

Why Can't Alison Sell her Drill? Evidence from eBay

(DRAFT- PLEASE DO NOT CITE)

Introduction:

Do products sold by women garner lower prices than products sold by men? If so, what are the mechanisms producing these disadvantages? Are women penalized for selling seemingly masculine products more than for selling seemingly feminine products? Are penalties greater when buyers have less information regarding the market price of the product?

Answers to these questions are important because while scholars have produced an impressive body of theoretical and empirical evidence about the extent and causes of gender inequality in labor markets, we know surprisingly little about how gender operates in product markets. If women experience similar disadvantages in product markets as they do in labor markets, the negative effects of gender in economic life are greater than previously understood. Moreover, understanding the mechanisms that produce gender price gaps in product markets, would enable us to better understand and address, gender inequality in them but also in other arenas like the labor market. Finally, it should be noted that unlike with labor-markets, peer to peer product markets are not regulated by law, so that whereas employment sex discrimination is prohibited, peer to peer sex discrimination in product markets is not.

In this paper we first review the literature on the effects of gender in product markets, describing in some detail the one recent study - on which we build and further explore - that shows that women are, in fact, disadvantaged as sellers compared to men, when selling the exact same new product. On average, women sellers received about 80 cents for every dollar a man received when selling the identical new product and 97 cents when selling the same used product. We then develop testable predictions about the mechanisms producing lower prices for women sellers. In short, we hypothesize that when there is a lack of fit between the gender of the actual owner and beliefs about who the prototypical owner should be, penalties arise. Thus for example, women are penalized more when selling seemingly masculine products. Moreover, we hypothesize that lack of information regarding the market prices of products increases the gender price gap because when uncertain about the market price, buyers use their cultural beliefs about gender as information

about how much they should pay for the product. We then combine the results from the previous eBay study on women in product markets with original data from three original experiments and use them all to explore our predictions; the three experiments document people's perceptions about traits of products: their seemingly masculinity and femininity, the availability of information about their market prices; and the traits of their prototypical owners. Responses of the participants in the experiments are then used in regression models to explain the gender price gaps obtained in the eBay data. By using this innovative research design, we therefore able to explain *real world price gaps* obtained in the eBay data, with *experimental results*.

Do women sellers receive lower prices than men?

To date, there has been comparatively little evidence regarding the effect of gender in product markets. There is some evidence, based on the findings of field experiments, which suggests discrimination against women buyers in product markets (11, 12). It has been hypothesized that women are discriminated against in the rental and insurance markets because they are believed to have a lesser ability to meet payments (12) or because gender is used as a proxy for reservation price (12). A couple of recent studies on women entrepreneurs suggest that women are disadvantaged compared to their male counterparts (Thebaud YEAR; Tinker et al YEAR). But how do women *sellers* fare? A recent study analyzing data on all the eBay auction transactions for 420 of the top selling products from 2009-2012 found evidence that women sellers are disadvantaged in product markets. Specifically, the study found that women sellers receive a smaller number of bids and lower final prices than do equally qualified men sellers of the exact same product in auctions on eBay. Women sellers received, on average, about 80 cents for every dollar a man received when selling the exact same new product and 97 cents when selling the same used product (1) (see Table S1). Note that the classification of products in the data is highly refined – a new blue iPod shuffle, second generation, for instance -- so that comparisons between women and men are for the exact same products. Hence, results regarding new products are particularly informative, as products are identical and new, so that quality-related explanations for the gender price gap become irrelevant.

As a policy, eBay does not reveal the gender of its users. Nonetheless, the authors were able to show in a separate experiment, that people were quite accurate in

guessing the gender of the seller from contextual cues in the product listing. Yet, while this paper was able to show that women sellers were disadvantaged, the data could not answer the question of *why* female disadvantage obtained. The study did find that the gender price gap was larger for some products than others.

In this project, we utilize the data on differences in the gender price gap across products from the eBay study, and pair it with original experimental data to test predictions about mechanisms generating the price differences between women and men sellers. We wish to compare for example, what happens when a man and a woman are selling the exact same new drill with what happens when they are selling a brand new sewing machine. Would the gender price gap be smaller for seemingly “feminine” compared to seemingly “masculine” products? And what if the buyer is uncertain about how much a sewing machine or a drill should cost? What happens if the product is used? Would buyers assume that women sellers have used female-typed products, such a sewing machine, more extensively than seemingly masculine products, such as a drill, such that they believe the former is more used or more worn out? Or would buyers, instead assume that female-typed products are of higher quality when sold by women compared to men, because women are more familiar with such products and therefore more able to accurately describe its condition and quality? Finally, would the buyer’s level of uncertainty be affected by the gender of seller and the gender-typing of the product? Is there less uncertainty, for example, when a man is selling a male-typed product than when a woman is selling the same product?

We first test whether people assign a likely gender to the owners-sellers of products. People may think, for example, that the prototypical owner of a drill is a man and that the prototypical owner of a sewing machine is a woman. We then test what happens when the actual owner-seller of a product is different than the prototypical owner, for example when a woman owns and sells a new drill. We predict that when a seller’s gender does not align with the gender of the prototypical seller, buyers will assign a lower value to the drill, compared to when the seller’s gender is prototypical. In other words, we predict the perceived characteristics of prototypical owners-sellers of products (being seemingly “feminine” or “masculine”, for example) interact with the perceived characteristics of the actual owner-seller (being a female or a male owner-seller) to generate prices and gender price gaps in product markets.

Why would people hold beliefs about the characteristics of prototypical owners of products? We draw on literature in social psychology which shows that people hold very defined and consensual cultural beliefs about social groups of people along two dimensions: competence and warmth (Fiske) and that these beliefs shape their social interactions with people from these social groups. Women for example, are typically perceived to be warmer, more emotional and sensitive than men, whereas men are perceived to be more competent, assertive and confident than women. In labor markets, this matters because higher competence ratings of individual employees lead to higher salaries, but higher warmth ratings do not (Cuddy et al 2004). If people tend to evaluate members of social groups along competence and warmth dimensions, they may similarly evaluate prototypical product owners along these same dimensions. People may think, for example, that an owner of a male-typed product, such as a drill is extremely competent but not very warm, and that the female-typed product, such as a Disney movie, is warm but not very competent.

If indeed people hold consistent and defined cultural beliefs about prototypical owners of products, we would expect such beliefs to interact with beliefs about actual sellers, affecting the prices people offer. For example, if prototypical owners of golf clubs are believed to be highly competent, but not very warm, then when women who are believed to be more warm than competent are selling such products, a perception of lack of fit between the product and the concrete seller may arise. Likewise a perception of lack of fit may arise, if the prototypical owners of golf clubs are perceived to be men, and the concrete seller of a golf club is a woman. As a result, women sellers of such golf clubs may receive lower prices, compared to men sellers. A perceived lack of fit may also arise when a man is selling a product that is perceived to be associated with warm or women owners, like a Disney movie for example. Nonetheless, because warmth and femininity in general are associated with lower value compared to competence and masculinity, lack of fit would have lower implications for men selling products associated with warmth compared to when women sell products associated with competence.

We borrow the term “lack of fit” from the literature in social psychology that suggests that when the job requirements contradict with stereotypes about job candidates, a failure will be predicted, whereas when there’s a fit between the stereotypes about the job candidates and the job requirements, a success will be

predicted (Heilman, 1983, 1995). In the context of product markets, a good fit will be perceived, when the perceived traits of prototypical owners, fit with the perceived traits of concrete owners-sellers.

One reason to expect that buyers will rely on cultural beliefs about the characteristics of prototypical and actual owners of products in their market interactions is that market interactions tend to involve uncertainty about value and prices. Under conditions of uncertainty, people are more likely to draw on widely shared beliefs, such as stereotypes and cultural schemas (Podolony XX; Correll et al 2016). In the case of products, people likely look for “clues” – like the type of owner, product and the fit between them - that would help them decide how much they are willing to pay for a product. When there appears to be a lack of fit between who the owner “should be” (the prototypical owner) and who the owner really is (the actual owner), uncertainty increases, which likely leads to increased scrutiny of the product, which can negatively affect judgments about product quality. Indeed research shows evaluators engage in greater scrutiny of women and black men when making hiring decisions for jobs in which they are not typical, which results in these groups receiving lower evaluations. If a similar process happens in product markets, products sold by atypical sellers may be judged as more flawed, problematic or less valuable compared with identical products sold by more prototypical sellers. Buyers therefore, may form their evaluations and expectations, on the basis of shared cultural beliefs or scripts about types of owners, products and the fit between them. Since products vary by the level of uncertainty associated with their value (e.g. there is less uncertainty in the value of a \$100 gift card compared with a painting), they vary in the need for buyers’ to search for “clues” contained in shared cultural beliefs. Thus, we predict that the more uncertain buyers are about the value of a product, the more they would rely on cultural beliefs about the specific owner, the prototypical owner of the product and the fit between them, when deciding how much to pay for a product.

Data and Methods

Overview

We created a dataset that combines data on the gender price gap of 306 of the bestselling products from the eBay study described above with original data from three original experiments. The first experiment assesses participants’ perceptions of

the sellers of a randomly drawn set of new or used products from the eBay data on the dimensions of competence and warmth as well as on their perceived gender. These perceptions are then used to predict the gender price gap of the products. The second experiment assesses participants' uncertainty about the price they would pay for a randomly drawn set of new or used eBay products. We also evaluate whether they indicate a desire to reduce their uncertainty and what effect uncertainty reducing preferences have on the gender price gap. The third experiment provides a deeper exploration of how gender affects the judgments of used products and how those judgments affect the gender price gap.

eBay Gender Price Gap Data

The original eBay data included 420 bestselling products, and each product appeared in the data both as a new product and as a used product. These data revealed that on average, women received 80 cents for every dollar men receive for selling the same new products, and about 97 cents for selling the same used product. Important for the current paper, the gender price gap varied by the type of product being sold.

For our analysis, we use only the products – either new or used - that were uniquely identifiable by a photo, resulting in a dataset of 306 bestselling products—99 of which are new and 207 that are used. For each we there is a unique gender price gap, controlling for all other transaction characteristics (such as seller's reputation, experience, number and type of pictures, time and duration of auction, etc. The dataset also contains the average final price received on eBay for each product. The average gender price gap is the difference between what sellers paid for products sold by women minus what they paid for the very same product sold by men, controlling for other transaction characteristics. For 306 products used in this study, the average range of the gender price gap is from negative 285% to 75% and the average is negative 2.26%. We combine the variables from the eBay dataset with the data from the three experiments below.

Experiments and Measures

2596 participants from Amazon Mechanical Turk were randomly assigned to one of the three experiments and to condition within each experiment. We randomly

assigned participants to the three experiments to ensure that participants participated in only one experiment. In each experiment the participants rated products on different dimensions, described below. These perceptions from the experimental data were then used to predict the price gap between actual women and men sellers obtained in the previous eBay study.

(1) Experiment 1: Perceptions of Warmth and Competence of Owners

Participants were presented with photos of five different products, randomly drawn from the larger dataset. The products were described as either used or new, depending on condition. Participants were then asked to rate, on 7-point scales, on how pleasant, sincere, tolerant, warm, competent, confident, intelligent, and high status they believed the owner of the product to be. They were then asked whether the prototypical owner-seller of the product was likely to be female on a 7-point scale ranging from “much more likely to be a man” to “much more likely to be a woman.”

Exploratory factor analysis revealed that the eight items loaded on two distinct dimensions: competence and warmth. Using the factor weights, we created a competence scale comprised of participants’ ratings of how competent, confident, intelligent, and high status they believed the owner to be ($\alpha=0.874$). The warmth scale was similarly constructed from participants’ ratings of the following items: pleasant, sincere, tolerant, warm ($\alpha=0.927$).

Participants tended to be consistent in their ratings of competence and warmth. They were also consistent in their perceptions of the seller’s gender. Finally, participants answered a series of demographic questions and were then directed to webpage where they entered payment information.

(2) Experiment 2: Uncertainty Experiment

Participants were presented with photos of five different products, randomly drawn from the larger dataset. The products were described as either used or new,

depending on condition. Participants were first asked how much they would be willing to pay for the new or used product in dollars. They were then asked on 7-point scales ranging from “not at all” to “very much:” 1) how uncertain they were about the price they would be willing to pay, 2) how interested they would be in the prices other people are willing to pay, and 3) how interested they would be in knowing who the current owner is. Whereas the first item captures their uncertainty, the latter two measure their desire to reduce uncertainty (and the ways in which they would want to do so).

(3) Experiment 3: What Does It Mean to be Used?

In this experiment participants learn the gender of the product owner. We varied the gender of the owner using first names (Alison/Brad) to signal gender. Participants were then asked to assess: 1) how extensively she or he has used the product on a 7-point scale ranging from “not at all” to “very extensively and 2) how familiar she or he is with products of this type on a 7-point scale ranging from “not at all” to “very much”. Finally, participants answered a series of demographic questions and were then directed to webpage where they entered payment information. Using the results of these survey, for each product in the dataset, we constructed two variables: one that captures the perceived gender gap in the usage of the product and the other that captures the perceived gender gap in the familiarity with the product.

Results

The data from all three experiment were used to construct variables capturing the perceived characteristics of products and of prototypical sellers-owners. For each product, the average of survey responses for each question was calculated. Thus, for example, the "competence" score for each product is the average score given by all raters of this specific product.

[Table 1 about here]

Table 1 reports the descriptive statistics for the variables used in the analysis. Note that the average gender price gap for the products we use in the analysis is -2.26%.

Note that here, unlike in the original eBay analysis, each product is represented only once (regardless of the number of sales on eBay) and there are about two times more used than new products.

Perceived Characteristics of Owners-Sellers

Survey participants tended to be consistent and defined in attributing characteristics to owners-sellers of products along the dimensions of competence and warmth. The standard deviation of participants' ratings of each of the owner-seller characteristics per product did not exceed 1.16 (answers were given on a scale of 1 to 7). Participants were also consistent as to whether the likely owner of a product is a woman or a man. The standard deviation of participants' assessment of the likelihood that the product owner-seller is a woman was 1.2 (here also, answers were given on a scale of 1 to 7).

A factor analysis of all eight characteristics revealed that they can be reduced to two factors that load distinctly on competence and warmth dimension. Traits that loaded consistently across products on the competence dimension were competence, intelligence, high status and confidence. Traits that loaded consistently across products on the warmth dimension were warmth, pleasance, tolerance and sincerity (See table S3). This suggests that competence and warmth differentiate prototypical owners-sellers of products.

For each product, competence and warmth composites of its owner-seller were calculated based on the factor analysis. Figure 1 plots products which received extreme high or low warmth and competence scores in the warmth-competence space. We can see, for example, that owners of a ping G15 Driver Golf Club are perceived to be highly competent but not very warm, whereas owners of a Susam Boyles' CD are perceived to be warm but not competent. Owners of a Kitchen Aid Mixed are perceived to be both warm and competent, whereas owners of an \$8 Lowe's Gift Card are perceived to be not very competent and not very warm.

[Figure 1 about here]

Perceptions of owners as being warm were correlated with believing the owner was likely to be a woman whereas perception of owners as being competent were correlated with believing the owner was a man. In a logistic regression model

predicting whether the owner of the product was perceived to be a woman, the warmth associated with the owner-seller had a positive significant effect. A one unit increase in warmth, generated a 0.92 increase in the probability of perceiving the owner-seller of the product to be a woman. On the other hand, a one unit increase in competence, decreased the probability that the owner-seller was perceived to be a woman by 0.5 (see Table S4).

The Gender Price Gap and the Femininity of Prototypical Owners-Sellers

The gender price gap is affected by the perceived gender of the prototypical owner of the product: when the prototypical owner of a product is perceived to be a man, the price gap between women and men sellers increases. when the prototypical owner of a product is perceived to be a woman, the price gap between women and men sellers decreases. Thus for example, the price gap between women and men sellers is larger when they sell drills – which are believed to be typically owned by men – compared to sewing machines – which are believed to be typically owned by men.

The results from OLS regression models predicting the gender price gap on eBay, are presented in Table 2. When the perceived prototypical owner of the product is a woman, the gender price gap for women seller decreases by 0.072 ($p < 0.1$, $N=306$). In other words, when the perceived prototypical owner of a given product is a man, women who sell it receive an additional 7% of a price penalty, compared to when the prototypical owner of the product is a woman. When the prototypical owner is perceived to be competent, the gender price gap increases by 0.10 for new products ($p < 0.05$, $N=306$). The perceived warmth does not affect the gender price gap.

It is worth noting that the perceived characteristics of prototypical owners-sellers predict not only the gender price gap, but also the final prices of products: products whose prototypical owners are perceived to be 1 unit warmer than others, receive prices that are \$2.9 lower, on average (final prices here are obtained from the previous eBay analysis, after controlling for auction and seller characteristics). Similarly, products prototypically owned and sold by women, receive prices that are \$2.2 lower ($p < 0.05$, $N=280$, see Table S5).

Uncertainty and the Gender Price Gap

People care more about the identity of the owner-seller of a product, when they are uncertain about the market price of the product or about how much they themselves are willing to pay for it. In an OLS regression model predicting how interested participants were in knowing who the owner-seller was, a one unit increase in their uncertainty about the market price of the product, increased their interest in who the owner is by 0.23 ($p < 0.05$, $N = 306$, see Table 3).

[Table 3, about here]

When people want to know who the owner-seller is, the price gap between women and men sellers increases. In an OLS regression model predicting the gender price gap, a unit increase in the interest in who the owner is, increases the gender price gap by 0.06 ($p < 0.1$, $N = 306$).

An alternative way for people, who are uncertain about the market price, to learn how much they themselves are willing to pay for a product, is to look at how much other people are willing to pay for the product. For some products ("public value products"), such as art and precious coins, such information is more valuable than for others. In an OLS regression model predicting interest in who the owner is, a one unit increase in being a "public value product" increases the interest in who the owner is by 0.46 ($p < 0.001$, $N = 306$).

Gender and Used products

With used products, two conflicting gender related mechanisms may operate; buyers may believe that women owners-sellers are more familiar with products typically owned by women, and therefore would be willing to pay for them *more* (especially when they themselves are uncertain about the price of the product); and, buyers may believe that women owners-sellers have used products typically owned by women more extensively than men owners-sellers and therefore would be willing to pay for them *less* (especially when the product is one that tends to be worn out). In Table 4, we distinguish between the two mechanisms operating in relation to used products and show that they both affect the price gap between women and men sellers.

[Table 4 about here]

We see in model 1, that women's familiarity with a product (survey participants' evaluations of the familiarity of women owners-sellers compared to men owners-sellers with the product) does not significantly affect the gender price gap. Nonetheless, because the familiarity of women with products (compared to men's) indicates not only that women know more about the product than men do, but also that women tend to use it more we distinguish between the two possible mechanisms in Model 2. We see that when a product was described by survey participants as a product that tends to wear out when used, women's familiarity with the product *increases* the price gap between women and men sellers (the interaction between women's familiarity and the tendency of the product to wear out is negative). Nonetheless, when the product is one that people tend to be uncertain about its price, women's familiarity with it *decreases* the price gap between women and men sellers (the interaction between women's familiarity and the uncertainty about the price of the product is positive). Table S7 presents the results of a robustness test where we interact the tendency of the product to wear out not with the familiarity of women with the product but with the tendency of women to use the product more (compare to men), and obtain similar results. In sum, we see that with used products, buyers are willing to pay women less, when products are typically used by women and the product tends to wear out, and at the same time, buyers are willing to pay women more, when women are more familiar with the products and they themselves are uncertain about the price of the product.

Discussion:

Bibliography:

- (1) Kricheli-Katz and Regev, "How Many Cents for a dollar? Women and Men in Product Markets", Science Advances (2), Feb
- (2) Fiske....