.
Case V: Consumers are indifferent to poverty but averse to relative deprivation ($\theta_p = 0$ and $\theta_r < 0$).
In this case the Fair Trade premium would also exhibit an "inverted U" shape based on the relative strengths of the composition and individual level impact of a change in q on non-participating producers. For small values of q the fair trade premium would increase as the individual level impact dominates while for larger values of q the composition impact would tend to reduce the Fair Trade price premium.
Case VI: Consumers are indifferent to poverty but show affinity towards relative deprivation ($\theta_p = 0$ and $\theta_r > 0$).
This is just the converse of Case V above, and the Fair Trade price premium would exhibit a regular "U" shape.
Case VII: Consumers are poverty averse but indifferent to relative deprivation ($\theta_p < 0$ and $\theta_r = 0$).
In this case, only the composition effect is in play for the non-participating producers and as a result, the Fair Trade price premium would be strictly decreasing in q.
Case VIII: Consumers are poverty loving but indifferent to relative deprivation ($\theta_p > 0$ and $\theta_r = 0$).
This is the opposite of Case VII above, and hence the price premium for the Fair Trade product would be strictly increasing with an increase in the scale of the program.

In sum, the above eight cases point to the fact that aversion to relative deprivation is *sufficient* for the Fair Trade price premium to exhibit an inverted "U" shape as the scale of the program increases. On the other hand, if relative deprivation is a non-issue then aversion to poverty is sufficient to show a strictly increasing price premium as the scale of the Fair Trade program increases.

2.5 Change in the Scope of a Fair Trade Program

Turning to the impact of the scope of a fair trade program on the price premium, we have by differentiating equation (3) with respect to δ,

$$\frac{\partial(p_l^p - p_{nl}^p)}{\partial\delta} = \theta_p \frac{\partial P_l}{\partial\delta} + \theta_r \left[\frac{\partial R_l}{\partial\delta} - \frac{n(1-q)}{(1-nq)}\frac{\partial R_u}{\partial\delta}\right] > (<)\ 0 \tag{5}$$

Note that in equation (5) above the relationship between the Fair Trade price premium and the scope of the program depends on the interplay of the following: (i) consumers' attitude towards poverty and relative deprivation; (ii) the impact of an increase in the scope of the program on the absolute poverty of the participating poor producers, (iii) the impact of the scope on the relative deprivation of the participating as well as the non-participating poor farmers. First, absolute poverty and relative deprivation of the participating poor producers is decreasing in δ (or $\frac{\partial P_l}{\partial\delta} < 0$ and $\frac{\partial R_l}{\partial\delta} < 0$) while relative deprivation for the non-participating poor producers is increasing in δ ($\frac{\partial R_u}{\partial\delta} > 0$)). We capture these three effects under the following cases.

Case I: Consumers are averse to poverty and relative deprivation ($\theta_p < 0$ and $\theta_r < 0$).
In this case the Fair Trade price premium is strictly increasing as the scope of the program increases. The intuition is as follows: an increase in the scope of the program reduces absolute poverty of the participating producers as well as their relative deprivation vis-à-vis the rich producers. If consumers are averse to poverty and relative deprivation (i.e., their indirectly utility decreases with increases in the poverty and relative deprivation index) then both these effects tilt their choice towards the fair trade labeled product. Further, since relative deprivation of the non-participating poor, vis-à-vis the rich and the participating poor farmers, increases with the scope of the fair trade program, consumers averse to relative deprivation shy away from purchasing the unlabeled product in favor of the labeled one. Thus, the impact of an increase in δ on the absolute poverty and relative deprivation of the the participating poor, and on the relative deprivation of the non-participating poor reinforce each other to increase the Fair Trade price premium.
Case II: Consumers are poverty loving and show affinity towards relative deprivation ($\theta_p > 0$ and $\theta_r > 0$).
This is the opposite of Case I above and the Fair Trade price premium would be strictly decreasing as the scope of the Fair Trade program increases.
Case III: Consumers are poverty loving but averse to relative deprivation ($\theta_p > 0$ and $\theta_r < 0$).
If consumers are poverty loving but averse to relative deprivation the Fair Trade price premium *may* exhibit an "inverted U" shape when the scope of the program increases depending

on (i) the magnitudes of θ_p and θ_r and (ii) the size of the income guarantee increase. As the scope of the program rises and the poverty index for participating producers fall, consumers would prefer to shy away from purchasing the Fair Trade product in favor of the unlabeled one. On the other hand, aversion to relative deprivation would have the opposite effect and induce consumers to purchase the Fair Trade product. Specifically, for small increases in income guarantees the impact on the poverty gap between the rich and the participating poor producers would be negligible. Under such circumstances, the Fair Trade price premium might show an increase as the positive effect on utility engendered through aversion to relative deprivation overweighs the negative impact on utility for poverty loving consumers. However, as the size of the income guarantee increases and the poverty gap between the rich and the participating poor producers becomes smaller the negative impact on utility through the weight given to the poverty index can overweigh the positive impact on utility due to aversion to relative deprivation thereby shifting consumers' choice towards the unlabeled product. These two effects generate an "inverted U" shape relationship between the scope and the Fair Trade price premium depending on the relative magnitude of the weights given to poverty and relative deprivation by a consumer.

Case IV: Consumers are poverty averse but show affinity towards relative deprivation ($\theta_p < 0$ and $\theta_r > 0$).

This is the opposite of Case III above and again depending on (i) the magnitudes of θ_p and θ_r and (ii) the size of the income guarantee increase, the Fair Trade price premium *may* exhibit a regular "U" shape as the scope of the program increases.

Case V: Consumers are indifferent to poverty but averse to relative deprivation ($\theta_p = 0$ and $\theta_r < 0$).

In this case the Fair Trade price premium is strictly increasing with the scope of the program since relative deprivation of the non-participating poor, vis-à-vis the rich and the participating poor farmers, increases with the scope of the Fair Trade program which induces consumers averse to relative deprivation to favor the labeled product.

Case VI: Consumers are indifferent to poverty but exhibit affinity for relative deprivation ($\theta_p = 0$ and $\theta_r > 0$).

In this case the Fair Trade price premium is strictly decreasing with the scope of the program since an increase in relative deprivation of the non-participating poor producers induces consumers to purchase the unlabeled product (the opposite of Case V above).

Case VII: Consumers are averse to poverty but indifferent to relative deprivation ($\theta_p < 0$ and $\theta_r = 0$).

In this case the Fair Trade price premium is strictly increasing with the scope of the program as consumers opt for the purchase of the Fair Trade labeled product as absolute poverty of the participating poor producers fall with increased income guarantee.
Case VIII: Consumers are poverty loving but indifferent to relative deprivation ($\theta_p > 0$ and $\theta_r = 0$).
In this case the Fair Trade price premium is strictly decreasing with the scope of the program as a decrease in absolute poverty of the participating poor producers leads consumers to shy away from purchasing the labeled product (the opposite of Case VII above).

Unlike the case with an increase in the scale of a Fair Trade program, consumers must exhibit affinity for poverty and aversion to relative deprivation for an inverted "U" relationship between the price premium and the scope of a Fair Trade program to be observed. However, for the price premium to be strictly increasing in the scope of the program consumers must exhibit *either* aversion to poverty *or* aversion to relative deprivation *or* both. These results are interesting, and reiterate the fact that the different components of a Fair Trade label in fact contribute to the price premium in very different ways. To investigate the testable implications of our model we develop an empirical approach for measuring how attitudes towards relative deprivation and poverty impact consumer willingness to pay for labeled coffee in the next section.

3 Empirical Approach

As we have seen from the theory in the preceding section, consumers may consider a number of attributes found on products widely available today including product labels conveying information on production practices and methods related to environmental or social issues. In moving beyond what is evaluated and chosen in today's product marketplace is the issue of what consumers value exactly when purchasing a labeled product. In particular, the valuation of a labeled product may be a function of what the consumer believes about how the labeled product impacts the underlying social or environmental issue. Yet today's products say nothing about label performance and merely leave it to the consumer to base decisions on expectations regarding a label's performance.

We therefore move beyond observable actual consumer choices and investigate instead the stated preferences over products having label performance attributes that do not exist in the product marketplace today. Following from the theory section, we design a choice experiment to examine the sensitivity of how varying degrees of information impact the valuation of Fair

Trade attributes by employing stated preference techniques for hypothetical Fair Trade coffees (for an introduction, see Louviere et al, (2000)). This methodology has been used for a wide variety of consumer products and has been shown to yield reliable information about market choices of respondents (e.g. see Ben-Akiva and Morikawa (1990) and Adamowicz et al. (1994)).

In our choice experiment we investigate consumer preferences for Fair Trade coffee in six separate information treatments. Table 1 broadly outlines the empirical strategy we employ in the paper. We vary the label information given to respondents in two ways. First, some respondents are offered choices having traditional Fair Trade coffees - similar to what is available in today's marketplace - while others are offered Fair Trade labeled coffees having label performance attributes. The details of these differences can be found in Table 2. Given a label type, we vary the description of Fair Trade in three ways: termed 'Full Information', 'Benefits Only', and 'No Information'. Table 3 shows the exact wording of the information differences found in the survey instruments. Our intent with altering information about Fair Trade is to examine if the response to label attributes may differ for respondents having different baseline information about the general features of Fair Trade programs. Taken together, we have six choice experiments explored in the empirics.

As outlined in Table 2 for traditional coffees, we present the usual coffee attributes such as coffee price, the country of origin, whether the coffee is organic or not, a Fair Trade insignia, and the certifying agency. Since this information mimics what is found on current Fair Trade labels, we term this the traditional label. Respondents then choose among two traditional Fair Trade coffees and a non-labeled coffee alternative (see Figure 1 for a sample question). In the second experiment, we offer respondents an opportunity to choose among Fair Trade labeled coffees with performance information and a non-labeled coffee (see Figure 2). In addition to the attributes presented on the traditional label, the performance label includes two criteria related to the performance of the Fair Trade labeling program: the increased revenue accruing to program participants, and the growth rate of program participation (both in percentage terms). Both of these attributes describe the degree to which the labeled coffee is lifting producers out of poverty and provide information far beyond the binary indicator of the traditional Fair Trade label.

Blocked experimental design techniques were used to select the fifteen sets of 5 questions *for*

each experiment that maximize the tradeoffs of the coffee choice experiments[11]. Although the levels and attributes of the two experiments are identical (except for performance information), the actual levels of the attributes chosen by the experimental design algorithm differ by question, block, and experiment. Details on the design and attribute levels are presented in Table 2. For each of the two experiments, respondents are randomly assigned to one of the fifteen blocks. Respondents consisted of students taking large introductory freshmen classes at the College of William and Mary during academic year 2008 which allows us to control for heterogeneity with respect to age, income and education levels.

The indirect utility functions for each individual n and coffee i, V_{ni}, are defined as follows:

$$V_{ni}(\beta) = \begin{cases} X_L\beta_L^{'} + X\beta^{'} + +\beta_{org} + \epsilon_{ni} & \text{if } i = A \text{ or } i = B \\ X\beta^{'} + \beta_{org}Org + \epsilon_{ni} & \text{if } i = C \\ 0 + \epsilon_{ni} & \text{if } i = \text{No Coffee} \end{cases}$$

The researcher is assumed to be able to observe the indirect utility function up to an alternative-specific error term, ϵ_{ni}. Also, since the fair trade coffee definition ensures organic, an organic effect (denoted by β_{org}) is estimated relative to the non-fair trade coffee, which is not always organic. In addition to the certifying agency found on the traditional fair trade label, the performance label includes additional items in the vector X_L: performance attributes on increased revenues and participation growth of the program (as well as quadratic terms on these performance attributes).

The theory of the preceding section outlines a variety of consumer preference types related to Fair Trade coffee. The experiments are designed to present different combinations of coffee choices associated with more product variation than could be observed in the market. To model consumer choice, we estimate the mixed logit model described by Train (2003) that allows for consumer preference heterogeneity across label attributes. The model also relaxes the IIA restrictions common in standard multinomial logit applications. In our application of the model, we assume that random parameters are distributed with a normal distribution. The probability of individual n choosing coffee i can be written as

$$P_{ni} = \int \left(\frac{e^{V_{ni}(\beta)}}{\sum_{j=1}^{J} e^{V_{nj}(\beta)}} \right) \phi(\beta \mid b, W) d\beta$$

where $\phi(\beta \mid b, W)$ is the normal density with mean b and covariance W. Notice that the J coffee choices over which the model is defined correspond to the four coffee options available to

[11] In practice, this means maximizing D efficiency or the determinant $|X'X^{-1}|$.

respondents, Coffee A or B (both labeled), Coffee C (not labeled), and the so-called "opt-out" choice if the respondent chooses not to buy any coffee. In addition, we estimate a model that estimates fixed effects for each performance level. This approach has the advantage that it does not impose a particular functional form on preferences but comes with a loss in degrees of freedom.

4 Results

The goal of our analysis is to empirically explore what underpins consumer preferences for labeled products and examine if consumer willingness to pay differs as label performance differs. We approached the data using two different functional forms for capturing consumer preferences for performance label attributes. In the first approach, we make the stronger assumption that preferences for the performance label attributes are represented by a quadratic functional form. However, by allowing for random parameters in performance attributes we are able to estimate a variety of preference shapes that are not as strait-jacketed as normal quadratic specifications. While a strong assumption, this specification facilitates the identification of curvature in consumer preferences over label attributes. In the second approach, we estimate a levels model. For each performance level, we estimate a dummy variable which captures the contribution of performance to the individual's utility.

4.1 Preferences Quadratic in Performance

Using the experimental design presented above, we were able to estimate a mixed logit model of product choice for each coffee experiment. The mixed logit specification allows for consumer heterogeneity, and we feel this is especially important when dealing with Fair Trade label attributes. For coffee attributes familiar to respondents (i.e., price and country of origin effects), we estimate fixed coefficients (which is akin to standard multinomial logit estimation). The results for each of our six experiments can be found in Tables 4, 5, and 6. In addition this table also contains information about the number of respondents and average number of questions for each respondent. To organize the discussion across the treatments, we first focus on the commonalities. Both sets of findings are similar for the fixed coefficients. Price, negative and significant, indicates that as the price of a coffee increases the probability of choosing the coffee decreases, *ceterus paribus*. The country of origin effects do not show marked differences amongst each other, however their strong positive and significant effects are largely driven by respondents choosing the coffee alternatives over the opt-out alternative which is normalized

to zero.

The random coefficients in the model associated with fair trade attributes are differentiated in the table by the mean and standard deviation estimates (denoted by μ and σ) for each attribute. Higher levels of σ imply increased consumer heterogeneity since preferences for each of these attribute are estimated as $N(\mu, \sigma)$ using simulation methods. Consider the effect of certifying agency for the Fair Trade coffees. In all of the experiments, the Fair Trade Labeling Organization (FLO) had the strongest effect relative to the USDA and local growers associations. The mean effect of certification was larger for the traditional label, indicating consumers attached more significance to the certifying agency. Recall that for traditional Fair Trade coffees, the certifying agency is the only attribute that differentiates fair trade coffee. The performance label has performance attributes in addition to the certifying agency, and the effect of certification is much smaller with much lower consumer heterogeneity. Because the performance attributes effectively adds credibility to any certifying agency, consumers appear to place less weight on it *ceterus paribus*. Only one of the certifying agency results differed markedly across the information treatments offered to respondents. Under the "Full Information" scenario, we see that consumers put no significant weight on any of the Fair Trade certifying agencies (for the performance label) and only prefer the FLO label under the traditional label experiment. Adding information about the positives and negatives of a Fair Trade labeling program seems to lessen the impact of the label, but our results indicate that this effect may be ameliorated via the performance of the program if conveyed on the label itself. Furthermore, when looking across all of our experiments, our results indicate that adding performance information decreases the importance of certifying agency.

We next turn to the performance attributes, denoted by Revenue and Participation in Tables 4, 5, and 6. Notice for each we have a linear term and a second order term. These attributes are instructive as to how the performance of labeling programs, and the associated change in poverty, impacts consumer willingness to pay for Fair Trade labeled products. First, we see that as both revenues and participation increases, the probability of choosing a labeled coffee increases. The quadratic specification allows for this effect to be non-linear[12] The estimated quadratic terms show that at the mean, there is a peak level beyond which increases in performance lowers the probability of choosing the labeled coffee. The estimates of σ for each of the first order terms shows that there is significant heterogeneity associated with both

[12]Further, the random parameters allow for significant heterogeneity in the curvature estimated with the quadratic terms.

revenue and participation increases. The second order terms do not exhibit significant degrees of heterogeneity.

We use the estimated coffee-specific indirect utility functions to calculate the price premia consumers are willing to pay over and above \$1 for a Fair Trade Columbian grown coffee, relative to a \$1 non-labeled Columbian coffee that is not organic. We perform this calculation for the traditional label (yielding a mean price premium of \$0.17[13], \$1.27, and \$0.89 for the Full Information, Benefits Only, and No Information treatments respectively). We also calculate WTP under the performance label and vary the levels of performance. In Figures 3, 4, and 5, we see similar patterns of how the willingness to pay for the performance label changes with respect to program performance. In these figures, we hold participation increases at 10% and examine how the WTP varies for different levels of grower revenue increases. The peak premium relative to the traditional coffee varies across the information treatments, occurring at 60% for the 'Full Information' label and as high as 71% for the 'Benefits Only' label. Beyond these critical values the WTP declines[14]. We summarize some of the critical values underlying these figures in Table 7. In this table, we calculate the performance level providing the maximum positive difference between the performance and traditional labels (WTP_{max}). We also provide the critical values WTP_0 and WTP_1. If the label performance attribute falls below WTP_0 or is above WTP_1 the willingness to pay for the traditional label exceeds the performance label. Note that these "inverted-U" shaped curves are calculated at the mean parameter estimates presented in the results tables.

Figures 3, 4, and 5 were calculated using the mean parameter estimates and do not account for the significant consumer heterogeneity found in our mixed logit specification. We therefore conduct a simulation study using a technique proposed by Krinsky and Robb (1986, 1991) to examine the distribution of preferences found in our experiment. As a result, we draw 500,000 parameter vectors from the estimated distribution of the model parameters and calculate the shape of the indirect utility function with respect to the performance parameters[15]. Given the quadratic specification of the model, there are four distinct classes of consumer heterogeneity (see Figure 6)[16]. The vertical intercept for each consumer class is a simple transformation of

[13]The FLO label parameter is not significant, so caution is warranted with this estimate.

[14]Similar figures can be generated for the participation performance attribute but are excluded for brevity.

[15]Note that this requires doing a primary draw of model parameters including the mean draws for the random parameters (μ and σ), followed by secondary draws for the individual-specific random parameter

[16]Because experimental evidence does not support other preference classes, we ignore linear WTP functions for performance and WTP curves whose minimum or maximum exist where the performance level is zero.

parameters on FLO certification and coffee price $\left(\frac{-\beta_{FLO}}{\beta_{price}}\right)$[17]. Notice, the most predominant consumer type exhibits the "inverted-U" shaped WTP curve (which we term consumer type II). This class alone accounts for over 80% (and as high as 90% for three of the six experiments) of all draws from the parameter distribution regardless of information treatment. As for the other classes, depending on the performance attribute examined either Type I or Type III preferences were the next most likely. Based on our specification, Type IV preferences can almost be ruled out as it accounts for a very small fraction of our sampled preference distribution.

4.2 Fixed Levels in Performance

The levels model estimates a separate fixed constant for each performance level after normalizing on one level. We normalize on the lowest performance level offered on the performance model (3%) for each performance attribute. The first thing to note about estimating the fixed levels model is that it increases the numbers of parameters in the model, and this increase is especially felt when trying to estimate random parameters for label attributes. Comparing to the quadratic specification using a mixed logit model, introducing random parameters for each fixed level performance coefficient requires an additional eight parameters. While identification such a model was possible, we found no significant heterogeneity amongst respondents with respect to performance attributes[18].

Table 8 contains three sets of results for each of the three information treatments. Since the results are very similar to the quadratic specification of the preceding section, we focus only on differences. First, in all but one case we find that the utility function increases as performance metrics increase. It is only for the 'Benefits Only' model do the level coefficients for participation increases, peak, and then decrease. Figures 7 and 8 shows the impact of changes in performance (relative to 3%) on consumer utility. Of particular interest is the pattern check of the contribution to the indirect utility function as performance increases (relative to 3%). As revenue increases, utility increases in near linear fashion across all information treatments. For participation increases, utility increases are flatter and only one information treatment shows an "inverted-U" shape.

[17]The WTP intercept term was positive for 95.2% of the draws from the parameter vector.

[18]Since we were not able to estimate random effects on the performance attributes, we restrict all parameters in the model to be non-random and estimate a simple multinomial logit model for this set of results.

5 Summary and Conclusion

Using Fair Trade labeled coffee as an example, this paper has explored how consumers' willingness to pay for a product with socially responsible attributes is dependent on both the overall information regarding the program in question as well as the performance metrics. Succinctly put, we test how consumer attitudes towards poverty and relative deprivation facing producers included as well as excluded by a Fair Trade program affects their willingness to pay a Fair Trade price premium as the scale and scope of the program varies. It bears emphasis that our theoretical model and experimental design conform as closely as possible to what is observed in the "real-world". As a result, our theoretical model does not account for how the current willingness to pay for a Fair Trade product might affect the scale and scope of the program at a future date. Along the same vein, we do not emphasize the certification issue, i.e., the entry and exit decisions of poor producers into or from a Fair Trade cooperative since these issues are not advertized by Fair Trade programs

Turning to the experimental analysis we find that respondents' willingness to pay for Fair Trade labeled coffee first rises and then falls for both increases in producer participation and increased income guarantee, and more importantly for all information treatments regarding the Fair Trade program. However, there are differences in the willingness to pay curvature across information treatments. For example, when respondents are presented with the "Benefits only" information about the Fair Trade program the observed willingness to pay occurs at 71% of increased revenues for producers participating in the Fair Trade program. On the other hand, the decline in willingness to pay occurs for 60% of increased revenue when respondents are presented with "Full Information" about the Fair Trade program. This alludes to the possibility that respondents are more aware of the relative deprivation deepening impact of an increase in scope of a Fair Trade program on non-participating poor producers when both the costs and benefits of a Fair Trade program is advertised.

Second, the empirical results allow for a link to our theory regarding consumer preferences towards poverty and relative deprivation facing producers both included and excluded by a Fair Trade program. As our quadratic specification shows, with respect to increases in scale (number of participating poor producers) respondents' willingness to pay exhibits an "inverted-U" shape, and this pattern is consistent across all information treatments regarding the costs and benefits of a Fair Trade program. Our theory suggests that the sufficient condition for such a pattern to be observed is that consumers are averse to relative deprivation (Section 2.4). Turn-

ing to the impact of increases in scope on the respondents' willingness to pay we also observe an "inverted-U" pattern. However, in this case our theory suggests that this pattern is consistent with consumers who are poverty loving but averse to relative deprivation (Section 2.5). This raises an intriguing possibility: do consumer attitudes towards poverty and relative deprivation depend on whether the scale or the scope of a Fair Trade program is of central focus? In other words, our results open up the question that consumers' attitudes towards poverty and relative deprivation facing poor producers in a given country may depend on which performance metric is being advertised as the objective of the Fair Trade program — an analysis of which is beyond the scope of this paper. Nevertheless, the illustration here shows that there is a limit to which consumption preferences that exhibit aversion to relative deprivation can effectively bring about an equal world, if market-based incentives (through performance based labeling) is used as a tool.

It is worth pointing out, however, that while relative deprivation of poor producers excluded from a Fair Trade program is a serious issue as highlighted by our theory, it requires respondents to have a nuanced understanding of the interplay between poverty and relative deprivation that a Fair Trade program engenders. In this respect, our levels specification shows that as revenue increases, utility increases in near linear fashion across all information treatments. For participation increases, utility increases are flatter and only one information treatment shows an "inverted-U" shape. Thus, an alternate explanation may well be that consumers are simply averse to poverty. As the discussions in Section 2.4 and 2.5 show, aversion to poverty is sufficient in exhibiting an increasing Fair Trade price premium with an increase in either the scale or the scope of the program.

References

Adamowicz W., J. Louviere and M. Williams; (1994). "Combining Revealed and Stated Preference Methods for Valuing Environmental Amenities". *Journal of Environmental Economics and Management*, 26(3): 271-292.

Basu, A. K. and R. L. Hicks (2008); "Label Performance and Willingness to pay for Fair Trade coffee: A Cross-National Perspective"; International Journal of Consumer Studies, 32, pp 470-478.

Basu, A. K., N. H. Chau and U. Grote; (2003). "Eco-Labeling and Stages of Development". *Review of Development Economics*, 7: 228-247.

Basu, A. K., N. H. Chau and U. Grote; (2004). "On Export Rivalry and the Greening of Agriculture: The Role of Eco-labels". *Agricultural Economics*, 31(2-3): 135-147.

Basu, A. K., N. H. Chau and U. Grote; (2006). "Guaranteed Manufactured without Child Labor: The Economics of Consumer Boycotts, Trade Sanctions and Social Labeling". *Review of Development Economics*, 10(3): 466-491.

Becchetti, L. and F. Camillo Rosati (2005); "The Demand for Socially Responsible Products: Empirical Evidence from a Pilot Study on Fair Trade Consumers," Working Papers 04, ECINEQ, Society for the Study of Economic Inequality.

Ben-Akiva, M. and T. Morikawa; (1990). "Estimation of Switching Models from Revealed Preferences and Stated Intentions". *Transportation Research*, 24A(6): 485-495.

Benebou, R. and J. Tirole (2009); "Individual and Social Responsibility"; IZA Discussion Paper No. 4570, Bonn, Germany.

Bjorner T.B., L.G. Hansen and C.S. Russell; (2004) "Environmental Labeling and Consumers' Choice - An Empirical Analysis of the Effect of the Nordic Swan". *Journal of Environmental Economics and Management*, 47(3): 411-434.

de Janvry, Alain; C. McIntosh and E. Sadoulet (2009); "What's Fair in Fair Trade? Generating Long-Term Benefits in a Disequilibrium Market"; working paper, Agricultural & Resource Economics, University of California at Berkeley.

DeShazo, J. R. and G. Fermo. (2002). "Designing Choice Sets for Stated Preference Methods: The Effects of Complexity on Choice Consistency", *Journal of Environmental Economics and Management*, 44, 123-143.

Elliott, K. and R. B. Freeman; (2003). "Can Labor Standards Improve Under Globalization?" Washington, DC: Institute for International Economics.

Forsyth K., Haley D., Kozak R. (1999) Will Consumers Pay More for Certified Wood

Products?. *Journal of Forestry*, 9(2), 19-22.

Hainmueller, J., M. J. Hiscox and S. Sequeira (2011) "Consumer Demand for the Fair Trade Label: Evidence from a Field Experiment". Unpublished Manuscript. Harvard University.

Hicks, R. L.; (2006). "Performance-Based Labeling" in Ulrike Grote, Arnab K. Basu and Nancy H. Chau (eds) New Frontiers in Eco and Social Labeling, Physica-Verlag: Springer Publishing, Berlin and New York.

Hicks, R. L.; (2002). Stated Preference Methods For Environmental Management: Recreatitional Summer Flounder Angling in the Northeastern United States (U. S. D. o. C. National Marine Fisheries Service, Ed.).

Hiscox, M. J. and N. F. B. Smyth; (2005). "Is There Consumer Demand for Improved Labor Standards? Evidence from Field Experiments in Social Labeling". Unpublished Manuscript. Harvard University.
http : *//www.people.fas.harvard.edu/ hiscox/SocialLabeling.pdf*

Hiscox, M. J., M. Broukhim, and C. S. Lewis (2011) "Consumer Demand for Fair Trade: New Evidence from a Field Experiment Using EBAY Auctions of Fresh Roasted Coffee". Unpublished Manuscript. Harvard University.

James, D; (2000). "Justice and Java: Coffee in a Fair Trade Market". NACLA Report to the Americas 34:11-42.

Kimeldorf, H., R. Meyer, M. Prasad and I. Robinson. (2004). "Consumers of the World Unite: A Market-based Response to Sweatshops". Unpublished paper. Department of Sociology, University of Michigan.

Koppel, H. and G. Schulze (2008). "Inefficient but Effective: A Field Experiment on the Effectiveness of Direct and Indirect Transfer Mechanisms"; Department of International Economic policy, Discussion paper Nr. 2, University of Freiburg, Germany.

Krinsky, I and A. Robb; (1986). "On Approximating the Statistical Properties of Elasticities", *Review of Economics and Statistics*, 68, pp 715-719.

Krinsky, I and A. Robb; (1991). "Three Methods for Calculating the Statistical Properties of Elasticities: A Comparison", *Empirical Economics*, 16, pp 199-209.

Louviere, J. J., D. A. Hensher and J. D Swait; (2000). Stated Choice Methods: Analysis and Applications with a Contribution by Wiktor Adamowicz. Cambridge University Press, New York and Melbourne.

Milford, A. (2004); "Coffee, Co-operatives and Competition: The Impact of Fair Trade"; CMI Report R2006:4, Chr. Michelson Institute, Bergen, Norway.

Milkman, R.; (2004). "Sales Results from UCSB Pants Experiment". Unpublished paper. Department of Sociology, University of California, Santa Barbara.

Nimon, W. and J. C. Beghin; (1999). "Are Eco-Labels Valuable? Evidence from the Apparel Industry", *American Journal of Agricultural Economics*, 81, 801-811.

O'Brien, K. A. and M. F. Teisl. (2004). "Eco-information and its effect on Consumer Values for Environmentally Certified Forest Products." *Journal of Forest Economics*, 10(2), 75-96.

Program on International Policy Attitudes (PIPA); (2000). Americans on Globalization. University of Maryland.

http : *//www.pipa.org/onlinereports/globalization/global_rep.html*

Reinstein, D. and J. Song (2009); "Efficient Consumer altruism and Fair Trade Products"; working paper, Department of Economics, University of Essex, UK.

Richardson, M. and R. Stahler (2007); "Fair Trade"; Economics Discussion Papers, No. 0709, University of Otago, Dunedin, New Zealand.

Ronchi, L. (2006); ""Fairtrade" and market failures in Agricultural Commodity Markets"; World Bank Policy Research Working Paper No 4011. Washington D.C.

Rode, J., R. M. Hogarth, and M. Le Menestrel. Ethical differentiation and market behavior: An Experimental Approach. *Journal of Economic Behavior and Organization*, Forthcoming.

Sen, A. K.; (1970). "Interpersonal Aggregation and Partial Comparability," *Econometrica*, 38, 393-409.

Taylor, P.; (2005). "In the Market But Not of It: Fair Trade Coffee and Forest Stewardship Council Certification as Market-Based Social Change". *World Development*, 33, 129-147.

Teisl, M., B. Roe and R. L. Hicks; (2002). "Can Eco-Labels Tune a Market", *Journal of Environmental Economics and Management*, 43, 339-359.

Train, K. E.; (2003). Discrete Choice Methods with Simulation. Cambridge University Press, Cambridge, U.K.

TransFair USA; (2005). 2005 Fair Trade Facts and Figures. TransFair USA, Oakland.

Yitzhaki, S.; (1979). "Relative Deprivation and the Gini Coefficient," *Quarterly Journal of Economics*, 93, 321-324.

Tables and Figures

Table 1: Summary of Empirical Strategy

Label Type	Fair Trade Label Narrative
Performance	Full Information
	Benefits Only
	No Information
Traditional Label	Full Information
	Benefits Only
	No Information

Table 2: Performance and Traditional Label Variable Descriptions and Attribute Levels.

Attribute	Description	Performance Label Range	Traditional Label Range
Price	Price per Cup of Coffee	Labeled: ($1.00,$1.25,$1.50,$1.75) Non-labeled:($0.30,$0.50,$0.70,$0.90)	Labeled: ($1.00,$1.25,$1.50,$1.75) Non-labeled:($0.30,$0.50,$0.70,$0.90)
Brazil	Grown in Brazil	(0,1)	(0,1)
Costa Rica	Grown in Costa Rica	(0,1)	(0,1)
Colombia	Grown in Colombia	(0,1)	(0,1)
Organic	Non-fair trade labeled coffee labeled as organic	(0,1)	(0,1)
USDA	Certified by the USDA	(0,1)	(0,1)
FLO	Certified by the Fairtrade Labeling Organisation International	(0,1)	(0,1)
CGA	Certified by the Country of Origin's Coffee Growers Association	(0,1)	(0,1)
Revenue	Increased Grower Revenue for Program Participants	(3%,25%,50%,75%)	
Participation	Increased Grower Participation in Fair Trade Program	(3%,25%,50%,75%)	

Table 3: Information Treatments used in Choice Experiments.

Information Treatment	Narrative
Full Information	Advocates argue that Fair Trade certified products ensure that farmers, workers, and artisans are paid a fair price for their products or labor, don't use child labor or forced labor, have healthy and safe working conditions, use sustainable and environmentally-friendly production methods, and have long-term and direct relationships with buyers. Others feel that fair trade is discriminatory against growers and countries that don't have the resources to institute a Fair Trade program.
Benefits Only	Advocates argue that Fair Trade certified products ensure that farmers, workers, and artisans are paid a fair price for their products or labor, don't use child labor or forced labor, have healthy and safe working conditions, use sustainable and environmentally-friendly production methods, and have long-term and direct relationships with buyers.
No Information	No narrative given to respondents.

Table 4: Results for the Full Information Experiment

	Performance Label		Traditional Label	
Variable	Estimate	S.E.	Estimate	S.E.
Price	-2.3137**	0.4010	-1.2045**	0.3135
Brazil	4.0065**	0.4522	2.4209**	0.3716
Colombia	4.0288**	0.4536	2.1798**	0.3660
Costa Rica	4.0860**	0.4499	2.2757**	0.3672
USDA (μ)	0.1049	0.3222	0.7257*	0.3926
USDA (σ)	0.0161	0.4893	0.1644	2.3113
FLO (μ)	0.4054	0.3505	0.9384**	0.2909
FLO (σ)	0.6200	0.8798	0.1242	1.4766
CGA (μ)	-0.1157	0.3312	0.4520	0.2800
CGA (σ)	0.5119	0.9197	0.2564	1.4840
ORG (μ)	0.2212	0.2592	-0.0069	0.2507
ORG (σ)	1.0828*	0.6155	0.8638	0.6909
Revenue (μ)	4.3259**	1.2481		
Revenue (σ)	3.5078**	1.2738		
$Revenue^2$ (μ)	-3.6361**	1.5573		
$Revenue^2$ (σ)	0.5219	1.6193		
Participation (μ)	3.4932**	1.2618		
Participation (σ)	1.7161	1.5792		
$Participation^2$ (μ)	-3.3019**	1.6145		
$Participation^2$ (σ)	2.3578	1.9213		
Observations	885		420	
Respondents	177		84	
Avg. No. of Questions	5.00		5.00	

Table 5: Results for the Benefits Only Experiment

	Performance Label		Traditional Label	
Variable	Estimate	S.E.	Estimate	S.E.
Price	-2.0274**	0.4231	-1.4890**	0.2788
Brazil	2.8296**	0.4296	3.1671**	0.4780
Colombia	2.8323**	0.4382	3.4847**	0.4703
Costa Rica	2.8114**	0.4224	3.3164**	0.4757
USDA (μ)	0.6002*	0.3554	1.2189**	0.4277
USDA (σ)	0.0087	0.5346	0.3101	1.7602
FLO (μ)	1.0360**	0.3946	1.8874**	0.2719
FLO (σ)	0.6691	0.8409	0.0232	0.6841
CGA (μ)	0.6133*	0.3655	1.3747**	0.2199
CGA (σ)	0.4317	1.0812	0.1060	2.4294
ORG (μ)	0.5117	0.5048	0.4914	0.4418
ORG (σ)	1.6094**	0.7784	1.4086**	0.7306
Revenue (μ)	4.9509**	1.5839		
Revenue (σ)	3.3307**	1.5347		
$Revenue^2$ (μ)	-3.4999**	1.6820		
$Revenue^2$ (σ)	0.3643	1.5386		
Participation (μ)	5.5743**	1.5628		
Participation (σ)	3.8456**	1.4629		
$Participation^2$ (μ)	-4.8006**	1.6904		
$Participation^2$ (σ)	0.8230	1.8128		
Observations	895		570	
Respondents	179		114	
Avg. No. of Questions	5.00		5.00	

Table 6: Results for the No Information Experiment

	Performance Label		Traditional Label	
Variable	Estimate	S.E.	Estimate	S.E.
Price	-2.7413**	0.3393	-1.3544**	0.2292
Brazil	3.3683**	0.3625	2.9388**	0.3041
Colombia	3.2547**	0.3597	2.9365**	0.3117
Costa Rica	3.3185**	0.3574	2.9007**	0.3127
USDA (μ)	1.2818**	0.3360	0.7472**	0.2568
USDA (σ)	-0.0325	0.6334	0.0188	0.6837
FLO (μ)	1.7137**	0.3547	1.2110**	0.2047
FLO (σ)	-0.0606	0.4725	0.0248	0.6461
CGA (μ)	1.1364**	0.3244	0.7323**	0.2518
CGA (σ)	0.0435	0.6566	0.9096	1.1752
ORG (μ)	0.1192	0.3596	0.3173	0.2299
ORG (σ)	1.2794*	0.6521	1.2850**	0.4662
Revenue (μ)	3.9177**	1.0019		
Revenue (σ)	2.0246**	0.9735		
$Revenue^2$ (μ)	-2.8656**	1.2168		
$Revenue^2$ (σ)	0.0520	1.4689		
Participation (μ)	3.2670**	1.0463		
Participation (σ)	2.4586**	1.0392		
$Participation^2$ (μ)	-2.6167**	1.2876		
$Participation^2$ (σ)	0.6129	1.5338		
Observations	864		900	
Respondents	174		180	
Avg. No. of Questions	4.97		5.00	

Table 7: Critical Performance Values and the Distribution of WTP

Information Treatment	WTP_0	WTP_1	WTP_{max}
Full	0.357	0.833	0.595
Benefits	0.252	1.163	0.707
None	0.123	1.245	0.684

Table 8: Robustness checks for Inverted U WTP

	Full Information		Benefits Only		No Information	
Variable	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
Price	-1.5481**	0.1772	-1.3230**	0.1830	-2.1314**	0.1964
Brazil	3.3085**	0.2591	2.1717**	0.2508	2.7888**	0.2408
Colombia	3.2489**	0.2544	2.0951**	0.2459	2.6680**	0.2368
Costa Rica	3.3830**	0.2555	2.1824**	0.2448	2.7553**	0.2352
USDA	-0.1029	0.2252	0.4843**	0.2349	1.0135**	0.2370
FLO	0.0752	0.2201	0.6681**	0.2286	1.3362**	0.2309
CGA	-0.2421	0.2237	0.3944**	0.2315	0.9169**	0.2314
Org	0.1235	0.1485	0.0736	0.1831	-0.1128	0.1747
Revenues 25%	0.5839**	0.1632	0.7395**	0.1519	0.5901**	0.1590
Revenues 50%	0.9075**	0.1571	0.8475**	0.1472	0.9371**	0.1554
Revenues 75%	1.0534**	0.1604	1.1415**	0.1513	1.0376**	0.1596
Participation 25%	0.4885**	0.1651	0.6089**	0.1568	0.5911**	0.1641
Participation 50%	0.5909**	0.1549	0.9468**	0.1513	0.6532**	0.1544
Participation 75%	0.6427**	0.1580	0.8675**	0.1546	0.8196**	0.1591
Observations	885		895		864	

In this section we would like for you to imagine that you are in your favorite campus coffee shop and are looking to purchase a cup of coffee. There are three different brands available for you to purchase. We will ask you to repeat the brand choice several times.

Please assume that the brand attributes are identical except for price, and any information given on the labels. **For example, please assume that product quality and size is the same across the three different brands**. If the fair trade or organic label is blank, then there is no information regarding whether that product meets standards or not.

6. Assume that you are going to buy a cup of coffee. If you could only choose from the following three choices, which one would you choose? (*Please check only **one** of the boxes at the bottom of this page).*

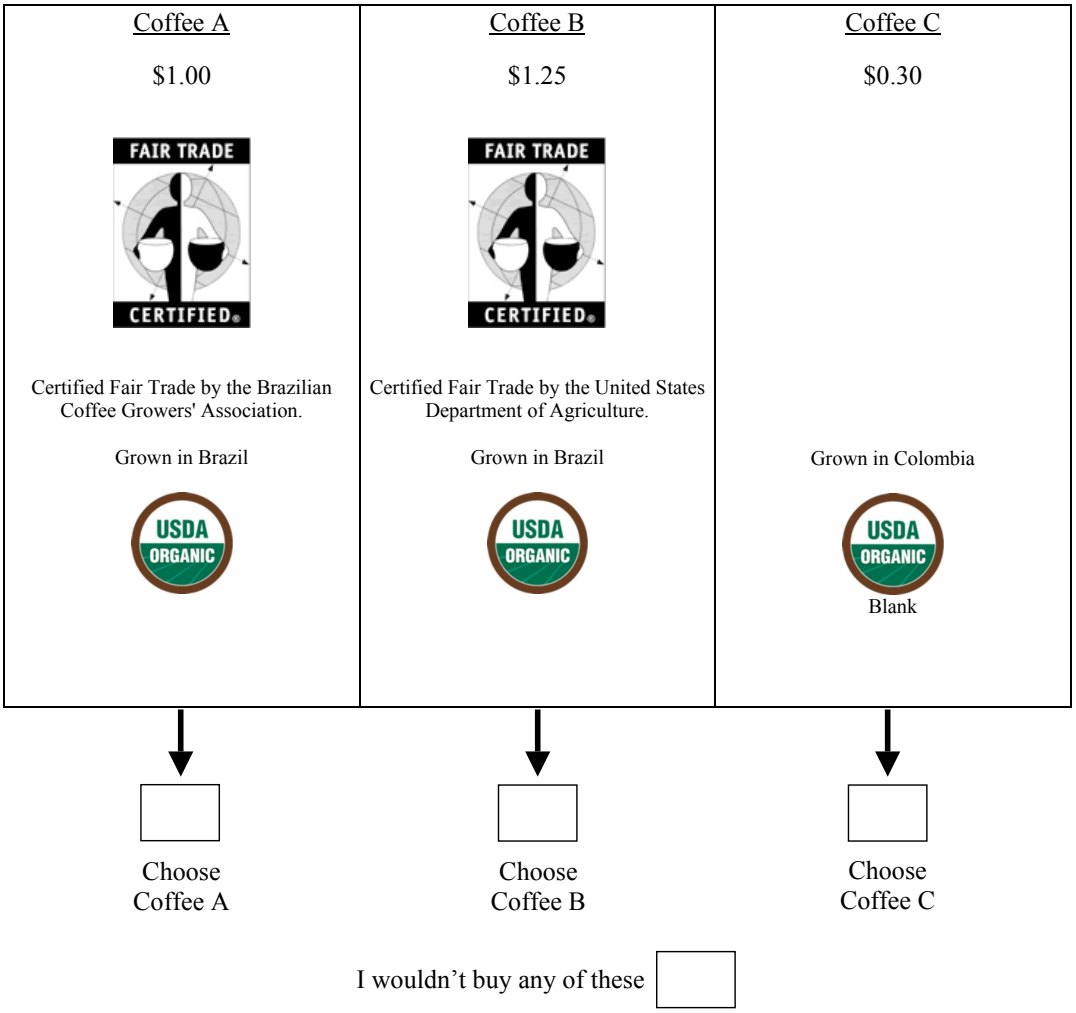

Figure 1: Traditional Label Survey Example

In this section we would like for you to imagine that you are in your favorite campus coffee shop and are looking to purchase a cup of coffee. There are three different brands available for you to purchase. We will ask you to repeat the brand choice several times.

Please assume that the brand attributes are identical except for price, and any information given on the labels. **For example, please assume that product quality and size are the same across the three different brands**. If the fair trade or organic label is blank, then there is no information regarding whether that product meets standards or not.

Any fair trade information describes the benefits and extent of the fair trade program relative to other coffee producers in the grower's home country.

6. Assume that you are going to buy a cup of coffee. If you could only choose from the following three choices, which one would you choose? (*Please check only **one** of the boxes at the bottom of this page).*

Coffee A	Coffee B	Coffee C
$1.25	$1.00	$0.50
This Brand's Fair Trade Performance for Program Participants	This Brand's Fair Trade Performance for Program Participants	No Information Available
Grower Revenue Increases Per Pound of Coffee: 50% Coffee Growers Participating: 25%	Grower Revenue Increases Per Pound of Coffee:: 25% Increased Grower Participation: 75%	
Certified Fair Trade by the United States Department of Agriculture. Grown in Brazil USDA ORGANIC	Certified Fair Trade by the United States Department of Agriculture. Grown in Colombia USDA ORGANIC	Grown in Costa Rica USDA ORGANIC Organic
☐ Choose Coffee A	☐ Choose Coffee B	☐ Choose Coffee C

I wouldn't buy any of these ☐

Figure 2: Traditional Label Survey Example

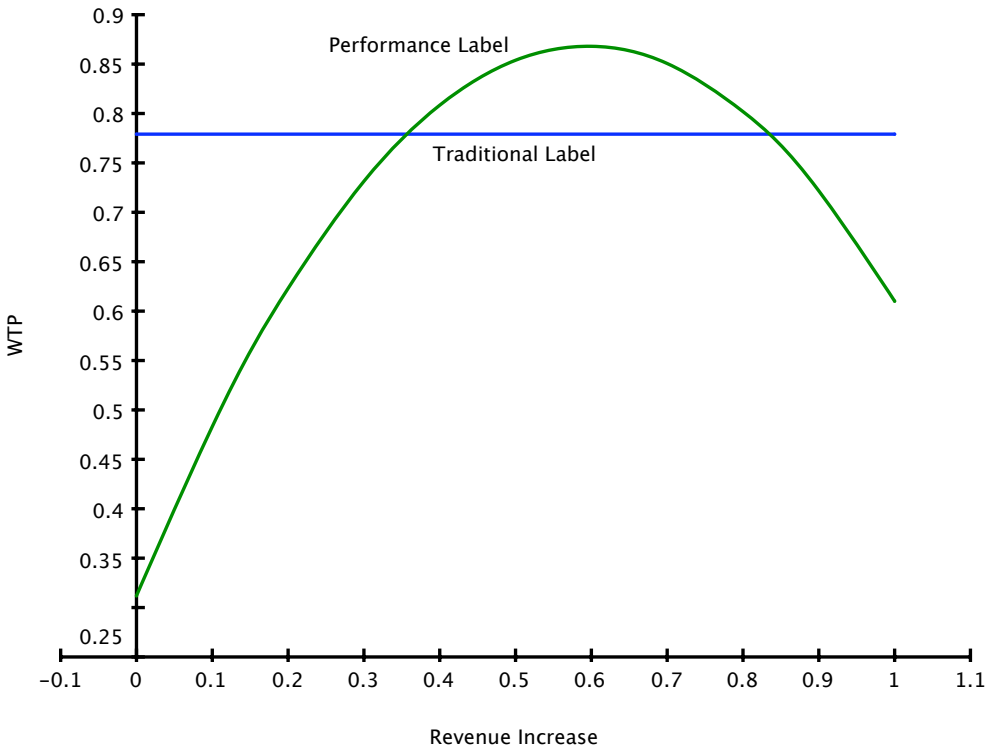

Figure 3: WTP Results for the Full Information Model

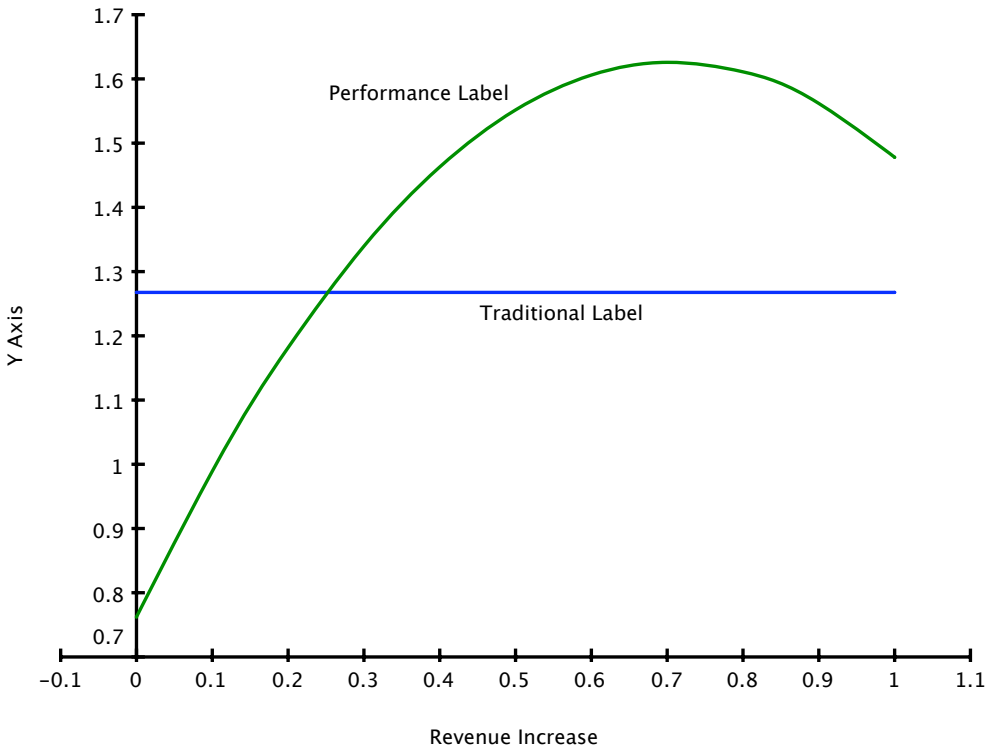

Figure 4: WTP Results for the Benefits Only Model

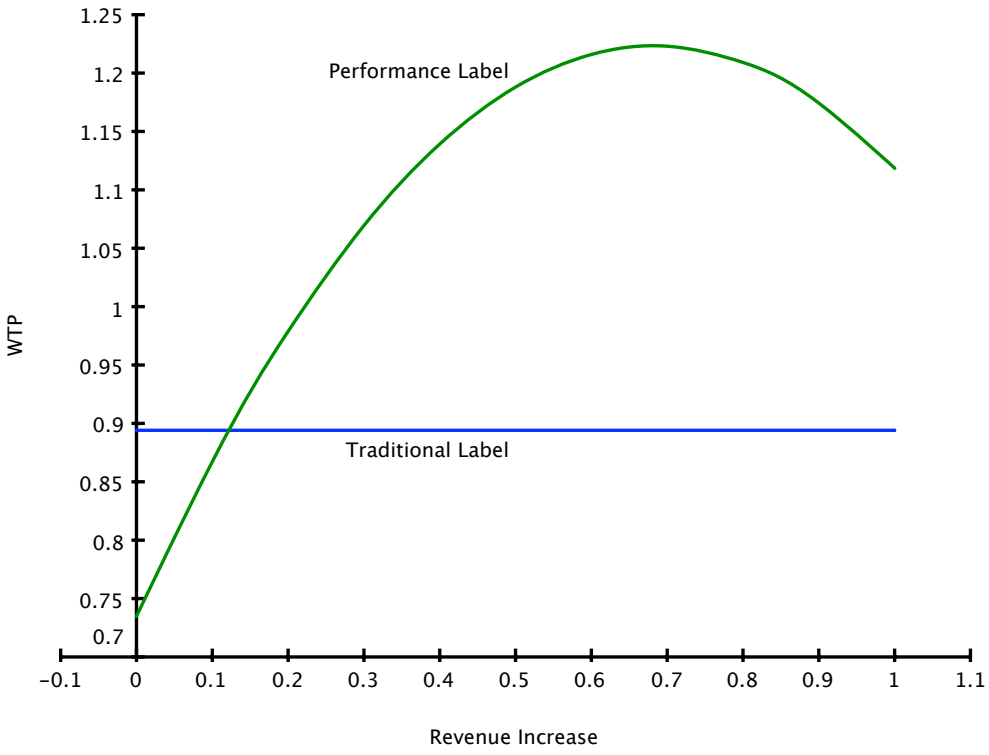

Figure 5: WTP Results for the No Information Model

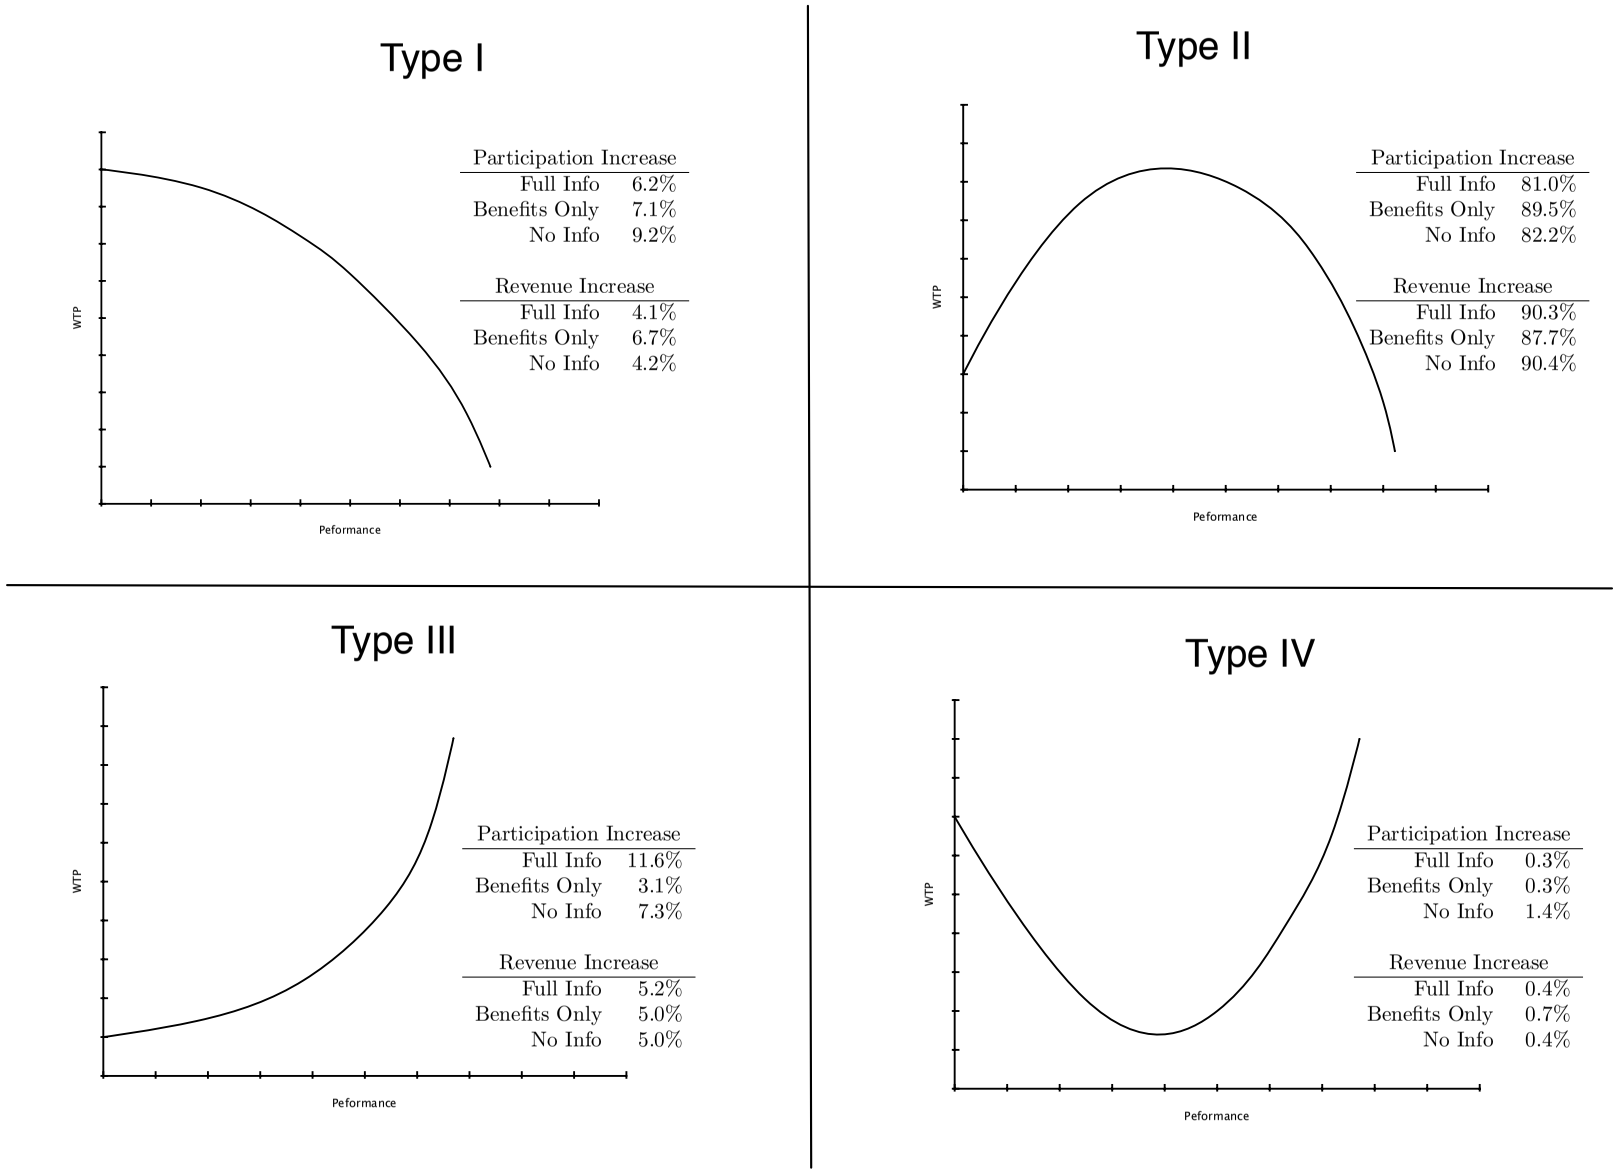

Figure 6: Distribution of Estimated Preferences

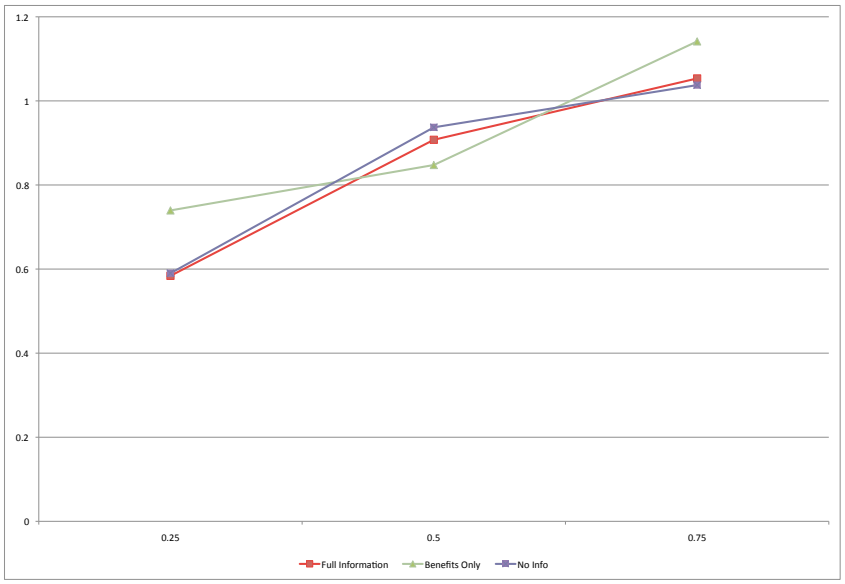

Figure 7: Curvature Plots from Level Model: Revenue Increases Relative to 3%

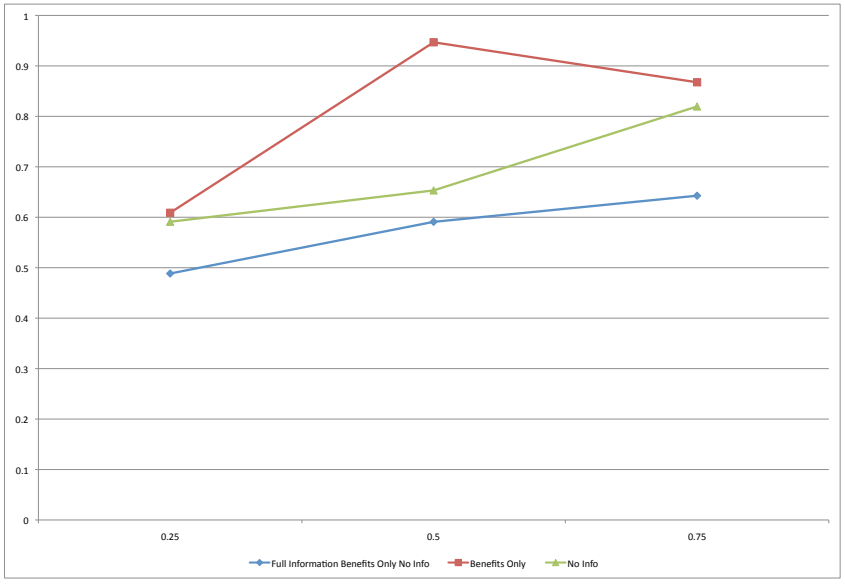

Figure 8: Curvature Plots from Level Model: Participation Increases Relative to 3%