E-Commerce Integration and Economic Development: Evidence from China*

Victor Couture†, Benjamin Faber‡, Yizhen Gu§ and Lizhi Liu¶

July 2017
–Preliminary and incomplete–

Abstract

The number of people buying and selling products online in China has grown from practically zero in 2000 to more than 400 million by 2015. Most of this growth has occurred in cities. In this context, the Chinese government recently announced the expansion of e-commerce to the countryside as a policy priority with the objective to close the rural-urban economic divide. As part of this agenda, the government entered a partnership with a large Chinese e-commerce firm. The program invests in the necessary transport logistics to ship products to and sell products from tens of thousands of villages that were largely unconnected to e-commerce. The firm also installs an e-commerce terminal at a central village location, where a terminal manager assists households in buying and selling products through the firm's e-commerce platform. This paper combines a new collection of survey and transaction microdata with a randomized control trial (RCT) across villages that we implement in collaboration with the e-commerce firm. We use this empirical setting to provide evidence on the potential of e-commerce integration to foster economic development in the countryside, the underlying channels and the distribution of the gains from e-commerce across households and villages.

Keywords: E-commerce, trade integration, economic development, rural-urban divide
JEL Classification: F63, O12, R13

*We are grateful to CEGA, the Clausen Center, the Haas School of Business, the Weiss Family Fund and the Bill and Melinda Gates Foundation for providing funding for this study. We have benefited from outstanding research assistance by Hero Ashman, Wenwei Peng, Jose Vasquez-Carvajal and Yi Wei. We are grateful to Lijun Sun, Wei Wang, Wei Zheng and Fang Ye for their tireless support during the implementation of this project. We also thank Hongbin Gao, Zhenzhong Sheng, Liang Chen, Wentao Zhang and Zhengwei Jiang for providing us access to and assistance with the transaction database. We are grateful to David Atkin, Jie Bai, Lauren Falcao-Bergquist, Lorenzo Casaburi, Thibault Fally, Frederico Finan, Cecile Gaubert, Jonas Hjort, Supreet Kaur, David Levine, Ted Miguel, Ralph Ossa, Andres Rodriguez-Clare, Daniel Sturm, Eric Verhoogen, Michael Walker, Noam Yuchtman, Fabrizio Zilibotti and participants at several seminars and conferences for very helpful comments. All views expressed and errors are our own.

†Real Estate Group, Haas School of Business, UC Berkeley
‡Department of Economics, UC Berkeley and NBER
§Institute for Economic and Social Research, Jinan University
¶Department of Political Science, Stanford University

1 Introduction

Since 2000, the number of Chinese buying and selling products online has grown from practically zero to more than 400 million by 2015, surpassing the US as the largest e-commerce market in terms of both users and total sales.[1] Outside of China, a growing number of developing countries, especially in Asia, Eastern Europe, Latin America and the Middle East, are also experiencing fast growth in e-commerce activity (WTO, 2013; UNCTAD, 2016b). Most of this growth to date has taken place in cities. In this context, the Chinese government recently announced the expansion of e-commerce to the countryside as a policy priority to foster rural economic development and reduce the rural-urban economic divide.[2] Other developing countries, such as Egypt, India and Vietnam, have recently announced similar policies to invest in the expansion of e-commerce outside of urban centers.[3]

These policies have been motivated mainly from the point of view of rural producers: by lowering trade and information costs to urban markets, the arrival of e-commerce is meant to increase rural incomes through higher demand for agricultural production and giving incentives for rural entrepreneurship. This motivation has been buttressed by a number of prominent case studies of successful "e-commerce villages" in China that have experienced rapid output growth by selling both agricultural and non-agricultural merchandise to urban markets through e-commerce.[4] On the consumption side, policy motivations rest on descriptive evidence of stronger demand for e-commerce in smaller cities with less pre-existing product variety on offer and higher tradable prices.[5]

Despite the fast growth of e-commerce in the developing world, and much recent hype and policy attention around the concept of "e-commerce villages", we currently have very limited empirical evidence on the economic consequences of e-commerce integration in developing countries. This paper combines a randomized control trial (RCT) that we implement across villages in collaboration with a large Chinese e-commerce firm with a new collection of household and store price survey microdata and the universe of transaction records from the firm's internal database. We use this empirical setting to provide evidence on the potential of e-commerce integration to foster economic development in the countryside, the underlying economic channels, and the distribution of the gains from e-commerce across households and villages. Our findings can serve as a first step towards building a rigorous evidence base on the economic consequences of the rapid

[1] See e.g. PFSweb (2016) and Statista (2016).

[2] The so-called "Number One Central Document" sets out yearly strategic priorities of the Chinese central government (the Central Committee of the Communist Party of China and the State Council in particular). The expansion of e-commerce to the countryside has featured in this document each January since 2014.

[3] As part of "Digital India", a collaboration between the Ministry of Electronics and IT and India Post have been tasked to expand online buying and selling in rural India (MEITY, 2016). Other recent examples include Egypt's National E-Commerce Strategy (MCIT, 2016) and Vietnam's E-Commerce Development Masterplan (PM, 2016). Following this policy interest, UNCTAD recently announced the launch of a new policy platform, "eTrade For All: Unlocking the Potential of E-Commerce in Developing Countries", to provide technical assistance and funding for e-commerce expansions in the developing world (UNCTAD, 2016a).

[4] These case studies have also received much press coverage. See e.g. "China's Number One E-Commerce Village" (BBC Global Business, 01 May 2013), "Inside China's Tech Villages" (The Telegraph, 05 Nov 2016), "Once Poverty-Stricken, China's "Taobao Villages" Have Found a Lifeline Making Trinkets for the Internet" (QZ, 01 Feb 2017), "Chinese 'Taobao Villages' Are Turning Poor Communities into Huge Online Retail Hubs" (Business Insider, 27 Feb 2017).

[5] In the US, the share of e-commerce in total US retail sales is estimated to be about 10 percent (e.g. FRED (2016)). In China, McKinsey (2016) reports this share to be as high as 20-30 percent in smaller cities, and Fan et al. (2016) find that this share increases by on average 1.2 percent as the city population is reduced by 10 percent.

growth of e-commerce in developing countries.

E-commerce is the ability to sell and buy through online transactions coupled with logistics for local parcel delivery and pickup from the producer. Bringing e-commerce to the countryside in developing countries requires more than internet access. In fact, as in many emerging countries, the internet has already spread rapidly to most parts of the Chinese countryside due to both smartphones and expanding broadband access. Instead, there are two currently binding barriers to e-commerce access in the Chinese countryside. First, commercial logistics operators in China do not offer parcel delivery or pickup services in most parts of the countryside. Servicing small and remote pockets of populations is subject to high per-unit costs related to the so called "last mile" between urban logistical hubs and villages. Second, many rural residents are unfamiliar with online platforms and lack access to online payment methods. In addition, villagers do not necessarily trust paperless transactions that occur before inspecting the product or without interacting with buyers in person. We refer to the first of these as the logistical barrier and to the second as the transactional barrier.

To overcome these barriers, the Chinese government recently partnered with a large e-commerce firm that operates a popular Chinese e-commerce platform. The program aims to invest in the necessary transport logistics to offer e-commerce in rural villages at the same price, convenience and service quality that buyers and producers face in their county's main city center. To this end, the e-commerce firm builds warehouses as logistical nodes for rural parcel delivery and fully subsidizes transport between the county's city center to and from the participating villages. To address the transactional barrier, the program installs an e-commerce terminal in a central village location. A terminal manager assists residents in buying and selling products through the firm's e-commerce platform, and villagers can pay upon receipt of their products or get paid upon pickup of their shipments in cash at the terminal location. From the end of 2014 to June 2016, approximately 16,500 Chinese villages in 333 counties and 27 provinces had been connected to e-commerce through the program. This expansion continues at the time of this writing, with an internal goal of 40,000 villages in 600 counties by the end of 2017. The RCT that we further describe below aims to evaluate the economic effects of this e-commerce expansion program.

Theoretically, the program can be rationalized as a reduction in trade and information costs between participating villages and the rest of urban China that is already connected to e-commerce. Importantly, this reduction in trade frictions is not confounded by first-time internet connections more broadly: villages were already connected to the internet and no investment is made on this front. Furthermore, the program only affects trading partners through e-commerce, while leaving other trade costs, e.g. to control villages, remain unchanged. The simplest way to approach this counterfactual is to model villages as small open economies whose market access, and thus tradables' prices, are predominantly determined by trade with larger urban centers that in turn are connected to other metropolitan regions and the rest of the world. In this scenario, control villages remain unaffected by the program, as rural-to-rural general equilibrium (GE) effects are negligible, and the comparison of outcomes in treatment vs control villages identifies the full GE effect of the arrival of e-commerce.[6]

[6]Another possible source of spillovers are rural-to-rural migration flows or access to the program's terminals in nearby villages, both of which we can observe directly.

Relaxing the assumption that village prices mainly depend on access to urban markets gives rise to potentially interesting GE spillover effects across villages. The presence of such spillovers would matter in our empirical analysis for three related reasons. First, spillovers on the control group would confound our estimates of the program's effect on treated villages. Second, spillovers on surrounding villages would have to be accounted for when quantifying the program's aggregate effect in the countryside. And third, knowing the magnitude of GE forces is required to evaluate the program's impact if it were scaled up to the entire Chinese countryside, instead of a subset of villages per county. In our empirical analysis, we begin by estimating simple differences in outcomes between treated and control villages under the baseline assumption that rural-to-rural GE effects are negligible. We then use additional empirical moments to investigate the presence of spillovers across rural communities.

Our analysis proceeds in four steps. In the first step, we derive a general expression for the program's effect on household economic welfare that guides the survey data collection and empirical analysis. Since the treatment we are interest in evaluating can affect not just individual behavior and the nominal earnings of households, that we can in principle record directly as part of the survey data collection, but also local household price indices in the denominator of real incomes, the evaluation of the welfare impact requires theoretical structure on the demand side. This is especially the case since some of the potential effects on household cost of living likely occur at the extensive margin of consumer choice, such as the arrival of a new e-commerce shopping option or local store exit. For such changes in the availability of local consumption options effective changes in consumer prices are unobserved since no information exists at either baseline (new options) or endline (disappearing options). Following a revealed-preferences approach, we can instead use observed substitution of household spending into new options or away from disappearing ones to infer the effective change in value-for-money over time across different product groups. More generally, the welfare expression allows us to break down the overall effect of e-commerce integration into several distinct subcomponents that we can link to the microdata.

In the second and third steps, we estimate the empirical moments of this welfare expression using household and village survey microdata as well as the firm's internal transaction database. The RCT takes place in 8 counties located in three provinces, Anhui, Henan and Guizhou, that have a large share of their population living in the countryside. The choice of these 8 counties was determined by the timing of the program's roll-out across different provinces and counties, and the internal evaluation of the senior management as to whether the provincial and county managers in question would be willing to cooperate with our research protocol. For each county, we were given authorization to randomly select control villages from a list of candidates that had been extended by 5 villages per county for the purpose of this research. Upon receipt of this extended list of village candidates, we randomly select 5 control villages and 7-8 treatment villages from each county's list. The remaining villages on the list also receive a program terminal as planned. Our sample thus includes 40 control villages and 60 treatment villages across the 8 counties, which we selected from a total number of 432 candidate villages (on average 54 villages per county). After obtaining the candidate list for each county, we were granted 2-3 weeks by the local operations team to run the randomization and send in our survey teams for data collection in the 12-13 sample villages before the terminal installations take place and e-commerce begins in

the treatment villages.

We complement this experimental design with survey data that we collect from households and local retail establishments in addition to the administrative transaction records from the firm's database. We collect baseline data in different counties in December 2015, January 2016 and April-May 2016 for 2800 households (roughly 8600 individuals) in the 100 villages. Half of these households are randomly selected within a 300 m radius of the planned terminal location ("inner village zone") and the other half is randomly selected from other parts of the village ("outer village zone"). For the endline, we collect data one calendar year after the baseline from the original household sample[7], and also extend the number of households living in the inner zone by 10 randomly selected households (leading to an endline sample of 3800 households).[8] For each household, we collect detailed information about e-commerce purchases, sales, other consumption expenditures, expenditures on production inputs, economic activities and incomes. We also collect baseline and endline information on 115 local retail price quotes for each village at the barcode-equivalent level across 9 consumer product groups as well as for business/production inputs. These survey data allow us to quantify the effect on household real incomes, and to distinguish between a number of underlying channels for both consumption gains (the denominator) and production-side effects (the numerator).

In the third step, we complement the RCT analysis using additional data from the firm's internal database that provide us with the universe of village transactions in 5 provinces (including the three above) over the period between November 2015 until April 2017. This database covers roughly 27.8 million purchasing and selling transaction records for about 12 thousand village terminals over the 18-month period. We use these data to answer four important questions that would be outside the scope of the RCT's survey data collection: i) to what extent are the RCT sample villages representative of the program villages in the Chinese countryside more broadly?; ii) to what extent are the results from the endline survey data sensitive to monthly seasonality?; iii) what is the time path of adjustment for e-commerce buying and selling each month since program entry, and do the effects increase for periods more than one year after the implementation?; and iv) to what extent are the survey data missing very successful, but rare, tail events on the production side that could affect the mean impact on household incomes per capita?

In the final step, we use the empirical estimates from steps 2 and 3 in combination with the theoretical framework in step 1 to quantify the impact of the program on average household economic welfare, the underlying channels and the distribution across households and villages.

[7]The fast pace of the program's expansion places bounds on the timing of the endline data collection. First, after our baseline the firm started planning additional waves of implementation in the 8 sample counties. Given that our control villages were selected from a list of promising candidate villages, they ranked high among candidates for additional installations. The additional waves did not take place within the first year after implementation. Second, in parallel the firm plans to roll out additional interventions through the village terminals, such as online banking, credit and health services. While interesting in their own right, these additional interventions would confound the impact of e-commerce integration that we are able to study cleanly in our sample. As discussed below, we are able to complement the RCT with additional evidence from the firm's internal database to investigate the time path of adjustments, including in periods after one year post-entry.

[8]In one of the 8 counties, the local government suspended further activity by our teams after we had completed endline data collection for 8 of the 12 sample villages in the county. This was unrelated to our operation, which followed the same protocol as elsewhere. As a result, we have endline data for 96 instead of the 100 villages. As the timing of data collection within the county was random, the 4 missing villages are not particular in any way. They included 2 control villages and 2 treatment villages.

The main findings can be summarized as follows. We find that the program leads to sizable gains in real incomes among rural households who are induced to use the e-commerce terminal. These users represent about 14 percent of the rural household sample and 13 percent of the village population after adjusting for sampling weights. For the average rural household, including non-users, these gains are statistically significant but more muted. Underlying these effects, we find strong heterogeneity across households and villages. The beneficiaries are on average significantly younger, richer, live in closer proximity to the e-commerce terminal and in villages that are relatively more remote. Conditional on these characteristics, we do not find evidence that household education and occupational status or the characteristics of the terminal managers affect the extent of household gains from e-commerce.

In terms of channels, we find significantly stronger economic gains among villages that were not previously serviced by commercial parcel delivery, suggesting that most of the program's gains are due to overcoming the logistical barrier, rather than the transactional one. On the consumption side, we find that the e-commerce terminals offer on average lower prices, higher convenience and increased product variety compared to pre-existing local retail choices, both within the village and in nearby towns. The gains in household purchasing power are strongest for durable product groups, such as electronics and appliances. We also find suggestive evidence that the program led to additional product variety in pre-existing local stores, as their managers source new products through e-commerce. We find no evidence of significant pro-competitive effects on local retailer prices, on the other hand. On the production side, we find no evidence for significant effects on the local economy: online selling activity, purchases of production inputs, household incomes and entrepreneurship are not significantly affected by the arrival of the program. Overall, we find that the gains from e-commerce are driven by a reduction in local household cost of living that is mainly due to the direct gains from access to the new e-commerce shopping option for local households.

Using the firm's database, we find little evidence on the consumption side suggesting that the adjustment takes longer than one year: the consumption-side uptake materializes within 2-4 months of entry and then remain mostly constant over time. On the production side, we find evidence that village-level out-shipments significantly increase over time beyond the 12-month window. However, the effect on total out-shipments remains relatively minor after more than two years post program entry, with a small upper-bound effect on local household incomes. Related to this, we do not find evidence that the survey data fails to pick up highly successful but rare tail events on the production side that could in principle shift the mean effect on local household outcomes.

Overall, our findings suggest that e-commerce expansions offer significant economic gains to certain groups of the rural population, rather than being broad-based. Compared to the recent case studies highlighting a set of highly successful rural e-commerce production hubs, our analysis reveals that a quite particular mix of local factors must be underlying these prominent success stories. In the absence of complementary interventions, such as for example business training, access to credit or targeted online promotions, large and significant production-side effects appear unlikely to materialize for the average rural market place. In this light, future work aimed at better understanding the factors under which the arrival of e-commerce can have transforma-

tive impacts on the production side of the rural economy seems a promising agenda for future research in this area.

This paper relates and contributes to the recent literature on globalization and development using within-country empirical variation (e.g. (Topalova, 2010; Donaldson, in press; Atkin et al., in press)). Given the empirical context, we also relate to the recent literature on the consequences of transport cost reductions within China (e.g. (Banerjee et al., 2012; Baum-Snow et al., 2016; Faber, 2014)). Instead of focusing on trade liberalization, transport infrastructure or the effects of FDI, we set focus on the economic consequences of e-commerce, a recent but fast-growing channel of market integration in developing countries that has so far received relatively little attention in the literature.

Our findings also relate to the literature on the consumer gains from e-commerce (e.g. (Brynjolfsson et al., 2003; Goldmanis et al., 2010)), and cost of living as a function of city size and urban density (e.g. (Handbury, 2013; Handbury & Weinstein, 2015; Couture, 2016; Fan et al., 2016)). In this literature, we most closely relate to recent work by Fan et al. (2016) who use data on e-commerce sales on the Taobao platform across 315 prefectures in China for the year 2013 to document a decreasing relationship between prefecture population and online expenditure shares in the cross-section. These findings would suggest that the consumer gains from e-commerce are expected to be the largest among small and remote market places. Relative to the existing literature, this paper uses experimental variation in the arrival of e-commerce to the countryside to quantify the effects on both consumption and production. Our findings suggest that the relationship between e-commerce usage and city size does not appear to hold monotonically as we move from relatively large urban centers further to the left tail of the population size distribution in the countryside.

The paper is also related to the recent literature on the effects of the internet in developing countries. Goyal (2010) studies the consequences of the introduction of internet kiosks with wholesale price information in the Indian state of Madhya Pradesh, and finds a positive effect on local soy prices and area under soy cultivation. More recently, Hjort & Poulsen (2016) use the geography of existing terrestrial communication networks in Africa in combination with the timing of submarine internet cable connections to study the effect of fast-speed internet on labor markets in several African countries. They find a positive effect on overall employment that is mainly driven by an expansion of higher-skill employment. Relative to the existing literature, this paper sets focus on a different question of policy interest. Rather than estimating the effects of the internet more broadly, we explore the consequences of the arrival of e-commerce among rural Chinese markets. Since the expansion of e-commerce requires specific investments to overcome both logistical and transactional barriers beyond the provision of internet access, our analysis can serve as a first step to inform the current wave of policy interest in the promise of e-commerce as a driver of rural development.

Finally, our findings relate to recent literature on the sources of the rural-urban economic divide in developing countries (e.g. (Young, 2013; Lagakos et al., 2016; Hamory et al., 2016)). A central question in this literature has been to what extent features of locations, rather than the selection of people across space, can explain the observed rural-urban gap in economic development. Our findings suggest, perhaps surprisingly, that in the Chinese case a lack of urban market

access –a characteristic that differs between rural and urban locations– does not by itself appear to be a strong factor explaining observed disparities between the countryside and urban centers, at least in the short to medium-run. In this respect, our findings complement existing evidence suggesting that selection plays an important role in rationalizing observed differences in rural and urban economic development.

The remainder of the paper is structured as follows. Section 2 describes the context, experimental design and data. Section 3 presents the theoretical framework. Section 4 presents the empirical analysis based on the RCT in combination with the survey data. Section 5 provides additional evidence using the firm's internal database. Section 6 presents the welfare quantification. Section 7 concludes.

2 Context, Experimental Design and Data

2.1 Context and Program Description

Following the announcement of the policy objective to expand e-commerce to the Chinese countryside as part of the so-called Number One Central Document in January 2014, the Chinese government entered a partnership with a large firm that operates a popular Chinese e-commerce platform. The program makes two main types of investments to enable villagers to buy and sell online through the firm's platform. The first is that the program invests in the local transport logistics, which the firm sees as a necessary condition to provide e-commerce access in rural areas. Before the arrival of the program, most villages were not serviced by commercial parcel delivery operators who had not solved the problem of the "last mile" transportation between dispersed rural households and urban county centers.[9]

The program sets out to change this lack of service for the purpose of e-commerce. In particular, the firm oversees the construction of warehouses in the county that serve as logistical nodes to pool all e-commerce related transportation requests to and from the participating villages. These warehouses are located close to the urban centers of the county with good cross-county transport access. The program also fully subsidizes the transportation cost between these county centers and the villages so that rural households face the same delivery costs and prices as households in the urban parts of the county. The rationale for this subsidy is that village deliveries and pickups start from a low basis, which due to economies of scale in rural transportation makes the starting phase of e-commerce prohibitively costly for village customers despite the investments in warehouses. The calculation of the government and the firm is that as the scale of rural e-commerce grows, per unit transport costs will decline enough to remove the need for a subsidy.

The second investment is the installation of a program terminal in a central village location. The e-commerce terminal is a flat-screen monitor mounted on the wall of shop space displaying the firm's website. On the screen, consumers and producers can choose their purchases or see their sales requests on the platform. There is also a keyboard and a mouse for villagers able to use the terminal without the store manager's help. Through this terminal, villagers can buy and sell products through the store manager's account, and they can pay in cash when the products arrive at the store for pickup, or they get paid upon delivery of their products for pickup at the

[9] To receive packages via mail in absence of commercial parcel delivery services, rural households have to travel to the county or township center to pick up the package after receiving notification by mail that it has arrived.

store location if selling online. The firm views the village terminals as instrumental in overcoming three challenges. First, local households may not be used to or comfortable with browsing products on online platforms. Second, they mostly do not have access to online payment methods. And third, they do not trust online purchases or sales before inspecting the goods in person or having received the payment. The terminal manager receives a reward of about 3-5 percent for each transaction completed through the terminal. Before deciding on terminal installations, the firm solicits applications from potential local store operators and schedules an exam for the applicants. The score of this exam is one of the criteria that the firm uses to determine whether a village is a candidate.

2.2 Sample, Design and Data

Selection of Provinces and Counties

There are two main factors determining our survey location in Anhui, Henan and Guizhou and the 8 counties within these provinces. First, our survey location depended on the timing of the program's roll-out across different provinces and counties, which had been decided before our collaboration with the firm. Second, we were guided by the internal evaluation of the program's senior managers as to whether the provincial and county managers in question would be willing to take part in our study and cooperate with our research protocol. These counties are: Huoqiu (Anhui), Linying (Henan), Linzhou (Henan), Minquan (Henan), Suixi (Anhui), Tianchang (Anhui), Xifeng (Guizhou) and Zhenning (Guizhou). In Section 5, we are also able to test the representativeness of our sample villages relative to all participating villages using the firm's internal transaction data in 5 provinces over this period.

Experimental Design

The unit of randomization is the village. For each county, we obtain a list of candidates that had been extended by 5 promising village candidates that would have not been part of the list in absence of our research. The two main factors determining the village selection within a county from the firm's operational perspective are i) a sufficient level of local population, and ii) the presence of a capable store applicant (as measured by the applicant's test score).

Upon receipt of this extended list of village candidates for each county, we randomly select 5 control villages and 7-8 treatment villages. The remaining villages on the extended list receive program terminals as planned. The full sample thus includes 40 control villages and 60 treatment villages across the 8 counties, which we selected from a total number of candidates of 432 villages (on average 54 villages per county).

We restrict the list of villages entering the stratification and randomization to villages with at least 2.5 km distance to the nearest village on the county list, wherever possible. We then stratify treatment and control villages along four dimensions. First, we balance the selection of treatment and control to both have a ratio of 85:15 with respect to pre-existing availability of commercial package delivery (85% not available, 15% available). We obtain this information for each village on the candidate list from the program's local county manager (who is not aware what we require that piece of information for). Having both villages with and without pre-existing commercial delivery services in principle allows us to disentangle the effect of the program that is driven by the terminal access point (i.e. providing more convenient access to e-commerce in

an environment where online buying/selling is already available through the internet), relative to the effect of providing both the terminal access point and the necessary transport logistics to make local online buying or selling possible regardless of the access point.

We further stratify the selection of treatment and control villages on the basis of the equally-weighted average of the z-scores for three village variables: the local store applicants' test score, the village population, and the ratio of non-agricultural employment over the local population. We obtain the last variable from the establishment-level data of the Chinese Economic Census of 2008 which surveys every non-agricultural establishment in the county.

After obtaining the candidate list for each county, we have about 2-3 weeks to run the randomization and send in the survey teams for data collection in 5 control villages and the 7-8 treatment villages before the terminal installations take place and e-commerce begins in the treatment villages. Compliance with our assignments to treatment and control villages is not perfect: the program was rolled out in 38 of the 60 villages assigned to treatment, and it was present in 5 out of the 40 control villages. We therefore report both intent-to-treat and treatment-on-treated effects. The main reason for imperfect compliance is the gap between the candidate list and actual terminal installation and operation that is part of our experimental design. In particular, not all candidates that apply and make it to the list of viable candidates (from which we randomize) end up accepting the offer and signing the contract after we send the results of our randomization back to the local operation team. As a result, not all of the 60 chosen candidate villages end up with effective program roll-out. In addition, our data collection imposed a 2-3 week delay relative to the standard roll-out of the program. Terminal installation normally happens almost immediately after the completion of the candidate list. In some cases, this delay meant that engineers had to come back to complete installations after having already moved on to different parts of the countryside, which due to logistical constraints they could not always commit to.[10]

Tables 1-4 present descriptive statistics of the baseline data at the individual level, the household level and for local retail prices. The experimental design appears to have been successful in creating treatment and control groups that are on average balanced in terms of pre-existing outcomes. As discussed in Section 4, our empirical analysis will also condition on the baseline values of the outcomes to be tested.

Sampling of Households and Household Survey Data

For the first round of data collection (December 2015 and January 2016 in Anhui and Henan, and April and May 2016 in Guizhou), we collect data from 28 households per village. 14 of those households are randomly sampled within a 300 meter radius of the planned terminal location ("inner zone"), and 14 households are randomly sampled from other parts of the village ("outer zone"). The household survey respondent is the member with the most knowledge of household consumption expenditures and income. Each respondent receives a gift to thank them for their participation in the survey (e.g., box of premium sweets, soaps, hand towels, etc). The value of the gift is about 4.5 USD. If the most knowledgeable respondent is not present at the time of the visit, then the surveyor schedules a follow-up visit.

The second round of data collection occurs one year after the first round in each county. We

[10] In our estimation sample, we only have data on 96 of the 100 villages for the endline survey. The treatment proportions for the sample of 96 villages are 37/58 and 4/38 respectively.

collect data from the same households as in the first round, plus 10 randomly sampled households within the inner zone around the planned terminal location.[11] If either the survey respondent or the primary earner of the initially surveyed household no longer resides at the same address, we record this in our data and replace the household with another randomly sampled household within the same sampling zone (inner or outer). In our welfare analysis, we report results both before and after weighting each sampled household in proportion to the share of the village population in its sampling zone.

We collect detailed information about household consumption expenditures across 9 household consumption categories for retail products (food and beverages, tobacco and alcohol, medicine and health, clothing and accessories, other every-day products, fuel and gas, furniture and appliances, electronics, transport equipment) as well as for expenditures on business inputs. We also collect information about household incomes, hours worked and economic activities of different members (occupation (e.g., farmer, manual worker, etc) and sector (agricultural, manufacturing, services)), in addition to data on asset ownership, financial accounts, internet use and migration.

In one of the 8 counties, the local government suspended further activity by our teams after we had completed endline data collection for 8 out of 12 villages in that county. This was unrelated to our operation, which followed the same protocol as elsewhere. As a result, we have endline data for 96 instead of the 100 villages. As the timing of data collection within the county was random, the 4 missing villages are not particular in any way. They include two control villages and two treatment villages.

Tables 1-3 present descriptive statistics of the baseline data from the household survey at the individual level and at the household level. The tables also present descriptive statistics for the same outcomes among the control group at the endline data collection. The median age of all household members in the baseline survey is 44 and the median household size is 3. 60 percent of households report the primary earner to be a peasant, and 82 percent of households report that the primary earner completed at least primary school. In terms of demographics, these statistics are very similar to nationally representative rural household samples from the China Family Panel Study. The same holds for household economic characteristics: median monthly income per capita and retail expenditure per capita are about RMB350 and RMB380 respectively, which makes these households significantly poorer than households living in urban centers. These households spend on average half of their retail expenditure outside the village, which requires significant travel as their main shopping destination outside the village is generally an urban center at a median two-way distance of 40 minutes. In terms of incomes, 80 percent of primary earners work inside the village.

As discussed in the introduction, many households report using the internet via smartphones or other devices: close to 40 percent report having used the internet, more than 50 percent own smartphones and close to 30 percent report owning a laptop or a PC (by comparison, almost all households own a TV). At the same time, e-commerce penetration is very limited compared to urban regions: the average share of household retail expenditure on local e-commerce deliveries is less than 1 percent, and this does not change over time for the endline survey in the control group

[11] This extended sample was possible due to a small remaining positive balance on the project account that we decided to invest in expanding the local household samples.

of villages. Similarly, the share of revenues from online selling in monthly household income is less than 0.5 percent, and again this does not change over time for the endline data collection in the control group of villages. By comparison, a recent survey conducted by McKinsey (2016) has found that urban households in Chinese cities spend on average up to 20-30 percent of total retail consumption on e-commerce deliveries.

Local Retail Price Survey Data

We aim to collect data on 115 price quotes for each village. 100 of these prices are from the same 9 household consumption categories for retail products as in our household survey (food and beverages, tobacco and alcohol, medicine and health, clothing and accessories, other everyday products, fuel and gas, furniture and appliances, electronics, transport equipment), and 15 price quotes are for local production/business inputs. Our protocol for the price data collection closely follows the IMF/ILO standards for store price surveys that central banks collect to compute the CPI statistics. The sampling of products across consumption categories is based on budget shares of rural households in Anhui and Henan that we observe in the microdata of the China Family Panel Study (CFPS) for 2012. The sampling across stores is aimed to provide a representative sample of local retail outlets (stores and market stalls). In villages with few stores we sampled all of them. The sampling of products within stores is aimed at capturing a representative selection of locally purchased items within that outlet and product group. Each price quote is at the barcode-equivalent level where possible (recording brand, product name, packaging type, size, flavor if applicable).

In the second round of data collection (one year after the first round), we aim to collect the price quotes of the identical products in the identical retail outlet where this is possible. Where this is not possible, due to either store closure or absence of product in the store, we record the reason for the absence and include a new price quote within the same product category that is sampled in the same way as in the first round.

Table 4 presents descriptive statistics of the baseline data from the local retail price surveys. Unsurprisingly durable goods categories (furniture and appliances, electronics and transport equipment) are an order of magnitude more expensive than goods in non-durable categories. The median number of sampled stores is 3 (40% of villages have 3 or fewer stores in total). These stores are small with a median square footage of $50m^2$ and the median store has not added any new product within the last month. For 67% of goods in these stores, there were no price labels and our surveyor was required to ask for a price quote.

Firm's Transaction Database

We complement the collected survey data with administrative records from two of the firm's internal datasets that we access through a remote server. The first database covers 5 provinces (our three study provinces plus two additional provinces with high shares of rural population: Guangxi and Yunnan) over the period from November 2015 (1 month prior to our survey data collection) to April 2017.[12] As summarized in Table 5, this database covers approximately 27.3

[12]To our knowledge, it is the first time that external researchers have obtained access to the firm's internal transaction database. Establishing this access took more than 1.5 years of negotiation and numerous rounds of approval procedures vis-a-vis different departments within the firm's organization.

million transaction records across 12,000 village terminals over the 18-month period.[13] Given that many terminals had already been in operation for several months prior to November 2015, these data cover adjustment periods significantly beyond the 12-months window that we are able to capture as part of the RCT. Terminals are observed up to two years and 4 months after the installation in these data. This first database covers all purchases made through the e-commerce terminals.

The second database covers the universe of sales transactions, i.e. of out-shipments from the villages, through the firms distribution network for the same universe of roughly 12 k village terminals in the 5 provinces over the period January 2016 to April 2017. The total number of e-commerce out-shipments over this period is roughly 500 thousand, as depicted in Table 5 that provides descriptive statistics for both datasets that we use in the analysis reported in Section 5.

3 Theoretical Framework

This section proceeds in three parts. We first describe the channels through which the program can affect the local economy. We then rationalize the empirical counterfactual in the light of potential GE spillovers across villages in the countryside. Finally, we derive a general expression of the program's effect on household economic welfare that guides the survey data collection and the empirical analysis in the following sections.

3.1 Channels at Work

What type of economic shock does the program imply for the local economies? The program makes no investment in internet accessibility for villagers, and the terminal cannot be used to browse the internet except for the e-commerce platform. This, together with the fact that roughly one third of village households report using the internet before the arrival of the program and more than 50 percent own smartphones (Table 2), indicate that the shock is particular to the arrival of e-commerce, rather than providing internet access more broadly. Being able to separate the effects of e-commerce from internet access more broadly (such as email, online search, social media, communication or news) is one of the strengths of this empirical setting.

The program has two central elements that are aimed at removing the logistical and transactional barriers to rural e-commerce. The first is that the program aims to bring e-commerce-related shipping costs to and from the village to the same level as those present in the county's main urban centers. To this end, the program builds warehouses as logistical hubs for village deliveries and pickups, and fully subsidizes transport costs between the county's city center and the villages. The second element is that the program installs an e-commerce terminal in a central village location, where a terminal manager assists villagers to buy and sell products through the firm's e-commerce platform using traditional offline payments.

For the local economy, both of these interventions affect the degree of trade integration between the village and the rest of urban China that is already connected to e-commerce. The logistical element reduces the physical trade costs to and from the village for bilateral connections that are connected to e-commerce. At the same time, the program does not directly affect the transport costs of non-participating villages, or trade flows of program villages outside of e-commerce. The

[13] We are able to identify 40 of the 43 e-commerce terminals in our RCT sample villages based on the Chinese county and village names that we have access to in the firm's transaction database.

transactional element (terminal installation), on the other hand, potentially reduces information and transactional frictions for trade flows to and from the village: e-commerce enables villagers to observe products and prices from other regions that are connected to e-commerce far beyond the local economy and, in turn, other regions can learn about the products and prices from local producers. To the extent that villagers were already aware of the online information offered by the e-commerce platform in absence of the program (e.g. through smartphones), the terminal installation may still alleviate transactional barriers by making it easier to buy from or sell to trade partners outside the village.

By overcoming both logistical and transactional barriers to e-commerce integration, the program provides villages with essentially urban market access through e-commerce. In villages that were not previously served by commercial parcel delivery services, the effect we observe will be driven by the removal of both of the barriers to e-commerce integration that we discuss above. In villages that already were serviced by parcel delivery, logistics were already operating at the same level as in urban areas and the effect we observe will be driven only by the additional provision of the terminal interface (removal of the second barrier to rural e-commerce).[14]

3.2 General Equilibrium Spillovers

Our estimation exploits differences in outcomes between program villages and comparable control villages. This raises the question to what extent these differences may reflect spillover effects from treated villages on nearby control villages. Theoretically, the presence and strength of local indirect effects is a priori unclear, and will depend on the degree of trade integration between villages in rural regions. If we think of Chinese villages as small open economies whose market access is mainly determined by trade with larger urban areas in the county that constitute the "rest of the world", rather than trading with other small rural markets, then the extent of such spillover effects could be limited. On the other hand, if the villages' trade market access is to a significant extent determined by trading with other villages in the countryside, rather than urban centers, then general equilibrium effects among rural regions could play an important role. In addition to spillovers driven by trade linkages between villages, it could also be the case that households in control villages use program terminals in nearby villages to access e-commerce.[15]

The extent of such spillover effects is an interesting empirical question for three main reasons. First, in principle we are interested to understand the effect of the program on the level of household welfare among treated villages. To the extent that the control group was indirectly affected, the difference in treated vs non-treated outcomes does no longer directly speak to this moment of interest. Second, even after correctly adjusting for indirect effects on the control group, the presence of spillovers would also have implications for the generalizability of our conclusions. In our current empirical setting, only a fraction of the Chinese countryside in any given county is currently part of the program. If we wanted to inform policy making on the welfare consequences

[14] According to the e-commerce firm, the transport cost subsidy does not affect villages that were previously serviced by parcel delivery services. The logistics operators essentially picked all areas in the county at which e-commerce-related deliveries and pickups could be offered at the same rate, i.e. there were no delivery cost differences within counties. The big difference was that logistics operators did not service most of the countryside in absence of the program's investments in logistics (warehouses and transport subsidies).

[15] Another possible source of spillovers in this setting are rural-to-rural migration flows for which we can test directly.

of scaling up e-commerce access in rural China to a larger fraction of the countryside, then the presence of spillovers on the control group would indicate that the implications for treated households could significantly differ from our treatment effects. Third, the presence of spillover effects would also change our understanding of the aggregate implications of the program –either in its current form, or when evaluating a scaled-up version of the program. That is, rather than focusing on the welfare effects on treated communities, we are also interested in the overall impact of the program among rural households as a whole. Here, knowing the extent of spillover effects allows us to compute the average effect of the program on rural households as a function of direct and indirect exposure to the program whose averages we can measure in the data (or simulate when scaling up).

In our empirical analysis, we begin by comparing economic outcomes in treatment and control villages, under the baseline assumption that rural-to-rural GE effects are negligible. We then proceed in two directions. First, we use a methodology close to Miguel & Kremer (2004) to investigate to what extent plausibly exogenous variation in exposure to nearby treatment villages affects local economic outcomes conditional on the local treatment status of the village in question. Second, we use trade theory (e.g. Eaton & Kortum (2002)) as a guidance and construct village-level measures of market access. Market access is the weighted sum of access to market expenditures across all rural and urban market places in China and beyond, where the weights are inversely related to the bilateral trade costs on each potential trading route. We can use information on the geographical position and market size of all rural and urban settlements in China during the baseline period in combination with measures of bilateral travel costs in order to investigate what fraction of theory-consistent trade market access in our village sample is driven by access to urban markets relative to other villages that participate in the e-commerce expansion program.

3.3 Quantification of Changes in Household Economic Welfare

As discussed above, the intervention that we are interested in evaluating has the potential to not just affect individual behavior and the nominal earnings of households, but also household cost of living in the denominator of real incomes. To empirically quantify the change in household price indices due to the arrival of the e-commerce terminal in the village, we require theoretical structure on the demand side.

In order to evaluate the change in economic welfare due to the arrival of e-commerce we consider the compensating variation (CV) for household h. This approach follows earlier work by Hausman (1996), Hausman & Leonard (2002) and more recently Atkin et al. (in press). The CV captures the change in exogenous income required to maintain utility when e-commerce arrives between period 1 and period 0, with periods denoted by superscripts:

$$CV_h = \underbrace{\left[e(\mathbf{P}^1, u_h^0) - e(\mathbf{P}^0, u_h^0)\right]}_{\text{Cost of living effect } (CLE)} - \underbrace{\left[y_h^1 - y_h^0\right]}_{\text{Nominal income effect } (IE)}, \quad (1)$$

where $\mathbf{P}^t$ is the vector of prices faced by the household in period t, u_h^t is the household's utility and y_h^t is its nominal income.

The first term is the cost of living effect, the welfare change due to the price changes induced by the arrival of e-commerce. The second term is the nominal income effect, the welfare change

due to any changes in household income that result from the arrival of e-commerce. While, at least in principle, we can record the effect on nominal household earnings and labor supply directly as part of the survey data collection, this is not the case for the cost of living effect. The store price survey data described above allow us to observe the vector of price changes $\mathbf{P}_C^1 - \mathbf{P}_C^0$ for continuing products in continuing local retailers (i.e. stores, market stalls, etc) that are present both before and after the arrival of the program. We index such continuing product prices by C.

However, there are three sets of price changes that are inherently unobservable: the consumer price changes $\mathbf{P}_T^1 - \mathbf{P}_T^0$ due to the entering e-commerce terminal indexed by T, the price changes $\mathbf{P}_X^1 - \mathbf{P}_X^0$ of potentially exiting local retailers or varieties within continuing stores indexed by X, and the price changes $\mathbf{P}_E^1 - \mathbf{P}_E^0$ due to local store entry or new product additions in pre-existing local retailers. For example, the price of shopping at the e-commerce terminal is prohibitively high prior to its installation, and exiting local retailers' prices are not observed ex post. As first noted by Hicks (1940), we can replace these three unobserved price vectors with 'virtual' price vectors, the price vectors that would set demand for these shopping options equal to zero given the vector of consumer prices for other goods and services.

In the following, we denote such price vectors with an asterix (the implicit prices that would set consumption equal to zero in a given period), and break up the total consumption price vector in expression (1), into the four different components of potential consumer price changes. This leads to a decomposition of the program's total cost of living effect (CLE) into different channels that we can map to observable moments in the survey microdata:

$$CLE = \underbrace{e(\mathbf{P}_T^1, \mathbf{P}_C^1, \mathbf{P}_X^{1*}, P_E^1, u_h^0) - e(\mathbf{P}_T^{1*}, \mathbf{P}_C^1, \mathbf{P}_X^{1*}, P_E^1, u_h^0)}_{\text{(1) Direct price-index effect } (DE)} + \underbrace{e(\mathbf{P}_T^{1*}, \mathbf{P}_C^1, \mathbf{P}_X^{1*}, P_E^{1*}, u_h^0) - e(\mathbf{P}_T^{0*}, \mathbf{P}_C^0, \mathbf{P}_X^{0*}, P_E^{0*}, u_h^0)}_{\text{(2) Pro-competitive price effect } (PP)}$$

$$+ \underbrace{e(\mathbf{P}_T^{1*}, \mathbf{P}_C^1, \mathbf{P}_X^{1*}, P_E^1, u_h^0) - e(\mathbf{P}_T^{1*}, \mathbf{P}_C^1, \mathbf{P}_X^{1*}, P_E^{1*}, u_h^0)}_{\text{(3) Entry effect } (EE)} + \underbrace{e(\mathbf{P}_T^{0*}, \mathbf{P}_C^0, \mathbf{P}_X^{0*}, P_E^{0*}, u_h^0) - e(\mathbf{P}_T^{0*}, \mathbf{P}_C^0, \mathbf{P}_X^0, P_E^{0*}, u_h^0)}_{\text{(4) Exit effect } (XE)} \tag{2}$$

The first term of the first bracket and the second term of the final bracket of this decomposition represent the same difference in expenditure functions as in the first term of (1): the amount of expenditure one would have to pay household h in order to obtain the pre-terminal level of wellbeing, but evaluated at the post-intervention consumption prices. The terms in the middle between these two terms cancel out one another, so that this decomposition yields the total gains or losses due to effective consumption price changes, including changes at the extensive margin of consumer choice (e.g. new shopping options).

The first bracket, the direct price index effect, captures the consumer gains due to the arrival of the new terminal shopping option holding all other prices fixed. These gains can arise from three distinct channels that are all captured in the quantification of the bracket: e-commerce can provide existing products at cheaper prices, it can offer new product variety that was not previously available, and it can offer different shopping amenities (e.g. convenience, saving trips outside the village, etc).

The second bracket captures changes in household cost of living due to price changes among pre-existing retailers and their products. Following Atkin et al. (in press), we label this the pro-

competitive price effect. Existing retailers could lower their markups due to increased competition from e-commerce. The third and fourth terms, that we refer to as the entry and exit effects, capture changes in product availability in the local retail environment. For example, the arrival of the e-commerce terminal could lead some local retailers to exit, it could in principle lead to store entry (e.g. stores sourcing online), and it could lead to disappearing or new product variety among pre-existing stores (for example due to local retailers starting to source their products online).

Up to this point, the welfare expressions in (1) and (2) are fully general without imposing functional forms. However, three of the four cost of living effects involve price index effects that are due to extensive margin changes in consumer choice: except for the pro-competitive price effect, we cannot collect data on effective consumer price changes due to new or disappearing retailer or product availability. To quantify the first, third and fourth welfare effects, one needs to impose theoretical structure on the expenditure functions. That same specification of consumer preferences will also provide a specific price index formula for the pro-competitive price effect, of course.

The logic behind this approach is as follows: once we know the shape of the demand curve that governs consumer substitution across different retailer options within a given product group, we can use the observed changes in the expenditure shares across different shopping options before and after the program intervention in order to infer the unobserved effective consumer price changes that underlie this observed substitution. Once we know the elasticity subject to which consumers usually substitute across retailers as a function of differences in effective value-for-money, then one can back out the implied effective price index change for consumption of a given product group that is consistent with the observed substitution into the e-commerce terminal for that product group. Again, these price index changes could be driven by price differences, different product availability and/or different shopping amenities. For the welfare evaluation, what matters is that we evaluate the observed changes in consumption expenditure in combination with knowledge about the consumer demand curve across retail outlets. A very similar logic follows when evaluating the price index movements due to disappearing stores, entering stores or product additions and disappearances within continuing stores.

In the appendix, we outline one such approach to guide the empirical estimation that follows a nested CES preference specification commonly used in international trade and macroeconomics. In particular, local households are assumed to have Cobb Douglas perferences across broad product groups in retail consumption (durables and non-durables). Within these nests, groups of different household types have CES preferences across retailers (e.g. e-commerce terminal, stores or stalls in village, stores outside village, etc). Within stores, households choose across varieties on offer within product groups as a function of quality-adjusted product prices. This structure follows recent work by Atkin et al. (in press) on Mexican households, as we describe further below and in the appendix.

Regardless of the particular demand specification one imposes, the raw empirical moments that are required to quantify the welfare impact of the intervention fall into three different types. The first set of empirical moments are estimates of the causal effects of the intervention on a number of observable economic outcomes, such as the fraction of total retail expenditure substituted

to the new e-commerce terminal across different product groups and by household types (for heterogeneity), the effects on household nominal incomes to capture the second term in (1), the effect on the price changes among pre-existing retailers, and the effect on the propensity of store exit or product introductions among the local retail environment.

The second type of required estimates are empirical moments from the baseline data collection, such as consumption shares across product groups and the shares of incomes across different activities. The third type of empirical moments are estimates of demand parameters that govern the degree of consumer substitution across retailers and products. This latter set of parameters differ across different functional form assumptions on the demand side. In this context, a benefit of our approach is related to the fact that the rich time series of consumer scanner data needed to obtain estimates of demand functions are not available in the Chinese empirical context. In this light, we can use existing recent estimates for households of very similar income ranges reported in Atkin et al. (in press), that to the best of our knowledge are the closest empirical estimates on the nature of retail demand and consumer substitution in an emerging market environment, such as China. In addition to tying our hands to existing estimates from the literature, we also report quantification results across a range of alternative demand parameters to document the sensitivity of the welfare estimates across a range of assumptions.

4 Empirical Analysis Using Survey Data

In this section, we estimate the program's effect on a number of economic outcomes related to: household consumption, incomes, economic activity, and local retail prices, that we observe in the survey microdata. In addition to being of interest in their own right, these empirical moments enter the quantification of changes in household economic welfare in Section 6.

4.1 Average Program Effects

As in e.g. McKenzie (2012), we run the following regressions:

$$y_{hv}^{Post} = \alpha + \beta_1 Treat_v + \gamma y_{hv}^{Pre} + \epsilon_{hv}, \tag{3}$$

where y_{hv} is an outcome of interest for household h living in village v. For outcomes from the retail price data, h indexes individual price quotes or store-level outcomes instead. $Treat_v$ is an indicator of intended treatment according to our randomization, so that β_1 captures the intent-to-treat effect (ITT). We also estimate the treatment-on-the-treated (TOT) after instrumenting for the actual treatment status using $Treat_v$. Finally, we run (3) after replacing the binary treatment indicator with a continuous measure of the log of household residential distance to the nearest program terminal, again using $Treat_v$ as an IV.

For households who were either replaced or added as part of our extended sample in the second round (from 28 to 38 households), we define y_{hv}^{Pre} as the mean pre-intervention outcome of households living in the same zone (inner or outer) in the same village. The implicit assumption is that households were not induced to in- or outmigrate of villages as a result of the program.[16] We cluster standard errors at the level of the treatment (village level).

[16] We find no evidence that households in treated villages are more or less likely to reside at the same address in the post-treatment survey. We also find no treatment effect on migration decisions of members within households.

Tables 6-8 present the estimation results for the average effects on household consumption, incomes and local retail prices. Our discussion focuses on the TOT results, while the tables display the three types of effects discussed above (ITT, TOT and log distance). We report average effects on households in our sample as described in Section 2. For the welfare quantification in Section 6, we will also report results after re-weighting village zones according to their local population shares.

Consumption

In Table 6, we find that the program on average leads to an uptake of 9 percent of households who report to have ever used the terminal for making purchases, relative to control villages. The effect is about 5 percent of local households who report using the terminal over the past month at the time of the endline survey, relative to the control group.[17] These effects on uptake on the consumption side may in part mask additional uptake from households in nearby villages. We return to this issue when investigating spillover effects further below.

The effect on the terminal share in total household retail expenditure is 1.24 percent for the average household in our survey data. Thus, households that report ever having used the terminal spent on average 1.24/8.9=14.1 percent of their retail consumption at the terminal during the past month. For those who bought over the past month, this share rises to 1.24/4.9=25.3 percent.

Looking at retail consumption across product groups, we find stronger effects on durables compared to non-durables. For durables, the treatment effect on the terminal share of household expenditure is 6.7 percent for the average local household in our sample, indicating a 44 percent shift in durable consumption to the e-commerce terminal among households who report to ever having used the terminal for consumption.[18] For non-durables, the treatment effect on the terminal share in household retail expenditure is 1 percent for the average household, indicating that ever-users spend on average about 11 percent of total non-durables expenditure at the terminal.[19] In contrast, we find no significant effect on household expenditures on production and business inputs (e.g. fertilizer, tools, machinery, materials, etc). Finally, note that although households do shift a small but significant share of their expenditures to the terminal, there are no significant treatment effects on total monthly retail expenditures. This result is consistent with the lack of income effect of the program that we discuss next.

To summarize, the program leads a minority of local households to take up the new e-commerce shopping option. Among the users, we find sizable effects on the substitution of total household retail expenditure to the e-commerce terminal, especially for durable consumption. These results are indicative of significant direct consumption gains for certain groups of local households. We return to the welfare computations based on these moments in the final section below.

[17] Following standard protocol, we construct monthly consumption based on the last two weeks of expenditures for non-durables (multiplied by two), and on the past three months for durables (divided by three). Usage over the past month is thus defined as either having purchased non-durables over the past two weeks, or as having purchased durables over the past three months.

[18] To compute durable consumption shares, the sample is restricted to households who buy any durables over the past three months. In this sample, the treatment effect on ever using the terminal is 15.3 percent instead of 9 percent. This yields an effect on the average durables consumption share among uptakers of 6.7/15.3=44 percent.

[19] Since all households consume non-durables, the treatment effect on uptake is as reported in Table 6, so that the average non-durables terminal share among ever users is 1/8.9=11 percent.

Incomes

The income effect of e-commerce on local producers could be both positive due to the possibility of selling online, and negative due to increased competition from the new terminal shopping option. In Table 7, we find no treatment effect on household incomes or labor supply in terms of hours worked by the primary and secondary earner. The point estimates on incomes per capita are close to zero and negative, and not statistically significant. We find no effects on either annual or monthly income, from agricultural or non-agricultural sources. In contrast to our consumption results, we find no treatment effect on online selling activity, online revenues or business creation offline or online. The point estimate on "any member of the household has ever sold online" is negative and not statistically significant. Given that the control group experienced no increase in income shares from online selling activity relative to its tiny level at baseline over this period (Table 3), these results suggest new e-commerce connections due to the program had no significant effect on the uptake of online selling activity or revenues.

We are cautious in drawing conclusions on the absence of production treatment effects from our household survey only. The 12-month period between baseline and endline surveys may be too short for local households to grow their online selling activities, and our survey sample may fail to capture rare but highly successful tail events of online businesses that could shift the local mean effect on incomes, for example. In Section 5 we use the firm's internal database to corroborate the analysis here that is based on the survey data. These data allow us to observe the universe of buying and selling transactions, and to estimate the monthly time path of adjustment within the 12 months after program installation and beyond.

Local Retail Prices

Table 8 shows the average program effects using the retail price survey data. We find no significant reduction in local store prices for identical continuing products that we observe in the same local retailer in both baseline and endline data. The point estimate is close to zero and positive and not statistically significant. Given our sampling framework, the unweighted average effect on local retail prices resembles the effect on a Laspeyres price index for local retail consumption.[20]

There is one piece of evidence suggesting potential pro-competitive effects on pre-existing local stores. The treatment effect on the number of new products per store over the past month is 4 goods, significant at the 10% level. This positive effect is large relative to the mean number of new goods of 1.4 in the baseline (5), but still small relative to the expected stock of goods in stores.[21] Furthermore, we find a negative but statistically insignificant effect for durables products. Given the small sample size of durable products observed in the villages, this could be consistent with the more pronounced treatment effect on household durable consumption above. We re-visit the plausibility and robustness of these pro-competitive forces in the heterogeneity analysis that follows below. Finally, we should again point out a limitation to the scope of the analysis that is based on our survey data collection: while we are able to estimate pro-competitive effects on

[20]Our price survey data collection mimics the data collection protocol of the IMF data dissemination standard for CPI analysis across countries. For example, this is same way that the BLS in the US or INEGI in Mexcio estimate the Laspeyres price index across product groups.

[21]We find no significant effect on store online sourcing, but this average masks significant heterogeneity with respect to the initial availability of commercial parcel delivery. We return to this result in the next section.

the local retail price environment within the villages, potential effects on retail prices in nearby towns, where households source part of their consumption, would be outside the scope of these data. We return to this possibility in the discussion of the welfare analysis in Section 6 below.

4.2 Heterogeneity Across Households and Villages

We now investigate to what extent the average effects mask significant heterogeneity across households and villages. To this end, we estimate regressions of the following form:

$$y_{hv}^{Post} = \alpha + \beta_1 Treat_v + \beta_2 X_{hv} + \beta_3 Treat_v \times X_{hv} + \gamma y_{hv}^{Pre} + \epsilon_{hv}, \tag{4}$$

where X_{hv} indicates different pre-existing household or village characteristics. As before, we report the results of specification (4) for both intent-to-treat and treatment-on-the-treated, and after replacing the binary treatments with log household residential distance to the nearest terminal location (again instrumenting with planned treatment status). We are particularly interested in the heterogeneous effect of the program with respect to pre-existing availability of commercial parcel delivery at the village level. Villages that were already serviced by commercial parcel delivery services at baseline, were essentially already connected to the same e-commerce network as city centers in the same county prior to the program's arrival.[22] This comparison thus allows us to distinguish the effect of removing both the transaction and logistical barriers to rural e-commerce (among villages without pre-existing parcel delivery), from the effect of removing only the transactional barrier (in villages with pre-existing parcel delivery). In a second step, we then investigate the heterogeneity across a basic set of pre-existing household and village characteristics: respondent age, education, household income per capita, residential distance to the planned terminal location, and a measure of the remoteness of the village (road distance to the nearest township center).

Table 9 reports the heterogeneous impact of the program with respect to pre-existing commercial parcel delivery across a number of economic outcomes. On the consumption side, we find that the average treatment effects are driven by villages that were not initially connected to commercial parcel delivery services. The average effects among the previously connected villages are precise zeroes on all outcomes that showed significant average effects in the pooled sample. This somewhat magnifies the previously reported average treatment effects on terminal consumption in the 85 percent of the village sample not previously connected to commercial parcel delivery. In these villages, slightly more than 10.5 percent of local households are induced to ever use the terminal relative to the control villages, and the average household in our sample spends 1.5% of their total retail expenditure on the e-commerce terminal over the past month. On the production side of the local economy, however, we find no significant effects in either group of villages, confirming the earlier pooled results. Considering the local retail outcomes, we now find a significant treatment effect on the number of stores sourcing their products online in villages without pre-existing delivery, and again find a treatment effect on new product varieties that is significant only in these villages. The treatment effect on local durable prices increases to minus 14.4 percent in villages without pre-existing delivery, but remains statistically insignificant at conven-

[22] According to the local program managers, e-commerce deliveries were not priced differently to the previously connected villages compared to cities.

tional levels. These results suggest that the removal of the logistical barrier to e-commerce is the main driver of the program's local economic effects, rather than the provision of an easy-to-use physical store interface in villages that were already connected to e-commerce logistically.

Table 10 extends the analysis of heterogeneous treatment effects to other household and village characteristics. We find that younger, richer households who live in closer proximity to the planned terminal, and in villages at larger distances from the nearest city center experience the most positive treatment effects on the consumption side. In particular, the average effect on uptake (roughly 10 percent of sample households relative to the control villages) is driven by, and more sizable among these groups of households. Somewhat surprisingly, we find no significant heterogeneity in household usage of the terminal with respect to the education (years of schooling) of the household respondent. We again find no significant heterogeneity in the treatment effect on the production side of the local economy.

4.3 Role of Program Implementation and GE Spillovers

So far, we find that the program has led to significant consumption-side effects among households who are induced to take up the e-commerce terminal. These households, however, do not represent a majority of the local village population, and we find no evidence of significant production-side effects. Before providing additional evidence from the firm's internal database in the next section, we use the survey data to investigate two potential explanations for the relatively muted effects for the average rural household: i) poor service and program implementation, and ii) general equilibrium spillovers on control villages.

The Role of Features of the Program Implementation

Poor terminal service compared to other shopping options could be an explanation for why the program did not attract a broader cross-section of the local population. The summary statistics in Table 2 and Table 11 point against this explanation. The new e-commerce terminal compares favorably with pre-existing shopping options because it is close, cheap and offers a wide variety of goods. To illustrate the lack of local shopping options in our survey villages, recall from Table 2 that households source more than half their retail consumption outside their village, which increases to 68 percent for durable goods. The need to travel outside the village to shop is unsurprising, given that our surveyors could not find any durable goods in local stores for about half of our sample villages as shown in Table 11. The household's main reported shopping destination outside the village is at a median distance of 10km return trip, representing a 40 minute round trip at a median cost of 4 RMB (Table 11).

In comparison, the terminal is much closer to our survey households, with a median distance to the planned terminal of 230m (Table 2), and it offers a variety of goods unavailable in local stores. 62 percent of goods bought through the terminal were not available in the village (Table 11), which rises to 84 percent for durable goods. When goods are available in both the terminal and the village, the terminal is cheaper by 15 percent (median).[23] The main shopping destination outside the village, generally the nearest township, is more competitive in terms of varieties offered (in 80 percent of cases terminal goods are available there), but the terminal is still much

[23]Our survey directly asks households, for each terminal purchase, whether the good was available in their village. If the good is available, we ask how much it would have cost.

cheaper by a median of 18 percent even before accounting for transport costs. Finally, delivery times in participating villages are close to identical to urban regions in China. These results support existing descriptive evidence on the popularity of e-commerce among urban regions of China (e.g. McKinsey (2016) and Fan et al. (2016)), and suggest that a lack of attractiveness of e-commerce is unlikely to account for the relatively muted average takeup of the new option as a result of the program.

A related question is to what extent poor planning and project implementation of the program by the e-commerce firm could account for the muted uptake of local households. This seems a priori unlikely, given the firm's high degree of professionalism, profit motive, institutional capacity and expertise, especially when compared to the resources generally available to implement public policies in developing countries. To further investigate this possibility, Table 12 presents regression results when estimating expression (4) with interaction terms for observable features of the program implementation. In particular, we test for heterogeneity in the program's effectiveness of reaching local households as a function of the terminal manager application test score, and a dummy for delay in the terminal installation with respect to the planned (and agreed upon) due date in our implementation schedule. We find that neither of these features affect take up of the terminal in a significant way. These results and the general context of the intervention suggest that a botched program implementation is unlikely the explanation behind limited household takeup.

The Role of Spillovers

We next investigate whether GE spillovers on surrounding villages could in part account for the small average effects. For example, if trade linkages with other nearby villages are an essential driver of the local economy, then it could be the case that the comparison between treated and control villages misses average income effects. If these villages are well integrated with one another, it could also be the case that store prices in surrounding villages respond to pro-competitive effects, potentially biasing toward zero the comparison between treatment and control villages. To investigate these mechanisms, we follow an approach similar to Miguel & Kremer (2004):

$$y_{hv}^{Post} = \alpha + \beta_1 Treat_v + \beta_2 Exposure_v^{treat} + \beta_3 Exposure_v^{all} + \gamma y_{hv}^{Pre} + \epsilon_{hv}, \tag{5}$$

where $Exposure_{vk}^{treat}$ measures the proximity of village v to other program villages, and $Exposure_{vk}^{all}$ measures proximity to all villages on the candidate list from which we randomly selected our control villages. Even though exposure to other program villages is not randomly assigned, our randomization means that conditional on exposure to all candidate villages, exposure to other treatment villages is plausibly exogenous. In turn, β_2 is an estimate of the the strength of cross-village spillovers.

We measure exposure as the number of treated villages within 3 or 10 km distance bins of a given village. Table 13 reports the estimation results. We find some evidence of positive spillover effects of nearby terminals within 3km of the village. These effects imply a slight increase of the total average effect of the program installation on household uptake that we estimated above. This increases from 9 percent in Table 6 to about 14 percent once we take into account positive spillovers from nearby villages, and about 13 percent in the village population when adjusted

for sampling weights. In contrast, we find no evidence of cross-village spillovers on local retail stores, or on the production side of the economy.

Summary of Findings from the Survey Data

We can summarize the results of this section as follows. On the consumption side, we find that the program leads to sizable substitution of retail expenditure among households who are induced to use the new e-commerce terminal shopping option. These households represent about 14 percent of the rural household sample and about 13 percent of the village population after adjusting for sampling weights. We find that the program's effect is subject to significant heterogeneity. The beneficiaries are on average younger, richer, live in closer proximity to the program's terminal and in villages that are more remotely located. Conditional on these characteristics, we do not find evidence that household education and occupation or the characteristics of the terminal manager are significant determinants of the program's impact. The consumption response is driven by the removal of the logistical barrier in villages with no pre-existing commercial parcel delivery, rather than by lifting additional transactional hurdles through the terminal interface. The terminals offer on average cheaper prices, more product variety and higher convenience/less travel costs. We find that the consumption effects are particularly pronounced for durable product groups, such as electronics and appliances. We also find some evidence of pro-competitive effects on the local retail environment: local store owners report significantly higher numbers of new product variety, and a higher likelihood of sourcing their products online in treated villages who did not initially have commercial parcel delivery. We do not find significant price reductions among local stores. On the production side, we find no evidence of significant effects on the local economy in terms of online selling activity, purchases of business inputs, household incomes, labor supply or entrepreneurship.

5 Additional Evidence Using the Firm's Transaction Database

In this section, we use the firm's internal transaction database to provide additional evidence on four remaining questions that are outside the scope and budget of our household survey data collection. First, are the villages in our RCT sample representative of the program's targeted villages across the Chinese countryside more broadly? Second, to what extent does seasonality and the timing of our endline data collection affect the estimation results? Third, what is the time path of adjustments on the consumption and production sides, and is terminal take-up increasing beyond our survey's 12-month post-treatment time window? And fourth, is our survey data missing rare but highly successful tail events on the production side that could shift the average effect on local household income per capita?

A described in Section 2, we have access to the universe of purchase transaction records over the period November 2015-January 2017, across roughly 11,900 village terminals that existed over this period in 5 provinces. This transaction database includes our treated RCT villages, and covers both user and terminal-level outcomes including terminal installation dates. To capture household sales through the e-commerce terminals, we also obtained access to the universe of village out-shipments for the same terminals over the period between January 2016 to April 2017. This second database provides us with the number of out-shipments and their weight in kilograms.

Are the RCT Sample Villages Representative?

One concern is that the 8 counties that our RCT study has been based on may not be representative of the Chinese countryside more broadly. To assess whether our RCT villages are representative of the population of program villages, we use the 5-province transaction database on both purchases and sales transactions to estimate regressions of the following form:

$$y_{vm} = \theta_m + \beta \text{RCTSample}_v + \gamma \text{MonthsSinceEntry}_{vm} + \epsilon_{vm},$$

where v indexes village terminals and θ_m is a set of monthly dummies indexed by m for the 15 months of operation from November 2015 to January 2017. y_{vm} is one of five terminal-level outcomes (monthly number of registered buyers, number of purchase transactions, total terminal sales, number of out-shipments and total weight of outshipments in kg), *RCTSample* is a dummy for whether the terminal is in our RCT sample, and *MonthsSinceEntry* controls for the number of months that terminal v has been in operation as of month m. The standard errors ϵ_{vm} are clustered at the terminal level.[24]

The results in Table 14 show no remarkable differences between our RCT villages and the entire set of villages in these 5 provinces. The same is true if we compare our RCT villages to all villages in our 3 survey provinces. The RCT sample seems marginally more successful on the out-shipment side, but the magnitudes are tiny. These results provide some reassurance against the potential concern that the e-commerce firm directed our team towards 8 counties that systematically differ from the program's target locations in the Chinese countryside.

Did We Collect Endline Data During Particular Months?

The timeline of pre-treatment data collection was determined by the roll-out schedule of the e-commerce firm, and we could not finance more than a single post-treatment round. As a result of these constraints, our survey cannot measure the impact of seasonality on treatment effects. We therefore use the transaction database to study seasonality effects by estimating:

$$y_{vm} = \theta_v + \beta \text{RCTMonth}_m + \gamma \text{MonthsSinceEntry}_{vm} + \epsilon_{vm},$$

where *RCTMonth* is a dummy for our survey months i.e., a dummy equal to 1 if month m is either in December, January, April or May, which are the four calendar months during which we conducted our survey. We again cluster standard errors ϵ_{vm} at the terminal level. The results are in Table 15. We find slightly higher numbers of terminal buyers during survey months relative to the rest of the calendar year, and slightly lower numbers of purchase transactions and outshipment volumes on the other hand. In both cases, the point estimates are very small: about one additional buyer per month, 4-5 less monthly transactions, and a reduction of less than one outshipment on the selling side. We conclude that seasonality is unlikely to be a significant driver of the conclusions from the RCT.

What Is the Time Path of Adjustments for Consumption and Production?

The program's objective to introduce e-commerce to all promising Chinese villages and continuous roll-out in our RCT counties imply that we cannot keep our control group untreated for

[24] With very rare exceptions there is only one terminal per village.

more than one year. However, we can use the firm's transaction data to see beyond this one-year horizon, and plot the time pattern of monthly terminal usage for both purchasing and selling starting from program installation. These plots tell us whether we can expect stronger impacts of the e-commerce terminals over time, either on the consumption or production sides.

We estimate the following event study specification:

$$y_{vm} = \theta_v + \delta_m + \sum_{j=0}^{18} \beta_j MonthsSinceEntry_{jvm} + \epsilon_{vm} \tag{6}$$

We describe the data construction to estimate this specification using the transaction dataset, but the methodology is exactly similar using shipment dataset. Each observation in equation 6 is a terminal in a month. A negative index j denotes the number of months prior to installation for terminal v and in this case the outcome y_{vm} will always be 0. A positive value of j indexes the number of month since terminal v started operation, so that β_0 is a measure of average outcomes for terminals during the month of their installation, β_1 captures averages one month after installation, and so on. We assign an index of j =18 to all observations equal or beyond 18 months after the first month of program entry, so that β_{18} captures average outcomes of terminals that have been in operation for more than 18 months. Since we have terminal and month fixed effects, each of the β_0-β_{18} are estimated relative to the omitted category that are periods pre-installation (zeros by nature of the data since ther terminals did not exist).

To estimate (6), we create a balanced panel in the sense that each of the 12 k village terminals ever observed in the raw data appears once per month in the panel, for each of the 18 months for which we have data. This panel starts in November 2015 and therefore spans terminal observations of up to 17 months pre-installation for villages connected in April 2017, to 28 months post-installation for the earliest terminals connected 10 month prior to the beginning of our data in November 2015.

Figures 1 and 2 present the event-study plots for terminal-level outcomes on the consumption and production sides. On the consumptions side, we find little evidence of increasing uptake past our survey's one-year timeline. Broadly, terminal usage appears to increase rapidly for about 2-4 months after opening, and then plateaus or declines over time. Interestingly, villagers appear to make the highest-value purchases almost immediately after the arrival of e-commerce and then switch to buying lesser-value products online.

On the production, we find evidence that village-level out-shipments increase smoothly over time after program entry, and that this increase continues beyond the 12-month window that we cover in our survey data collection. The effect increases by roughly one third when comparing the point estimate on the total weight of out-shipments 12 months post-entry to the point estimate for more than 1.5 years post-entry (including periods up to 2 years and 4 months post-entry). These results suggest that production-side adjustments may take longer to fully materialize than the 1-year horizon covered in the survey data. Having said this, the estimated effects at the village level remain relatively minor even two year post implementation. The average number of monthly out-shipments increases to about 9 in periods more than 1.5 years after the arrival of e-commerce. In turn, the combined weight of all village-level out-shipments increases to about 23 kg on average.

Are the Survey Data Missing Successful Tail Events on the Production Side?

Our survey sampling of 38 households per village may be insufficient to capture rare but very successful events on the production side. If neglected, such tail events of high-volume online businesses enabled by the terminal could in principle shift the average effect of the terminal on household incomes that we estimate as part of the RCT analysis. To investigate this issue, we use the universe of e-commerce shipments from 5 provinces over the period January 2016 to April 2017. As discussed above, we observe total shipment weight in kg, but not revenues. Figure 2 shows that the mean monthly number of e-commerce shipments out of the villages peaks around 9 with a mean total weight of less than 25 kg for the entire village.

To obtain a conservative upper-bound for these shipments' value to the local village economy, we assume that 1) all of these shipments are pure local value-added and thus 1:1 equivalent with local incomes, and 2) that the average value per kg of these shipments is as high as that of global Chinese exports (i.e. on average RMB66.5 per kg in 2015 and 2016).[25] Under these assumptions, we find that e-commerce out-shipments account for at most a 0.14 percent increase in local income per capita one-and-a-half to 2-and-a-quarter years after the arrival of e-commerce. Note that this estimate is also abstracting from the fact that we do not find a significant treatment effect of the program on out-shipments in the RCT, suggesting that the small amount of out-shipments through the terminals is on average roughly matched by other e-commerce related out-shipments among control villages that occur in absence of the program.

Summary of Findings from Transaction Database

When comparing our RCT villages to the roughly 12 thousand other villages in the transaction data, we find that they are broadly representative of the Chinese village population that is being considered to be part of the e-commerce expansion program. The periods during which we collected endline data appear to be slightly above-average for some outcomes related to terminal purchasing use, and slightly below-average for some outcomes related to purchasing price tags and village-level out-shipments. However, the point estimates are very small in magnitude, suggesting that seasonality is unlikely to confound the RCT analysis. In terms of time path of adjustment, we find little evidence on the consumption side that the program's effect takes longer to materialize than the one-year period covered by our survey. The effects occur within 2-4 months after installation and remain roughly stable afterward. On the production side, we find evidence that village-level out-shipments are increasing steadily over time after installation. The effects remain small, however, in terms of total out-shipment weights, suggesting a minor upper-bound effect on village income per capita more than 2 years post installation. Related to this, we find no evidence that our survey data collection missed rare but highly successful tail events on the production side that could have in principle shifted the village-level average effect on economic outcomes.

6 Quantification

This section combines the empirical results from the previous sections with the theoretical framework in Section 3 and the appendix to quantify the program's effect on average household welfare, decompose the underlying channels and estimate the distribution of the gains from e-

[25] From World Bank's WITS database that provides total value of Chinese exports and total weight.

commerce integration across households and villages.

Average Effect

The most robust evidence of significant treatment effects that we find in the previous sections is related to the substitution of local households' retail expenditure to the e-commerce terminal after the program is implemented. As discussed in Section 3, these treatment effects enter the direct price index effect as part of the consumer gains due to the program. Even though it is impossible to directly observe the implicit price index changes due to the arrival of a new retail shopping option (that includes differences in prices, product variety as well as shopping amenities), we can use existing knowledge about the slope of household demand across retail shopping options in order to quantify the change in consumption value that is consistent with the observed changes in household expenditure on the ground.

As derived in the appendix, the expression for the direct consumer gains from the arrival of the e-commerce terminal, expressed as a percentage of initial household expenditure, is given by:

$$\frac{DE}{e(\mathbf{P}_T^{0*}, \mathbf{P}_C^0, \mathbf{P}_E^{0*}, \mathbf{P}_X^0, u_h^0)} = \prod_{g \in G} \left((\sum_{s \in S_g^C} \phi_{gsh}^1)^{\frac{1}{\sigma_g - 1}} \right)^{\alpha_{gh}} - 1 \tag{7}$$

where $\sum_{s \in S_g^C} \phi_{gsh}^1$ is the share of retail expenditure that is not spent on the new e-commerce terminal post-intervention, σ_g is the elasticity of substitution across retail options to source consumption in product group g, and α_{gh} is the Cobb-Douglas expenditure share on that product group for household group h.

To estimate this expression empirically, we require information on the program's effect on $\sum_{s \in S_g^C} \phi_{gsh}^1$ as well as σ_g and α_{gh}. For the α_{gh}, we use our baseline survey data on household expenditure shares across product group. For ex-post expenditure shares on the new e-commerce option, we use the treatment effects among the 85 percent of villages without pre-existing parcel delivery connections reported in Table 9. These villages experienced the removal of both logistical and transactional barriers to e-commerce integration, which is the counterfactual that we focus on for the quantification exercise. We include the intercept among control villages in these treatment effects to account for positive spillovers (Table 13).

We perform this welfare computation for three different groups of local households: for the average sample household (treatment effect of 1.6 percentage points for terminal share of total retail consumption), for households who report ever having used the terminal for purchases (treatment effect of 14 percentage points), and for households who have used the terminal for purchases during the consumption period covered by our survey data (treatment effect of 25.5 percentage points). Given the important discrepancy in treatment effects across categories of goods documented in the previous sections, we estimate welfare effects separately for durables and non-durable retail consumption.

Note that these treatment effects give equal weight to all households in our endline data. To obtain welfare estimates that are representative at the village-level, we also re-estimate the treatment effects after weighting each household in our sample according to the fraction of the village population that resides within its sampling zone (inner or outer) in our endline data. These estimates are slightly smaller, but very similar (1.5, 11 and 24 percentage points respectively),

suggesting that our sampling procedure did not significantly distort the average household in the village by much. For exposition, we provide welfare estimates both with and without re-weighting households.

For σ_g, the final set of required moments in (7), we use the closest existing estimate of consumer demand across retailer choices in an emerging market context from recent work by Atkin et al. (in press) in Mexico. In particular, we use demand parameter estimates for households in Mexico with incomes comparable to that of rural Chinese households in our survey. For non-durables consumption, the baseline parameter is $\sigma_N = 3.87$, and for durables consumption the baseline parameter is $\sigma_D = 3.85$.

To obtain standard errors for the welfare evaluation, we take into account that the treatment effects on ex-post e-commerce consumption shares are point estimates, not data points. We bootstrap the computation of expression 7 across 1000 iterations with random re-sampling. Each iteration uses the mean and standard deviation of the estimated treatment effect on terminal share of retail consumption for durables and non-durables, and for each of the three household groups discussed above, and draws from a normal distribution around the mean of the respective point estimate.

Table 16 reports the estimation results. The average reduction in retail cost of living among households who experienced the lifting of both logistical and transactional barriers is 0.81 percent. This effect increases to 5.5 percent among the roughly 14 percent of households who were induced to have ever used the terminal for purchases, and to 13.2 percent for the roughly 7 percent of households who used the terminal during the consumption period covered by our endline data collection. These effects are slightly lower at 0.71, 4.8 and 11.3 percent respectively when weighting our sample households to represent the village population as a whole. Underlying these effects are strong consumer gains in durable consumption: 2.9 percent for the average village household, 16.6 percent for ever-users and 45.8 percent for households who used the terminal in the past month.

Distribution of the Gains from E-Commerce Integration

We now investigate the distribution of the gains from the arrival of e-commerce across households and villages. We use treatment effects from the heterogeneity specification in the last rows of Table 10, which includes all interactions with program treatment. We estimate this specification with the dependent variable being either household terminal share in durable retail consumption or in non-durable retail consumption. For each sample household living in treatment villages without pre-existing parcel delivery, we compute a fitted value of the treatment effect on terminal retail consumption shares based on the primary earner's age, education, income per capita, residential distance to the planned terminal as well as distance to the nearest township center (remoteness).

We use these estimated effects for $\sum_{s \in S_g^C} \phi_{gsh}^{t1}$ in expression (7), and then plot the effect on household retail price indices flexibly across all sample households in treated villages. Figure 3 shows these plots for household income per capita (upper left), respondent age (upper right), distance to terminal (lower left) and distance to the nearest township center (lower right). These plots quantify the distribution of the gains to the average household, without restricting attention to uptakers. The confidence intervals in these figures are based on sampling variation in

household characteristic on the x-axis after clustering standard errors at the village-level.

The income plot shows that households in the 5th percentile of the income distribution on average experience a 0.25 percent reduction in cost of living due to terminal arrival, which roughly quadruples to more than 1 percent for households at the 95th income percentile. A household with a 20 year old primary earner on average experiences gains of close to 2 percent, which drops below 1 percent past the age of 40. The gains are close to 1.5 percent on average in close residential proximity to the terminal and decrease to on average less than half a percent toward the largest distances in the sample. In contrast, villages in close proximity to the nearest township center experience small average gains that more than quadruple as the distance becomes larger in the sample.

Overall, these figures confirm the heterogeneity of the program's impact on consumption uptake that we discussed above and report in Table 10. In particular, we find that richer, younger households, who live in closer proximity to the e-commerce terminal and in villages that are farther away from existing shopping options in nearby township centers benefit significantly more from the arrival of e-commerce..

Quantification Across Alternative Parameter Values

To account for uncertainty in the demand parameters, we compute results across alternative values of σ_N and σ_D, relative to the baseline parameterization ($\sigma_N = 3.87$ and $\sigma_D = 3.85$). In particular, we allow for household shopping demand to be both more or less price elastic across retailer options. Intuitively, the less price sensitive households are across retailers (i.e. the lower σ_g), the higher will be the implied consumer gains that are consistent with the observed household substitution to the new shopping option. The low-elasticity scenario is $\sigma_N = 2.87$ and $\sigma_D = 2.85$ while the high-elasticity scenario is $\sigma_N = 4.87$ and $\sigma_D = 4.85$. A priori, it is unclear which scenario is more likely in our current empirical context, relative to the baseline parameters estimated in Atkin et al. (in press) for similarly poor Mexican households living in urban areas. Rural Chinese households may be less sensitive to effective price differences across retailers due to higher shopping travel costs to a nearby town compared to urban Mexicans. Conversely, rural Chinese households may be intrinsically more price sensitive than Mexicans with similar real incomes.

Table 17 reports the estimation results. As discussed above, assuming that rural Chinese shopping demand is somewhat less price elastic across retailer options yields significantly larger estimated welfare gains in retail consumption: 1.25 percent for the average household in our sample, 8.5 percent for households that were ever terminal consumers, and 21 percent for households that are current terminal consumers. The baseline estimates were 0.81, 5.5 and 13.2 percent respectively. Conversely, assuming more price elastic shopping demand yields slightly smaller welfare effects of 0.6, 4 and 9.6 percent.

Summary of Results

We find that the program leads to sizable gains in real income among households who are induced to use the e-commerce terminal, who represent about 14 percent of the rural household sample and about 13 percent of the village population adjusting for sampling weights. The welfare gains for the average rural household are more muted, suggesting strong heterogeneity in the effect of the arrival of e-commerce rather than broad-based welfare gains. The beneficiaries

are on average richer and younger, live in closer proximity to the e-commerce terminal, and in villages that are farther away from the nearest township center. The welfare gains are driven by a significant reduction in household cost of living due to greater product variety, cheaper prices and a reduction in travel costs. These gains are strongest for durable product groups such as electronics and appliances.

7 Conclusion

The potential of e-commerce integration as a driver of economic development has featured prominently in recent policy reports and the popular press. In this context, the Chinese government has recently launched the first nation-wide e-commerce expansion program to remove the barriers to e-commerce development outside of cities. As internet access has already widely spread in the countryside, this program aims to invest in removing two main remaining barriers to e-commerce: the lack of modern transport logistics necessary for commercial parcel delivery and pickup (logistical barrier), and the transitioning to non-traditional online user interfaces and paperless payments (transactional barrier).

This paper uses this empirical context to study the economic consequences of e-commerce integration on the local economy, the underlying channels, and the distribution of the gains from e-commerce across households and villages. To this end, we combine an RCT that we implement across villages in collaboration with a large Chinese e-commerce firm with a new collection of microdata on household consumption, production and retail prices in the Chinese countryside.

The analysis provides several insights. We find that the program leads to sizable gains in real incomes among rural households who are induced to use the e-commerce terminal. These users represent about 14 percent of the rural household sample and 13 percent of the village population after adjusting for sampling weights. For the average rural household, including non-users, these gains are statistically significant but more muted. Underlying these effects, we find strong heterogeneity across households and villages. The beneficiaries are on average significantly younger, richer, live in closer proximity to the e-commerce terminal and in villages that are relatively more remote. Conditional on these characteristics, we do not find evidence that household education and occupational status or the characteristics of the terminal managers affect the extent of household gains from e-commerce.

In terms of channels, we find significantly stronger economic gains among villages that were not previously serviced by commercial parcel delivery, suggesting that most of the program's gains are due to overcoming the logistical barrier, rather than the transactional one. On the consumption side, we find that the e-commerce terminals offer on average lower prices, higher convenience and increased product variety compared to pre-existing local retail choices, both within the village and in nearby towns. The gains in household purchasing power are strongest for durable product groups, such as electronics and appliances. We also find suggestive evidence that the program led to additional product variety in pre-existing local stores, as their managers source new products through e-commerce. We find no evidence of significant pro-competitive effects on local retailer prices, on the other hand. On the production side, we find no evidence for significant effects on the local economy: online selling activity, purchases of production inputs, household incomes and entrepreneurship are not significantly affected by the arrival of the

program. Overall, we find that the gains from e-commerce are driven by a reduction in local household cost of living that is mainly due to the direct gains from access to the new e-commerce shopping option for local households.

Using the firm's database, we find little evidence on the consumption side suggesting that the adjustment takes longer than one year: the consumption-side uptake materializes within 2-4 months of entry and then remain mostly constant over time. On the production side, we find evidence that village-level out-shipments significantly increase over time beyond the 12-month window. However, the effect on total out-shipments remains relatively minor after more than two years post program entry, with a small upper-bound effect on local household incomes. Related to this, we do not find evidence that the survey data fails to pick up highly successful but rare tail events on the production side that could in principle shift the mean effect on local household outcomes.

Overall, our findings suggest that e-commerce expansions offer significant economic gains to certain groups of the rural population, rather than being broad-based. Compared to the recent case studies highlighting a set of highly successful rural e-commerce production hubs, our analysis reveals that a quite particular mix of local factors must be underlying these prominent success stories. In the absence of complementary interventions, such as for example business training, access to credit or targeted online promotions, large and significant production-side effects appear unlikely to materialize for the average rural market place. In this light, future work aimed at better understanding the factors under which the arrival of e-commerce can have transformative impacts on the production side of the rural economy seems a promising agenda for future research in this area.

References

Anderson, S. P., De Palma, A., & Thisse, J. F. (1992). *Discrete choice theory of product differentiation*. MIT press.

Atkin, D., Faber, B., & Gonzalez-Navarro, M. (in press). Retail globalization and household welfare: Evidence from mexico. *Journal of Political Economy*.

Banerjee, A., Duflo, E., & Qian, N. (2012). *On the road: Access to transportation infrastructure and economic growth in china* (Tech. Rep.). National Bureau of Economic Research.

Baum-Snow, N., Brandt, L., Henderson, J. V., Turner, M. A., & Zhang, Q. (2016). Roads, railroads and decentralization of Chinese cities. *Forthcoming, Review of Economics and Statistics*.

Brynjolfsson, E., Hu, Y., & Smith, M. D. (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science*, *49*(11), 1580–1596.

Couture, V. (2016). Valuing the consumption benefits of urban density. *Working Paper, UC Berkeley*.

Donaldson, D. (in press). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*.

Eaton, J., & Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 1741–1779.

Faber, B. (2014). Trade integration, market size, and industrialization: Evidence from China's national trunk highway system. *The Review of Economic Studies*, *81*(3), 1046–1070.

Fan, J., Tang, L., Zhu, W., & Zou, B. (2016). The alibaba effect: Spatial consumption inequality and the welfare gains from e-commerce. *Michigan State University mimeo*.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review*, *84*(1), 155–57.

FRED. (2016). E-commerce retail sales as a percent of total sales (ecompctsa). *Economic Data, Federal Reserve of St. Louis*.

Goldmanis, M., Hortaçsu, A., Syverson, C., & Emre, Ö. (2010). E-commerce and the market structure of retail industries. *The Economic Journal*, *120*(545), 651–682.

Goyal, A. (2010). Information, direct access to farmers, and rural market performance in central india. *American Economic Journal: Applied Economics*, *2*(3), 22–45.

Hamory, J., Kleemans, M., Li, N., & Miguel, E. (2016). Individual ability and selection into migration in Kenya. *Mimeo, UC Berkeley*.

Handbury, J. (2013). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across us cities. *Working Paper, Wharton*.

Handbury, J., & Weinstein, D. E. (2015). Goods prices and availability in cities. *The Review of Economic Studies*, *82*(1), 258–296.

Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In *The economics of new goods* (pp. 207–248). University of Chicago Press.

Hausman, J. A., & Leonard, G. K. (2002). The competitive effects of a new product introduction: A case study. *Journal of Industrial Economics*, *50*(3), 237–263.

Hicks, J. R. (1940). The valuation of the social income. *Economica*, 105–124.

Hjort, J., & Poulsen, J. (2016). The arrival of fast internet and skilled job creation in Africa. *Columbia University mimeo*.

Lagakos, D., Mobarak, M., & Waugh, M. (2016). Urban-rural wage gaps in developing countries: spatial misallocation or efficient sorting. *Mimeo, UC San Diego, Yale, and NYU*.

MCIT. (2016). Egypt's national e-commerce strategy. *Ministry of Communications and Information Technology Report*.

McKenzie, D. (2012). Beyond baseline and follow-up: The case for more t in experiments. *Journal of Development Economics*, *99*(2), 210–221.

McKinsey. (2016). China's e-tail revolution. *Report, McKinsey Global Institute*.

MEITY. (2016). Digital india: Driving e-commerce in rural and semi urban India. *Ministry of Electronics and IT Report*.

Miguel, E., & Kremer, M. (2004). Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, *72*(1), 159–217.

PFSweb. (2016). China's e-commerce market 2015. *Annual Market Report*.

PM. (2016). Vietnam's e-commerce development masterplan. *Vietnam's Office of the Prime Minister Report*.

Statista. (2016). Online retail statistics for china. *Market Statistics*.

Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics*, *2*(4), 1–41.

UNCTAD. (2016a). E-commerce opens new opportunities for developing countries. *UNCTAD/PRESS/U14/IN/2016/011*.

UNCTAD. (2016b). Unctad e-commerce index 2016. *Report, United Nations Conference on Trade and Development, Geneva*.

WTO. (2013). E-commerce in developing countries: Opportunities and challenges for small and medium-sized enterprises. *Policy Report, World Trade Organization, Geneva*.

Young, A. (2013). Inequality, the urban-rural gap and migration. *The Quarterly Journal of Economics*, 1727–1785.

8 Figures and Tables

Table 1: Baseline Data: Individual Level

		Full Sample at Baseline	Treatment Villages at Baseline	Control Villages at Baseline	P-Value (Treat-Control=0)	Control Villages at Endline
Age	Median	44.000	44.000	43.000		46.000
	Mean	38.950	39.329	38.407	0.208	39.943
	Standard Deviation	23.580	23.658	23.460		23.759
	Number of Obs	8491	5001	3490		4194
Gender (Female=1)	Median	1.000	1.000	1.000		1.000
	Mean	0.534	0.526	0.546	0.025	0.537
	Standard Deviation	0.499	0.499	0.498		0.499
	Number of Obs	8484	5001	3483		4188
Employed (for age>15) (Yes=1)	Median	1.000	1.000	1.000		1.000
	Mean	0.767	0.766	0.769	0.882	0.762
	Standard Deviation	0.423	0.424	0.422		0.426
	Number of Obs	6070	3590	2480		3015
Peasant (for age>15) (Yes=1)	Median	1.000	1.000	1.000		1.000
	Mean	0.527	0.527	0.526	0.971	0.513
	Standard Deviation	0.499	0.499	0.499		0.500
	Number of Obs	6369	3760	2609		3144
No Schooling (for age>15) (No School=1)	Median	0.000	0.000	0.000		0.000
	Mean	0.270	0.273	0.266	0.745	0.319
	Standard Deviation	0.444	0.446	0.442		0.466
	Number of Obs	6368	3758	2610		3132
Completed Junior High School (for age>15) (Yes=1)	Median	0.000	0.000	0.000		0.000
	Mean	0.437	0.429	0.449	0.419	0.422
	Standard Deviation	0.496	0.495	0.498		0.494
	Number of Obs	6368	3758	2610		3132
Completed Senior High School (for age>18) (Yes=1)	Median	0.000	0.000	0.000		0.000
	Mean	0.104	0.104	0.104	0.969	0.097
	Standard Deviation	0.305	0.305	0.305		0.296
	Number of Obs	6286	3719	2567		3096

Notes: See Section 2 for discussion.

Table 2: Baseline Data: Household Level

		Full Sample at Baseline	Treatment Villages at Baseline	Control Villages at Baseline	P-Value (Treat-Control=0)	Control Villages at Endline
Age of Primary Earner	Median	50.000	50.000	50.000		52.00
	Mean	49.824	49.953	49.631	0.634	51.395
	Standard Deviation	12.673	12.710	12.621		13.55
	Number of Obs	2548	1530	1018		1348
Gender of Primary Earner (Female=1)	Median	0.000	0.000	0.000		0.00
	Mean	0.288	0.295	0.276	0.457	0.295
	Standard Deviation	0.453	0.456	0.447		0.46
	Number of Obs	2547	1530	1017		1348
Primary Earner Went to School (Yes=1)	Median	1.000	1.000	1.000		1.00
	Mean	0.815	0.814	0.817	0.874	0.750
	Standard Deviation	0.388	0.389	0.386		0.43
	Number of Obs	2550	1531	1019		1342
Primary Earner Is Peasant (Yes=1)	Median	1.000	1.000	1.000		1.00
	Mean	0.590	0.600	0.577	0.620	0.587
	Standard Deviation	0.492	0.490	0.494		0.49
	Number of Obs	2549	1531	1018		1348
Primary Earner Self-Employed (Yes=1)	Median	0.000	0.000	0.000		0.00
	Mean	0.073	0.087	0.053	0.036	0.072
	Standard Deviation	0.261	0.282	0.224		0.26
	Number of Obs	2549	1531	1018		1348
Household Size	Median	3.000	3.000	3.000		3.00
	Mean	3.114	3.053	3.205	0.075	2.987
	Standard Deviation	1.422	1.420	1.421		1.40
	Number of Obs	2740	1647	1093		1405
Household Monthly Income Per Capita in RMB	Median	350.000	339.000	375.000		466.67
	Mean	876.412	841.198	929.473	0.365	1028.960
	Standard Deviation	1717.456	1687.169	1761.560		2005.31
	Number of Obs	2740	1647	1093		1405
Household Monthly Retail Expenditure Per Capita in RMB	Median	381.000	372.833	400.500		364.00
	Mean	732.017	663.034	835.966	0.135	686.616
	Standard Deviation	2304.540	1139.788	3368.220		1512.06
	Number of Obs	2735	1644	1091		1405
Household Monthly Expenditure on Business Inputs Per Capita in RMB	Median	0.000	0.000	0.000		0.00
	Mean	123.417	123.007	124.033	0.981	128.464
	Standard Deviation	1033.757	1076.656	966.070		1069.52
	Number of Obs	2736	1644	1092		1405
Any Member of the Household Has Ever Used the Internet (Yes=1)	Median	0.000	0.000	0.000		0.00
	Mean	0.368	0.354	0.390	0.249	0.427
	Standard Deviation	0.482	0.478	0.488		0.49
	Number of Obs	2739	1646	1093		1402
Household Owns a Smartphone (Yes=1)	Median	1.000	1.000	1.000		1.00
	Mean	0.526	0.509	0.552	0.153	0.551
	Standard Deviation	0.499	0.500	0.498		0.50
	Number of Obs	2731	1642	1089		1400

Notes: See Section 2 for discussion.

Table 3: Baseline Data: Household Level – Continued

		Full Sample at Baseline	Treatment Villages at Baseline	Control Villages at Baseline	P-Value (Treat-Control=0)	Control Villages at Endline
Share of Household Monthly Expenditure on E-Commerce Deliveries	Median	0.000	0.000	0.000		0.00
	Mean	0.007	0.006	0.007	0.693	0.008
	Standard Deviation	0.050	0.046	0.057		0.05
	Number of Obs	2720	1637	1083		1397
Share of E-Commerce Sales in Household Monthly Income	Median	0.000	0.000	0.000		0.00
	Mean	0.003	0.001	0.006	0.103	0.003
	Standard Deviation	0.052	0.030	0.074		0.05
	Number of Obs	2055	1244	811		1161
Distance in Meters to Planned Terminal Location	Median	231.556	232.891	231.454		203.63
	Mean	290.346	293.364	285.797	0.789	286.631
	Standard Deviation	243.450	247.778	236.820		267.06
	Number of Obs	2740	1647	1093		1405
Share of Retail Expenditure Outside of Village	Median	0.553	0.489	0.623		0.60
	Mean	0.500	0.470	0.545	0.193	0.531
	Standard Deviation	0.395	0.402	0.379		0.38
	Number of Obs	2720	1637	1083		1397
Share of Business Input Expenditure Outside of Village	Median	1.000	1.000	1.000		1.00
	Mean	0.613	0.610	0.618	0.916	0.633
	Standard Deviation	0.465	0.470	0.457		0.46
	Number of Obs	926	558	368		544
Travel Time One-Way to Main Shopping Destination Outside Village (minutes)	Median	20.000	20.000	20.000		20.00
	Mean	29.892	29.941	29.826	0.962	28.862
	Standard Deviation	27.825	27.380	28.429		26.19
	Number of Obs	2234	1284	950		1188
Travel Cost One-Way to Main Shopping Destination Outside Village (RMB)	Median	2.000	2.000	1.500		1.00
	Mean	3.739	3.847	3.591	0.715	4.236
	Standard Deviation	10.092	11.774	7.196		16.78
	Number of Obs	2216	1278	938		1185
Household Owns a PC or Laptop (Yes=1)	Median	0.000	0.000	0.000		0.00
	Mean	0.283	0.276	0.295	0.631	0.284
	Standard Deviation	0.451	0.447	0.456		0.45
	Number of Obs	2731	1642	1089		1400
Household Owns a Car (Yes=1)	Median	0.000	0.000	0.000		0.00
	Mean	0.108	0.107	0.110	0.851	0.131
	Standard Deviation	0.311	0.309	0.313		0.34
	Number of Obs	2731	1642	1089		1400
Household Owns a Motorcycle (Yes=1)	Median	0.000	0.000	1.000		0.00
	Mean	0.486	0.456	0.532	0.031	0.467
	Standard Deviation	0.500	0.498	0.499		0.50
	Number of Obs	2731	1642	1089		1400
Household Owns a TV (Yes=1)	Median	1.000	1.000	1.000		1.00
	Mean	0.977	0.977	0.977	0.953	0.977
	Standard Deviation	0.149	0.148	0.150		0.15
	Number of Obs	2731	1642	1089		1400

Notes: See Section 2 for discussion.

Table 4: Baseline Data: Local Retail Prices

		Full Sample at Baseline	Treatment Villages at Baseline	Control Villages at Baseline	P-Value (Treat-Control=0)	Control Villages at Endline
Number of Stores at Village Level	Median	3.00	3.00	2.00		2.00
	Mean	4.15	4.38	3.79	0.33	3.61
	Standard Deviation	2.94	2.91	2.98		2.99
	Number of Obs	99	60	39		38
Establishment Space in Square Meters	Median	50.00	50.00	40.00		50.00
	Mean	99.07	74.42	146.76	0.35	121.33
	Standard Deviation	320.38	89.60	532.73		375.35
	Number of Obs	361	238	123		126
Number of Establishment's New Products Added Over Last Month	Median	0.00	0.00	0.00		0.00
	Mean	1.43	1.56	1.17	0.57	0.63
	Standard Deviation	7.44	8.88	3.42		2.26
	Number of Obs	330	215	115		126
Prices of All Retail Consumption (9 Product Groups) in RMB	Median	7.00	7.00	6.00		6.00
	Mean	71.03	76.74	61.43	0.47	71.23
	Standard Deviation	411.24	433.67	370.33		390.31
	Number of Obs	9382	5884	3498		3259
Price Was Not Displayed on Label (Needed to Ask=1)	Median	1.00	1.00	1.00		1.00
	Mean	0.67	0.66	0.67	0.97	0.73
	Standard Deviation	0.47	0.47	0.47		0.44
	Number of Obs	8977	5597	3380		3370
Prices of Business or Production Input in RMB	Median	10.00	10.00	8.80		9.00
	Mean	45.63	42.88	49.78	0.76	43.84
	Standard Deviation	195.09	206.23	177.46		97.92
	Number of Obs	444	267	177		111
(1) Prices of Food and Beverages in RMB	Median	4.38	4.60	4.00		4.00
	Mean	11.58	11.81	11.21	0.73	10.05
	Standard Deviation	24.35	23.31	25.99		17.75
	Number of Obs	4853	3021	1832		1834
(2) Prices of Tobacco and Alcohol in RMB	Median	12.00	13.00	12.00		13.00
	Mean	28.81	30.35	26.36	0.46	29.32
	Standard Deviation	53.97	59.45	43.77		55.16
	Number of Obs	1331	818	513		531
(3) Prices of Medicine and Health Products in RMB	Median	10.00	10.00	9.98		8.40
	Mean	26.13	24.40	29.31	0.66	18.50
	Standard Deviation	43.35	38.46	51.11		33.77
	Number of Obs	399	258	141		90
(4) Prices of Clothing and Accessories in RMB	Median	15.00	12.00	20.00		22.00
	Mean	46.31	45.69	47.79	0.90	57.00
	Standard Deviation	74.71	71.49	82.13		85.66
	Number of Obs	401	282	119		65
(5) Prices of Other Everyday Products in RMB	Median	10.00	10.00	9.00		9.00
	Mean	14.68	14.53	14.93	0.93	13.10
	Standard Deviation	31.03	32.69	28.06		18.17
	Number of Obs	1462	916	546		626
(6) Prices of Fuel and Gas in RMB	Median	5.00	5.00	5.00		5.83
	Mean	11.65	15.36	8.08	0.26	5.82
	Standard Deviation	21.46	28.88	9.59		0.23
	Number of Obs	53	26	27		4
(7) Prices of Furniture and Appliances in RMB	Median	110.00	85.00	187.00		398.00
	Mean	1009.49	1001.66	1026.34	0.95	1167.30
	Standard Deviation	1504.81	1583.03	1333.52		1350.70
	Number of Obs	183	125	58		43
(8) Prices of Electronics in RMB	Median	449.00	609.50	17.50		1799.00
	Mean	917.05	976.41	782.14	0.59	1782.71
	Standard Deviation	1224.37	1242.82	1184.20		871.58
	Number of Obs	144	100	44		45
(9) Prices of Transport Equipment in RMB	Median	1440.00	1980.00	30.00		2800.00
	Mean	1700.66	1794.74	1534.21	0.71	2578.24
	Standard Deviation	1822.07	1770.33	1922.34		1697.82
	Number of Obs	108	69	39		21

Notes: See Section 2 for discussion.

Table 5: Firm's Transaction Data

group	N_Transactions	N_Users	N_Shipments	N_Terminals	N_Counties
Full Sample	27270532	3785019	500743	11941	175
3 Provinces	20647373	2832872	442319	8561	116
8 Counties	1835897	216529	44148	706	8
RCT Villages	130769	15099	3158	40	8

group	N_Provinces	N_Days	N_Months	N_ProductGroups	Sum_Payments	Sum_Weight
Full Sample	5	547	18	42	4480424896	1169673
3 Provinces	3	547	18	42	3409227245	1019373
8 Counties	3	503	17	42	330930097	95908
RCT Villages	3	482	16	41	17618900	7817

Notes: The table provides information from both the purchasing and sales transaction databases.
N_Transactions is the number of village purchase transactions. N_Users is the number of village users making purchases. N_Shipments is the number of village out-shipments through the e-commerce terminals. Sum_Payments is the sum of village purchases. Sum_Weight is the sum of the weight of all village out-shipments through e-commerce. See Section 2 for discussion.

Table 6: Average Effects: Consumption

Dependent Variables		Intent to Treat	Treatment on Treated	Log Distance (IV using
Monthly Total Retail Expenditure Per Capita	Treat or Log Dist	-21.93	-40.92	11.72
		(31.96)	(60.19)	(17.09)
	R-Squared	0.038		
	First Stage F-Stat		43.92	38.31
	Number of Obs	3,434	3,434	3,434
Household Has Ever Bought Something at Terminal (Yes=1)	Treat or Log Dist	0.0480***	0.0886***	-0.0253***
		(0.0169)	(0.0271)	(0.00801)
	R-Squared	0.008		
	First Stage F-Stat		45.56	39.22
	Number of Obs	3,518	3,518	3,518
Household Has Bought Something at Terminal in Past Month (Yes=1)	Treat or Log Dist	0.0263***	0.0490***	-0.0140***
		(0.00981)	(0.0171)	(0.00484)
	R-Squared	0.009		
	First Stage F-Stat		43.93	37.70
	Number of Obs	3,482	3,482	3,482
Share of Online Terminal in Total Monthly Retail Expenditure	Treat or Log Dist	0.00666***	0.0124***	-0.00355***
		(0.00239)	(0.00434)	(0.00126)
	R-Squared	0.006		
	First Stage F-Stat		44.03	37.84
	Number of Obs	3,434	3,434	3,434
Share of Terminal in Monthly Business Inputs	Treat or Log Dist	-0.00715	-0.0154	0.00481
		(0.00778)	(0.0191)	(0.00615)
	R-Squared	0.003		
	First Stage F-Stat		16.46	11.70
	Number of Obs	1,207	1,207	1,207
Share of Terminal in Monthly Non-Durables	Treat or Log Dist	0.00536***	0.00999***	-0.00286***
		(0.00195)	(0.00355)	(0.00103)
	R-Squared	0.003		
	First Stage F-Stat		44.11	37.83
	Number of Obs	3,433	3,433	3,433
Share of Terminal in Monthly Durables	Treat or Log Dist	0.0398**	0.0669**	-0.0220**
		(0.0159)	(0.0261)	(0.00870)
	R-Squared	0.011		
	First Stage F-Stat		52.64	25.14
	Number of Obs	768	768	768
Share of Terminal in Monthly Food and Beverages (1)	Treat or Log Dist	0.00121	0.00223	-0.000634
		(0.000823)	(0.00152)	(0.000436)
	R-Squared	0.001		
	First Stage F-Stat		45.63	39.31
	Number of Obs	3,359	3,359	3,359
Share of Terminal in Monthly Tobacco and Alcohol (2)	Treat or Log Dist	0.000608	0.00123	-0.000373
		(0.000515)	(0.00109)	(0.000326)
	R-Squared	0.001		
	First Stage F-Stat		33.02	23.77
	Number of Obs	1,653	1,653	1,653
Share of Terminal in Monthly Medicine and Health Products (3)	Treat or Log Dist	0.000693	0.00126	-0.000358
		(0.000689)	(0.00124)	(0.000352)
	R-Squared	0.000		
	First Stage F-Stat		51.06	43.31
	Number of Obs	2,416	2,416	2,416
Share of Terminal in Monthly Clothing and Accessories (4)	Treat or Log Dist	0.0465***	0.0734***	-0.0227***
		(0.0140)	(0.0216)	(0.00678)
	R-Squared	0.019		
	First Stage F-Stat		70.69	41.66
	Number of Obs	1,269	1,269	1,269
Share of Terminal in Monthly Other Household Products (5)	Treat or Log Dist	0.00430	0.00804	-0.00241
		(0.00395)	(0.00713)	(0.00215)
	R-Squared	0.001		
	First Stage F-Stat		43.87	32.53
	Number of Obs	2,336	2,336	2,336
Share of Terminal in Monthly Heating, Fuel and Gas (6)	Treat or Log Dist	0	0	0
		(0)	(0)	(0)
	R-Squared	.		
	First Stage F-Stat		.	.
	Number of Obs	1,463	1,463	1,463
Share of Terminal in Monthly Furniture and Appliances (7)	Treat or Log Dist	0.0546**	0.0908**	-0.0302**
		(0.0217)	(0.0368)	(0.0119)
	R-Squared	0.019		
	First Stage F-Stat		47.51	22.10
	Number of Obs	380	380	380
Share of Terminal in Monthly Electronics (8)	Treat or Log Dist	0.0697**	0.110**	-0.0446**
		(0.0345)	(0.0522)	(0.0225)
	R-Squared	0.024		
	First Stage F-Stat		43.20	11.48
	Number of Obs	232	232	232
Share of Terminal in Monthly Transport Equipment (9)	Treat or Log Dist	0.0353*	0.0554*	-0.0162*
		(0.0201)	(0.0313)	(0.00942)
	R-Squared	0.014		
	First Stage F-Stat		43.07	36.75
	Number of Obs	141	141	141

Notes: See Section 4 for discussion. Standard errors are clustered at the level of villages. * 10%, ** 5%, *** 1% significance levels.

Table 7: Average Effects: Incomes

Dependent Variables		Intent to Treat	Treatment on Treated	Log Distance (IV using Treat)
Monthly Income Per Capita in RMB	Treat or Log Dist	-7.838	-14.48	4.190
		(70.78)	(129.9)	(37.55)
	R-Squared	0.038		
	First Stage F-Stat		45.33	37.69
	Number of Obs	3,437	3,437	3,437
Monthly Income Per Capita Net of Transfers in RMB	Treat or Log Dist	-12.55	-23.21	6.711
		(72.18)	(132.4)	(38.24)
	R-Squared	0.051		
	First Stage F-Stat		45.16	37.66
	Number of Obs	3,445	3,445	3,445
Annual Income Per Capita in RMB	Treat or Log Dist	-45.95	-85.08	24.74
		(586.9)	(1,080)	(314.2)
	R-Squared	0.046		
	First Stage F-Stat		44.77	36.69
	Number of Obs	3,388	3,388	3,388
Monthly Agricultural Income Per Capita	Treat or Log Dist	-70.23	-130.3	37.53
		(140.3)	(257.7)	(74.07)
	R-Squared	0.033		
	First Stage F-Stat		44.23	37.43
	Number of Obs	3,448	3,448	3,448
Monthly Non-Agricultural Income Per Capita	Treat or Log Dist	-46.65	-86.06	24.86
		(137.3)	(249.6)	(72.25)
	R-Squared	0.157		
	First Stage F-Stat		45.74	38.23
	Number of Obs	3,441	3,441	3,441
Weekly Hours Worked Primary Earner in Hours	Treat or Log Dist	1.295	2.418	-0.696
		(3.429)	(6.392)	(1.844)
	R-Squared	0.000		
	First Stage F-Stat		43.71	36.73
	Number of Obs	3,324	3,324	3,324
Weekly Hours Worked Secondary Earner in Hours	Treat or Log Dist	0.334	0.609	-0.185
		(3.969)	(7.212)	(2.194)
	R-Squared	0.000		
	First Stage F-Stat		44.28	34.71
	Number of Obs	1,878	1,878	1,878

Dependent Variables		Intent to Treat	Treatment on Treated	Log Distance (IV using Treat)
Member of Household Has Ever Sold Online (Yes=1)	Treat or Log Dist	-0.00700	-0.0129	0.00370
		(0.00562)	(0.0104)	(0.00300)
	R-Squared	0.347		
	First Stage F-Stat		45.30	38.28
	Number of Obs	3,504	3,504	3,504
Member of Household Has Sold Online In Past Two Weeks (Yes=1)	Treat or Log Dist	-0.00132	-0.00244	0.000700
		(0.00237)	(0.00438)	(0.00126)
	R-Squared	0.038		
	First Stage F-Stat		44.30	37.80
	Number of Obs	3,498	3,498	3,498
Online Sales in Past Month in RMB	Treat or Log Dist	-10.09	-18.75	5.366
		(12.89)	(23.94)	(6.855)
	R-Squared	0.012		
	First Stage F-Stat		44.26	37.82
	Number of Obs	3,498	3,498	3,498
Share of Online Sales in Household Monthly Income	Treat or Log Dist	-0.00120	-0.00224	0.000652
		(0.00176)	(0.00330)	(0.000960)
	R-Squared	0.032		
	First Stage F-Stat		41.62	33.64
	Number of Obs	2,830	2,830	2,830
Primary Earner Working As Peasant (Yes=1)	Treat or Log Dist	-0.0229	-0.0425	0.0122
		(0.0319)	(0.0597)	(0.0172)
	R-Squared	0.140		
	First Stage F-Stat		44.42	37.25
	Number of Obs	3,327	3,327	3,327
Member of Household Started a Business Over Last 6 Months (Yes=1)	Treat or Log Dist	-0.00802	-0.0149	0.00428
		(0.00631)	(0.0120)	(0.00344)
	R-Squared	0.001		
	First Stage F-Stat		44.37	37.74
	Number of Obs	3,468	3,468	3,468
New Business Selling in Part Online (Yes=1)	Treat or Log Dist	0.000212	0.000394	-0.000113
		(0.00159)	(0.00294)	(0.000846)
	R-Squared	0.000		
	First Stage F-Stat		44.33	37.68
	Number of Obs	3,468	3,468	3,468

Notes: See Section 4 for discussion. Standard errors are clustered at the level of villages. * 10%, ** 5%, *** 1% significance levels.

Table 8: Average Effects: Local Retail Prices

Dependent Variables		Intent to Treat	Treatment on Treated
Log Prices (All)	Treat	0.0189	0.0352
		(0.0142)	(0.0263)
	R-Squared	0.893	0.893
	First Stage F-Stat		41.66
	Number of Obs	6,877	6,877
Product Replacement Dummy (Not Counting Store Closures) (Yes=1)	Treat	-0.00516	-0.00983
		(0.00947)	(0.0181)
	R-Squared	0.000	-0.002
	First Stage F-Stat		39.82
	Number of Obs	8,956	8,956
Store Closure (at Product Level) (Yes=1)	Treat	0.00124	0.00236
		(0.0294)	(0.0556)
	R-Squared	0.000	0.000
	First Stage F-Stat		39.82
	Number of Obs	8,956	8,956
Number of New Products Per Store	Treat	2.194**	4.020*
		(1.073)	(2.278)
	R-Squared	0.277	0.212
	First Stage F-Stat		19.69
	Number of Obs	312	312
Store Owner Sources Products Online (Yes=1)	Treat	-0.00145	-0.00261
		(0.0258)	(0.0461)
	R-Squared	0.000	-0.001
	First Stage F-Stat		23.76
	Number of Obs	341	341
Log Prices of Business Inputs	Treat	0.00229	0.00337
		(0.129)	(0.186)
	R-Squared	0.811	0.811
	First Stage F-Stat		24.86
	Number of Obs	237	237
Log Prices of Non-Durables	Treat	0.0211	0.0398
		(0.0146)	(0.0276)
	R-Squared	0.860	0.860
	First Stage F-Stat		40.36
	Number of Obs	6,455	6,455
Log Prices of Durables	Treat	-0.0320	-0.0522
		(0.0711)	(0.115)
	R-Squared	0.951	0.952
	First Stage F-Stat		9.753
	Number of Obs	185	185

Dependent Variables		Intent to Treat	Treatment on Treated
Log Prices of Food and Beverages (1)	Treat	0.0368**	0.0706*
		(0.0185)	(0.0375)
	R-Squared	0.870	0.870
	First Stage F-Stat		39.37
	Number of Obs	3,686	3,686
Log Prices of Tobacco and Alcohol (2)	Treat	0.0212	0.0421
		(0.0340)	(0.0662)
	R-Squared	0.809	0.810
	First Stage F-Stat		32.39
	Number of Obs	1,071	1,071
Log Prices of Medicine and Health Products (3)	Treat	-0.0474	-0.0756
		(0.0741)	(0.122)
	R-Squared	0.794	0.795
	First Stage F-Stat		19.18
	Number of Obs	266	266
Log Prices of Clothing and Accessories (4)	Treat	0.0809	0.115
		(0.111)	(0.158)
	R-Squared	0.845	0.842
	First Stage F-Stat		42.80
	Number of Obs	152	152
Log Prices of Other Household Products (5)	Treat	-0.0328	-0.0619
		(0.0382)	(0.0744)
	R-Squared	0.756	0.755
	First Stage F-Stat		28.85
	Number of Obs	1,268	1,268
Log Prices of Heating, Fuel and Gas (6)	Treat	-0.0115	-0.0440
		(0.0955)	(0.332)
	R-Squared	0.007	-0.095
	First Stage F-Stat		0.795
	Number of Obs	12	12
Log Prices of Furniture and Appliances (7)	Treat	-0.0347	-0.0617
		(0.0881)	(0.156)
	R-Squared	0.952	0.953
	First Stage F-Stat		6.757
	Number of Obs	109	109
Log Prices of Electronics (8)	Treat	-0.0892	-0.163
		(0.305)	(0.570)
	R-Squared	0.884	0.890
	First Stage F-Stat		3.180
	Number of Obs	23	23
Log Prices of Transport Equipment (9)	Treat	0.0297	0.0398
		(0.0840)	(0.110)
	R-Squared	0.946	0.946
	First Stage F-Stat		22.67
	Number of Obs	53	53

Notes: See Section 4 for discussion. Standard errors are clustered at the level of villages. * 10%, ** 5%, *** 1% significance levels.

Table 9: Effects of Logistical vs Transactional Barriers

Effects on Consumption				
Dept Variables		Intent to Treat	Treatment on the Treated	Log Distance (IV Using Treat)
Monthly Total Retail Expenditure Per Capita	Treat or Log Dist	-26.72	-49.03	15.10
		(36.25)	(67.96)	(20.73)
	Treat or Log Dist * Delivery	31.42	58.59	-16.79
		(69.33)	(140.5)	(31.84)
	First Stage F-Stat		2.388	14.91
	Number of Obs	3,434	3,434	3,434
Household Has Ever Bought Something at Terminal (Yes=1)	Treat or Log Dist	0.0573***	0.105***	-0.0323***
		(0.0190)	(0.0288)	(0.00922)
	Treat or Log Dist * Delivery	-0.0603**	-0.110**	0.0335***
		(0.0251)	(0.0438)	(0.0113)
	First Stage F-Stat		2.683	14.88
	Number of Obs	3,518	3,518	3,518
Household Has Bought Something at Terminal in Last Month (Yes=1)	Treat or Log Dist	0.0329***	0.0604***	-0.0187***
		(0.0111)	(0.0189)	(0.00576)
	Treat or Log Dist * Delivery	-0.0422***	-0.0790**	0.0224***
		(0.0155)	(0.0329)	(0.00774)
	First Stage F-Stat		2.513	14.48
	Number of Obs	3,482	3,482	3,482
Share of Expenditure at Terminal in Total Monthly Retail Expenditure	Treat or Log Dist	0.00796***	0.0146***	-0.00451***
		(0.00274)	(0.00488)	(0.00153)
	Treat or Log Dist * Delivery	-0.00833***	-0.0153***	0.00466***
		(0.00294)	(0.00542)	(0.00160)
	First Stage F-Stat		2.413	14.66
	Number of Obs	3,434	3,434	3,434
Share of Expenditure at Terminal in Total Monthly Business Inputs	Treat or Log Dist	-0.00830	-0.0190	0.00621
		(0.00827)	(0.0222)	(0.00764)
	Treat or Log Dist * Delivery	0.0158	0.0296	-0.00865
		(0.0105)	(0.0241)	(0.00789)
	First Stage F-Stat		6.346	4.150
	Number of Obs	1,207	1,207	1,207
Share of Expenditure at Terminal in Total Monthly Non-Durables	Treat or Log Dist	0.00637***	0.0117***	-0.00361***
		(0.00225)	(0.00400)	(0.00126)
	Treat or Log Dist * Delivery	-0.00646**	-0.0119***	0.00365***
		(0.00246)	(0.00452)	(0.00133)
	First Stage F-Stat		2.413	14.66
	Number of Obs	3,433	3,433	3,433
Share of Expenditure at Terminal in Total Monthly Durables	Treat or Log Dist	0.0486***	0.0807***	-0.0284***
		(0.0177)	(0.0284)	(0.0102)
	Treat or Log Dist * Delivery	-0.0694***	-0.118***	0.0367***
		(0.0258)	(0.0442)	(0.0129)
	First Stage F-Stat		3.150	9.741
	Number of Obs	768	768	768

Effects on Incomes				
Dept Variables		Intent to Treat	Treatment on the Treated	Log Distance (IV Using Treat)
Monthly Income Per Capita Net of Transfers in RMB	Treat or Log Dist	-13.87	-25.27	7.906
		(77.86)	(140.7)	(43.96)
	Treat or Log Dist * Delivery	12.70	23.04	-7.437
		(188.3)	(367.2)	(80.54)
	First Stage F-Stat		2.635	14.44
	Number of Obs	3,445	3,445	3,445
Monthly Income Per Capita in RMB	Treat or Log Dist	-14.99	-27.14	8.513
		(77.55)	(140.1)	(43.82)
	Treat or Log Dist * Delivery	50.29	97.16	-22.44
		(171.2)	(339.1)	(75.42)
	First Stage F-Stat		2.694	14.42
	Number of Obs	3,437	3,437	3,437
Annual Income Per Capita in RMB	Treat or Log Dist	70.33	124.2	-39.47
		(645.0)	(1,168)	(366.8)
	Treat or Log Dist * Delivery	-734.1	-1,462	306.9
		(1,484)	(2,755)	(622.5)
	First Stage F-Stat		2.501	14.22
	Number of Obs	3,388	3,388	3,388
Member of Household Has Ever Sold Online (Yes=1)	Treat or Log Dist	-0.00857	-0.0156	0.00483
		(0.00608)	(0.0111)	(0.00348)
	Treat or Log Dist * Delivery	0.0102	0.0188	-0.00546
		(0.0141)	(0.0280)	(0.00614)
	First Stage F-Stat		2.561	14.66
	Number of Obs	3,504	3,504	3,504
Share of Online Sales in Household Monthly Income	Treat or Log Dist	-0.00172	-0.00316	0.00101
		(0.00210)	(0.00387)	(0.00124)
	Treat or Log Dist * Delivery	0.00282	0.00540	-0.00145
		(0.00233)	(0.00441)	(0.00131)
	First Stage F-Stat		2.402	12.70
	Number of Obs	2,830	2,830	2,830
Primary Earner Working as Peasant (Yes=1)	Treat or Log Dist	-0.0192	-0.0352	0.0109
		(0.0341)	(0.0624)	(0.0196)
	Treat or Log Dist * Delivery	-0.0284	-0.0609	0.00806
		(0.0813)	(0.185)	(0.0372)
	First Stage F-Stat		2.503	14.27
	Number of Obs	3,327	3,327	3,327
Member of Household Has Started a Business Over Last 6 Months (Yes=1)	Treat or Log Dist	-0.00328	-0.00601	0.00187
		(0.00635)	(0.0116)	(0.00358)
	Treat or Log Dist * Delivery	-0.0297	-0.0604	0.0113
		(0.0183)	(0.0536)	(0.0104)
	First Stage F-Stat		2.517	14.48
	Number of Obs	3,468	3,468	3,468

Effects on Retail Prices			
Dept Variables		Intent to Treat	Treatment on the Treated
Log Prices (All)	Treat	0.0114	0.0215
		(0.0144)	(0.0273)
	Treat * Delivery	0.0417	0.0739
		(0.0377)	(0.0572)
	First Stage F-Stat		17.26
	Number of Obs	6,877	6,877
Product Replacement Dummy (Not Counting Store Closures) (Yes=1)	Treat	-0.00680	-0.0129
		(0.0108)	(0.0206)
	Treat * Delivery	0.00907	0.0173
		(0.0213)	(0.0415)
	First Stage F-Stat		2.648
	Number of Obs	8,956	8,956
Store Closure (at Product Level) (Yes=1)	Treat	0.00111	0.00209
		(0.0355)	(0.0668)
	Treat * Delivery	0.000779	0.00162
		(0.0423)	(0.0805)
	First Stage F-Stat		2.648
	Number of Obs	8,956	8,956
Number of New Products Per Store	Treat	1.403*	2.352*
		(0.828)	(1.354)
	Treat * Delivery	3.403	7.993
		(3.876)	(12.77)
	First Stage F-Stat		1.247
	Number of Obs	312	312
Store Owner Sources Products Online (Yes=1)	Treat	0.0250**	0.0416**
		(0.0122)	(0.0201)
	Treat * Delivery	-0.0911	-0.185
		(0.0814)	(0.166)
	First Stage F-Stat		1.320
	Number of Obs	341	341
Log Price of Business Inputs	Treat	-0.0858	-0.108
		(0.134)	(0.182)
	Treat * Delivery	0.289	0.473
		(0.273)	(0.447)
	First Stage F-Stat		1.972
	Number of Obs	237	237
Log Price of Non-Durables	Treat	0.0192	0.0366
		(0.0157)	(0.0308)
	Treat * Delivery	0.0137	0.0214
		(0.0362)	(0.0585)
	First Stage F-Stat		16.09
	Number of Obs	6,455	6,455
Log Prices of Durables	Treat	-0.118	-0.144
		(0.0880)	(0.104)
	Treat * Delivery	0.164	0.288
		(0.134)	(0.366)
	First Stage F-Stat		0.488
	Number of Obs	185	185

Notes: See Section 4 for discussion. Standard errors are clustered at the level of villages. * 10%, ** 5%, *** 1% significance levels.

Table 10: Heterogeneity Across Households and Villages

Type of Heterogeneity		Intent to Treat	Treatment on the Treated	Log Dist (IV)	Intent to Treat	Treatment on the Treated	Log Distance (IV)	Intent to Treat	Treatment on the Treated
	Dependent Variables:	Household Has Ever Bought Something at Terminal (Yes=1)			Monthly Income Per Capita (RMB)			Log Local Retail Prices	
Average Effect	Treat or Log Dist	0.0480***	0.0886***	-0.0253***	-7.838	-14.48	4.190	0.0189	0.0352
		(0.0169)	(0.0271)	(0.00801)	(70.78)	(129.9)	(37.55)	(0.0142)	(0.0263)
	R-Squared	0.008			0.038			0.893	0.893
	First Stage F-Stat		45.56	39.22		45.33	37.69		41.66
	Number of Obs	3,518	3,518	3,518	3,437	3,437	3,437	6,877	6,877
Village Was Previously Connected to Parcel Delivery (Yes=1)	Treat or Log Dist	0.0573***	0.105***	-0.0323***	-14.99	-27.14	8.513	0.0114	0.0215
		(0.0190)	(0.0288)	(0.00922)	(77.55)	(140.1)	(43.82)	(0.0144)	(0.0273)
	Treat or Log Dist *	-0.0603**	-0.110**	0.0335***	50.29	97.16	-22.44	0.0417	0.0739
	Delivery	(0.0251)	(0.0438)	(0.0113)	(171.2)	(339.1)	(75.42)	(0.0377)	(0.0572)
	R-Squared	0.016			0.040			0.894	
	First Stage F-Stat		2.683	14.88		2.694	14.42		17.26
	Number of Obs	3,518	3,518	3,518	3,437	3,437	3,437	6,877	6,877
Village Distance to Township Center	Treat or Log Dist	-0.0156	-0.00882	-0.00268	-23.53	-43.67	14.71	-0.0219	-0.0322
		(0.0288)	(0.0429)	(0.0126)	(181.7)	(289.2)	(84.33)	(0.0375)	(0.0632)
	Treat or Log Dist *	0.0388**	0.0612***	-0.0138**	0.389	0.371	-1.272	0.0216	0.0358
	Log Dist Planned	(0.0162)	(0.0227)	(0.00570)	(97.50)	(152.0)	(40.55)	(0.0198)	(0.0336)
	R-Squared	0.014			0.040			0.893	
	First Stage F-Stat		15.63	11.79		15.66	10.98		16.96
	Number of Obs	3,518	3,518	3,518	3,437	3,437	3,437	6,877	6,877
Primary Earner's Age	Treat or Log Dist	0.140***	0.223***	-0.0669***	-136.4	-237.8	70.34		
		(0.0506)	(0.0778)	(0.0230)	(172.5)	(286.5)	(84.03)		
	Treat or Log Dist *	-0.00172**	-0.00251*	0.000778**	2.561	4.551	-1.341		
	Age	(0.000774)	(0.00129)	(0.000370)	(2.734)	(4.825)	(1.404)		
	R-Squared	0.023			0.049				
	First Stage F-Stat		16.07	15.63		16.34	15.65		
	Number of Obs	3,304	3,304	3,304	3,292	3,292	3,292		
Primary Earner's Educaction	Treat or Log Dist	0.0407*	0.0977**	-0.0266**	52.80	119.7	-33.46		
		(0.0206)	(0.0412)	(0.0115)	(83.52)	(195.0)	(53.92)		
	Treat or Log Dist *	0.00161	-0.000469	-5.85e-05	-8.666	-17.79	5.057		
	Years of Education	(0.00267)	(0.00506)	(0.00141)	(12.14)	(24.03)	(6.774)		
	R-Squared	0.014			0.063				
	First Stage F-Stat		8.462	10.62		8.662	10.78		
	Number of Obs	3,296	3,296	3,296	3,284	3,284	3,284		
Household Income Per Capita	Treat or Log Dist	0.00806	0.0209	-0.00505	35.86	59.51	-16.75		
		(0.0213)	(0.0375)	(0.00998)	(96.83)	(165.5)	(45.62)		
	Treat or Log Dist *	0.00712**	0.0121**	-0.00370**	-9.204	-15.79	4.564		
	Log Income PC	(0.00326)	(0.00541)	(0.00162)	(21.22)	(36.31)	(10.39)		
	R-Squared	0.011			0.355				
	First Stage F-Stat		22.78	17.96		22.57	17.62		
	Number of Obs	3,416	3,416	3,416	3,437	3,437	3,437		
Household Distance to Planned Terminal	Treat or Log Dist	0.144**	0.231**	-0.0636**	185.9	400.1	-108.9		
		(0.0591)	(0.109)	(0.0315)	(350.6)	(697.5)	(188.3)		
	Treat or Log Dist *	-0.0181*	-0.0274	0.00739	-36.54	-79.67	21.85		
	Log Dist Planned	(0.00981)	(0.0193)	(0.00587)	(61.53)	(128.5)	(34.90)		
	R-Squared	0.012			0.039				
	First Stage F-Stat		9.905	11.64		9.325	14.15		
	Number of Obs	3,518	3,518	3,518	3,437	3,437	3,437		
Combined	Treat or Log Dist	0.154*	0.289**	-0.0838*	108.5	213.4	-57.26	-0.0398	-0.0435
		(0.0805)	(0.140)	(0.0438)	(333.8)	(619.5)	(184.7)	(0.0362)	(0.0531)
	Treat or Log Dist *	-0.0400	-0.106	0.0342**	98.21	229.2	-53.30	0.0413	0.0517
	Delivery	(0.0285)	(0.0687)	(0.0149)	(137.1)	(336.0)	(69.69)	(0.0361)	(0.0622)
	Treat or Log Dist *	0.0458***	0.0813***	-0.0178***	-37.85	-81.46	18.11	0.0284	0.0380
	Log Dist Township	(0.0174)	(0.0298)	(0.00688)	(62.90)	(134.2)	(31.65)	(0.0188)	(0.0312)
	Treat or Log Dist *	-0.00181**	-0.00314**	0.000964**	0.929	1.742	-0.511		
	Age	(0.000775)	(0.00129)	(0.000390)	(2.567)	(4.664)	(1.378)		
	Treat or Log Dist *	0.000370	-0.00380	0.000671	-2.778	-1.854	1.218		
	Years of Education	(0.00268)	(0.00499)	(0.00144)	(10.22)	(21.43)	(6.086)		
	Treat or Log Dist *	0.00908***	0.0162***	-0.00544***	-12.43	-21.38	6.717		
	Log Income PC	(0.00339)	(0.00555)	(0.00174)	(22.39)	(38.60)	(11.50)		
	Treat or Log Dist *	-0.0249**	-0.0417*	0.0109	-8.134	-20.40	5.556		
	Log Dist Planned	(0.0107)	(0.0218)	(0.00671)	(45.46)	(96.39)	(26.75)		
	R-Squared	0.051			0.353			0.894	
	First Stage F-Stat		0.474	2.991		0.420	2.938		1.579
	Number of Obs	3,261	3,261	3,261	3,282	3,282	3,282	6,877	6,877

Notes: See Section 4 for discussion. Standard errors are clustered at the level of villages. * 10%, ** 5%, *** 1% significance levels.

Table 11: How Do the E-Commerce Terminals Compare?

Could You Have Purchased This Product in Your Village? (Yes=1)	Sample Fraction	0.380	Household Living in Village Without Any Durables on Sale (Yes=1)	Sample Fraction	0.547
	Number Obs	255		Number Obs	3,508
Log Price Difference between Terminal and Village Retail	Sample Mean	-0.166	Travel Cost to Main Shopping Destination Outside Village (RMB)	Sample Mean	11.85
	Sample Median	-0.154		Sample Median	4
	Number Obs	95		Number Obs	2,766
Could You Have Purchased This Product in the Nearby Town? (Yes=1)	Sample Fraction	0.836	Travel Time to Main Shopping Destination Outside Village and Back (Minutes)	Sample Mean	58.14
				Sample Median	40
	Number Obs	238		Number Obs	2,366
Log Price Difference between Terminal and Nearby Town Retail	Sample Mean	-0.227	Travel Distance to Main Shopping Destination Outside Village and Back (Km)	Sample Mean	15.38
	Sample Median	-0.182		Sample Median	9.045
	Number Obs	197		Number Obs	2,773

Notes: See Section 4 for discussion.

Table 12: Role of Program Implementation

Type of Heterogeneity		Intent to Treat	Treatment on the Treated	Log Distance (IV Using Treat)
	Dependent Variable: Household Has Ever Bought Something at Terminal (Yes=1)			
Average Effects	Treat or Log Dist	0.0480***	0.0886***	-0.0253***
		(0.0169)	(0.0271)	(0.00801)
	R-Squared	0.008		
	First Stage F-Stat		45.56	39.22
	Number of Obs	3,518	3,518	3,518
Terminal Manager Test Score	Treat or Log Dist	0.0594	0.104	8.91e-05
		(0.147)	(0.242)	(0.0576)
	Treat or Log Dist * Score	-0.000214	-0.000384	-0.000238
		(0.00164)	(0.00270)	(0.000667)
	R-Squared	0.006		
	First Stage F-Stat		8.786	8.162
	Number of Obs	3,042	3,042	3,042
Terminal Manager Test Score Above the Median	Treat or Log Dist	0.0314	0.0616	-0.0148
		(0.0295)	(0.0501)	(0.0120)
	Treat or Log Dist * Above Median	0.0191	0.0182	-0.0109
		(0.0347)	(0.0583)	(0.0153)
	R-Squared	0.006		
	First Stage F-Stat		8.654	9.503
	Number of Obs	3,042	3,042	3,042
Terminal Installation Delayed	Treat or Log Dist	0.0392	0.0656*	-0.0188*
		(0.0247)	(0.0357)	(0.00978)
	Treat or Log Dist * Delay Dummy	0.0167	0.0486	-0.0138
		(0.0335)	(0.0554)	(0.0169)
	R-Squared	0.009		
	First Stage F-Stat		10.93	10.98
	Number of Obs	3,518	3,518	3,518

Notes: See Section 4 for discussion. Standard errors are clustered at the level of villages. * 10%, ** 5%, *** 1% significance levels.

Table 13: Role of GE Spillovers

Dependent Variables		Treatment on Treated without Spillovers	ToT with Spillovers: Number of Terminals within 3 km Outside of Village	ToT with Spillovers: Number of Terminals within 10 km Outside of Village
Monthly Income Per Capita (RMB)	Treat Dummy	-14.48	-5.367	-30.73
		(129.9)	(139.4)	(122.8)
	Exposure to Terminals Outside the Village		-145.6	-0.452
			(185.1)	(27.19)
	Exposure to Other Villages		-33.89**	-12.75***
			(16.19)	(3.939)
	First Stage F-Stat	45.33	47.59	44.94
	Number of Obs	3,437	3,437	3,437
Any Member of Household Has Ever Sold Online (Yes=1)	Treat Dummy	-0.0129	-0.0139	-0.0155
		(0.0104)	(0.0101)	(0.0101)
	Exposure to Terminals Outside the Village		-0.00112	-0.00291
			(0.0102)	(0.00198)
	Exposure to Other Villages		-0.00356***	-0.000358
			(0.00107)	(0.000369)
	First Stage F-Stat	45.30	47.43	45.04
	Number of Obs	3,504	3,504	3,504
Household Has Ever Bought Something at Terminal (Yes=1)	Treat Dummy	0.0886***	0.0780***	0.0846***
		(0.0271)	(0.0265)	(0.0263)
	Exposure to Terminals Outside the Village		0.0665**	-0.00877
			(0.0311)	(0.00579)
	Exposure to Other Villages		-0.00326	0.00223*
			(0.00540)	(0.00115)
	First Stage F-Stat	45.56	47.90	45.34
	Number of Obs	3,518	3,518	3,518
Share of Terminal in Total Retail Expenditure	Treat Dummy	0.0124***	0.0100**	0.0116***
		(0.00434)	(0.00398)	(0.00418)
	Exposure to Terminals Outside the Village		0.0159*	-0.00185**
			(0.00835)	(0.000943)
	Exposure to Other Villages		-0.000581	0.000517**
			(0.000519)	(0.000232)
	First Stage F-Stat	44.03	46.38	43.96
	Number of Obs	3,434	3,434	3,434
Log Local Retail Prices (All Prices)	Treat Dummy	0.0352	0.0341	0.0391
		(0.0263)	(0.0259)	(0.0249)
	Exposure to Terminals Outside the Village		0.00284	0.00476
			(0.0315)	(0.00566)
	Exposure to Other Villages		-0.00207	-0.00106
			(0.00304)	(0.000969)
	First Stage F-Stat	41.66	43.72	44.90
	Number of Obs	6,877	6,877	6,877

Notes: See Section 4 for discussion. Standard errors are clustered at the level of villages. * 10%, ** 5%, *** 1% significance levels.

Table 14: Are Sample Villages Representative?

	(1)	(2)	(3)	(4)	(5)	(6)
	Full Sample	Full Sample	Full Sample	3 Provinces	3 Provinces	3 Provinces
VARIABLES	N Users	N Transactions	Sales (RMB)	N Users	N Transactions	Sales (RMB)
RCT_Sample	-4.110	0.0605	-6,034	0.149	12.65	-3,747
	(7.751)	(25.33)	(4,061)	(7.734)	(25.32)	(4,066)
Observations	125,204	125,204	125,204	100,098	100,098	100,098
R-squared	0.037	0.047	0.029	0.031	0.046	0.030
N Cluster	11731	11731	11731	8471	8471	8471

	(1)	(2)	(3)	(4)
	Full Sample	Full Sample	3 Provinces	3 Provinces
VARIABLES	N Transactions	Weight (kg)	N Transactions	Weight (kg)
RCT_Sample	1.712**	5.154	1.364*	4.680
	(0.753)	(4.332)	(0.752)	(4.333)
Observations	120,483	120,483	95,744	95,744
R-squared	0.060	0.023	0.067	0.026
N Cluster	11904	11904	8591	8591

Notes: The upper panel presents point estimates from regressions based on the purchase transaction database. The lower panel presents point estimates from regressions based on the sales transaction database. See Section 5 for discussion. Standard errors are clustered at the level of village terminals. * 10%, ** 5%, *** 1% significance levels.

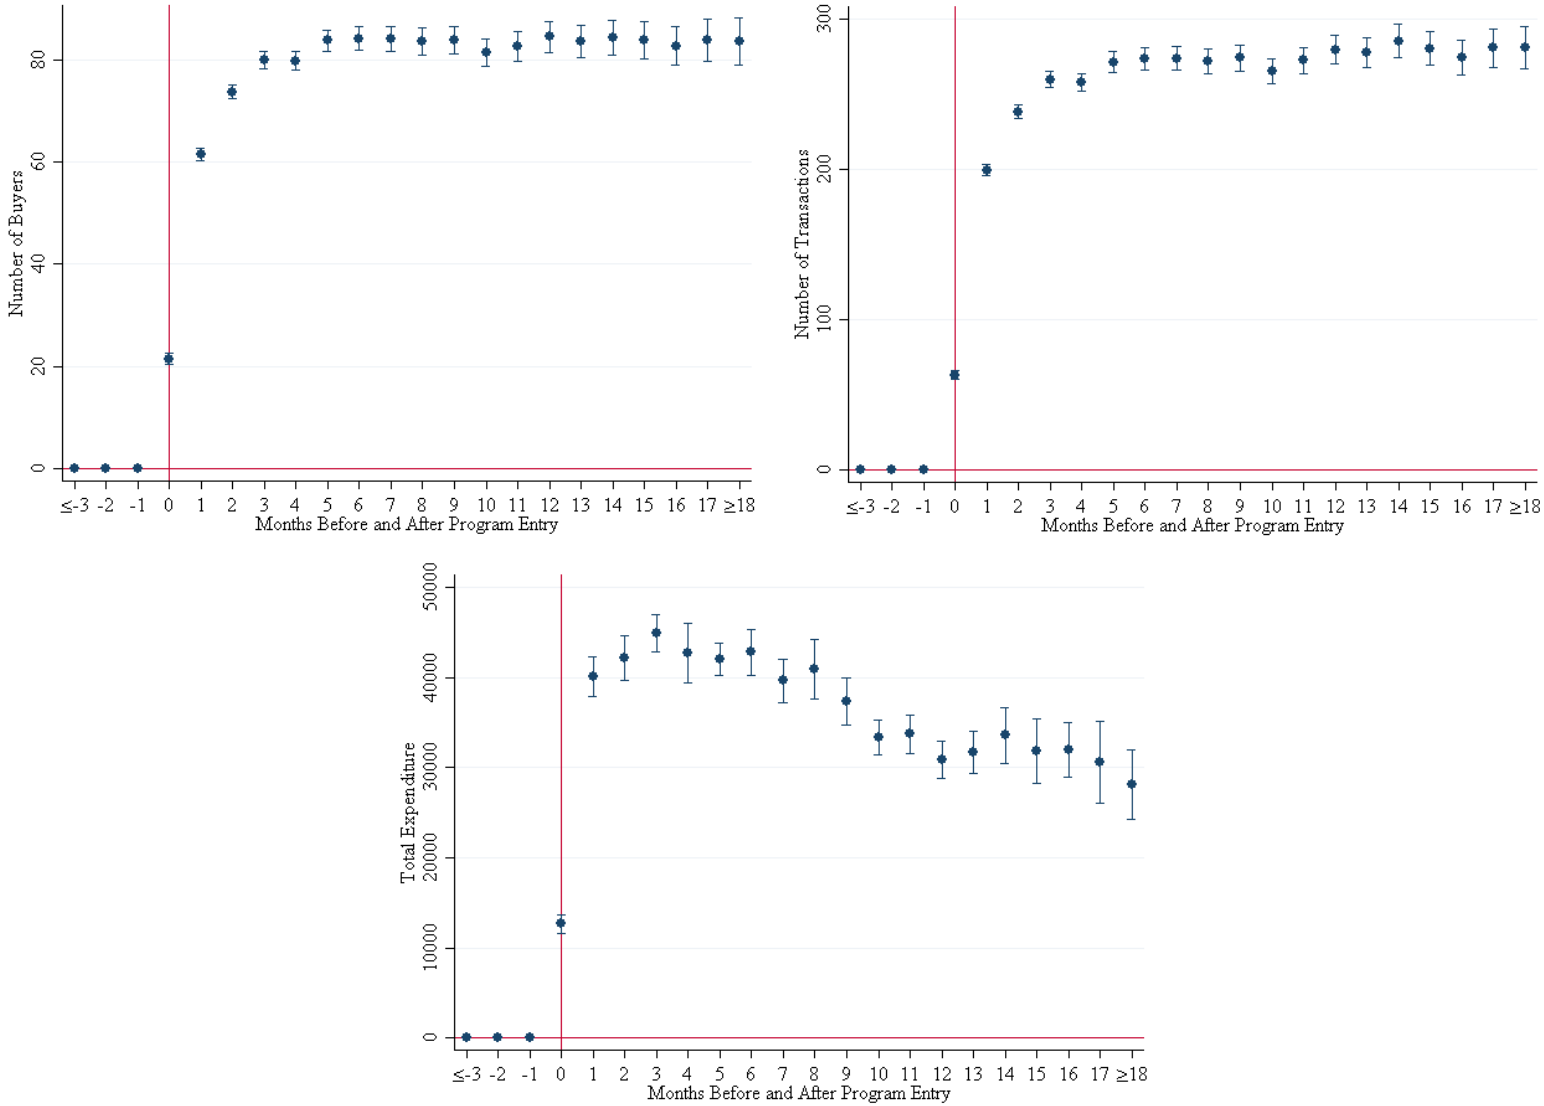
Table 15: Role of Seasonality

	(1)	(2)	(3)	(4)	(5)	(6)
	Full Sample	Full Sample	Full Sample	3 Provinces	3 Provinces	3 Provinces
VARIABLES	N Users	N Transactions	Sales (RMB)	N Users	N Transactions	Sales (RMB)
RCT_Month	0.893***	-4.671***	-1,565***	0.568**	-5.290***	-585.9
	(0.255)	(0.818)	(451.5)	(0.274)	(0.863)	(458.0)
Observations	125,204	125,204	125,204	100,098	100,098	100,098
R-squared	0.694	0.680	0.219	0.679	0.667	0.227
N Cluster	11731	11731	11731	8471	8471	8471

	(1)	(2)	(3)	(4)
	Full Sample	Full Sample	3 Provinces	3 Provinces
VARIABLES	N Transactions	Weight (kg)	N Transactions	Weight (kg)
RCT_Month	-0.387***	-1.256***	-0.498***	-1.407***
	(0.0225)	(0.125)	(0.0261)	(0.138)
Observations	120,483	120,483	95,744	95,744
R-squared	0.592	0.432	0.570	0.422
N Cluster	11904	11904	8591	8591

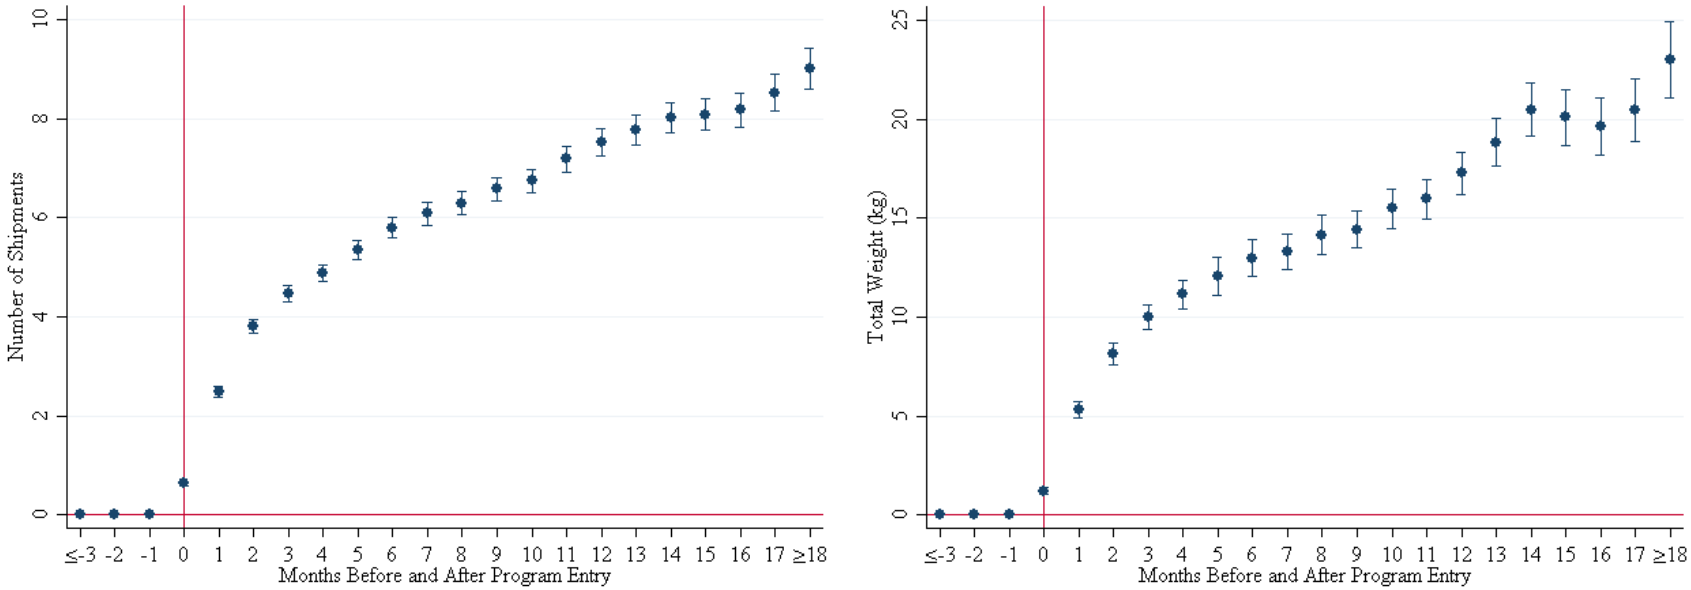
Notes: The upper panel presents point estimates from regressions based on the purchase transaction database. The lower panel presents point estimates from regressions based on the sales transaction database. See Section 5 for discussion. Standard errors are clustered at the level of village terminals. * 10%, ** 5%, *** 1% significance levels.

Figure 1: Timeline of Adjustment: Consumption (Terminal-Level)



Notes: See Section 5 for discussion. Standard errors are clustered at the level of village terminals. * 10%, ** 5%, *** 1% significance levels.

Figure 2: Timeline of Adjustment: Selling (Terminal-Level)



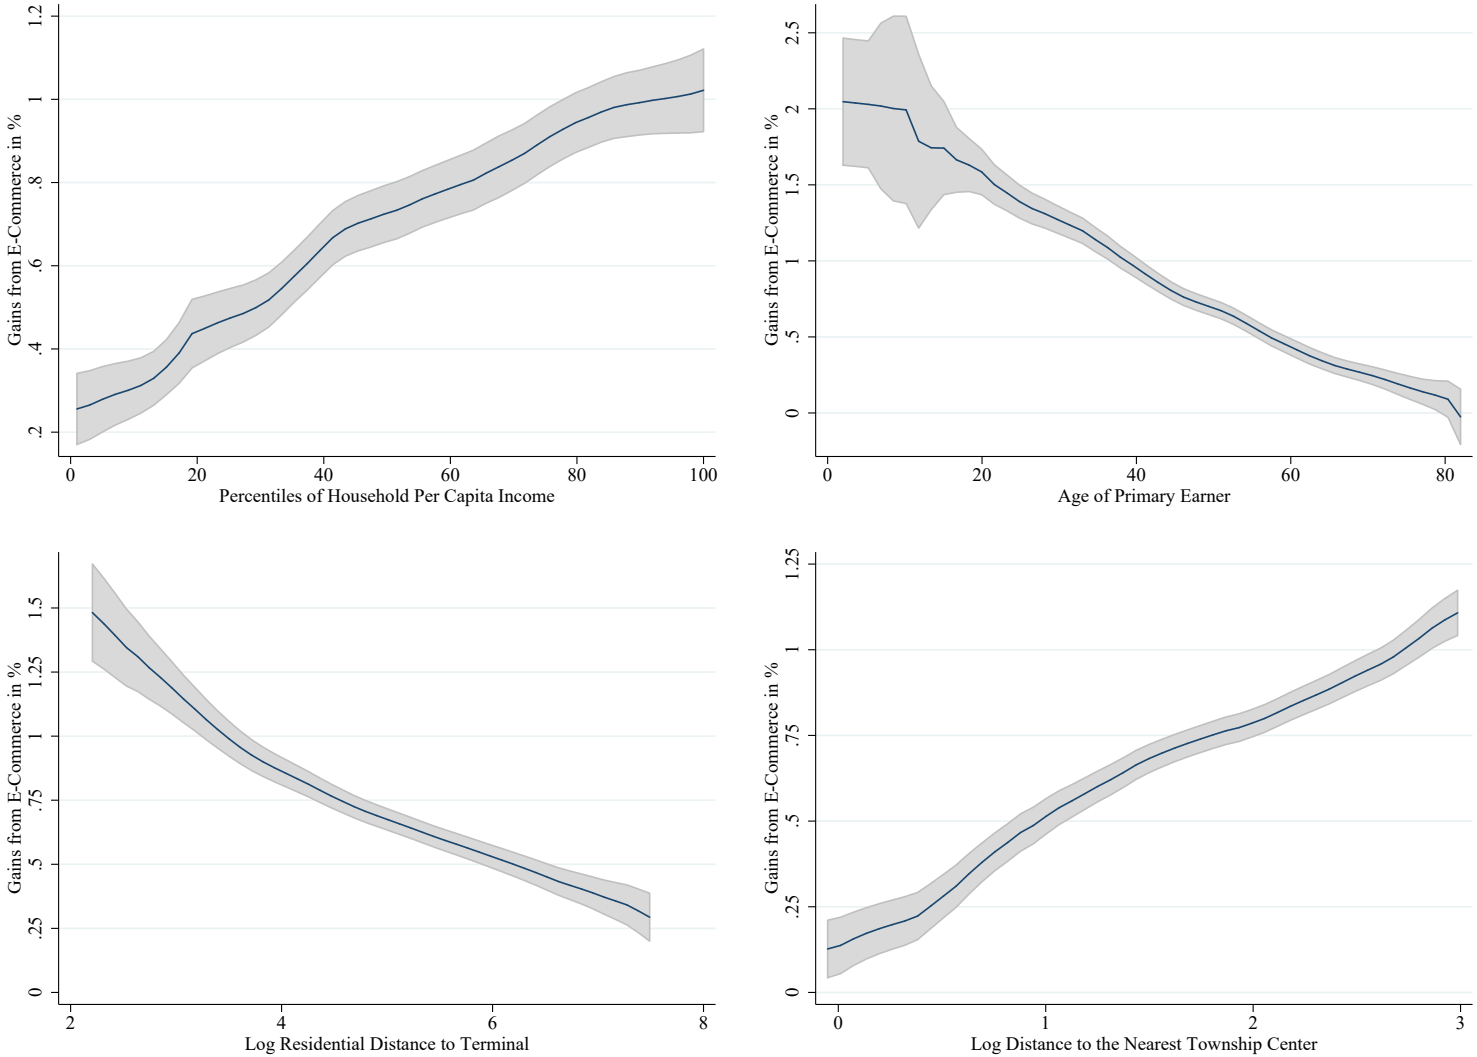
Notes: See Section 5 for discussion. Standard errors are clustered at the level of village terminals. * 10%, ** 5%, *** 1% significance levels.

Table 16: Average Effects On Household Economic Welfare

	Unweighted (Effects in Sample)			Weighted (Effects in Village Population)		
	Durables Consumption	Non-Durables Consumption	Total Retail Consumption	Durables Consumption	Non-Durables Consumption	Total Retail Consumption
Reduction in Retail Cost of Living for All Households	3.298% (0.027)	0.478% (0.004)	0.812% (0.005)	2.908% (0.031)	0.419% (0.003)	0.714% (0.005)
Reduction in Retail Cost of Living Among Users	19.331% (0.215)	3.722% (0.029)	5.464% (0.035)	16.599% (0.215)	3.267% (0.024)	4.764% (0.032)
Reduction in Retail Cost of Living Among Users Last Month	57.1% (2.634)	8.862% (0.073)	13.183% (0.131)	45.802% (2.411)	7.797% (0.064)	11.259% (0.124)

Notes: See Section 6 for discussion. Standard errors are bootstrapped across 1000 iterations with random re-sampling.

Figure 3: Heterogeneity of Welfare Effect



Notes: See Section 6 for discussion. Gains are expressed in terms of percentage point reductions of retail cost of living. Confidence intervals are based on standard errors that are clustered at the level of villages.

Table 17: Quantification Under Alternative Demand Parameters

	$\sigma_D = 2.87, \sigma_N = 2.85$			$\sigma_D = 3.87, \sigma_N = 3.85$			$\sigma_D = 4.87, \sigma_N = 4.85$		
	Durables Consumption	Non-Durables Consumption	Total Retail Consumption	Durables Consumption	Non-Durables Consumption	Total Retail Consumption	Durables Consumption	Non-Durables Consumption	Total Retail Consumption
Reduction in Retail Cost of Living for All Households	5.129% (0.043)	0.735% (0.005)	1.252% (0.007)	3.298% (0.027)	0.478% (0.004)	0.812% (0.005)	2.431% (0.02)	0.355% (0.003)	0.601% (0.003)
Reduction in Retail Cost of Living Among Users	31.47% (0.368)	5.773% (0.046)	8.526% (0.056)	19.331% (0.215)	3.722% (0.029)	5.464% (0.035)	13.942% (0.151)	2.747% (0.022)	4.02% (0.026)
Reduction in Retail Cost of Living Among Users Last Month	116.98% (10.082)	13.94% (0.118)	21.032% (0.222)	57.1% (2.634)	8.862% (0.073)	13.183% (0.131)	38.053% (1.331)	6.495% (0.053)	9.598% (0.092)

Notes: See Section 6 for discussion. Standard errors are bootstrapped across 1000 iterations with random re-sampling.

9 Appendix

Additional Figures and Tables

[Work in progress.]

Theoretical Framework for Welfare Evaluation

Following recent work in international trade (Atkin et al., in press), we propose a three-tier demand system. In the upper tier there are Cobb-Douglas preferences over broad product groups $g \in G$ (durables and non-durables) in total consumption, in the middle tier there are asymmetric CES preferences over local retailers selling that product group $s \in S$ (e.g. local stores, market stalls or the e-commerce terminal), and in the final tier there are preferences over the individual products within the product groups $b \in B_g$ that we can leave unspecified for now:

$$U_h = \prod_{g \in G} [Q_{gh}]^{\alpha_{gh}} \tag{A.1}$$

$$Q_{gh} = \left(\sum_{s \in S_g} \beta_{gsh} q_{gsh}^{\frac{\sigma_g - 1}{\sigma_g}}\right)^{\frac{\sigma_g}{\sigma_g - 1}} \tag{A.2}$$

where α_{gh} and β_{gsh} are (potentially household- or income-group-specific) preference parameters that are fixed across periods. Q_{gh} and q_{gsh} are product-group and store-product-group consumption aggregates with associated price indices P_{gh} and r_{gsh} respectively, and σ_g is the elasticity of substitution across local retail outlets. For each broad product group, consumers choose how much they are going to spend at different retail outlets based on the store-level price index r_{gsh} (which itself depends on the product mix and product-level prices on offer across outlets).

While the demand system is homothetic, we capture potential heterogeneity across the income distribution by allowing households of different incomes to differ in their expenditure shares across product groups (α_{gh}) and their preferences for consumption bundles at different stores within those product groups (β_{gsh} and the preference parameters that generate q_{gsh}). As shown by Anderson et al. (1992), these preferences can generate the same demands as would be obtained from aggregating many consumers who make discrete choices over which store to shop in. Building on Feenstra (1994), the following expression provides the exact proportional cost of living effect under this demand system:

$$\frac{CLE}{e(\mathbf{P}_T^{0*}, \mathbf{P}_C^0, \mathbf{P}_E^{0*}, \mathbf{P}_X^0, u_h^0)} = \frac{e(\mathbf{P}_T^1, \mathbf{P}_C^1, \mathbf{P}_E^1, \mathbf{P}_X^{1*}, u_h^0)}{e(\mathbf{P}_T^{0*}, \mathbf{P}_C^0, \mathbf{P}_E^{0*}, \mathbf{P}_X^0, u_h^0)} - 1 = \prod_{g \in G} \left(\left(\frac{\sum_{s \in S_g^C} \phi_{gsh}^1}{\sum_{s \in S_g^C} \phi_{gsh}^0}\right)^{\frac{1}{\sigma_g - 1}} \prod_{s \in S_g^C} \left(\frac{r_{gsh}^1}{r_{gsh}^0}\right)^{\omega_{gsh}} \right)^{\alpha_{gh}} \tag{A.3}$$

where S_g^C denotes the set of continuing local retailers within product group g, $\phi_{gsh}^t = r_{gsh}^t q_{gsh}^t / \sum_{s \in S_g} r_{gsh}^t q_{gsh}^t$ is the expenditure share for a particular retailer of product group g, and the ω_{gsh}s are ideal log-change weights.[1]

For each product group g, the expression has two components. The $\prod_{s \in S_g^C} \left(\frac{r_{gsh}^1}{r_{gsh}^0}\right)^{\omega_{gsh}}$ term is a Sato-Vartia (i.e. CES) price-index for price changes in continuing local stores that forms the *pro-competitive price effect*.[2] The price terms r_{gsh}^t are themselves price indices of product-specific prices p_{gsb}^t within local continuing stores which, in principle, could also account for new product

[1] In particular, $\omega_{gsh} = \left(\frac{\tilde{\phi}_{gsh}^1 - \tilde{\phi}_{gsh}^0}{\ln \tilde{\phi}_{gsh}^1 - \ln \tilde{\phi}_{gsh}^0}\right) / \sum_{s \in S_g^C} \left(\frac{\tilde{\phi}_{gsh}^1 - \tilde{\phi}_{gsh}^0}{\ln \tilde{\phi}_{gsh}^1 - \ln \tilde{\phi}_{gsh}^0}\right)$, which in turn contain expenditure shares of different retailers within product groups where the shares consider only expenditure at continuing retailers $\tilde{\phi}_{gsh}^t = r_{gsh}^t q_{gsh}^t / \sum_{s \in S_g^C} r_{gsh}^t q_{gsh}^t$.

[2] Notice that the assumption of CES preferences does not imply the absence of pro-competitive effects as we do not impose additional assumptions about market structure (e.g. monopolistic competition).

varieties or exiting product varieties using the same methodology. While we name these price changes pro-competitive, they may derive from either reductions in markups or increases in productivity at local stores (distinctions that do not matter on the cost-of-living side, but would generate different magnitudes of profit and income effects that we capture on the nominal income side).

The $(\frac{\sum_{s\in S_g^C}\phi_{gsh}^1}{\sum_{s\in S_g^C}\phi_{gsh}^0})^{\frac{1}{\eta_{gh}-1}}$ term captures the gains to customers of the e-commerce terminal in the numerator, from both the *direct price index effect* and the *entry effect*, and local store exit in the denominator, i.e. the *exit effect*. As in expression (2) of Section 3, we can decompose the total cost of living effect in equation (A.3) into four different types of effective consumer price changes by adding and subtracting terms.

For expositional purposes, consider the simple case where the program's only effect on cost of living is driven by the direct price index effect. In that case, the expenditure share spent on continuing local retailers $(\sum_{s\in S_g^C}\phi_{gsh}^1)$ is lower than unity only due substitution to the new e-commerce terminal (abstracting from a potential effect on store entry). The gains from the program as a proportion of initial household spending are then:

$$\frac{DE}{e(\mathbf{P}_T^{0*},\mathbf{P}_C^0,\mathbf{P}_E^{0*},\mathbf{P}_X^0,u_h^0)}=\prod_{g\in G}\left((\sum_{s\in S_g^C}\phi_{gsh}^1)^{\frac{1}{\sigma_g-1}}\right)^{\alpha_{gh}}-1 \tag{A.4}$$

The welfare gain from a new shopping option is a function of the market share of that outlet post-entry and the elasticity of substitution across stores. The revealed preference nature of this approach is clear. If consumers greatly value the arrival of the new option—be it because it offers low prices p_{gsb}^1, more product variety that reduces r_{gsh}^1 or better amenities β_{gsh}—the market share is higher and the welfare gain greater. Hence, these market share changes capture all the potential consumer benefits of shopping through the e-commerce terminal. How much greater depends on the elasticity of substitution. Large terminal market shares will imply small welfare changes if consumers substitute between local shopping options very elastically, and large welfare changes if they are inelastic. A similar logic would apply to effects on the entry of local retailers, or on the exit of local stores (where a large period 0 market share means large welfare losses, again tempered by the elasticity of substitution).

Data Appendix

[Work in progress.]