Targeted Incentives, Broad Impacts: Evidence from an E-commerce Platform*

Xiang Hui† Washington University

Meng Liu‡ Washington University

Tat Chan§ Washington University

October 11, 2020

Abstract

Digital platforms sometimes offer incentives to a subset of sellers to nudge behaviors, but how do these targeted incentives affect *all* sellers? In this paper, we study a policy change on a large e-commerce platform that offers financial incentives only to platform-certified sellers when they provide fast handling and generous return policies on their listings. We find that both the targeted and non-targeted sellers become more likely to adopt the promoted behavior after the policy change. Exploiting a large number of markets on the platform, we find that in markets with a larger proportion of the targeted population—hence more affected by the policy change—non-targeted sellers are more likely to adopt the promoted behavior and experience a larger increase in sales with little price changes. This finding is consistent with our key insight that a targeted incentive may *increase* demand for non-targeted sellers if both the targeted type and the promoted behavior are observed and valued by consumers. Our results have managerial implications for digital platforms that use targeted incentives.

Keywords: targeted incentives, quality provision, signalling, demand expansion

*We thank Seth Benzell, Avi Collis, Oren Reshef, Song Yao, and seminar participants at Washington University for their helpful comments. We are grateful to eBay for providing access to the data.

†hui@wustl.edu

‡mengl@wustl.edu

§chan@wustl.edu

1 Introduction

Digital platform designers often use incentives to improve sellers' quality provision. For example, Amazon offers a 75% discount on storage fees to sellers who store their products in its warehouses, and Uber and Lyft give bonuses to drivers who accept consecutive rides during busy hours.[1] Giving universal incentives to all sellers can be costly for platforms, so to reduce costs, platforms sometimes offer incentives only to a subset of sellers (targeted incentives). The trade-off, however, is that targeted incentives may have a limited effect on nudging quality provision by sellers because these incentives do not directly affect non-targeted sellers. Additionally, targeted incentives offer a cost advantage to targeted sellers, and can therefore reduce demand for non-targeted sellers. In this paper, we ask the following questions: When platforms adopt targeted incentives, how does the target size affect the effectiveness of these incentives, and what are the underlying mechanisms at play? Answering these questions is critical for determining the optimal target size when designing targeted incentives.

Our study draws evidence from eBay, one of the largest e-commerce platforms in the world. We exploit a policy change in March 2012 when eBay started to exclusively provide its certified sellers, namely eBay Top Rated Sellers (eTRS hereafter), a 5% discount on selling fees on their listings if they offered 1-day handling and 14-day returns (premium service, or PS hereafter). The setting is ideal for studying our research questions for three reasons. First, eBay has more than 400 distinct product markets that differ in the number of targeted sellers (i.e., eTRS), which creates essential variations in targeting sizes for our research questions. Second, the incentive's target is chosen based on sellers' certification status, which is determined before the policy change and is relatively fixed in the short run. Lastly, the incentive promotes a specific behavior, so we can directly measure quality provision in terms of this behavior.

We start by discussing our conceptual framework of targeted incentives. A seller's targeted type (i.e., eTRS status) is fixed in the short run, and the seller's quality provision (i.e., adopting the promoted PS or not) is endogenous. Consumers observe and value both sellers' eTRS status and whether they offer PS, and they make purchase decisions by comparing four substitute seller types — eTRS–PS, eTRS–non-PS, non-eTRS–PS, and non-eTRS–non-PS. The targeted incentive immediately encourages targeted sellers to adopt the promoted behavior, causing an increase in supply among eTRS–PS sellers and a decrease in supply among eTRS–non-PS sellers. The reallocation

[1] https://www.cnet.com/news/to-power-prime-one-day-shipping-amazon-asks-sellers-to-send-it-more-stuff/; https://www.uber.com/blog/los-angeles/consecutive-trips-earnings/; (9/14/2020)

in supply decreases the equilibrium price among eTRS–PS sellers but *increases* the equilibrium price among eTRS–non-PS sellers. How does this affect demand for non-eTRS sellers? Because the prices of their two substitutes from eTRS sellers move in opposite directions, the resulting change in demand is ambiguous. Specifically, demand for non-eTRS sellers can *increase* based on how consumers substitute across these four seller types.[2] Consequently, demand for PS, or PS premium, among non-eTRS sellers changes, causing them to re-optimize their decision on quality provision. This framework guides our analyses on how target size affects the impacts of targeted incentives, and predicts more salient equilibrium changes in markets with a larger targeted population.

The potential increase in demand for the non-targeted sellers may seem prima facie counterintuitive: in standard models of *targeted subsidies*, which are monetary transfers to sellers of a given type regardless of their behavior, the demand for non-subsidized firms will be unambiguously lower because the prices of their subsidized competitors will decrease.[3] However, our key insight is that this no longer holds for *targeted incentives*, which are targeted subsidies contingent on behavior, if the targeted types and the promoted behavior are observed and valued by consumers. Specifically, if consumers regard non-eTRS–PS and eTRS–non-PS sellers as closer substitutes than non-eTRS–PS and eTRS–PS sellers, then by steering eTRS sellers towards offering PS, the targeted incentive lessens competition for non-eTRS sellers, leading to a larger residual demand for them.

Our empirical analyses show that the share of PS listings of eTRS sellers increases from 22% to 43% after the policy change. The share of PS listings of non-targeted (i.e., non-eTRS) sellers also increases, from 10% to 17%. These findings suggest that the targeted incentive has spillover effects on the non-targeted group, whose quality provision also increases as a result.

To study how targeting size affects the impacts of targeted incentives, we leverage more than 400 product subcategories on eBay, which vary in the size of the targeted group (i.e., eTRS sellers) relative to the non-targeted group (i.e., non-eTRS sellers) prior to the policy change. We adopt a continuous difference-in-differences approach that exploits this cross-market variation and compare the temporal changes in outcomes in "more affected" markets where a larger share of listings comes from the targeted eTRS sellers before the policy change with the temporal changes in outcomes in "less affected" markets.[4] Additionally, we estimate the PS premium by comparing almost iden-

[2] For example, if consumers view eTRS and PS as similarly good signals for seller quality, the substitutability between eTRS–non-PS and non-eTRS–PS sellers will be higher because both have one quality signal, and a price increase among the former group will *increase* demand for the latter group.

[3] Examples of targeted subsidies are government subsidies for small firms and tax cuts for a sector of the economy.

[4] We define the policy exposure measure this way because both PS and the incentives are offered at the listing level. Using the share of eTRS sellers as the policy exposure measure does not qualitatively change the results.

tical listings that only differ in whether PS is offered and study how offering PS affects the sales probability of the product conditional on price.

We find that in more affected markets, there is a larger supply reallocation from eTRS–non-PS to eTRS–PS sellers, where supply is proxied by the number of listings (regardless of whether they sell). As expected, the equilibrium price among eTRS–non-PS sellers increases more in more affected markets. Given this result, our model predicts a larger demand increase among non-eTRS–PS sellers in more affected markets, and if in addition the PS premium is also larger in these markets, we should also expect non-eTRS sellers to offer more PS. The empirical results corroborate these predictions: the PS premium for non-eTRS sellers increases more in more affected markets, and consistently, we find a larger supply reallocation from non-eTRS–non-PS to non-eTRS–PS sellers in these markets. Additionally, the changes in equilibrium price and quantity sold are consistent with a larger increase in demand for non-eTRS sellers in more affected markets. Lastly, we find no change in seller exit or quality provision beyond offering PS from non-eTRS sellers.

Our findings have managerial implications for the use of targeted incentives on digital platforms. In terms of the effectiveness of nudging behavior, targeted incentives can affect the quality provision from both targeted and non-targeted sellers through market forces. A larger targeting size increases the spillover effect on the non-targeted group, causing them to improve quality further. Additionally, targeted incentives may increase the demand for non-targeted sellers, especially those who offer the promoted behavior, provided that the targeted type and the promoted behavior are quality signals valued by consumers. This finding suggests that offering targeted incentives does not necessarily give an unfair advantage to a selected group of sellers—a concern that some platform designers may have.

1.1 Related Literature

Our paper contributes to three strands of literature. First, it contributes to the literature that studies the effect of supply-side incentives. Researchers have shown that financial incentives effectively encourage user contribution in posting reviews (e.g., Cabral and Li (2015), Fradkin et al. (2018), Sun et al. (2017), and Burtch et al. (2018)), in knowledge sharing (e.g., Kuang et al. (2019)), and in participating in open source software communities (e.g., Roberts et al. (2006)). Besides financial incentives, social incentives can also be effective motivators for user behaviors. For example, users post more reviews after being informed of the social norm (e.g., Chen et al. (2010) and Burtch et al. (2018)), and they generate more content when the size of their social network is larger (e.g., Zhang and Zhu (2011) and Shriver et al. (2013)). Ahn et al. (2011) and Kumar et al. (2014) build structural models of content generation in which the social network enters users' utility of content creation. Our paper contributes to this literature by studying the consequences of targeting incentives to only a subset of sellers.[5]

Second, our paper is related to a large strand of literature that analyzes the effects of targeted firm subsidies (e.g., Rotemberg (2019)). These subsidies can be in the form of directed lending (e.g., Banerjee and Duflo (2014)), capital subsidies (e.g., Bergström (2000)), access to finance (e.g., Krishnan et al. (2015)), export facilitation (e.g., Hui (2019)), and procurement subsidies (e.g., Marion (2007)), among others. In this paper, we show that unlike targeted subsidies, targeted incentives may expand demand for the non-targeted sellers.

Lastly, our paper contributes to the literature that studies the demand expansion effect alongside increases in competition on platforms. In this literature, the increase in competition typically comes from new sellers' entry to the platform. For example, Cennamo et al. (2016) show that in the home-video-game market, competition can induce firms to create new product niches, expanding the product market for future consumers. Li and Agarwal (2017) show that the integration of Instagram on Facebook benefits large third-party applications on Facebook due to a larger customer base, while hurting the small ones due to higher competition. Similarly, Reshef (2019) documents that the entry of new restaurants on a food delivery platform increases the performance of high-quality incumbent businesses, because consumers have more options to choose from on the platform, while hurting low-quality incumbent businesses. Furthermore, Cao et al. (2018) show that entrants expand market demand for incumbents due to the network effect in the bike-sharing industry in China. Lastly, Shen and Xiao (2014) and Yang (2019) study the learning effect of observing competitors' entry in the fast food industry. In this paper, we also show a demand expansion effect, but which is due to targeted incentives and does not involve entry of new market participants, network effects, or learning.

2 Background and Data

The eBay Top Rated Seller (eTRS) is eBay's flagship certification program to reduce buyers' asymmetric information about seller quality. Sellers are evaluated on the 20th of each month and are eTRS if they pass a set of requirements, which are based on past sales (at least 100 items and

[5] Several papers have studied the effect of consumer coupons on non-promoted products, including Bawa and Shoemaker (1987), Venkatesan and Farris (2012), and Sahni et al. (2017).

$3,000 in sales in the previous year) and past quality (98% or higher positive consumer feedback, less than 0.5% of 1s and 2s on the 5-point Detailed Seller Ratings, and less than 0.5% or two buyer claims) as of 2012. Sellers who obtain the eTRS status enjoy several benefits. First, certified sellers get an eTRS badge, which is prominently displayed on every listing of the seller. Second, eTRS sellers get a 20% discount on the final value fee (i.e., eBay's commission fee). Lastly, the listings from eTRS sellers can appear higher on the product search results page.

On February 28, 2012, eBay initiated a platform-wide campaign to encourage sellers to offer fast shipping and generous return policies, the campaign being in effect from March 1, 2012, to May 31, 2012. During this three-month period, eBay offered eTRS sellers an additional 5% discount of the final value fee, besides the usual 20% discount, when an eTRS seller offered a 14-day (or more) money-back return policy *and* same-day or one-day handling for a listing (Premium Service, or PS henceforth). Note that the benefit is listing-specific: eTRS sellers only get the policy benefit for listings for which they offer the above-mentioned service. Additionally, the incentive only applies to eTRS sellers, and not to non-eTRS sellers. To offer PS, sellers choose the qualified handling days and return option when they list an item, and eBay automatically detects PS listings and applies the commission discount when these listings sell.

Consumers can learn about whether PS is offered by reading the information of the return option and handling time on the item listing page. The return specifics are shown underneath the price and shipping information, as shown in Figure A1. To get the information on handling days, consumers need to scroll down the listing page and click on the "Shipping and payments" tab. The seller-specified handling time is listed at the bottom of this section.

To study the targeted incentive, we use internal data from eBay, which includes detailed listing attributes, transaction outcomes, product characteristics, buyer history, seller history, and feedback and reputation. Our main dataset covers the period from 12 weeks before and 12 weeks after the date of the policy change, March 1, 2012. We choose the 24-week window because afterwards eBay implemented another policy to further encourage eTRS sellers to offer PS.[6] To study the effects of this policy change on commercial sellers, we focus our attention on sellers who had sold at least $5,000 in the year before the beginning of our sample period.

A key feature of the data that enables our identification is the large number of subcategories on eBay. There are more than 400 subcategories, such as "Household Supplies & Cleaning", "DVDs &

[6] Specifically, eTRS sellers would lose the 20% discount and higher position on the search results page for listings without PS.

Blu-ray Discs", "Men's Clothing", and "Cell Phones & Smartphones". Similar to Hui et al. (2017), we treat each subcategory as a separate market.[7] Additionally, our data allows us to directly observe the quality provision in terms of the promoted behavior: for every listing, we observe whether PS is offered. This direct, listing-level measurement contrasts with previous literature which proxies quality provision with feedback and assumes that consumers' rating behavior is unaffected by the policy change. Henceforth, we refer to the listings that offer PS as "PS listings", and to the other listings as "non-PS listings".

In Figure 1, we plot the time series of share of PS listings in the 12 weeks before and 12 weeks after the policy change. On the X-axis, "0" refers to the policy week and all the other weeks are normalized to this week. We can make two main observations based on this graph. First, the share of PS listings is consistently higher from eTRS sellers than from non-eTRS sellers, consistent with the fact that eTRS sellers are of higher quality on average. Second, immediately after the introduction of the incentives targeted to eTRS sellers, both eTRS and non-eTRS sellers increase their PS offerings in their listings. The fact that non-eTRS sellers are also more likely to offer PS after the policy change suggests that the targeted incentive has spillover effects on the non-targeted group.

We report the summary statistics of our sample in Table 1. The sample we use for our main analyses contains data from the eight weeks before and the eight weeks after the policy change. We chose this sample duration because we use data from the first four weeks (i.e., from 12 weeks before to 9 weeks before the policy change) to construct the policy exposure measure, which is defined in the next section. We therefore exclude this period to remove endogeneity concerns in the main analyses. We have repeated the analyses using the entire sample and the qualitative results do not change. All the values in Table 1 except for Share of PS Listings are normalized with respect to the value of eTRS sellers in the first week of our sample.

In the eight weeks after the policy change, the number of listings decreases by about 8%, i.e., $(1.38 - 1.5)/1.5$, for eTRS sellers, and by about 3% for non-eTRS sellers. Due to seasonality in sales, we should be cautious of interpreting these as supply decreases due to the policy change. However, assuming seasonality affects sellers in similar ways, this result suggests that non-eTRS sellers experience a supply increase because of the policy change. In terms of offering the promoted

[7] Note that eBay has a finer catalog, namely "leaf categories". For example, a leaf category within "Cell Phones & Smartphones" may identify a phone brand. However, this catalog can be too fine for defining markets, as it is unclear whether similar products such as Samsung phones and Google Pixel phones, for example, belong to two separate markets.

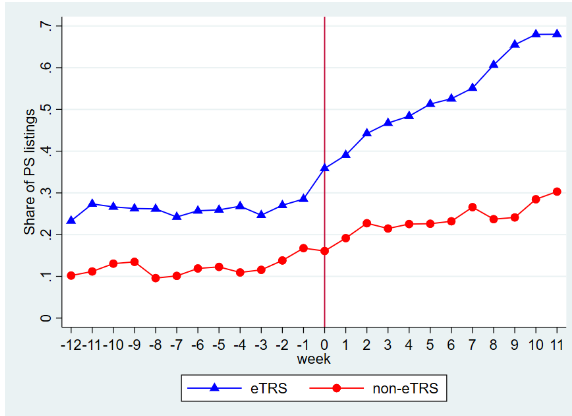

Figure 1: Share of Listings with Premium Service, eTRS vs. Non-eTRS

Notes: Week 0 is the policy week of March 1, 2012. The data consists of sellers who had sold at least $5,000 in the year before the beginning of our sample period.

Table 1: Summary Statistics

	eTRS Sellers		Non-eTRS Sellers	
	8 Weeks Before	8 Weeks After	8 Weeks Before	8 Weeks After
Num. of Listings	1.50	1.38	0.66	0.64
Share of PS Listings	0.22	0.43	0.10	0.17
Quantity Sold	1.11	1.03	0.41	0.41
Revenue	1.22	1.23	0.56	0.57

Notes: Averages for eTRS and non-eTRS sellers. One observation is a seller–market pair. All values except for Share of PS Listings are normalized with respect to the value of eTRS sellers in the first week of our sample.

behavior, we see a 96% increase among eTRS sellers and a 70% increase among non-eTRS sellers, consistent with the patterns in Figure 1. Next, eTRS sellers experience an about 7% drop in their sales volume in the eight weeks after the policy change, but non-eTRS sellers do not experience a decrease in their sales volume. Lastly, we see small changes in revenue for both eTRS and non-eTRS sellers. Given these changes, this result implies a price increase among eTRS sellers and little price change among non-eTRS sellers, suggesting that sellers do not engage in price competition after the policy change. In Section 4, we will estimate the effect of the policy change on these outcome variables using our identification strategy to account for seasonality, and will also interpret the results through the lens of our conceptual framework.

3 Conceptual Framework and Empirical Strategy

3.1 Conceptual Framework

We discuss the intuition of our conceptual framework to guide our empirical analyses, and leave further details for the appendix. In our framework, a seller's targeted type (i.e., eTRS status) is fixed and the seller's quality provision (i.e., offering PS or not) is endogenous.[8] Consumers differentially value seller quality but do not observe it. Instead, they rely on signals of seller quality, namely sellers' eTRS status and whether they offer PS, to make purchase decisions among four types of sellers—eTRS–PS, non-eTRS–PS, eTRS–non-PS, and non-eTRS–non-PS.

For clarity of illustration, we develop the framework in three steps. First, we provide a benchmark model of targeted subsidies. Next, we modify it to study targeted incentives, in which targeted sellers receive compensation *only if* they adopt the promoted behavior. Lastly, we further modify the model to study how the size of the targeted population affects the impact of targeted incentives on quality provision and market outcomes.

Starting with the first step, suppose eBay offers a targeted subsidy to eTRS sellers. This decreases their production cost and therefore increases their supply (i.e., an outward shift of their supply curve). Demand for non-eTRS sellers decreases (i.e., an inward shift of their demand curve) because the price of their subsidized competitors decreases. Therefore, targeted subsidies unambiguously reduce demand for non-targeted sellers.

In the second step, consider the case of a targeted incentive and suppose that eBay compensates eTRS sellers *only if* they adopt PS. This targeted incentive immediately encourages eTRS sellers to adopt PS, leading to an increase in supply among eTRS–PS sellers and a decrease in supply among eTRS–non-PS sellers. This reallocation in eTRS sellers' supply decreases the equilibrium price among eTRS–PS sellers but *increases* the equilibrium price among eTRS–non-PS sellers.[9] How does demand for non-eTRS sellers respond? There can be two cases. First, if consumers treat the eTRS status and PS as two independent seller characteristics, as in the standard choice model, then non-eTRS–PS sellers are more substitutable with eTRS–PS sellers than with eTRS–non-PS sellers. In this case, the residual demand for non-eTRS–PS sellers will decrease because the price of their closest substitute decreases, leading to lower price and quantity sold. The residual demand for non-eTRS–non-PS sellers will increase, since they are more substitutable with eTRS–

[8] We do not allow for entry and exit in this model, to keep it tractable. Empirically, we find little changes in the exit rate of sellers, in Section 5.2.

[9] It is assumed that the policy does not directly affect the demand curves for eTRS sellers.

non-PS sellers than with eTRS–PS sellers. In the second case, assuming buyers have asymmetric information about seller quality in markets, they could view eTRS and PS as similarly good signals for seller quality. Then the substitutability between the eTRS–non-PS and non-eTRS–PS sellers will be higher because both have one quality signal. In this case, the residual demand for non-eTRS–PS sellers will *increase* because the price of their closest substitute increases, leading to higher price and quantity sold. The residual demand for non-eTRS–non-PS sellers also increases but to a lesser extent, because their substitutability with eTRS–non-PS sellers is smaller than that between non-eTRS–PS and eTRS–non-PS sellers.[10]

To tell apart the two cases, we define "PS premium" as the difference in the sales probability between a product with PS and an otherwise identical product without PS. In the first case discussed above, demand for non-eTRS–PS sellers decreases while demand for non-eTRS–non-PS sellers increases, which leads to a decreasing PS premium among non-eTRS sellers. Given this, non-eTRS sellers will shift their supply from PS to non-PS. In the second case, however, demand for non-eTRS–PS sellers increases while for non-eTRS–non-PS sellers it increases less. This implies an increasing PS premium among non-eTRS sellers, and as a result, non-eTRS sellers will shift their supply from non-PS to PS. In this case, the targeted incentive creates a positive spillover effect on the non-targeted sellers both in terms of incentivizing their quality provision and increasing their residual demand.

In the last step, we expand this framework to study how the impacts of the targeted incentive on market equilibrium vary across markets that differ in the targeting size. We find that the equilibrium changes in quality provision and market outcomes are more salient in more affected markets. Assuming we are in the second case, targeted incentives will be more effective in markets with a larger targeted population because more non-eTRS will also offer PS. Additionally, in these markets non-eTRS sellers, especially those that offer PS, will benefit from a larger demand increase.

3.2 Empirical Strategy

Recall that our core research question is, How does the impact of targeted incentives vary by the size of the targeted population? To answer this question, the experimental ideal is to randomly assign different targeting sizes across replicas of the market and see how market outcomes differ

[10] This can also be interpreted through a standard vertical differentiation model. Suppose there are three quality levels: high (eTRS–PS), medium (eTRS–non-PS, non-eTRS–PS), and low (non-eTRS–non-PS). When there are more eTRS–PS sellers and fewer eTRS–non-PS sellers, competition becomes less fierce for both types of non-eTRS sellers, but its extent is larger for non-eTRS–PS sellers as they are closer to the eTRS–non-PS sellers on the vertical line.

across versions of the treatment. Motivated by the experimental ideal, we take advantage of more than 400 markets on eBay, which differ in the *ex ante* share of sellers who are eligible for the incentives (i.e., eTRS sellers). Our identification strategy compares temporal changes in outcomes across markets with different target sizes.

To be specific, for each market, we calculate the share of listings from eTRS sellers out of all listings in the first four weeks of our sample period (from Week -12 to Week -9), and use this measure as the policy exposure in each market. Additionally, we use the share of listings from eTRS sellers, instead of the share of eTRS sellers, as the policy exposure measure to capture the fact that the premium service is offered at the listing level, and so is the financial incentive. However, using the share of eTRS sellers as the policy exposure measure does not change the results qualitatively because the two measures are highly correlated.

This *ex ante* policy exposure measure across markets creates a continuum of treatment and control groups at the market level, and allows us to estimate a continuous difference-in-differences (DiD) model that compares temporal changes in the "more affected" and "less affected" markets:

$$ln(Y_{mt}) = \beta Share_m \times Post_t + \eta_m + \xi_t + \epsilon_{mt}, \tag{1}$$

where Y_{mt} are outcome variables in market m and week t; $Share_m$ is the market-specific policy exposure measure defined previously; $Post_t$ is a dummy variable that equals 1 after the policy change; η_m are market fixed effects; ξ_t are week fixed effects; and ϵ_{mt} is the idiosyncratic error term. Our coefficient of interest β measures how the size of the targeted population affects the impact of incentives. Another benefit of our continuous DiD approach, as opposed to an event study approach, is that the ξ_t term allows us to control for common time trends on the platform. Throughout the analysis, the standard errors are clustered at the market level to account for serial correlations and the heteroskedasticity of outcome variables in a given market.

The identification assumption of equation 1 is that the policy exposure measure affects market outcomes only through differences in the policy intensity due to the size of the targeted population. This assumption would be violated if, for example, the markets with a larger share of eTRS listings had a higher level of competition. To account for market-level heterogeneity, we control for market fixed effects, η_m, in the regression. We also control for week dummies, ξ_t, to account for common time trends in outcomes that are the same across all markets. Like any difference-in-differences exercises, our approach cannot control for time-varying, market-specific error terms that could be correlated with the policy exposure variable (e.g., the markets with high policy exposure are those that experience a faster growth rate in the level of competition). We provide evidence in Section 5 that is consistent with the validity of our identification assumption, including a graph on the parallel trends and a leads–and–lags analysis.

For the continuous DiD to work, we also need enough variation in the policy exposure measure. To check this, we plot the distribution of the ex ante share of eTRS listings across more than 400 markets on eBay in Figure 2. We indeed observe a large amount of variation in this measure, with the lowest exposure close to 0 and the highest close to 1. In most markets, eTRS listings account for more than half of all listings.

Based on the conceptual framework, the changes in quality provision across markets should be consistent with the changes in PS premium. For example, if the non-eTRS sellers offer more PS listings in more affected markets, it must be because the PS premium for non-eTRS sellers becomes larger in these markets. To estimate the PS premium, we use the matched listing approach, which is necessary because PS listings and non-PS listings possibly differ in ways that correlate with both demand and the propensity of offering PS. For example, the decision of whether to offer PS may depend on the price and shipping cost of the item, which affect demand. Seller heterogeneity could also confound the estimate if high-quality sellers are more likely to offer the PS service than are lower-quality sellers, and demand is higher for high-quality sellers. Lastly, market conditions change over time and can also affect consumer demand.

To mitigate the above-mentioned concerns, we match the listings in several key components to control for product, seller, and market-level heterogeneity as much as we can. Following Elfenbein et al. (2012)) and Einav et al. (2015), we match the listings based on following variables: seller identity, item listing title, item listing subtitle, item's leaf category on eBay, sales price, and listing start week. In addition, these matched listings must have variation in whether PS is offered. Essentially, matched listings can be considered as identical listings except for whether the PS is offered. As argued by these authors, the matched listings can be thought of as instances where eBay sellers experiment with sales parameters (i.e., in this case, whether to offer PS) to understand consumers' preferences. Having constructed these matched sets of listings, we then exploit the within-set variation in whether PS is offered to identify its effect on the sales probability using the following equation:

$$Success_{ij} = \gamma PS_{ij} + \mu_i + \nu_{ij}, \tag{2}$$

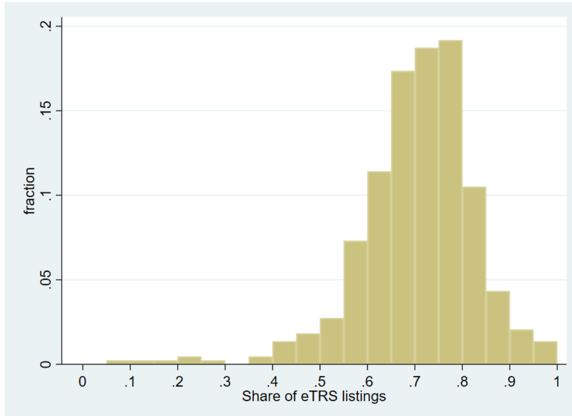

Figure 2: Distribution of Policy Exposure across Markets on eBay

Notes: The figure plots the distribution of policy exposure, i.e., the share of listings from eTRS sellers, across more than 400 markets. The sample contains listings from the 12^{th} week to the 9^{th} week before the policy change.

where $Success_{ij}$ is an indicator for whether listing j within the matched set of listings i results in at least one sale; PS_{ij} is the dummy variable for whether listing j in matched set i offers PS; μ_i is the fixed effects for the matched set of listings; and ν_{ij} is the random error. Our coefficient of interest is γ, which measures consumer demand for PS.

Lastly, to study how the PS premium varies across markets with different policy exposures, we modify equation 2 to

$$Success_{ijt} = \alpha PS_{ij} + \lambda PS_{ij} \times Share_m + \mu_i + \delta_t + \nu_{ijt}, \tag{3}$$

where $Share_m$ is as previously defined and δ represents the weekly time fixed effects. Our coefficient of interest is λ, which captures the difference in consumer demand for PS in more affected markets.

4 Results

We present our empirical results following the order of our theoretical reasoning. First, we show the effect of the targeted incentive on market outcomes and quality provision for eTRS sellers in Section 4.1. Next, we estimate the changes in the PS premium for non-eTRS sellers to study the changes in consumer demand for PS in Section 4.2. Lastly, we investigate the program's effect on non-eTRS sellers in Section 4.3. Overall, the empirical results highlight the findings that the targeted incentive is effective in our setting: in markets with a larger targeted population, non-targeted sellers are more likely to offer PS and experience a larger demand expansion when doing so.

4.1 Policy Effect on Targeted Sellers

We adopt equation 1 to study the policy's effect on targeted sellers and report the estimated β in Table 2. The sample we use is at the market–week level. Starting with eTRS–PS sellers, column (1) reports the estimation results on the differential changes in supply as measured by the log of number of new listings in a week. The estimate is statistically insignificant, indicating that the percentage change in supply is constant across markets with a different share of eTRS listings. Similarly, column (2) shows that the percentage change in equilibrium quantity is also constant across markets. The constant across-market percentage change implies that the *absolute increase* in supply and equilibrium quantity, as a share of total listings, are larger in markets that have a higher share of eTRS listings.[11] These results are consistent with our conceptual framework: in more affected markets, the targeted incentive causes a larger supply increase among eTRS–PS sellers, leading to a larger increase in the equilibrium quantity and a larger decrease in the equilibrium price. Although the values in column (3) are consistent with lower equilibrium prices in more affected markets, this estimate is not statistically significant at the 10% level, which could be due to a secondary demand increase in this market or to the fact that sellers prefer to enjoy a higher profit margin, both causes which are excluded from the model. However, a decrease in price among eTRS–PS sellers is not essential for our main results, as will be explained soon.

For eTRS–non-PS sellers, column (4) shows that in a market with a 10 percentage points (pp) larger share of eTRS listings, the supply decrease is 5.75% larger. This is consistent with the fact that more eTRS sellers start providing PS in markets with a larger targeted population. Similarly, column (5) shows that the equilibrium quantity also decreases more in more affected markets. Lastly, column (6) shows that the equilibrium price increases by 3.07% more in markets that have a 10 pp higher share of eTRS listings. The results on eTRS–non-PS sellers are consistent with a larger supply decrease, and hence a larger price increase in more affected markets. The price result among eTRS–non-PS sellers is critical, because it is a necessary condition for a potential demand

[11] To illustrate the idea for the result on listings, consider two markets with M_1 and M_2 total number of listings. Market 1 has 80% eTRS listings and Market 2 has 60% eTRS listings, i.e., Market 1 has a higher policy exposure. Suppose that prior to the policy change, 50% of the eTRS listings were offering PS in both markets. That is, there were $0.4M_1$ ($= 0.8\times0.5\times M_1$) eTRS–PS listings in Market 1 and $0.3M_2$ ($= 0.6\times0.5\times M_2$) eTRS–PS listings in Market 2. Suppose after the policy change, both markets experienced the same growth rate of PS listings from eTRS sellers, for example, by 40%. This is equivalent to an increase of eTRS–PS listings from $0.4M_1$ to $0.56M_1$ in Market 1 and $0.3M_2$ to $0.42M_2$ in Market 2. Therefore, while the growth rate is the same for both markets, the absolute increase in eTRS–PS listings as a share of total listings in Market 1, i.e., 16%, is larger than that of Market 2, i.e., 12%.

Table 2: Market Outcomes and Quality Provision: Targeted Sellers (i.e., eTRS)

	PS = 1			*PS = 0*		
	(1) ln(num. listings)	(2) ln(quantity sold)	(3) ln(sales price)	(4) ln(num. listings)	(5) ln(quantity sold)	(6) ln(sales price)
Share×Post	0.041	0.119	-0.154	-0.575***	-0.299**	0.307**
	(0.299)	(0.213)	(0.248)	(0.141)	(0.118)	(0.151)
R^2	0.956	0.963	0.800	0.976	0.980	0.869
Observations	10,248	10,248	10,248	10,248	10,248	10,248

Notes: This table shows the results of regressing outcome variables on policy exposure times the post dummy, controlling for market and week fixed effects. Standard errors are clustered at the market level.

*** indicates significance at p = 0.01; ** p = 0.05; * p = 0.1.

increase in the non-eTRS markets, which we find evidence on in Section 4.3.

Lastly, note that columns (1) and (4) collectively show that eTRS increases quality provision in terms of offering PS more in markets with a larger targeted population. In a separate analysis, we perform equation (1) on all eTRS sellers regardless of whether they offer PS and use log of number of PS listings as the outcome variable. The coefficient estimate is close to zero and has a t-value of 0.02. The essentially constant percentage change in PS listings across markets indicates that the *absolute increase* in supply of PS listings is larger in markets with a larger share of eTRS listings.

4.2 PS Premium

Given our results on eTRS sellers' supply reallocation towards offering PS, the next step is to study the changes in PS premium among non-eTRS sellers. There are two cases according to our conceptual framework. If eTRS and PS are two independent characteristics, as in the standard choice model, then there will be no increase in the PS premium for non-eTRS sellers.[12] Alternatively, if consumers view eTRS and PS as similar quality signals for sellers, then we would expect an increase in the PS premium for non-eTRS sellers. This indicates that the PS premium for non-eTRS sellers should be larger in markets with a larger targeted population in the second case.

To estimate the PS premium, we use the matched listings approach as in equation 2 on data from the eight weeks before and the eight weeks after the policy week. Column (1) in Table 3 shows that offering PS increases the sales probability by 3.4 pp for non-eTRS sellers, which is an

[12]This statement is true even though the price change does not vary across markets among eTRS–PS sellers.

Table 3: PS Premium for Non-Targeted Sellers (i.e., non-eTRS)

Dependent Variable: Success			
	(1)	(2)	(3)
	Entire Sample	8 Weeks Before	8 Weeks After
PS	0.034***	0.073	-0.144*
	(0.006)	(0.075)	(0.079)
Share×PS		-0.048	0.242**
		(0.105)	(0.110)
R^2	0.732	0.766	0.706
Observations	9,027	3,584	5,443

Notes: This table shows the results of regressing the dummy of whether a listing sells on the PS dummy and its interaction with policy exposure, controlling for matched listing fixed effects. The matching is based on seller ID, listing title, listing subtitle, leaf category, listing start week, and price. Standard errors are clustered at the matched listing level.

*** indicates significance at p = 0.01; ** p = 0.05; * p = 0.1.

18% increase. Next, we study how the PS premium differs across markets with varying policy exposure using equation 3. We perform this analysis separately for the eight weeks before and the eight weeks after the policy change and report the results in columns (2) and (3). Starting with column (3), we find a larger PS premium for non-eTRS–PS sellers in more affected markets, which indicates a larger demand expansion for these sellers after the policy change. However, the results in column (2) show no significant across-market difference in the PS premium before the policy change. The two columns together show that the across-market pattern exists only *after* the policy change, suggesting that the result is not driven by omitted variables that correlate with both the policy exposure and PS premium across markets.

As a robustness check, we repeat the matched listing analysis using logged sales quantity plus one as the dependent variable, and additionally control for the logarithm of the number of available quantities in the regression. The results, reported in Table A1 in the online appendix, show qualitatively similar results on changes in the PS premium, both overall and across markets.

4.3 Policy Effect on Non-Targeted Sellers

We now study the policy effects among non-eTRS sellers. Given the higher equilibrium price for eTRS–non-PS sellers and consequently larger PS premium for non-eTRS sellers in more affected markets, our conceptual framework predicts that non-eTRS sellers should offer more PS in these markets.

We estimate β in equation 1 and report the results in Table 4. Similar to our analyses in Section 4.1, our data is at the market–week level. Columns (1) and (2) show that among non-eTRS–PS sellers, supply and equilibrium quantity increase more in markets with a higher share of eTRS listings, consistent with a larger PS premium in these markets. Column (3) shows that the equilibrium price is constant across markets, which happens when the increase in demand for non-eTRS–PS sellers is also larger in more affected markets, for otherwise we should see a price decrease. The larger demand increase is consistent with the finding that the price of non-eTRS–PS sellers' closest substitute (eTRS–non-PS sellers) decreases more in more affected markets, as shown in column (6) in Table 2.

Next, we study the changes in equilibrium outcomes among non-eTRS–non-PS sellers. Column (4) and (5) show that supply (not significant) and equilibrium quantity (significant at the 5% level) decrease more in more affected markets. Column (6) shows that the price on average is larger in more affected markets, although this estimate is not statistically significant at the 10% level. While the demand change for these sellers is ambiguous given the insignificant price coefficient, it is clear that demand cannot decrease too much because otherwise one should expect a negative estimate on the changes in equilibrium price. These results are consistent with our conceptual framework.

Lastly, columns (1) and (4) combined show that non-eTRS sellers increase quality provision by offering more PS in markets with a larger targeted population. We have also performed equation (1) on all non-eTRS sellers regardless of whether they offer PS and used log of number of PS listings plus one as the dependent variable. The coefficient estimate is 0.803 and is statistically significant at the 1% level, indicating a larger increase in PS listings of non-eTRS sellers in more affected markets.

There are two main takeaways in this section. First, targeted incentives cause both targeted and non-targeted sellers to adopt the promoted behavior, and sellers do not seem to engage in price competition on a large scale. Second, the fact that non-eTRS sellers do not lower their price and they sell more when adopting the promoted behavior highlights a key insight: demand for non-targeted sellers, especially for those who adopt the promoted behavior, can in fact *increase* even though targeted sellers adopt the promoted behavior to a larger extent. This result critically depends on how consumers value the promoted behavior as a signal of seller quality vis-à-vis the existing certification. If consumers regard PS as a strong quality signal, then non-targeted sellers can attract more demand by adopting the promoted behavior. This finding is in contrast with an

Table 4: Market Outcomes and Quality Provision: Non-Targeted Sellers (i.e., non-eTRS)

	PS = 1			*PS = 0*		
	(1) ln(num. listings)	(2) ln(quantity sold)	(3) ln(sales price)	(4) ln(num. listings)	(5) ln(quantity sold)	(6) ln(sales price)
Share×Post	0.709**	0.610**	0.081	-0.153	-0.275**	0.374
	(0.304)	(0.260)	(0.245)	(0.150)	(0.114)	(0.294)
R^2	0.915	0.900	0.691	0.977	0.982	0.824
Observations	10,248	10,248	10,248	10,248	10,248	10,248

Notes: This table shows the results of regressing outcome variables on policy exposure times the post dummy, controlling for market and week fixed effects. Standard errors are clustered at the market level.

*** indicates significance at p = 0.01; ** p = 0.05; * p = 0.1.

unambiguous decrease in demand for non-targeted sellers in a targeted subsidy model.

5 Robustness

In this section, we start by providing evidence for the identification assumption of our continuous DiD approach by first graphically showing parallel trends across markets with different policy exposure and then modifying equation 1 to perform leads–and–lags analyses. Next, we test our key assumptions in our conceptual framework, and, in particular, whether sellers exit the market or respond to the targeted incentive in ways beyond offering the promoted behavior. Overall, the results suggest that the use of our empirical specification and conceptual framework is justified.

5.1 Parallel Trends Assumption

In a normal DiD exercise, researchers typically plot the time series of the outcome variable in the treatment and control groups and check for parallel pre-policy trends. Since we have many treatment and control groups, we divide them into markets with above- and below-median policy exposure, and check for the parallel trends assumption across the two groups. In Figure 3a, we plot the weekly share of PS listings for the two groups among eTRS sellers. The two series have similar time trends in the 12 weeks before the policy change, which is consistent with the parallel trends assumption. Additionally, the two curves remain parallel after the policy change, suggesting that while eTRS offer more PS after the policy change, as seen in Figure 1, this effect does not vary by

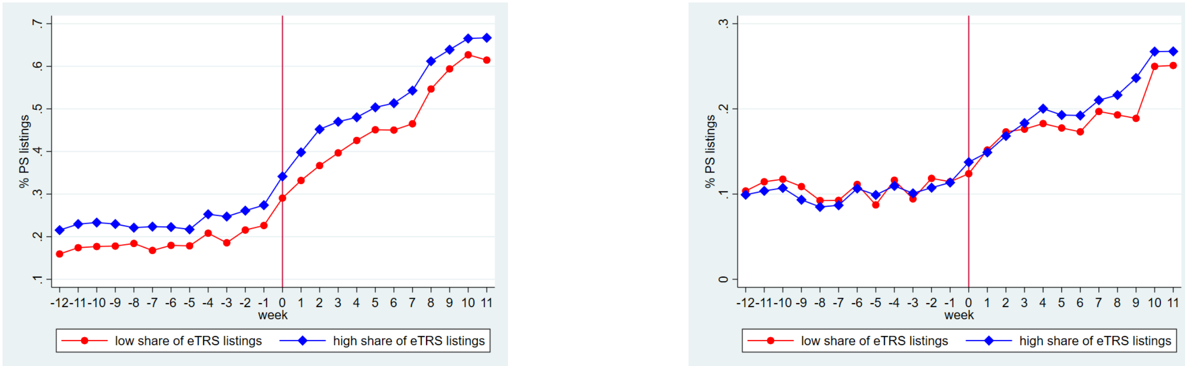

(a) eTRS sellers (b) Non-eTRS sellers

Figure 3: Parallel Trends

Notes: We divide markets on eBay into those with above- and below-median policy exposure, i.e., share of eTRS listings. We plot the average share of PS listings for these two groups, separately for eTRS and non-eTRS sellers. Week 0 is the policy week of March 1, 2012.

policy exposure. In Figure 3b, we plot the same graph for non-eTRS sellers. We observe parallel trends before the policy change, but afterwards non-eTRS sellers are visibly more likely to offer the PS service in their listings in markets with above-median exposure. This result is consistent with the parallel trends assumption in our continuous DiD specification.

Next, we modify equation 1 and perform a leads–and–lags analysis that allows us to incorporate the continuous nature of our treatment to test for parallel trends in key variables before the policy change. The specification is given as

$$ln(Y_{mt}) = \sum_{l=-8}^{7} \beta_l Share_m \times T_t(t = k + l) + \eta_m + \xi_t + \epsilon_{mt}, \tag{4}$$

where k is the week of the policy change, $T_t(t = k + l)$ are dummies for whether the current period is $k+l$, and l represents the l-th lead ($l > 0$) or lag ($l \leq 0$) relative to k. The coefficient of interest is β_l, which measures the difference in average outcome across markets with different policy exposure in each period. Since the policy exposure measure is defined based on the first four weeks of our sample, i.e., the four weeks before $Week = -7$, it is still an ex ante measure of the market and therefore is unlikely to bias the estimation results.

Overall, the leads-and-lags analyses show results consistent with the validity of the identification assumption of equation 1. For example, for eTRS sellers, we do not observe a statistically significant difference in the log of PS listings across markets in almost all the leading periods (except in the seventh week before the policy change). For non-eTRS sellers, we mostly do not observe significant across-market differences until after the policy change. We report the estimation results on more outcome variables in the online appendix.

5.2 Seller Exit and Quality Provision Beyond Offering PS

In our conceptual framework, we make the simplification assumption that sellers cannot exit the market. In reality, however, when sellers' profit becomes too low, they may choose to exit the market. If non-eTRS sellers are more likely to exit the market after the policy change, this could also lead to a larger PS premium in more affected markets.

To study seller exit, we define a proxy for seller engagement in a market. Specifically, we use the logged number of active sellers in the market, and a seller is considered active if she has at least one active listing in that week. We estimate equation 1 using the same sample as in our main analysis and report the results in column (1) in Table 5. Starting with non-eTRS sellers in panel B, we do not find evidence that they exit more in more affected markets. For eTRS sellers, however, panel A shows that they tend to engage less (or exit more) in markets with a higher policy exposure, although the results are only marginally significant. These results suggest that sellers, especially non-eTRS sellers, do not exit the market at large.

Another simplification assumption is that sellers cannot improve their quality provision in ways other than providing PS. However, if in reality sellers also change their quality provision beyond offering PS, then we would overestimate the PS premium and wrongly attribute the demand expansion for non-eTRS sellers to adopting the promoted behavior. To mitigate this concern, we use equation 1 to test whether sellers have changed their quality provision as measured by the number of photos in a listing, whether they include subtitles in a listing, title length, whether the title is in bold font, whether the listing is available for international shipping, and whether they allow for a greater-than-14-day return period (PS only requires a 14-day return). The results are reported in columns (2)–(7) in Table 5. Across specifications, we do not find statistically significant estimates for either eTRS or non-eTRS sellers. These estimates suggest that sellers mainly respond to the targeted incentive by adopting the promoted behavior, namely offering PS in listings, but do not seem to change their behaviors beyond that.

To summarize, in this section we find evidence consistent with our empirical specification and conceptual framework. These results add confidence to our earlier findings: when the targeted incentive is offered to more sellers in a market, more non-targeted sellers will respond by offering the promoted behavior, and therefore the incentives will be more effective in this market. Additionally,

Table 5: Seller Exit and Other Service Provision

	(1) log(Num. sellers)	(2) Num. Photos	(3) Share Subtitles	(4) Title Length	(5) Share Bold Titles	(6) Share Intl Available	(7) Share >14 day return
Panel A: eTRS							
Share×Post	-0.100*	-0.040	-0.030	-0.657	0.002	-0.083	-0.057
	(0.055)	(0.113)	(0.023)	(0.671)	(0.001)	(0.074)	(0.036)
R^2	0.997	0.872	0.695	0.778	0.620	0.828	0.805
Observations	6,816	6,816	6,816	6,816	6,816	6,816	6,816
Panel B: Non-eTRS							
	(1) log(Num. sellers)	(2) Num. Photos	(3) Share Subtitles	(4) Title Length	(5) Share Bold Titles	(6) Share Intl Available	(7) Share >14 day return
Share×Post	-0.003	0.101	0.020	0.332	0.001	0.023	-0.024
	(0.067)	(0.128)	(0.028)	(0.307)	(0.001)	(0.058)	(0.027)
R^2	0.995	0.927	0.556	0.660	0.780	0.798	0.728
Observations	6,816	6,816	6,816	6,816	6,816	6,816	6,816

Notes: This table shows the results of regressing outcome variables on policy exposure times the post dummy, controlling for market and week fixed effects. Standard errors are clustered at the market level.
*** indicates significance at p = 0.01; ** p = 0.05; * p = 0.1.

in markets with a larger targeted population, non-targeted sellers experience a larger demand expansion, especially when they offer the promoted behavior. The key mechanism, as we argue, is that consumers use and value both the targeted type and the promoted behavior as quality signals for sellers.

6 Conclusion

In this paper, we study the effects of a targeted incentive. We leverage a platform-wide campaign on eBay that targets financial incentives to platform-certified sellers who provide fast shipping and generous return policies. We find that besides the targeted sellers, non-targeted sellers are also more likely to adopt the promoted behavior. Importantly, when the targeted incentive is offered to more sellers in a market, non-targeted sellers will also respond by adopting the promoted behavior to a larger extent, and therefore the incentive will be more effective in this market. Additionally, non-targeted sellers experience a larger demand expansion in markets with a larger targeted population, especially when they adopt the promoted behavior. A key determinant of demand expansion is how consumers value the targeted type and the promoted behavior as quality signals for sellers.

One limitation of our work is that in our setting the targeted sellers are always platform-certified, who are of higher quality and are essentially market leaders. This fact could make the spillover effect more salient, because an average seller may closely observe the market leaders when they sell. We should keep this caveat in mind when interpreting the results: if an incentive is targeted to non-certified sellers instead, we may not observe a spillover effect of a similar magnitude on certified sellers. Additionally, we may not observe a demand expansion effect for certified sellers, because they already have the quality certification from the platform. Another limitation is that we cannot study the effects of targeted incentives in the long run, allowing for entry and exit of sellers, because of eBay's subsequent policy changes.

Our findings provide managerial implications for digital platforms that use or consider using targeted incentives. To determine the optimal targeting size, a platform may want to estimate both targeted and non-targeted sellers' elasticity of adopting the promoted behavior with respect to the incentive, using methods such as the one in this paper or field experiments. Based on these estimates, the platform can determine the optimal targeting size given its valuation on the promoted behavior and the cost of giving incentives. Additionally, to take advantage of the potential demand expansion effect for non-targeted sellers, a platform may want to make the promoted behavior a salient quality signal to consumers.

References

Ahn, D.-Y., J. A. Duan, and C. F. Mela (2011). An equilibrium model of user generated content. *Available at SSRN 1957989*.

Banerjee, A. V. and E. Duflo (2014). Do firms want to borrow more? testing credit constraints using a directed lending program. *Review of Economic Studies 81*(2), 572–607.

Bawa, K. and R. W. Shoemaker (1987). The coupon-prone consumer: some findings based on purchase behavior across product classes. *Journal of marketing 51*(4), 99–110.

Bergström, F. (2000). Capital subsidies and the performance of firms. *Small business economics 14*(3), 183–193.

Burtch, G., Y. Hong, R. Bapna, and V. Griskevicius (2018). Stimulating online reviews by combining financial incentives and social norms. *Management Science 64*(5), 2065–2082.

Cabral, L. and L. Li (2015). A dollar for your thoughts: Feedback-conditional rebates on ebay. *Management Science 61*(9), 2052–2063.

Cao, G., G. Z. Jin, X. Weng, and L.-A. Zhou (2018). Market expanding or market stealing? competition with network effects in bikesharing. Technical report, National Bureau of Economic Research.

Cennamo, C., Y. Gu, and F. Zhou (2016). Value co-creation and capture in the creative industry: The us home video game market. Technical report, Working paper.

Chen, Y., F. M. Harper, J. Konstan, and S. X. Li (2010). Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review 100*(4), 1358–98.

Einav, L., T. Kuchler, J. Levin, and N. Sundaresan (2015). Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics 7*(2), 215–247.

Elfenbein, D. W., R. Fisman, and B. McManus (2012). Charity as a substitute for reputation: Evidence from an online marketplace. *Review of Economic Studies 79*(4), 1441–1468.

Fradkin, A., E. Grewal, and D. Holtz (2018). The determinants of online review informativeness: Evidence from field experiments on airbnb. *Available at SSRN 2939064*.

Hui, X. (2019). Facilitating inclusive global trade: Evidence from a field experiment. *Management Science*.

Hui, X., M. Saeedi, G. Spagnolo, and S. Tadelis (2017). Certification, reputation and entry: An empirical analysis. *Unpublished Manuscript*.

Krishnan, K., D. K. Nandy, and M. Puri (2015). Does financing spur small business productivity? evidence from a natural experiment. *The Review of Financial Studies 28*(6), 1768–1809.

Kuang, L., N. Huang, Y. Hong, and Z. Yan (2019). Spillover effects of financial incentives on non-incentivized user engagement: Evidence from an online knowledge exchange platform. *Journal of Management Information Systems 36*(1), 289–320.

Kumar, V., B. Sun, and K. Srinivasan (2014). Why do consumers contribute to connected goods? *A Dynamic Model of Competition and Cooperation Working Paper*.

Li, Z. and A. Agarwal (2017). Platform integration and demand spillovers in complementary markets: Evidence from facebooks integration of instagram. *Management Science 63*(10), 3438–3458.

Marion, J. (2007). Are bid preferences benign? the effect of small business subsidies in highway procurement auctions. *Journal of Public Economics 91*(7-8), 1591–1624.

Reshef, O. (2019). Smaller slices of a growing pie: The effects of entry in platform markets.

Roberts, J. A., I.-H. Hann, and S. A. Slaughter (2006). Understanding the motivations, participation, and performance of open source software developers: A longitudinal study of the apache projects. *Management science 52*(7), 984–999.

Rotemberg, M. (2019). Equilibrium effects of firm subsidies. *American Economic Review 109*(10), 3475–3513.

Sahni, N. S., D. Zou, and P. K. Chintagunta (2017). Do targeted discount offers serve as advertising? evidence from 70 field experiments. *Management Science 63*(8), 2688–2705.

Shen, Q. and P. Xiao (2014). Mcdonald's and kfc in china: Competitors or companions? *Marketing Science 33*(2), 287–307.

Shriver, S. K., H. S. Nair, and R. Hofstetter (2013). Social ties and user-generated content: Evidence from an online social network. *Management Science 59*(6), 1425–1443.

Sun, Y., X. Dong, and S. McIntyre (2017). Motivation of user-generated content: Social connectedness moderates the effects of monetary rewards. *Marketing Science 36*(3), 329–337.

Venkatesan, R. and P. W. Farris (2012). Measuring and managing returns from retailer-customized coupon campaigns. *Journal of marketing 76*(1), 76–94.

Yang, N. (2019). Learning in retail entry. *International Journal of Research in Marketing*.

Zhang, X. M. and F. Zhu (2011). Group size and incentives to contribute: A natural experiment at chinese wikipedia. *American Economic Review 101*(4), 1601–15.

Online Appendix

A Listing Screenshot

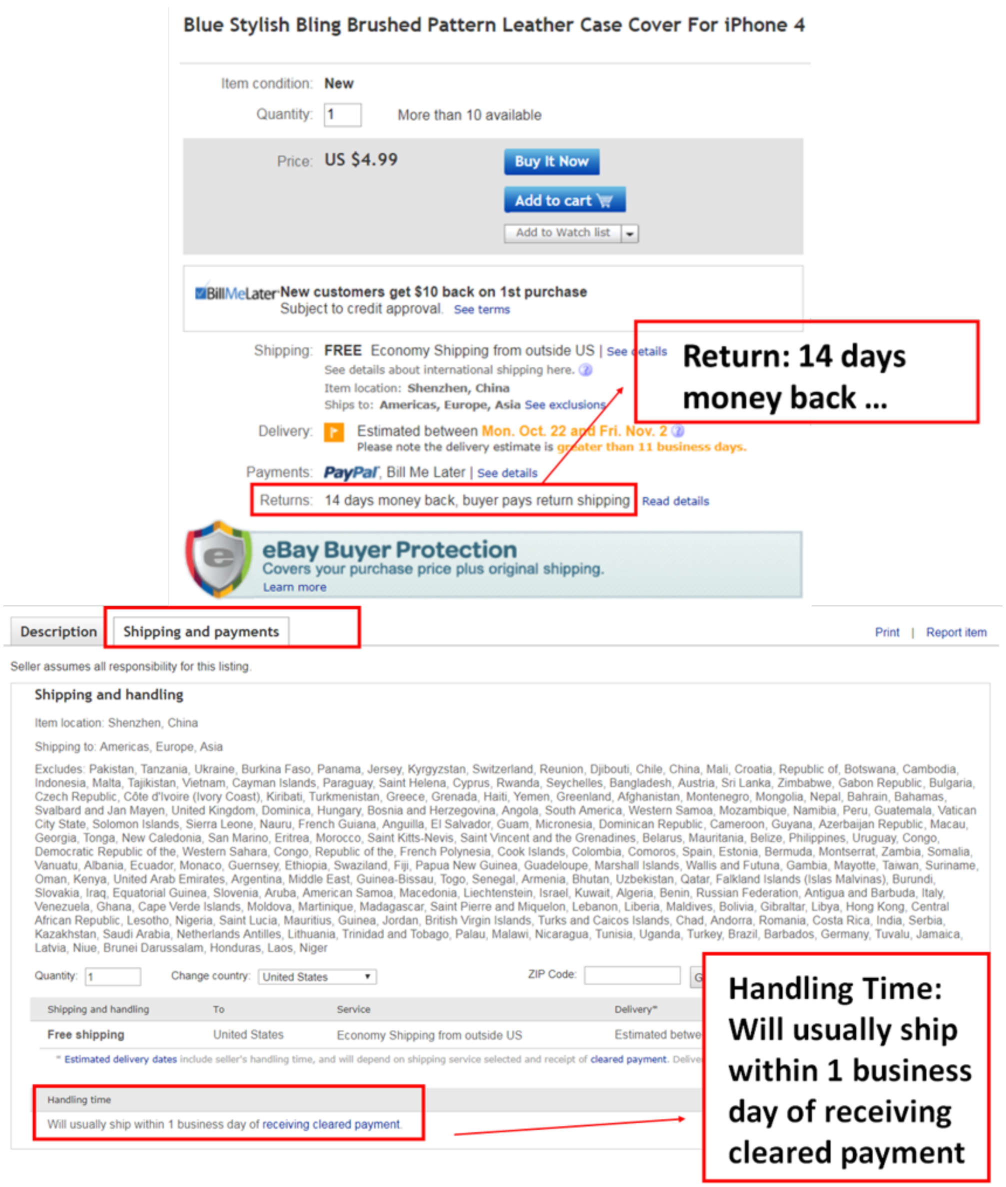

Figure A1: Information on Return Specifics and Handling Time

Figure A2: Targeted Subsidy

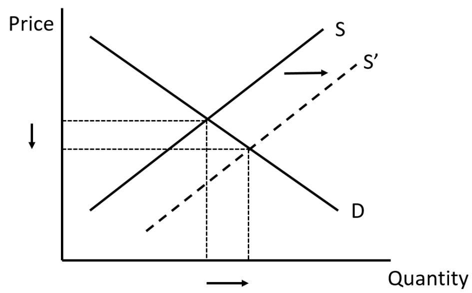

(a) eTRS

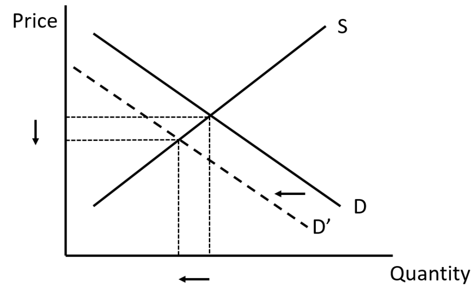

(b) non-eTRS

B Elaboration on the Conceptual Framework

In this section, we elaborate on our theoretical framework. Since the framework is based on the standard demand and supply framework, we show the essential forces using graphs. We follow the same logic as seen in Section 3.1: we start by offering a benchmark model of targeted subsidies, next, we modify the benchmark model to study targeted incentives, and finally, we further modify the model to study the effect of the size of the targeted population on market outcomes and quality provision. To proceed to step one, we need to make the following assumptions:

Market: There are many sellers and many consumers. Sellers differ in their quality and cost of providing PS. There is no entry and exit of sellers. Consumers differ in their reservation prices and valuation for quality.

Information: Buyers do not directly observe sellers' quality but see whether a seller has the eTRS certification. A seller is certified if her quality level is above a quality threshold set by eBay. Since the certification is determined based on a seller's performance in the past year, it is fixed in the short run.

It is useful to start the model construction by considering the benchmark case of targeted subsidies. Examples of these incentives include government subsidies for small firms and tax cuts for one sector of the economy. We plot the benchmark model in Figure A2. As illustrated in Figure A2a, a subsidy to firms reduces their cost of production, leading to an outward shift in supply and a lower price. Demand for the non-targeted firms in Figure A2b shifts inward because the price of their substitute decreases, leading to a reduction in price and quantity sold. This simple model highlights the insight that while targeted subsidies can increase the sales of subsidized firms, this comes at the cost of lower revenue for the non-targeted firms because the targeted subsidies *decrease* their residual demand.

Does this negative spillover result necessarily hold for targeted incentives? The answer is, no, if both the targeted characteristics and the promoted behavior are observed and valued by consumers. To see this, let us first define sellers' choice variable:

Quality Choice: Sellers choose whether to offer PS in their listings, where the cost of providing PS can vary across sellers. Besides the eTRS certification, buyers also observe whether PS is offered in a listing. Therefore, the following four products are substitutes in the eyes of consumers: PS products by eTRS sellers (eTRS–PS), non-PS products by eTRS sellers (eTRS–non-PS), PS products by non-eTRS sellers (non-eTRS–PS), and non-PS products by non-eTRS sellers (non-eTRS–non-PS).

In Figure A3, we illustrate the changes in market equilibrium after the implementation of an incentive that is targeted to eTRS sellers who provide PS. Figure A3a plots the immediate policy impact among eTRS–PS sellers, which is essentially the same as in Figure A2a: the financial incentive causes an outward shift in supply from S to S', reducing the equilibrium price and increasing the equilibrium quantity for them. Because the total number of eTRS sellers is fixed in the short run, the supply curve for the eTRS–non-PS sellers shifts inward, causing an increase in the equilibrium price and a decrease in the equilibrium quantity (Figure A3b). Note that the policy does not have a first-order effect on the demand curves for eTRS sellers. The quality provision of eTRS sellers increases as more of them offer PS afterwards.

Next, we study the changes in market equilibrium for non-eTRS sellers. Unlike in the benchmark model where demand decreases unambiguously (Figure A2b), in this case demand for non-eTRS–PS and non-eTRS–non-PS sellers could either increase or decrease, as plotted in Figures A3c and A3d. The reason is that while the price of one substitute decreases (the eTRS–PS product), the price of another substitute increases (eTRS–Non-PS). Therefore, the changes in demand for non-eTRS sellers depend on the substitution pattern across the four products including two important cases. Consider demand among non-eTRS–PS sellers: In the first case, consumers can treat eTRS and PS as two independent seller characteristics as in the standard choice model; for example, if consumers value PS very much because they highly value fast shipping, then the substitutability between non-eTRS–PS and eTRS–PS sellers will be higher than the substitutability between non-eTRS–PS and the eTRS–non-PS sellers. This means that all else equal, the price increase among the eTRS–PS sellers has a larger impact on demand for non-eTRS–PS sellers, which leads to a net inward shift in demand for them. Since non-eTRS–non-PS sellers are more substitutable with

Figure A3: Targeted Incentive

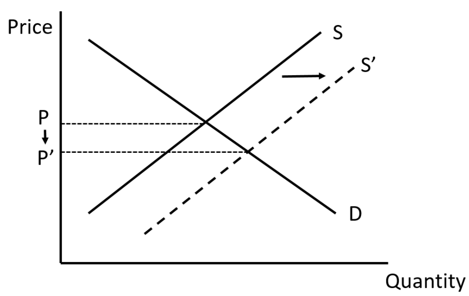

(a) eTRS–PS

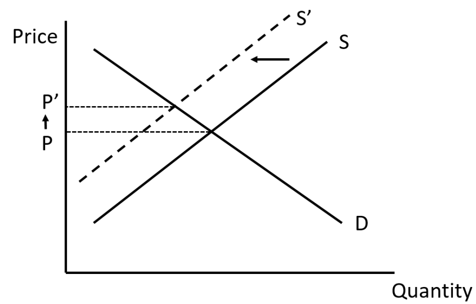

(b) eTRS–non-PS

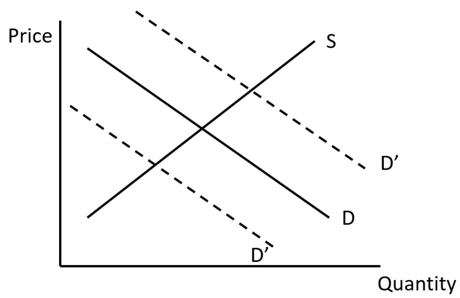

(c) non-eTRS–PS

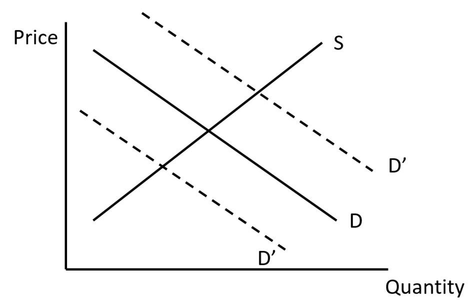

(d) non-eTRS–non-PS

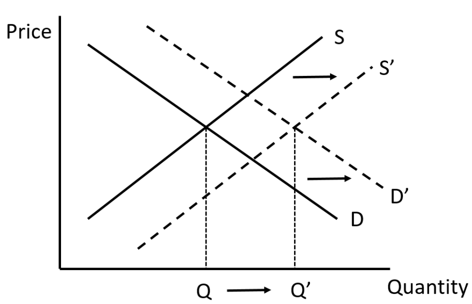

(e) noneTRS–PS

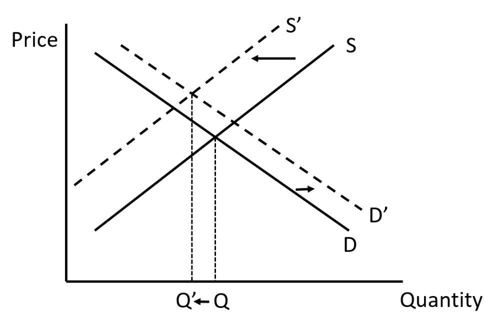

(f) non-eTRS–non-PS

eTRS–non-PS sellers than with eTRS–PS sellers, their residual demand will shift outward. In the second case, if consumers' preference is such that they value the eTRS and the PS signals similarly because they think both are equally good signals for seller quality, then the substitutability between the eTRS–non-PS and the focal products will be higher because both products have one quality signal. In this case, the price increase among eTRS–non-PS sellers has a larger effect on demand for non-eTRS–PS sellers, leading to an outward shift in demand for them. The residual demand for non-eTRS–non-PS sellers also increases but less than it does for non-eTRS–PS sellers. This is because the substitutability of non-eTRS–non-PS sellers with eTRS–non-PS sellers is smaller compared with that between non-eTRS–PS and eTRS–non-PS sellers. An alternative way of thinking of this is that non-eTRS–non-PS sellers are farther away from eTRS–non-PS sellers on the vertical quality line, and therefore the impact of the price increase on demand for these sellers is smaller than it is for non-eTRS–PS sellers.

To tell the two cases apart, note that in the first case, the difference in the residual demand between non-eTRS–PS and non-eTRS–non-PS is smaller, and therefore consumer demand for PS decreases among eTRS sellers. This means that the PS premium measured as the difference in the probability of sales conditional on price decreases. However, in the second case, this difference increases and hence the PS premium also increases.

To assess how supply for non-eTRS sellers would change in response to changes in demand, note that the supply responses must be in opposite directions because the number of non-eTRS sellers is fixed in the short run. Given this, it suffices to study how supply changes among non-eTRS–PS sellers. It will increase (resp., decrease) if the PS premium increases (resp., decreases). As a result, the changes in supply among non-eTRS sellers, and therefore quality provision, could go in either direction depending on the demand changes for them. For illustration purposes, let us assume demand changes as in the second case, which is consistent with our empirical findings: demand increases among non-eTRS–PS sellers and increases slightly among non-eTRS–non-PS sellers. These demand changes imply that offering PS becomes relatively more attractive for non-eTRS sellers in the sense of higher PS premium, leading to an increase in the supply curve among non-eTRS–PS sellers and a decrease in the supply curve among non-eTRS–non-PS sellers. As shown in Figure A3e, an increase in both the demand and the supply curves among non-eTRS–PS sellers increases the equilibrium quantity, and the change in the equilibrium price depends on the relative strength of demand and supply increase. Furthermore, a small increase in demand and a decrease in supply among non-eTRS–non-PS sellers decrease the equilibrium quantity and increase

Figure A4: Targeted Incentive: Different Target Sizes

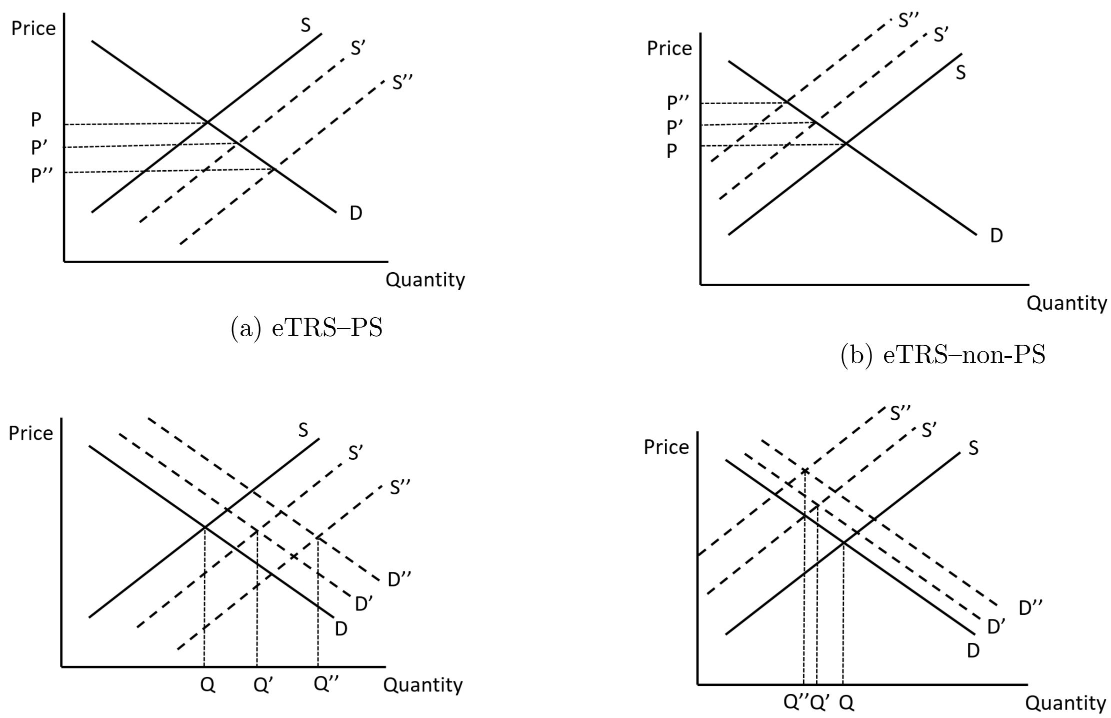

(a) eTRS–PS

(b) eTRS–non-PS

(c) non-eTRS–PS

(d) non-eTRS–non-PS

the equilibrium price (Figure A3f). Lastly, note that the quality provision of non-eTRS sellers increases in this case as more of them offer PS afterwards.

Having illustrated two intermediate models, we move to the model that addresses our core research question on designing targeted incentives, which is on how many people to target. To study this, we compare two markets that are differentially affected by the targeted incentives in terms of how many sellers are eligible for the incentive (eTRS sellers), i.e., how many sellers will receive the financial incentive when they provide PS. We graphically show the comparative statics in Figure A4, where market 2 is assumed to be more affected than market 1. Essentially, the model predicts that the equilibrium changes discussed previously will be more salient in market 2 than in market 1. Starting with Figure A4a, we see that after the policy change, the supply curve in market 2 (S") shifts out more than the supply curve in market 1 (S') among eTRS–PS sellers, which implies the equilibrium price decreases even more in market 2 than in market 1 ($P" < P'$). Similarly, Figure A4b shows that among eTRS–non-PS sellers, the supply curve shifts inward more in market 2, causing a larger increase in the equilibrium price in this market than in market 1 ($P" > P'$). Lastly, the quality provision in the form of offering PS from eTRS sellers should increase more in market 2 than in market 1 afterwards.

Next, we study the changes among non-eTRS sellers in Figures A4c and A4d. Again, we assume that demand increases among the non-eTRS–PS sellers and slightly increases among the non-eTRS–non-PS sellers based on empirical findings. Under this assumption, among the non-eTRS–PS (resp., non-eTRS–non-PS) sellers, market 2 will experience a larger outward (resp., inward) shift in the demand curve than in market 1. The changes in demand imply that, fixing the supply curve, the PS premium for non-eTRS sellers is larger in market 2 than in market 1, leading to a larger outward (resp., inward) shift in the supply curve among non-eTRS–PS sellers (resp., non-eTRS–non-PS). This combination of changes implies that the change in the equilibrium quantity is larger in market 2 than in market 1 among non-eTRS–PS sellers ($Q" > Q'$), and is the opposite among non-eTRS–non-PS sellers ($Q" < Q'$). Similar to before, the prices could either go higher or lower among non-eTRS–PS sellers depending on the relative changes in the demand and supply curves. Lastly, non-eTRS sellers increase their quality provision more in the more affected market 2 than in market 1.

We summarize the predictions of our stylized framework. Note that all the predictions below are about across-market comparisons given our research question. In more affected markets, i.e., markets with more eTRS sellers, we expect the following:

1. **Market outcomes among eTRS–PS sellers**: (1) supply shifts outward more, (2) the equilibrium quantity increases more, and (3) the equilibrium price decreases more.

2. **Market outcomes among eTRS–non-PS sellers**: (1) supply shifts inward more, (2) the equilibrium quantity decreases more, and (3) the equilibrium price increases more.

3. **Market outcomes among non-eTRS–PS sellers**: if the PS premium for non-eTRS sellers is larger in more affected markets, then (1) supply shifts outward more and (2) the equilibrium quantity increases more.

4. **Market outcomes among non-eTRS–non-PS sellers**: if the PS premium for non-eTRS sellers is larger in more affected markets, then (1) supply shifts inward more and (2) the equilibrium quantity decreases more.

5. **Quality provision of eTRS**: the increase in the total number of PS listings is larger.

6. **Quality provision of non-eTRS**: if the PS premium for non-eTRS sellers is larger in more affected markets, then the increase in the total number of PS listings is larger.

C PS Premium in Sales Quantity

In this section, we use equation 2 to estimate the PS premium in terms of units of sales in a listing. Since a listing may not sell, we use logged quantity plus one as the outcome variable. We also control for logged available units in a given listing on the right hand side of the equation. Column (1) shows that offering PS increases the sales quantity by 6.8% for non-eTRS sellers. Next, column (2) shows that there is no across-market difference in the PS premium before the policy change. However, column (3) shows that after the implementation of the targeted incentive, the PS premium is larger in markets with a higher policy exposure, which indicates a larger demand for the non-eTRS–PS sellers in more affected markets after the policy change. All the results in Table A1 are qualitatively similar to our previous results in Table 3 in the main text of the paper.

Table A1: Robustness: PS Premium for Non-Targeted Sellers

Dependent Variable: ln(quantity+1)			
	(1)	(2)	(3)
	Entire Sample	8 Weeks Before	8 Weeks After
PS	0.068***	0.125	-0.139
	(0.007)	(0.084)	(0.090)
Share×PS		-0.082	0.288**
		(0.117)	(0.124)
ln(quantity available)	0.189***	0.134***	0.230***
	(0.016)	(0.024)	(0.021)
constant	-0.057***	0.018	-0.110***
	(0.018)	(0.028)	(0.023)
R^2	0.705	0.718	0.697
Observations	9,027	3,584	5,443

Notes: This table shows the results of regressing logged sales quantity plus one on the PS dummy, its interaction with policy exposure, and the log of available quantity in a listing, controlling for matched listing FE. The matching is based on seller ID, listing title, listing subtitle, leaf category, listing start week, and price. Standard errors are clustered at the matched listing level.

*** indicates significance at p = 0.01; ** p = 0.05; * p = 0.1.

D Leads–and–Lags Analysis

We report the estimation results on the leads–and–lags analysis in Table A2 using specification 4 in the paper. The results are broadly consistent with the parallel trends assumption in our continuous DiD approach. Columns (1) – (3) report the results for eTRS sellers. Before the policy change, there is very little difference across markets in the number of PS listings, price, and number of active sellers. This is consistent with the parallel trends assumption across markets. After the policy change, there is no statistically significant change in the number of PS listings and price, and eTRS sellers seem to be more likely to exit in more affected markets five weeks after the policy change.

Turning to non-eTRS sellers, column (4) shows no statistically significant estimates at the 5% level before the policy change, except in the week right before it, which could be due to sellers' response to eBay's announcement of the change. After the policy change, non-eTRS sellers offer more PS listings in markets with a higher policy exposure, and the increase is rather persistent over time. In columns (5) and (6), we repeat the analysis for sales price and number of active sellers. Most of the estimates before the policy week are not statistically significant, which is consistent with the parallel trends assumption. Additionally, none of the coefficient estimates after the policy week are statistically significant at the 5% level, which is consistent with our previous results that non-eTRS sellers do not engage in price competition or exit the market at large.

Table A2: Leads-and-Lags Analysis

	eTRS Sellers			*Non-eTRS Sellers*		
	(1)	(2)	(3)	(4)	(5)	(6)
	ln(#PS)	ln(price)	ln(#seller)	ln(#PS)	ln(price)	ln(#seller)
Share×Week= −7	0.800***	0.006	0.049	-0.278	1.135*	0.077
	(0.305)	(0.857)	(0.086)	(0.293)	(0.626)	(0.079)
Share×Week= −6	-0.390	0.275	-0.119	0.516	0.444	0.003
	(0.321)	(0.344)	(0.084)	(0.472)	(0.557)	(0.083)
Share×Week= −5	0.433	-0.170	0.027	0.797*	0.135	0.124
	(0.353)	(0.319)	(0.080)	(0.462)	(0.715)	(0.107)
Share×Week= −4	-0.328	-0.142	-0.125*	0.391	-0.563	0.017
	(0.335)	(0.380)	(0.065)	(0.482)	(0.454)	(0.109)
Share×Week= −3	-0.228	-0.115	0.009	0.643	0.443	0.169
	(0.330)	(0.316)	(0.080)	(0.431)	(0.784)	(0.111)
Share×Week= −2	0.255	-0.116	-0.070	0.414	1.103	0.078
	(0.383)	(0.280)	(0.078)	(0.466)	(0.788)	(0.119)
Share×Week= −1	0.383	-0.310	-0.100	0.964**	0.407	0.072
	(0.402)	(0.318)	(0.087)	(0.486)	(0.583)	(0.170)
Share×Week= 0	-0.341	-0.440	-0.160*	0.938*	1.097	0.256
	(0.345)	(0.553)	(0.092)	(0.478)	(0.779)	(0.161)
Share×Week= 1	0.389	0.425	-0.118	1.232**	0.553	0.186
	(0.399)	(0.319)	(0.076)	(0.501)	(0.585)	(0.164)
Share×Week= 2	0.513	-0.548	-0.077	1.452***	1.370*	0.180
	(0.438)	(0.454)	(0.085)	(0.544)	(0.791)	(0.128)
Share×Week= 3	0.685	0.163	-0.156	1.674***	0.836	0.154
	(0.471)	(0.395)	(0.100)	(0.519)	(0.662)	(0.125)
Share×Week= 4	0.297	0.162	-0.081	1.634***	0.464	0.073
	(0.465)	(0.496)	(0.086)	(0.519)	(0.662)	(0.125)
Share×Week= 5	-0.245	0.250	-0.198**	0.973*	1.041	-0.058
	(0.460)	(0.309)	(0.083)	(0.507)	(0.723)	(0.132)
Share×Week= 6	-0.241	0.414	-0.132	1.007**	0.882	-0.140*
	(0.429)	(0.370)	(0.123)	(0.474)	(0.681)	(0.084)
Share×Week= 7	-0.169	-0.237	-0.204**	0.965**	0.203	-0.138
	(0.450)	(0.364)	(0.084)	(0.491)	(0.602)	(0.154)
R^2	0.960	0.855	0.997	0.927	0.796	0.995
Observations	6,816	6,778	6,816	6,816	6,771	6,816

Notes: This table shows the leads-and-lags regression results of outcome variables on policy exposure times a vector of week dummies, controlling for market and week fixed effects. Standard errors are clustered at the market level.

*** indicates significance at p = 0.01; ** p = 0.05; * p = 0.1.