Three Essays in Development Economics

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Stephen Michael Pitts, SJ

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Advisor: Marc F. Bellemare

June, 2025

Acknowledgements

I am grateful to many people who have accompanied me on this five-year journey through graduate school.

My professors: Thank you to my committee chair Mike Boland for encouraging me to come to Minnesota, for many breakfasts after daily Mass, and for visiting me in the field in Chiapas. Thank you to my advisor Marc Bellemare for connecting in unconventional ways during the pandemic, emphasizing the importance of econometric rigor at every stage, encouraging me to research what I wanted, and giving candid yet kind feedback. Thank you to Paul Glewwe for encouraging me to integrate my professional and personal life, especially for inviting me to lead a lectio divina group for faculty and students. Thank you to Joe Ritter for patiently answering many admissions questions, teaching me econometrics and R, and employing me as your teaching assistant for two years. Thank you to Raahil Madhok for accepting the invitation to serve on his first dissertation committee and for our biweekly meetings in the final semester. Thank you to my outside member Adriana Uscanga, who brought first-hand experience of coffee fields in Chiapas to my final defense. Finally, thank you to Jesse Anttila-Hughes for ten years of mentoring from my time at the University of San Francisco to the present. I started this journey as your student and ended up as your coauthor.

My classmates: Thank you to my Jimmy Campbell for reminding me ”you’ve got to want it” as we spent hours working problem sets during the cold Minnesota winter in the pandemic during the first year and to many visits to Crooked Pint and the Guthrie Theater in subsequent years. Thank you to members of BeRG (the Bellemare Research Group) and the UMN Applied Economics Student Seminar for giving me feedback on the enclosed projects and others in their various stages.

Our Chiapas research team: Thank you to Chris Boyd for helping develop the experimental design and run the experiment in the field. Thank you to Grant Storer for assistance throughout the research process. Thank you to Alejandro Rodriguez and the entire staff at the Yomol A'tel cooperative group in Chiapas, especially our research team in Summer 2022.

Professional colleagues: Thanks to Nicholas Magnan, Jacob Ricker-Gilbert, and three anonymous reviewers for *Agricultural Economics* for helping me turn Essay 1 into my first peer-reviewed publication. Along the way, I appreciate feedback from the Association for Economic Research of Indigenous Peoples seminar; participants in seminars at Fordham University, Boston College, and Loyola University Chicago; and participants in the invited poster session of the AAEA 2023 Annual Meetings in Washington, DC.

My Jesuit brothers: Thank you to the Markoe House Jesuit Community in Minneapolis for providing both the physical space in the form of an office and the spiritual space in the form of daily liturgies, delicious meals, and other community recreation that both supported me in my academic work and reminded me that I will always be more than an academic. Thank as well to my other Jesuit brothers in special studies as we have supported each other in our common mission. Thanks also to the Sagrada Familia Jesuit Community in Mexico City and the San Francisco Xavier Jesuit Community in Chiapas for hospitality during my research visits.

Other friends in the Twin Cities: Thanks to parishioners at St. Thomas More Catholic Community, St. Catherine's University, St. Stephen's in Minneapolis, Assumption in Richfield, and St. Francis de Sales in St. Paul for sharing the Eucharist and fellowship with me. Thanks also to Crossfit Kingfield for providing wonderful coaching and a physical outlet during these years.

My family: Thanks to my parents Frank and Janet for their love and support, and to my aunt Terre, uncle Jimmy, and cousin Andrea, all of whom visited me in the farthest north place that I have lived.

Funding: Fieldwork and conference travel for this dissertation were supported by the Department of Applied Economics at the University of Minnesota. In particular, I am grateful for the support of the grant "The Value of Relationships" which was funded by the Center for International Food and Agriculture Policy.

Dedication

To the Salvadoran martyrs – Ignacio Ellacuría, Ignacio Martín-Baró, Segundo Montes, Juan Ramón Moreno, Joaquín López y López, Amando López – who used their intellectual gifts in the service of *el pueblo* and paid the ultimate price for it.

To their housekeeper and daughter – Elba and Celina Ramos – and many innocent victims of injustice around the world.

To Jon Sobrino, who told me that it didn't matter what subject that I studied, so long as I stayed close to the reality of the people.

Abstract

Many smallholder agricultural producers around the world struggle to make a living in difficult market environments. These producers face important decisions: whether to enter producer cooperatives, whether to remain loyal to these cooperatives, and whether to stay in rural agriculture or migrate to cities. They are influenced by the decisions of their peers, the market environment, and climate change. My research aims to better understand these decisions at the micro- and macro-levels. With these insights in hand, I hope to work with producer organizations to improve their services and local governments to improve their policies, in both cases to better serve smallholder producers.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Essay 1. With the rise of market-led development, marketing cooperatives have emerged that offer smallholder producers a guaranteed minimum price for their cash crops. Their existence is threatened when members side-sell a part of their harvest to outside buyers. My first essay, entitled *Unpacking Side-Selling: Experimental Evidence from Rural Mexico*, describes a framed field experiment with indigenous coffee producers in southern Mexico to examine the effect of four factors on the marketing decision: additional income, the presence of microcredit and/or technical assistance, average outside buyer price, and harvest quantity. Our results show that participants allocate on average 82% of their harvest to the certain-price buyer. Changes in harvest quantity and outside-buyer price have minimal effects. The offer of complementary services has a null effect. Moreover, 22% of the participants always allocate their entire harvest to the certain-price buyer. Extra income increases this probability by 10%. Subgroup analysis reveals that this effect is limited to existing cooperative members.

Essay 2. Producer organizations can help smallholder producers adapt to climate shocks by insuring production and teaching them climate-resilient production techniques. However, information about the benefits of membership takes time to reach potential adopters and often decays before it reaches an entire population. My second essay, entitled *Information Decay and Cooperative Entry under Risk*, examines the entry of two different cooperatives by indigenous coffee producers: a coffee cooperative and a honey cooperative. Our analysis leverages a network graph of entry decisions that spans 22 years and includes the locations of the producers, who live in 124 villages grouped in ten regions. To characterize the temporal lags, we estimate two specifications of a linear-in-means model: one with with producer fixed effects and another with first differences. To characterize the spatial lags, we estimate three specifications of a spatial lag model with different weighting matrices. In both models, we interact the peer adoption rate with the number of periods of seasonal drought. The linear-in-means estimation results reveal a longer entry period for coffee than for honey. The spatial lag estimation results reveal more information decay for honey than coffee. In space, seasonal drought in other villages and regions increases the probability of entry into the coffee cooperative but not the honey cooperative. In both, we find that periods of seasonal drought counteract network effects for coffee and honey. Our results provide insight for policy makers to strengthen producer organizations in contexts that experience climate shocks.

Essay 3. Recent qualitative evidence suggests that social networks play an important role in potential migrants' decisions to migrate and their choice of destination. Yet even the latest literature employing microeconomics migration models with social networks often only estimates these models on small household panel data sets. My second essay, entitled *Where You Go Depends on Who You Know: Social Networks as Determinants of Mexican Internal Migration*, uses the Mexican population census to estimate a structural gravity model with social networks on internal migration flows from origin municipalities to destination states over three recent five-year periods at the intensive and extensive margin. To proxy for the social networks, I use internal migrant flows along the same corridor in a previous time period. My results show that social networks affect migration flows. At the extensive margin, a 1% increase in the size of the social network increases by 5%, 12%, and 13% the likelihood of a migration corridor; at the intensive margin, the equivalent social network elasticities are 19%, 30%, and 32%. I identify the effects using origin and destination characteristics as well as the presence of a migrant flow in 1960 to control for other factors that could drive migration along these corridors. These results contribute to both microeconomic and macroeconomic analysis of the determinants of migration.

Overview. Each of the three essays examines a different decision of smallholder producers: the decision of how to market their cash crop, the decision of whether to join a producer cooperative, and the decision of whether to exit agriculture and migrate. Taken together, these decisions are some of the most important decisions that smallholder producers make, and our results illuminate the important roles of factors that standard utility maximization models do not include. Liquidity constraints affect a producer's marketing decision, peer effects and drought shocks affect a producer's decision to join a cooperative, and peer effects affect an internal migrant's choice of destination. The relevance of these understudied factors underscores the importance of considering them in future research and policy work.

Chapter 2

Unpacking Side-Selling: Experimental Evidence from Rural Mexico

2.1 Introduction

Smallholder agricultural producers face a variety of market imperfections that reduce the welfare they receive from the sale of their cash crops: output price volatility, monopsony power by traders, and transaction costs.[1] In many developing countries, state-backed organizations, such as commodity boards, alleviate these market imperfections by providing price insurance and other services to producers. However, in recent years, governments have reduced or eliminated these agricultural support programs. As a result, market-based organizations such as producer cooperatives have emerged in their place. Since they lack state support, however, these producer cooperatives depend on the continued loyalty of their members to finance their services, which often improve welfare over the medium and long term. When members sell a portion of their harvest to outside traders in the short term, this side-selling threatens the economic viability of cooperatives.

[1] Thanks to an anonymous reviewer for suggesting that we frame the paper in this way

Empirical estimates of the incidence of side-selling vary widely: 12% (Keenan et al., 2024; Woldie, 2010; Wollni & Fischer, 2015), 20% (Ewusi Koomson et al., 2022), 30% (Alemu et al., 2021; Arana-Coronado et al., 2019), 40% (Gerard et al., 2021) or 55% (Fischer & Qaim, 2014; Geng et al., 2023). Moreover, the amount of side-selling varies both among producers in the same cooperative and within the same producer over different marketing years. Wollni and Fischer (2015) find that side-selling behavior follows the U-shaped pattern first reported by Fafchamps and Hill (2005) regarding producer marketing decisions. Farmers with a low or high production quantity are more loyal to a cooperative. The former cannot pay the fixed cost of side-selling, and the latter are not as affected by the liquidity constraints that often drive side-selling. In addition, production shocks (Keenan et al., 2024) and liquidity shocks (Geng et al., 2023) can also increase side-selling from one year to the next in the same producer. Finally, risk aversion (Binswanger, 1980), length of cooperative membership (Bhuyan, 2007), and the presence of complementary services such as microcredit or technical assistance (Mujawamariya et al., 2013) are also associated with side-selling.

In this essay, we use a framed field experiment to determine the effect of four factors on side-selling: production shocks, income shocks, transaction cost shocks, and nudge reminders of complementary services. Participants play 60 rounds of a game in which each round corresponds to a marketing year. In a given round, they must allocate their harvest across a certain-price and an uncertain-price buyer.[2] In order to estimate the value participants place on the services offered by the certain-price buyer, we vary its description: certain price; certain price and microcredit; certain price, microcredit, and technical assistance. Moreover, we vary the harvest quantity and the mean of the price offered by the uncertain-price buyer to estimate the effect of production shocks and transaction cost shocks, respectively, on marketing behavior. Finally, we give half of the participants additional income from another source to estimate the effect of an income shock. Our experiment integrates these four separate sources of variation that prior work has associated with side-selling. To our knowledge, we are the first to use an experiment to study side-selling.

Our results are as follows. First, price certainty matters at both the intensive and extensive margins. At the overall margin, producers allocate on average 82% of their harvest to the certain-price buyer. At the extensive margin, 22% of the producers (58 of 268) allocate their entire harvest

[2] Thanks to Marc Bellemare for pointing out that technically the uncertain-price buyer is a risky price buyer since the distribution of the outside buyer price is known.

to the certain-price buyer in each round. This estimate of an 18% incidence of side-selling approaches the lower bound of the empirical results above. It suggests that in cases where cooperatives offer a fixed price and outside traders a variable price, side-selling behavior, or its inverse, producer loyalty to cooperatives, is associated with producer risk preferences.

Second, additional income influences side-selling at the extensive margin but not at the intensive margin. At the extensive margin, it increases by 10% a producer's probability of selling the entire harvest to the certain-price buyer in each round. At the intensive margin, it does not affect round-level performance. When we estimate the extensive margin of the effect of the additional income separately for cooperative members and non-members, we find significant heterogeneity in the treatment effects: 17% for members and 2% for non-members. The former effect is significant at the 5% level and the latter is not significant. Two additional moderator analyses give additional information on the mechanisms behind the effect of additional income. First, for cooperative members, the treatment effect of additional income decreases with the number of years of cooperative membership: for a new member it is 42% and decreases by 3% for each year of membership. Second, the treatment effect varies depending on cooperative membership and risk aversion, as measured by a no-loss lottery based on that of Eckel and Grossman (2008). For the least risk-averse cooperative members (CRRA near 0), it is 7%. From there it increases to 31% for the most risk-averse cooperative members (CRRA near 2). For the least risk-averse cooperative non-members, it is 6%. From there it decreases to -29% for the most risk-averse cooperative non-members. None of these effects are statistically significant.

Third, production shocks affect the marketing decision by at most 3% in either direction. Thus we confirm the U-shaped behavior reported by Wollni and Fischer (2015) and Keenan et al. (2024). Though our point estimates are small, they are similar in magnitude to these results. Finally, nudge reminders of complementary services do not affect the marketing decision. This result differs from that of Mujawamariya et al. (2013) and suggests that behavioral economics may not offer a solution to side-selling (Wuepper et al., 2023).

Our results contribute to three distinct strands of literature. First, we contribute to the literature on marketing decisions of agricultural producers. Previous literature has examined the determinants of participation in cooperatives (Bernard & Spielman, 2009; Mojo et al., 2017) and intensity of participation in cooperatives (Bhuyan, 2007; Fischer & Qaim, 2014; Klein et al., 1997; Mujawamariya et al., 2013) using reduced-form models on cross-sectional data sources. Fafchamps and Hill (2005), Woldie (2010), and Wollni and Fischer (2015) propose structural models and test their predictions, once again on cross-sectional data. Our work goes deeper. Instead of the likelihood or intensity of cooperative participation, we examine the demand for the services that cooperatives typically provide. Our results provide insight into the mechanisms behind how much and under what conditions cooperative members market their agricultural production through cooperatives.

Second, we contribute to the literature on the use of experiments to understand producer decision making. Palm-Forster and Messer (2021) provides a recent review of the use of experiments to study the behavior of agricultural producers. Framed field experiments are not new, as the pioneering work of Binswanger (1980) demonstrates. However, they are still as relevant in 2025 as in 1980. They improve on the internal validity of the cross-sectional research above at a fraction of the cost of a Randomized Control Trial. Moreover, they allow for the study of more variation. Casaburi and Reed (2022) pay bonuses to a random subset of traders to examine effects further down the value chain. We too could have randomly subsidized coffee producers with additional income, but at the expense of losing the three other sources of variation in our experiment. The subsidies alone would have cost as much as the entire budget of our experiment.

Our experiment is most similar to three recent experiments. Bellemare et al. (2020) test the prediction of Sandmo (1971) that producers reduce production in situations of price risk and finds that this prediction does not hold. Boyd and Bellemare (2022) both corroborate this finding and also find that the provision of insurance causes producers to increase production in situations of price risk. Mattos and Zinn (2016) find evidence for the existence of producer reference prices in marketing decisions. These three experiments survey a mix of 119 college students and producers, 101 producers, and 75 producers, respectively. Our sample size of 268 producers improves on the external validity of all three experiments, especially since we confirm their findings in a different context.

Third, we contribute to the small literature on price risk (Boyd & Bellemare, 2020). In situations of output price risk, Newbery and Stiglitz (1981) propose methods for evaluating the welfare effects of commodity price stabilization programs. Their work and much of the following work focus on the differential effects of such programs depending on whether agricultural households are net buyers or sellers of the good in question (Barrett, 1996; Bellemare et al., 2013; Finkelshtain & Chalfant, 1991). Our situation differs for two reasons. First, coffee is a cash crop, not a staple, so we do not need to consider the producers' own welfare as consumers. Second, most smallholder producers do not have the infrastructure to store coffee from year to year. Thus, there is no opportunity for arbitrage between growing seasons, just as in the case of the Kenyan roses that Macchiavello and Morjaria (2015) study.

Instead, our work complements that of Bellemare et al. (2021), who consider producers that face output price risk and can allocate their production between a contract that pays a fixed price and an intermediary who pays the market price. They find that contract farming reduces participants' income variability, which they proxy with the residual from a cross-section regression. Moreover, they find that participation in contract farming schemes would be more beneficial for producers who do not participate than it is for those who do. Our sixty-round experiment extends these results, since we can observe producer behavior over many simulated growing seasons and producers can in effect choose the amount of income variability to which they are exposed. Our work also shows the fragility of informal price insurance schemes, especially in years when the market beats the contract and producers are liquidity-constrained.

Our essay proceeds as follows. Section 2 gives background on coffee production worldwide and in Mexico and describes the context where we conducted the experiment. Section 3 describes the design of the experiment and relates it to previous work. Section 4 describes our data and gives descriptive statistics. Section 5 presents the empirical strategy we use to test the effect of the four additional factors on the marketing decision. Section 6 presents and discusses the results, at the participant-round-game level, the participant-round level and the participant level. It also presents results of subgroup and moderation analysis. Section 7 gives policy implications and concludes.

2.2 Context

In this section, we first describe the situation of smallholder coffee producers in Chiapas, Mexico. Next we describe two different development strategies that have sought to improve their welfare and the welfare of other smallholder producers in the developing world: state-led development and market-led development.[3] We touch briefly on the macroeconomic factors that led to a transition from state-led development to market-led development in the early 1990s. Third, we describe the particular institutional features of our partner cooperative. Finally, we describe the challenge that side-selling poses to the cooperative.

Worldwide, coffee is cultivated on approximately 12.5 million farms. Ninety-five percent of coffee producers have farms no larger than five hectares, and eighty-four percent have farms of two hectares or less. For many producers, coffee is their primary cash crop, and therefore their annual income depends on two factors: the yield of their harvest and the world price of coffee. Mexico is the tenth largest coffee producer in the world.[4]

Although our analysis focuses on smallholder Mexican coffee producers, the issues here are not limited to Mexico or coffee. The picture we paint here is broadly similar to the situation of smallholder producers of specialty crops in Latin America and elsewhere who are members of cooperatives (Pitts, 2023). Side-selling has been observed with coffee producers in Peru (Keenan et al., 2024), coffee producers in Burundi (Gerard et al., 2021), coffee producers in Costa Rica (Wollni & Fischer, 2015), banana producers in Ethiopia (Woldie, 2010), dairy producers in Kenya (Geng et al., 2023), sorghum producers in Kenya (Nyamamba et al., 2022), and barley producers in Ethiopia (Alemu et al., 2021). In all of these contexts, cooperatives offer value chain integration, quality upgrading, and microcredit to smallholder agricultural producers, but the provision of these services is hindered by producer members who do not market their harvest through the cooperative.

2.2.1 Smallholder Mexican Coffee Production

Our setting is a group of indigenous Mexican coffee producers in the state of Chiapas in southern Mexico. Coffee is the primary cash crop for these producers. They typically produce 4 quintals (240 kilograms) on 1-2 hectares of land and sell their coffee for MXN 70-80 (approximately US$3.50) per kilogram. Thus, they earn around US$1000, which they use to purchase everything they do not

[3] We are grateful to the lead article in a special issue of *Food Policy* for this distinction (Markelova et al., 2009).

[4] Wright et al. (2024) provides a recent systematic review of the literature on coffee supply chains at the global level.

grow for themselves.[5] Typically, they grow corn and vegetables for their own consumption.

Coffee has been grown in Mexico since the nineteenth century (Bobrow-Strain, 2007). Initially, Mexican peasants worked as hired labor on large coffee plantations. In the early twentieth century, as a result of the land redistribution of the 1917 Mexican Constitution, these smallholder producers received their own plots of land, nearly all of which were less than 5 hectares.

Green coffee is only the first stage in the coffee value chain. Figure 2.1 provides a high-level overview of the entire coffee value chain. A smallholder producer sells to a local intermediary (either a village trader or in our case, a coffee cooperative). This local intermediary in turn sells to a national intermediary. Finally, that national intermediary sells to a multinational corporation.

Smallholder coffee producers face substantial output price volatility at the first level of the value chain. Because of this output price volatility, they do not produce an optimal amount of coffee. In addition, they do not make long-term investments in coffee production through quality upgrading that would allow them to increase the income they receive from coffee production.

2.2.2 State-Led Development for Mexican Coffee Producers

The past hundred years have seen two different approaches to improve the welfare that smallholder coffee producers receive from their harvest: state-led and market-led development. In the first approach, state actors provided increased support for smallholder coffee producers as coffee production developed in Mexico through the early and middle of the twentieth century. At the international level, in 1962 the coffee producing nations of the world formed the International Coffee Organization (ICO) in order to stabilize the world market for coffee after a series of boom-bust cycles. With the establishment of the ICO, the International Coffee Agreement (ICA) used periodically renewed export quotas to stabilize the international price of coffee. This agreement lasted until 1994.

At the national level, in 1973 the Mexican government founded a state agency to support coffee producers, the Mexican Coffee Institute (Renard & Breña, 2010). This agency provided direct support to coffee producers: subsidized inputs, technical assistance, and a guaranteed purchase price. In turn, it helped Mexican coffee producers sell their coffee internationally for almost 20 years.

[5] This profile comes from previous work in this region by Pitts (2019).

The life cycle of the Mexican Coffee Institute overlapped with the external debt crisis faced by Mexico and other Latin American countries during the 1980s. As part of the Baker Reforms in 1986, Mexico agreed to reduce the level of agricultural support for domestic producers in order to receive international financing to cover its external debt. As a result, beginning in 1990, the administration of Mexican president Carlos Salinas phased out the Mexican Coffee Institute as it implemented a larger series of market-based reforms. Since then, Mexican smallholder coffee producers have been exposed to the international price of green coffee as it is traded on commodity markets like the New York Mercantile Exchange. Figure 2.2 shows the price per kilogram of green Arabica coffee in US cents from 1990 to the present.

2.2.3 Market-Led Development for Mexican Coffee Producers

With the elimination of the Mexican Coffee Institute, producer cooperatives emerged in Mexico in the 1990s that provide the same services to smallholder producers: a guaranteed purchase price, technical assistance, and microcredit (Folch & Planas, 2019). These cooperatives are often associated with the fair trade movement (Dragusanu et al., 2014). In addition, they frequently promote organic farming practices. Typically, members have three years from joining the cooperative to adopt organic farming practices.

Producer cooperatives improve the welfare of their producer members by taking advantage of upstream contracts in the value chain. These contracts help in two ways. First, they allow cooperatives to offer a fixed price to their members, in contrast to the volatile price of the commodity market. Second, they spread the fixed marketing costs faced by smallholders over a larger marketed volume to reduce the cost per unit of marketing. Cooperatives use these savings to finance complementary services, such as microcredit and technical assistance. However, the effectiveness of cooperatives depends on a guaranteed volume of deliveries from members. As a condition of membership, these cooperatives often require that their members sell all their coffee through the cooperative. When members sell their coffee to outside buyers, it threatens the financial viability of the cooperative.

2.2.4 Our Partner Cooperative: Ts'umbal Xitalha'

Producer cooperatives provide a variety of services and operate in a variety of ways, so we describe the particular way our partner cooperative operates and the particular services it provides.

The producer cooperative Ts'umbal Xitalha' (TX) has existed since 2000. It has evolved to provide price insurance, emergency loans, and technical assistance to its producer members. In October, at the beginning of each marketing year, the executive board of the cooperative sets the purchase price for the coming year. The TX members agree to deliver their coffee to the cooperative during the harvest season. Unlike other cooperatives, which pay their members at the end of the marketing year, TX pays on delivery. Local intermediaries or traders are also active in the region. They buy coffee at the world price, which varies daily, as Figure 2.2 indicates, plus a small markup of MXN 5 to MXN 10 (US$0.50 or US$1.00). Although TX stipulates that its members market their entire coffee harvest through the cooperative, it cannot enforce this requirement. Thus, when local traders offer a higher price than TX, members face the temptation to market some or all of their coffee harvest through these local traders instead of the cooperative.

Figure 2.3 shows the TX price and the world price for the marketing years 2019 to 2025. Figure 2.4 summarizes TX administrative data to show the number of members who delivered their coffee to TX in each marketing year and the total volume of coffee that these members delivered.

In the marketing year after the pandemic, the world price of coffee (and thus the price offered by local traders) increased above the price offered by TX for an extended period of time. This situation emerged for two reasons. First, a decrease in demand among TX's customers left it with excess inventory and reduced the price it could offer the following year. Second, higher transaction costs and labor issues across the worldwide coffee industry caused an increase in the world price of coffee and thus the price offered by local intermediaries.

As a result, member coffee deliveries to TX decreased by half during 2021 and 2022, a phenomenon that affected the viability of TX. Figure 2.5 shows that more of the decline occurred at the intensive margin than at the extensive margin. Although the total number of members who delivered any of their coffee harvest to TX decreased, many members continued to deliver some of their coffee harvest to the cooperative, but substantially reduced the size of their deliveries. Because members do not disclose the total amount of their coffee harvest to TX, TX cannot know whether members are side-selling or how much they are side-selling.

In order to continue serving its members, TX sought external financing to increase the price it could offer its members. As Figure 2.4 indicates, by 2023, the world price of coffee decreased, so TX members no longer faced the temptation to side-sell. However, TX partnered with us to understand more deeply the causes of side-selling behavior and explore potential policy responses to prepare for a possible future scenario in which the world price of coffee could once again exceed the price offered by TX.

2.3 Experimental Design

2.3.1 Experiment Overview

In this section, we describe our experimental protocol that examines coffee producers' marketing decisions. We present participants with a simplified version of the marketing decision they face in real life, which we described in Section 2.2. Within the taxonomy of field experiments, our experiment is an framed field experiment (Harrison & List, 2004) or a lab-in-the-field experiment (Eckel & Londono, 2021) because we invite members of the target population to replicate a concrete task that they perform in their daily lives. We simplify the decision in four ways to better understand the core mechanism at work.

1. Ideally, any side-selling by members would be punished by expulsion from the cooperative. Thus the cooperative would be able to force its members to always deliver their entire harvest to the cooperative. This sort of punishment is infeasible for two reasons:

    (a) First, in this region, as in many regions with a substantial population of smallholder producers, nearly all of the cooperative members have social ties that stretch back for generations. Punishing members who side-sell would negatively affect these ties in ways that would spill over to religious, cultural, or other economic interactions.

    (b) Second, the cooperative does not record the total harvest of members, so it cannot verify the fraction of members' harvest that they are marketing through the cooperative. For this reason, we model side-selling as an isolated decision that producer members make independently each year.

2. Many estimates of side-selling in the literature come from contexts with variation in the timing of payment. Smallholder producers may choose between a local trader that pays them immediately and a cooperative that pays them at the end of the growing season. In this case, a producer's time preferences would influence the decision to side-sell. To eliminate this potential confounder, in our experiment participants are paid immediately by both the certain-price and the uncertain-price buyer

3. The presence of transaction costs also varies depending on the context. In some contexts, producers who side-sell to a local trader incur a fixed cost compared to selling to the cooperative. In other contexts, producers who sell to the cooperative incur a fixed cost compared to selling to the local trader. To consider both situations, we vary the mean of the price offered by the uncertain-price buyer as either above, below, or the same as price offered by the certain-price buyer. These three options correspond to contexts where there is a fixed cost to side-selling, no fixed cost to either marketing decision, or a fixed cost to selling to the cooperative.

4. Finally, institutional arrangements with respect to complementary services vary tremendously. In some contexts, local traders provide microcredit and possibly even technical assistance. In other contexts, only cooperatives provide these services. In addition, institutional arrangements vary in terms of eligibility for either of these services. The strictest possible arrangement would restrict complementary services to cooperative members. Spillover effects among neighbors, some of whom are cooperative members and others who are not, often prevent the enforcement of this sort of restriction. Thus, we provide nudge reminders to test for the effect of the provision of these services.

   Nudge reminders are often used in food economics to persuade consumers to make healthier food choices (Caputo & Just, 2022). For example, choosing healthier food or smaller portions will benefit consumers in the medium to long run. Instead of coercing consumers (a paternalistic approach), researchers have tried to randomly vary the labeling of foods or provide additional information on menus as a way to induce consumers to voluntarily make healthier choices. Depending on the context, these types of interventions can have a moderate effect, especially if they are directly relevant (salient) to the consumer. We approach a participant's marketing decision in a similar way, testing the effect of changing the description of the certain-price buyer on the participant's allocation decision. We incorporate salience into our descriptions: for example, we describe the certain-price buyer specifically as a buyer that "provided you (the participant) with microcredit in the past year" instead of generically as a buyer "who provides microcredit." These nudge reminders allow us to separately evaluate the appeal of the three services that cooperatives and other traders most commonly provide: price insurance, microcredit, and technical assistance.

By examining the effect (if any) of the nudge reminders on the allocation decision, we hope to estimate the participant's willingness to pay for these additional services. Moreover, using nudge reminders allows us to separate the potential service itself from the intermediary (cooperative or local trader) who provides it. For example, in the second game, we add microcredit to the description of the certain-price buyer but do not describe it as a cooperative provides it. In this way, our aim is to estimate separately participants' willingness to pay for microcredit and their preference for a cooperative.

In the experiment, participants market their coffee 60 times in three games of 20 rounds apiece. During these 60 rounds, we vary four factors to determine their effect on the marketing decision.

1. Half of the participants receive **additional income** at the beginning of the experiment that increases their earnings in each round of the three games they play.

2. In blocks of 20 rounds, we vary the **framing** of the certain-price buyer as a buyer who offers a certain price (Game 1); a certain price and microcredit (Game 2); a certain price, microcredit, and technical assistance (Game 3). All participants play all three games in random order.

3. By round, we vary the **harvest size of the participants**, **the mean of the price offered by the uncertain-price buyer**, and **the realized price of the uncertain-price buyer**.

In the subsequent sections, we describe each part of the experiment conceptually in detail: both the antecedents in the literature and the practical details in our experiment. In Section 2.4, we introduce the notation for the different pieces of the experiment and provide the payoff function. Appendix A gives the complete experimental protocol.

2.3.2 Preliminary Questions

Participants first answer two basic questions and three arithmetic and probability questions. We use questions similar to those in Boyd and Bellemare (2022).

Basic Questions

1. Have you ever sold coffee you or your family has produced?

2. Do you know how to read and write? Yes/No

Filter Questions

1. What is 40% of MXN 100?

2. If you produce 17 bags of coffee and sell 9, how many remain?

3. Imagine that there are 3 blue balls and 7 red balls. You pick a ball at random. Is it more probable that it is red or blue?

Descriptive statistics of the responses to four out of five of these questions are reported in Table 2.1. All participants report experience selling coffee, so we omit these responses. 74% of participants report some literacy.

The second set of questions allows us to determine whether side-selling behavior is associated with poor multiplication, subtraction, or probability skills. Originally, we intended to exclude (filter) participants who missed more than one of the questions. This criterion would have disqualified two of the participants. However, based on the guidance of our implementing partner, we did not exclude any participants.

Next, the order in which the three games and the lottery are played is randomized by a roll of a 12-sided die. Table 2.2 shows the results of this randomization. Half of the participants complete the lottery before the three games, and the other half complete it after the three games.

2.3.3 Additional Income Treatment

Next, half of the participants receive MXN 3,000 (US$150) in fake money that serves as additional income in each round of the three games and contributes to their overall earnings. The treated participants are selected based on their identification number within the sample: participants with odd numbers receive the additional income and participants with even numbers do not receive the additional income.

The additional income is meant to proxy for the real-world effect of income from another source such as the sale of another cash crop, income from off-farm labor, or support from a Mexican government program. We choose an amount (MXN 3,000) that is about half of what producers could conceivably earn from these sources in a month.

1. **Another cash crop**. The main alternative cash crop in the region is honey. According to records from a honey cooperative in the region, producer members earned on average MXN 20000 (US$1000) from honey sales during the three and a half months of the honey season the year before the experiment, or just under MXN 6000 (US$300) per month.

2. **Income from off-farm labor**. Similarly, weekly pay is MXN 1500 (US$150) in manufacturing plants on the US/Mexico border, where many producers report migrating seasonally. With one to two months of work, minus expenses, a producer could earn about MXN 6000 (US$300).

3. **Support from a Mexican government program**. Finally, participants in this region are eligible for a Mexican government agricultural support program (Sembrando Vida), in which smallholder producers can earn up to MXN 6000 (US$300) per month by planting trees on their land parcels (*Reglas de Operación Del Programa Sembrando Vida*, 2022).

Randomly assigning this treatment allows us to determine the effect of additional income on the marketing decisions of the participants who receive it. Previous work examines the effect of additional income on production decisions of cash crops (Pfeiffer et al., 2009) or marketing decisions of staple goods (Woldeyohanes et al., 2017). Pfeiffer et al. (2009) find that additional income causes producers to increase production in the presence of a credit market failure because they use it to finance the purchase of production inputs. Woldeyohanes et al. (2017) find that producers market less of staple goods in the presence of off-farm income in order to keep a food reserve and insure consumption. Our experiment does not allow participants to store coffee across years.

The closest study to the present is that of Wollni and Fischer (2015), who hypothesize that non-agricultural income will increase member deliveries to cooperatives. In their model, however, cooperatives deliver patronage refunds at the end of the marketing year, so the non-agricultural income merely allows for consumption smoothing across time periods.[6] In our context, any effect of the additional income will indicate deviation from purely profit-maximizing behavior. To our knowledge, we are the first to experimentally test the effect of additional income on the marketing decision of a cash crop.

2.3.4 Eckel-Grossman Lottery

Next participants roll Participants complete an Eckel-Grossman lottery to measure their risk preferences. Eckel and Grossman (2008) propose a simple task for measuring risk preferences similar to that of Binswanger (1980). Participants choose one of five gambles, each with a low payoff and a high payoff that occur with 50% probability. The gambles are increasing in both expected payoff and risk, as measured by the standard deviation between the two payoffs. After participants choose their preferred gamble, they roll a die and receive the corresponding payoff.

An advantage of the Eckel-Grossman lottery compared to other lotteries such as that of Holt and Laury (2002) is its simplicity (Charness et al., 2013). This simplicity allows its use in other settings in Latin America with a population similar to our indigenous coffee growers (Cardenas & Carpenter, 2013; Moya, 2018). Moreover, despite its simplicity, the participant's choice of gamble can be used to estimate his or her risk preferences in the form of a Constant Relative Risk Aversion (CRRA) parameter of the power utility function $U(x) = x^{(1-r)}/(1-r)$.

Table 2.3 shows the Eckel-Grossman lottery that we present to our participants. Eckel and Grossman (2008) provide two sets of gambles: one with negative payoffs (to test for loss aversion)

[6] In many cooperatives, additional profits above and beyond the price paid for members' production are distributed to members at the end of the fiscal year. These additional payments are called patronage refunds. This feature sets cooperatives apart from investor-owned firms (IOFs), which distribute profits to shareholders at the end of the fiscal year.

and one without. For simplicity, we use the no-loss lottery and scale the payoffs ($16 = MXN 10000 or US$500) so that the first gamble has a guaranteed payoff of MXN 10000 MXN. We choose MXN 10000 because it is the average payoff in a round of the game (4 quintals · 60 kilograms per quintal · MXN 50 per kilogram = MXN 10000).

2.3.5 The Presence of Complementary Services

After the preliminary activities, participants complete 10 rounds of Game 1 for practice. The results of this practice game are not recorded.[7] Next they complete Games 1-3 in random order. Game 2 and Game 3 vary the framing of the certain-price buyer by describing up to two complementary services that the participant received last year from the buyer. In addition, in Game 3, the certain-price buyer is described as a cooperative.

Game 1 certain-price buyer offers a fixed price of MXN 50 per kilogram.

Game 2 certain-price buyer offers a fixed price of MXN 50 per kilgram **and gave the participant microcredit in the past year**.

Game 3 **A cooperative** offers a fixed price of MXN 50 per kilogram **and gave the participant microcredit and technical assistance last year**.

As Section 2.2 describes, microcredit and technical assistance are provided by TX, the cooperative that operates in this region. The welfare-enhancing effects of both services are confirmed by a recent systematic review (Liverpool-Tasie et al., 2020). However, supplying these services imposes additional costs on the cooperative that lower the guaranteed minimum price it can offer members for their coffee. Moreover, not only the cooperative, but also local traders can provide fixed prices and complementary services. Here, we are interested in whether participants value these services enough to market at least a fraction of their coffee through a buyer that offers these services even if they could earn more by marketing it through a buyer that does not.

In all three games, the certain-price buyer provides a fixed price. Thus, the allocation decision tests the participant's preference for price certainty. In the second game, the certain-price buyer

[7] Because of enumerator error, 40 participants did not complete the practice game. We include it as a control in the regressions.

is described as a buyer who provided the participant with microcredit in the past year. This buyer could be a trader or a cooperative; the game does not specify. Rather, by comparing the participant's preference for the second over the first framing of the certain-price buyer, we hope to estimate the participant's willingness-to-pay for microcredit. Only in the third game is the certain-price buyer described as a cooperative that provides microcredit technical assistance. By comparing the participant's preference for the third over the second framing of the certain-price buyer, we hope to estimate the participant's willingness-to-pray for technical assistance provided by a cooperative.[8]

2.3.6 Harvest Quantity

Each round of the experiment corresponds to a marketing year. At the beginning of the round, the participant's harvest quantity for that year is determined randomly by the roll of a 12-sided die. Each of the four possibilities for the harvest quantity – 2, 4, 6, or 8 quintals – appears with the same probability (25%).[9] Once the harvest quantity is realized, participants receive a corresponding number of miniature burlap bags.

Under a profit-maximizing framework, harvest quantity should not impact the marketing decision. Profit-maximizing participants should sell their entire harvest to the buyer who gives them the best price. However, previous studies indicate that harvest quantity affects the marketing decision; moreover, they find that it affects the decision differently for poor producers and rich producers.

Fafchamps and Hill (2005) examine the decision to sell coffee at the farmgate or market by Ugandan coffee producers. They find a U-shaped relationship: the very poor and very rich are more likely to sell at the farmgate, because of lack of transportation to the market for the former and the higher opportunity cost of time for the trip to the market for the latter.

Wollni and Fischer (2015) examine determinants of how producers allocate their coffee harvest among two buyers. They also find a U-shaped relationship between farm size and coffee deliveries. Initially, the relative profitability of marketing to outside buyers increases with farm size, and

[8]An anonymous reviewer raised the concern that adding two elements at the same time in Game 3 to the framing in Game 2 conflates the preference for technical assistance and for cooperatives. In this context, only cooperatives provide technical assistance. However, future work could vary these attributes separately.

[9]A quintal is a local unit that corresponds to 60 kilograms of green coffee.

producers with medium-sized farms sell more to outside buyers than producers with smaller farms. However, as farm size continues to increase, however, producers' discount rate of patronage refunds decreases as well. The reason is that producers with larger farms have more access to other sources of income than producers with medium-sized farms to insure their consumption and deal with unexpected expenses. Thus, producers with larger farmers sell a smaller share of their harvest to outside buyers than producers with medium-sized farms. Based on this previous work, we expect to find a U-shaped relationship between harvest quantity and producer marketing decisions.

2.3.7 Certain-Price Buyer vs Uncertain-price Buyer

In each round, participants allocate their harvest between a certain-price and an uncertain-price buyer. The certain-price buyer always offers them MXN 50 (US$2.50) per kilogram for their coffee. The description of the certain-price buyer varies according to the presence of complementary services above. The uncertain-price buyer offers them a price whose mean varies: below the certain price (MXN 45 or US$2.25), the same as the certain price (MXN 50 or US$2.50), or above the certain price (MXN 55 or US$2.75). The price follows a multinomial distribution with five supports that is constructed to approximate a normal distribution. Figure 2.6 shows the three possible distributions. In each distribution, the mean appears four times, the two values MXN 5 above and below the mean appear three times, and the values MXN 10 above and below the mean appear once. Constructing the distribution in this way allows the roll of a 12-sided die to approximate a draw from a normal distribution.

Crucially, participants allocate their coffee harvest after learning the mean of the price of uncertain-price buyer (in other words, which of the three distributions the realized price will follow) but before learning the realized price. Next, they allocate their coffee harvest between the two buyers in increments of one quintal. To aid them in the allocation decision, a payoff table is shown that gives the revenue from all possible allocations conditional on the coffee harvest and distribution of the uncertain-price buyer in the current round. They must allocate the entire harvest and cannot store coffee for subsequent rounds. Figure 2.7 gives a representative table.[10]

Conditional on the mean of the uncertain-price buyer, expected utility theory predicts the

[10] Appendix A contains all 12 possible payoff tables.

behavior of a risk-neutral participant as follows.

1. **Uncertain-price buyer mean of MXN 45.** Allocate the entire harvest to the certain-price buyer.

2. **Uncertain-price buyer mean of MXN 50.** Be indifferent between the certain-price buyer and uncertain-price buyer.

3. **Uncertain-price buyer mean of MXN 55.** Allocate the entire harvest to the uncertain-pirce buyer.

Notably, in all three scenarios, depending on the realization of the price of the uncertain-price buyer, participants could potentially make more revenue by allocating some or all of their harvest to the uncertain-price buyer.

In practice, participants are not risk neutral. They are risk averse but vary in the degree of risk aversion. Thus, we can use their allocation decisions to recover their risk preferences as follows. Examining allocation decisions in the situation where the mean of the uncertain-price buyer is MXN 50, the same price offered by the certain-price buyer, allows us to determine participants' preferences for price certainty. Adding the other two possibilities for the uncertain-price buyer (mean that is MXN 5 higher or MXN 5 lower than the price offered by the certain-price buyer) allows for the estimation of the effect of small changes in the market environment on participants' allocation decisions. As we pointed out above, these slight variations in price could reflect differences in transaction costs or daily variation in the world price of coffee.

After the allocation decision, participants learn the realized price of the uncertain-price buyer and their revenue for the round. This revenue is added to their running total for the experiment. If they are in the treatment group for the additional income, then this income is added as well at the end of each round.

2.3.8 Final Activities

Recall that we randomly assigned half of the participants to complete the Eckel-Grossman lottery before the three games and half to complete it after the three games. The half of the participants who did not complete it before complete it now. All participants complete an exit survey with information about their household and agricultural production.[11]

2.3.9 Compensation

We compensate participants based on their performance in the experiment. On the advice of our implementing partner, we do not make cash payments to participants. In this way, we differentiate ourselves from the representatives of the Mexican government who distribute various support programs either directly in cash or via direct deposit. Rather, we provide vouchers redeemable on site for dry goods: a bottle of cooking oil, laundry detergent, a bag of sugar, a bag of salt, or a bag of rice. Each voucher corresponds to earnings of MXN 250,000 in the game. Participants can earn between three and six vouchers.

A potential concern is that participants who are assigned the additional income treatment could receive more compensation than those who are in the control group. Recall that treated participants receive MXN 3,000 additional income per round or MXN 180,000 of additional income over sixty rounds. At most, they receive one voucher more compared to a counterfactual scenario with identical performance in the game but without the treatment. Thus, we argue that the possible compensation is nearly the same for treatment and control participants and thus treatment assignment does not affect participants' behavior in the game.

This compensation satisfies the three criteria proposed by Eckel and Londono (2021). It is *monotonic* because participants who perform better in the game receive more compensation. It is *salient* because participants understand how their actions in the experiment translate into their level of compensation. It is *dominant* because the market value of these products corresponds to the opportunity cost of a day's wages that participants give up to participate in the game.

[11] Appendix A contains the entire survey.

2.4 Data and Descriptive Statistics

2.4.1 Sample Selection

Data come from a framed field experiment that we conducted with 268 indigenous coffee producers in northeast Chiapas, Mexico, in summer 2022. During this period, we scheduled eleven field visits to eight of the ten regional centers in the area served by the Ts'umbal Xitalha' (TX) coffee cooperative. For logistical reasons, we were unable to visit two of the regional centers. The field visit dates were announced and arranged through local churches and community centers, so cooperative members and non-members were equally aware of the opportunity to participate. At three regional centers, more participants volunteered than we could accommodate in a single day, so we returned for a second day to those sites to accommodate all participants. After all field visits were completed, we used the TX member list to determine which participants came from families that marketed their coffee through the cooperative and classified them accordingly. Table 2.4 gives an overview of the field visits and a breakdown of the number of cooperative members and non-members who participated in the experiment at each regional center.

We briefly discuss the external validity of the study. The external validity of our study refers to the extent to which the results are representative of those of the population under study, indigenous coffee producers. One potential threat to the external validity of our study could be selection bias. For example, Frijters et al. (2015) found selection bias in an artefactual field experiments in rural China. We argue that our sample does not suffer from selection bias for the following reasons:

1. Any coffee producer can participate in the experiment. We do not allow more than one individual from the same family to participate in the study due to the limited amount of dry goods we bring on the field visit for compensation.

2. Participation is not associated with on-farm economic opportunities. We conducted the experiments in the summer between the planting season and the harvest season. The coffee harvest of participants would not be affected if they neglected it for one day to participate in the experiment. Similarly, it is unlikely that their neighbors would ask them for their help with their coffee fields on the day of the field visit. Thus, there is no social or financial

opportunity cost to participating in the experiment.

3. Participation is not associated with off-farm economic opportunities. Although some indigenous people in this region internally migrate to work off-farm in the summer months, whole families do not. Thus, if one member of a family is away pursuing off-farm work, then a family can send another member to participate. In fact, some did.

4. Our sample of 268 producers is larger than the sample for similar experiments. It is slightly larger than that of Binswanger (1980), who surveyed 240 Indian smallholder producers, and it is considerably larger than that of Mattos and Zinn (2016), who surveyed 75 grain producers in Manitoba; the sample of Bellemare et al. (2020), who surveyed a combination of 119 US undergraduates and Peruvian potato farmers; and the sample of Boyd and Bellemare (2022), who surveyed 101 Peruvian potato farmers.

The external validity of our study also refers to the degree to which our results generalize to other populations. As we describe in Section 2.2, this population of coffee producers is demographically representative of smallholder coffee producers in other places. In addition, with the breakdown of the International Coffee Agreement, cooperatives similar to our partner cooperative have emerged in coffee-producing regions around the world. Like our partner cooperative, these cooperatives struggle to compete with local traders as they provide value-added services such as microcredit and technical assistance. Due to weak institutional arrangements, they also struggle with side-selling. Moreover, as we describe in the Introduction, the issue of side-selling extends beyond coffee to any number of other cash crops that smallholder producers market through agricultural producer cooperatives. We argue that our study also sheds light on the causes of side-selling in these contexts.

2.4.2 Descriptive Statistics at the Participant Level

Table 2.1 presents summary statistics at the participant level. The first group of characteristics comes from the exit survey that participants complete after the experiment. The sample has slightly more men than women. The mean age of the participants is 44 years with a standard deviation of 16 years. There are slightly less women (n=131) than men (n=137). In addition to gender, we also report on the educational level of participants. Mexico requires nine years of compulsory education:

six of primary school and three of secondary school. Most of the participants (75%) report only a primary school education. 14% report only a middle school (secondary school) education. 11% have also completed high school (preparatory school). All participants speak an indigenous language (Tseltal) as their first language and learn Spanish as their second language starting at primary school.

The second group of characteristics comes from administrative data from the cooperative. As we mentioned above, after completing all field visits, we matched participant names to the TX member list to label 126 participants as cooperative members. For 124 of these members, the cooperative could provide us with the number of years in the period 2013-24 that these members delivered coffee to the cooperative. We use this value to measure members' loyalty to the cooperative. Figure 2.11 displays a histogram of these values.

The third group of characteristics comes from the preliminary activities: filter questions, treatment assignment, and lottery. Participants answer five preliminary questions before participating in the experiment to assess their understanding of basic mathematical concepts. Section 2.3.2 gives more information. All participants grow coffee and 74% report being able to read and write. All 268 correctly answer the arithmetic question, 266 correctly answer the percentage question, and 200 correctly answer the probability question. After the preliminary questions, they are randomly assigned MXN 3,000 additional non-farm income. We see an equal number of treatment (n=134) and control (n=134) participants.

In the main part of the experiment, participants complete the practice game and three games of 20 rounds apiece. The games differ in how they frame the certain-price buyer. Section 2.3.5 gives more information. We randomize game order and lottery placement using a 12-sided die. Table 2.2 shows the results of this randomization. Both lottery placement and the order of the games are approximately randomized.

2.4.3 Eckel-Grossman Lottery

Participants complete an Eckel-Grossman risk preference elicitation lottery before or after the practice rounds and three games. Section 2.3.4 gives more information. Figure 2.8 shows gamble choices of participants broken down by gender. In our results, men and women show the same preferences with the highest preference for gamble 5. These results differ from those of Eckel and Grossman (2008), who find gender differences in lottery preferences. In their results, men's preferences are right-skewed with the highest preference for gamble 5, and women's preferences follow a normal distribution with the highest preference for gamble 3. Figure 2.9 shows gamble choices of participants broken down by cooperative membership status. As we discuss in Section 2.6, cooperative members are slightly more risk averse than cooperative non-members. Figure 2.10 shows the gamble choices of the participants broken down by lottery position. There does not appear to be an association between lottery position and gamble choice.

2.4.4 Descriptive Statistics at the Participant-Round Level

Table 2.5 presents summary statistics at the participant-round level. In each round, the size of the participant's harvest and the mean of the price offered by the uncertain-price buyer both vary randomly according to a roll of a 12-sided die. Section 2.3 gives more details. We code both of these experimental variables as dummy variables with four and three possibilities, respectively. Perfectly randomized experimental variables would exhibit sample probabilities of 0.25 for each possibility of the harvest and 0.33 for each possibility of the mean price of the uncertain-price buyer. Our sample slightly favors a harvest of 6 or 8 quintales and a mean price of the uncertain-price buyer of MXN 50 due to physical idiosyncrasies with the die.

2.4.5 Outcomes of Interest

The outcome of interest is the share of the harvest that participants allocate to the certain-price buyer in each round of the experiment. We compute it as follows. Let i denote the participant, $g \in \{1, 2, 3\}$ denote the game, and $t \in \{1, 2, \ldots, 20\}$ denote the round. In each round, participants learn the harvest quantity, $q^g_{i,t} \in \{2, 4, 6, 8\}$, and the mean price of the uncertain-price buyer $p^{pg}_{i,t} \in \{45, 50, 55\}$. They choose how many quintals $z^g_{i,t} \in \{0, 1, 2, 3, 4, 5, 6, 7, 8\}$ to allocate to the certain-price buyer. We compute the share as $\delta^g_{i,t} = z^g_{i,t}/q^g_{i,t}$.

When we pool allocation decisions for all three games for the same participant, the notation above changes slightly. Here we denote the round as $t \in \{1, 2, \ldots, 60\}$ and drop the g superscript from the harvest quantity and the mean price of the uncertain-price buyer, so they are $q_{i,t}$ and $p^p_{i,t}$, respectively. The participant's choice is $z_{i,t}$. We compute the share as $d_{i,t} = z_{i,t}/q_{i,t}$. For round-level regressions, our outcome of interest is precisely the game-level allocation $\delta^g_{i,t}$ or the pooled allocation $d_{i,t}$. The pooling of the allocations does not change their cardinal values. It just maps them from $\delta^g_{i,t}$ space where $g \in \{1, 2, 3\}$ and $t \in \{1, 2, \ldots, 20\}$ to $d_{i,t}$ space where $t \in \{1, 2, \ldots, 60\}$. Table 2.5 gives descriptive statistics for this outcome.

For the participant level regressions, we aggregate the pooled participant-round allocation $d_{i,t}$ across rounds as follows. Because one quarter of the sample (n=58) allocate their entire harvest to the certain-price buyer in every round, we separate the overall margin into the extensive and intensive margin so that we can analyze them separately. Table 2.1 gives descriptive statistics for these outcomes.

1. The **overall margin** is the average allocation for a participant over 60 rounds, or $d_i = \frac{1}{60}\sum_{t=1}^{60} d_{i,t}$.
2. The **extensive margin** is an indicator variable of whether a participant allocates his or her entire harvest to the certain-price buyer in all rounds, or $\bar{d}_i = I[d_i = 1]$.
3. The **intensive margin** is the average allocation of those participants who do not allocate their entire harvest to the certain-price buyer in all rounds.

Figure 2.12 presents a histogram of the overall margin broken down into participants who received the MXN 3,000 additional income treatment and those who did not. The left shift in the allocation of the additional income group suggests that the treatment is associated with a decrease in the overall margin.

Figure 2.13 presents a histogram of the overall margin broken down by cooperative membership status. The right shift in the allocation of the non-members suggests that cooperative membership status is associated with a decrease in the incidence of side-selling.

2.4.6 Payoff Function

We put the payoff function of the experiment below. We suppress the subscript i for each participant and consider the arrangement of the data in which the three games are pooled together into sixty rounds per individual. Each round is denoted by t. In round t, the harvest quantity is denoted by q_t, the realized price of the uncertain-price buyer by p_t^p, and the fraction of the allocation to the certain-price buyer by δ_t. The payoff of the Eckel-Grossman lottery is denoted by L. The indicator variable extra$_i$ is 1 if the participant receives the additional income treatment and 0 otherwise.

$$\Pi = L + \sum_{t=1}^{60} (3000 \cdot \text{extra} + \delta_t \cdot q_t \cdot 50 + (1 - \delta_t) \cdot q_t p_t^p) \tag{2.1}$$

2.5 Empirical Framework

We now describe our empirical framework. First, we discuss our estimation strategy at the participant-game-round level, the participant-round level, and the participant level. In particular, we use moderation analysis at the participant level to estimate the effect of the additional income treatment moderated by risk aversion as measured by the Eckel-Grossman lottery and moderated by loyalty to the cooperative as measured by years of participation. Next, we discuss our identification strategy. Finally, we discuss subgroup analysis among cooperative members and non-members to test for heterogeneous treatment effects.

2.5.1 Estimation Strategy

We estimate the effect of four sources of variation on the marketing decisions of participants: the presence of additional income, a change in the framing of the certain-price buyer, an increase or decrease in the harvest quantity, and an increase or decrease in the mean price offered by the uncertain-price buyer. Since these four sources of variation vary at three levels, our estimation strategy operates at three levels. First, we estimate the effect of the harvest quantity and the mean price offered by the uncertain-price buyer at the participant-game-round level. Next, we pool all three games and estimate the effect of harvest quantity, mean price offered by uncertain-price buyer, and game framing, this time at the participant-round level. Finally, we aggregate participants' allocations across all 60 rounds and estimate the effect of the additional income treatment at the participant level.

Estimation at the Participant-Game-Round Level

Recall from Section 2.4.5 that we denote round-level outcomes in two ways to distinguish between the estimation in this section, which separates allocations by game, and the estimation in the next section, which pools allocations across all three games. Table 2.5 gives descriptive statistics for both outcomes of interest.

1. The expression $\delta^g_{i,t}$ denotes the share that participant i allocates to the certain-price buyer in round t of game g. Here $g \in \{1, 2, 3\}$ and $t \in \{1, 2, \ldots, 20\}$.

2. The expression $d_{i,t}$ denotes the share that participant i allocates to the certain-price buyer in round t. Here $t \in \{1, 2, \ldots, 60\}$.

We estimate the following equation for each game:

$$\delta^g_{i,t} = \alpha^g_i + \sum_{s\in\{45,55\}} \beta^{pg}_s I[p^{pg}_{i,t} = s] + \sum_{h\in\{2,6,8\}} \beta^{qg}_h I[q^g_{i,t} = h] + \lambda^g t + \epsilon^g_{i,t} \tag{2.2}$$

To allow for non-linear effect of variation in the mean price offered by the uncertain-price buyer and the harvest quantity, we code both variables using dummy variables. First, we code the mean of the price offered by the uncertain-price buyer with the dummy variables $I[p^{pg}_{i,t} = s]$. The coefficients β^{pg}_s compare two alternative prices to a reference price of MXN 50 per kilogram. In the first case, the mean price of MXN 45 per kilogram is below the reference price, and in the second case the mean price of MXN 55 per kilogram is above the reference price. Recall that the reference price is the same as the price that is always offered by the certain-price buyer.

Similarly, we code the participant's harvest quantity with the dummy variables $I[q^g_{i,t} = h]$. The coefficients β^{qg}_h compare three alternative harvests to a reference harvest of 4 quintals. Recall that 1 quintal is 60 kilograms. In the first case, the harvest quantity is half the size of the reference harvest quantity (2 quintals), and in the second and third cases it is 50% larger (6 quintals) or double (8 quintals) the size of the reference harvest quantity. We use 4 quintals (240 kilograms) as a reference harvest quantity because this possibility is closest to the typical quantity of participants' harvests in real life. The exit survey indicates that the mean coffee harvest quantity of the sample is 371 kilograms and the median coffee harvest quantity is 270 kilograms.

In both this participant-game-round estimating equation and in the participant-round estimating equation below, we include a linear time trend to control for the effect of later rounds. The effect could be positive (participant learning) or negative (participant fatigue or boredom). Here, this time trend is denoted by λ^g. As we discuss in Section 2.5.2 below, we include participant fixed effects α_i^g to control for unobserved participant heterogeneity that does not vary by round. Following Boyd and Bellemare (2022), we cluster standard errors at the participant level to allow for correlation among unobservables within rounds played by the same participant.

Estimation at the Participant-Round Level

Next, we augment Equation 2.2 with additional dummy variables for Game 2 and Game 3 denoted by $I[g_{i,t} = c]$. The new equation appears as Equation 2.3 below. Recall that Game 2 and Game 3 vary the framing of the certain-price buyer. Section 2.3 gives more detail. The new coefficients c_g capture the effect of the variation in framing.

The remaining coefficients use Latin letters to refer to the same parameters denoted by Greek letters in Equation 2.2. The coefficients b_s^p capture variation in the mean price offered by the uncertain-price buyer, and the coefficients b_h^q capture variation in harvest quantity. As before, we include participant fixed effects a_i and a linear time trend l_1. We pool participant results across all three games and estimate this equation on the pooled sample.

$$\begin{aligned} d_{i,t} = a_i &+ \sum_{s \in \{45,55\}} b_s^p I[p_{i,t}^p = s] + \sum_{h \in \{2,6,8\}} b_h^q I[q_{i,t} = h] \\ &+ \sum_{g \in \{2,3\}} c_g I[g_{i,t} = c] + lt + e_{i,t} \end{aligned} \tag{2.3}$$

Estimation at the Participant Level

Finally, we consider the allocation decisions of participants in all three games. As Section 2.4.5 describes, we average participants' allocation to the certain-price buyer across all 60 rounds, which we denote by d_i below. Wollni and Fischer (2015) use a similar outcome of interest: the fraction of coffee harvest sold to one buyer. They note that this dependent variable is a fractional variable bounded between 0 and 1. For this reason, they use the quasi-likelihood estimator proposed by

Papke and Wooldridge (1996).

We do not follow their approach. Instead, we estimate equation 2.4 separately for the overall margin, the extensive margin, and the intensive margin. This method resembles the double-hurdle model used by Shumeta et al. (2018) with the added benefit that the point estimates are directly interpretable.

$$d_i = \theta_1 \text{extra}_i + \beta_1 X_i + \epsilon_{1i} \tag{2.4}$$

The coefficient of interest is θ_1, the effect of the additional income on these three outcomes. In addition, as controls, we include the following covariates: age, gender, education level, CRRA calculated based on the Eckel-Grossman lottery, completion of the practice game, reported literacy, correct answer on the probability filter question, game order, and lottery position. Since the unit of analysis is the participant and the treatment is at the participant level, we do not cluster the standard errors. We simply compute heteroskedasticity-robust standard errors.

We use an augmented version of Equation 2.4 to examine the effect of the additional income treatment moderated by two characteristics of the participants. First, we estimate the effect of the treatment moderated by CRRA for the full sample as well as for subsamples of cooperative members and non-members. Second, we estimate the effect of the additional income treatment moderated by cooperative loyalty as measured by the number of years that the participant sold to the cooperative for a subsample of cooperative members. In both cases, the covariate Z_i denotes the moderator.

$$d_i = \theta_2 \text{extra}_i + \gamma Z_i + \tau \text{extra}_i Z_i + \beta_2 X_i + \epsilon_{2i} \tag{2.5}$$

Here, there are three coefficients of interest. First, the coefficient θ_2 captures the overall effect of the additional income treatment. Second, the coefficient γ captures the effect of the moderator. Third, the coefficient τ captures the additional treatment effect of a one-unit increase in CRRA or a further year of loyalty to the cooperative. Once again β_2 captures the effect of the vector of controls.

2.5.2 Identification Strategy

First, we consider identification for the estimations at the participant-game-round level and the participant-round level. Within each game, at the round level, we randomize the harvest quantity and the mean of the price offered by the uncertain-price buyer, so the corresponding coefficients in Equations 2.2 and 2.3 are causally identified. Across the three games, the order is randomized and participants play all three games, so we argue that the additional coefficients for Game 2 and Game 3 in Equation 2.3 are also causally identified.

Two concerns remain for causal identification. First, we consider the potential correlation between the share allocated to the certain-price buyer in each round and unobservable characteristics at the participant level such as risk preference or skill at playing the game. We use participant fixed effects to control for these unobservable characteristics. Second, participants' allocation decisions in earlier rounds and later rounds might differ in unobservable ways, due to participant learning or fatigue. For this reason, all participants play ten rounds of a practice game that are not counted, either in their overall score in the game or in our analysis. The practice game controls for participants who learn the game faster than others. Moreover, we include a linear-time trend to control for additional learning, boredom, or fatigue.[12]

Next, we consider identification for the estimation at the participant level in Equation 2.4 and Equation 2.5. Here, the additional income treatment is randomized at the participant level, so the parameters θ_1 and θ_2 are causally identified. Moreover, CRRA and participant loyalty are considered exogenously fixed prior to the experiment, so we argue that the parameters γ for the direct effect and τ for the interaction effect of these moderators are also causally identified.

2.5.3 Subgroup Analysis

We would like to estimate the effect of the four factors above separately for cooperative members and non-members to uncover potentially heterogeneous treatment effects. Recall that 126 of our 268 participants are cooperative members. Cooperative membership is a time-invariant participant characteristic, so we cannot include a membership dummy in Equation 2.2 or 2.3 because it would

[12] At the request of an anonymous reviewer, we estimated two alternative specifications of Equation 2.3: one that omitted this time trend and another that replaced it with round fixed effects. The results of these alternative specifications were nearly identical to the results of our preferred specification.

be absorbed in the participant fixed effects. Moreover, it is a choice variable based on observed and unobserved characteristics, so we cannot add it to the vector of controls X in 2.4.

For this reason, we use subgroup analysis. We estimate Equations 2.3 and 2.4 separately for cooperative members and non-members to allow for a comparison of the estimated parameters. We argue that the parameters in these estimated results are causally identified for the reasons we discuss in the previous section. One drawback to this approach is the reduced sample size in the subsamples of 126 members and 142 non-members compared to the full sample of 268 members. This reduced sample size limits the statistical power of the associated hypothesis tests.

2.6 Results and Discussion

In this section, we first present estimation results at the participant-game-round level (Equation 2.2) and the participant-round level (Equation 2.3). Next, we present estimation results at the participant level on the full sample (Equation 2.4 and Equation 2.5). Finally, we present estimation results at the participant level on the subsamples of cooperative members and non-members.

2.6.1 Participant-Game-Round Level Results

Table 2.6 presents the results of the estimation of Equation 2.2 at the participant-game-round level. Recall from Table 2.5 that the baseline allocations to the certain-price buyer for Game 1, Game 2, and Game 3 are 82%, 83%, and 82% respectively. The strong preference of the participants for price certainty stands out as the most important result at the participant-game-round level and the participant-round level. These allocations reveal an 18% incidence of side-selling. This estimate is higher than the 12% incidence of side-selling reported by Keenan et al. (2024), Woldie (2010), and Wollni and Fischer (2015) and close to the 20% incidence of side-selling reported by Ewusi Koomson et al. (2022). It is lower than the estimates of the other studies we cite in the Introduction, which range from 30% to 55%.

Moreover, this high baseline provides context to the point estimates below. Our point estimates of the effect of varying the harvest quantity, varying the mean of the price offered by the uncertain-price buyer, and varying the framing of the certain-price buyer range between 1% and 4%. These effect sizes may seem small, but we argue that they are still important relative to the overall incidence of 18% of side-selling.

We first examine these effect of varying the harvest quantity. Reducing the harvest quantity by half from the reference of 4 quintals to 2 quintals increases the incidence of side-selling by 3%. Increasing it by 50% from the reference of 4 to 6 quintals does not affect side-selling. Doubling it from 4 to 8 quintals, however, increases the incidence of side-selling, this time by 2%. These point estimates of the effect of varying harvest quantity on the incidence of side-selling are comparable in magnitude to the effect sizes of Keenan et al. (2024), which range from 1% to 7%. Interestingly, their effect sizes are negative, while ours are positive. On the other hand, our effect sizes are the same sign and approximately the same magnitude that Wollni and Fischer (2015) find, although their use of nonlinear econometric methods makes a direct comparison of point estimates difficult.

To shed light on this puzzle, we use the general framework of Fafchamps and Hill (2005), who examine the distinction between selling at the farmgate and going to the market. These authors suggest that producers only travel to market when they have a sufficient quantity to justify the fixed cost of the trip. In other words, producers with a medium harvest tend to travel more to the market than producers with a small harvest. However, producers with large harvest do not travel to the market as frequently because the opportunity cost of time for them is too high. In our study context, local traders come to the farmgate while the cooperative recollection points are at a distance. Thus, in our study, participants will only deliver their harvest to the cooperative if they have enough to justify the trip, but not so much that the opportunity cost of time is too high. The context of Wollni and Fischer (2015) is the same and, for this reason, the sign of their results matches that of ours. In contrast, in the context of Keenan et al. (2024), the cooperative is near and the local traders are far away, so their results have the opposite sign. Producers with a small harvest do not side-sell as much because they cannot justify the fixed cost of the trip to the local traders. Producers with a large harvest have better things to do with their time.

Next, we examine the effect of varying the mean of the price offered by the uncertain-price buyer, which is a proxy for a change in market conditions or a change in transaction costs. We see that a MXN 5 reduction is associated with a 2% increase in side-selling. This result does not match profit-maximizing behavior, and we cannot find an easy explanation for it. Our hypothesis is that this result reflects a characteristic of the local context. Perhaps hearing about a reduction in the price of the local trader causes producers to think that the price will rise in the future. A MXN 5 increase in the mean of the price offered by the uncertain-price buyer does not affect the allocation decision.

When we compare the estimation results across the three games (columns 1, 2, and 3), we do not find much difference. The baseline allocations for all three games are very close. Moreover, so are the coefficients for the variation in harvest size and mean of the price offered by the uncertain-price buyer. These similarities suggest that varying the framing of the certain-price buyer does not make a difference in the allocation decision. These results contrast with those of Mujawamariya et al. (2013), which studies side-selling in a context where some local traders offer credit and others do not, so the provision of credit by some traders induces producers to market their production through these traders. Similarly, Ewusi Koomson et al. (2022) find that access to extension services provided by the cooperative (credit and technical assistance) reduces the incidence of side-selling.

2.6.2 Participant-Round Level Results

Table 2.7 presents results from estimating Equation 2.3, a specification that augments Equation 2.2 with dummies for Game 2 and Game 3, on a pooled sample that combines participant allocation decisions across all three games. The point estimates here do not differ meaningfully from those in the previous specification. The framing of the certain-price buyer in Game 2 (microcredit) appears not to affect the allocation decision. The framing of the certain-price buyer in Game 3 (cooperative with microcredit and technical assistance) causes participants to allocate 1% less coffee to the certain-price buyer, but the point estimate lacks statistical significance. As we mentioned above, these results differ from those of Ewusi Koomson et al. (2022) and Mujawamariya et al. (2013).

2.6.3 Participant-Level Results

In this section, we first examine the baseline direct effect of the additional income treatment on the average allocation of participants in all rounds of the game. Then we examine the indirect effect of risk aversion on this average allocation.

Table 2.8 presents the results of the estimation of Equation 2.4 at the participant level. Recall from Table 2.1 that 58 of 268 participants allocate the entire harvest to the certain-price buyer in every round. Thus we separate the overall margin into the extensive and intensive margin.

At the extensive margin, the presence of additional income increases the likelihood by 10.0% that a participant will not side-sell to the cooperative at all. This result differs from that of Keenan et al. (2024), who find that non-farm income only reduces side-selling within the same producer (variation in non-farm income over the three-year panel), but not between producers. Moreover, our treatment effect of 10.0% is much higher than theirs of 1.5%. In contrast, it matches that of Shumeta et al. (2018), who find a larger effect of off-farm income at the extensive margin than at the extensive margin. In their sample, 49 of 190 Ethiopian coffee farmers are completely loyal to the cooperative. 67% of the loyal farmers have off-farm income, while only 23% of the side-selling farmers do. We improve on their results by randomizing the presence of off-farm income. Our result also matches that of Geng et al. (2023), who find that an unexpected health shock (which they use as a proxy for an income shock) in a given week decreases the share of milk delivered to a dairy cooperative in the same or subsequent week by 2.5%.

The presence of additional income does not affect side-selling behavior at the overall margin. The effect at the intensive margin is also small (-1.9%) and not statistically significant.

Three covariates are associated with allocation decisions at the participant level: completed only middle school, understanding probability, and completing the practice game. All increase side-selling behavior. We present these as associations that warrant further study. Wollni and Fischer (2015) and Keenan et al. (2024) also find an association between an increase in the education level of producers and side-selling behavior.

Next, we turn to the moderating effect of risk aversion on side-selling. Table 2.9 presents the results of the estimation of Equation 2.5 at the participant level. Recall that Equation 2.5 augments Equation 2.4 with an interaction term of the participants' CRRA as measured by the Eckel-Grossman lottery. Half of participants completed the lottery before the experiment and half after the experiment. As Figure 2.10 shows, we find the same distribution of lottery choice for both groups, so we argue that lottery placement does not affect lottery choice.

In general, we find that increased risk aversion decreases side-selling, consistent with Woldie (2010). Interpreting the results, we find a baseline effect of the additional income of 6.7% at the extensive margin that increases by 7.5% with each one-unit increase in the CRRA. Table 2.3 shows the estimated CRRA range given by each lottery choice. The treatment effect for lottery choice 1 (CRRA = 2) is 13.5%. The treatment effect for lottery choice 5 (CRRA = 0.2) is 7.4%. These results imply that additional income reduces side-selling more for more risk-averse participants. They match those of Boyd and Bellemare (2022) and Bellemare et al. (2020), who both used estimated participant risk preferences from risk-elicitation lotteries and found differential effects of the provision of crop insurance. However, as in these two prior studies, our results here also lack statistical significance.

2.6.4 Subgroup Analysis by Cooperative Membership

Finally, we estimate the round-level outcomes and the participant-level outcomes separately for cooperative members and non-members. Recall that Figure 2.13 shows a histogram of the participant level outcomes broken down by cooperative membership status. Throughout this section, the smaller sample size (126 members and 142 non-members) of the two subgroups limits the statistical power of the hypothesis tests. However, we argue that the differences in the point estimates warrant the analysis.

Participant-Round Level Results

Table 2.10 presents estimation results at the participant-round level separately for cooperative members and non-members. We see differential effects for changes in harvest quantity and uncertain-price buyer between members and non-members. We consider first the case of an 8 quintal harvest relative to the reference harvest of 4 quintals. For non-members, the point estimates that we saw in the overall sample double (3.2% vs 1.7%). For members, doubling the harvest size does not affect the allocation decision. This difference indicates that non-members value profit maximization more than price certainty.

Next, we consider the case of a 2 quintal harvest relative to the reference harvest of 4 quintals. Recall the overall effect of 3.0% from Table 2.7. In the subgroup analysis, members side-sell 2.5% more of their harvest and non-members side-sell 3.6% respectively. In contrast to the situation above, here both groups choose profit maximization over loyalty to the cooperative. In interpreting these coefficients, we note that participants in the experiment only have three choices to allocate their harvest: 2 quintals, 1 quintal, or 0 quintal to the certain-price buyer. Thus, instead of an average increase in side-selling of 2.5%, a better interpretation could be that 1 in 50 participants changed their allocation decision.[13]

Unlike in the pooled results, we find an effect of the experiment framing here. Access to microcredit decreases side-selling by 1.3%, indicating that cooperative members value this service. In the same vein as above, a better interpretation might be that approximately 1 in 100 cooperative member participants change their behavior when reminded of access to microcredit. In contrast, when the certain-price buyer is described to non-members as a cooperative, the framing reduces their allocation to the certain-price buyer by 3.2% (or 1 in 33). This result possibly indicates a dislike for cooperatives.

Participant Level Results

Tables 2.11 and 2.12 present results for participant-level outcomes on subgroups of cooperative members and non-members respectively. The smaller sample size (126 members and 142 non-members) limits the statistical power of the hypothesis tests. Nevertheless, we see a sharp contract in the point estimate of the additional income treatment at the extensive margin. For cooperative members, it is 16.3%, while for non-members it is 2.5%. This difference suggests that the additional income may relieve a budget constraint that allows cooperative members who already prefer price certainty to pursue it even more.

Tables 2.13 and 2.14 present the results for the moderating effect of risk-aversion on the treatment effect of the additional income for cooperative members and non-members, respectively. Recall that Figure 2.9 shows a breakdown of lottery choice by members and non-members. Members are slightly more risk-averse than non-members.

In these estimation results, the baseline treatment effect of additional income at the extensive margin for cooperative members is comparable to the baseline treatment effect in Table 2.9 (4% vs. 6% reduction in likelihood of side-selling). However, we find that the differential effect by unit of CRRA is double for members compared to the overall sample (18.2% vs 9.2% reduction in

[13]Thanks to David Rosenkranz for pointing this out

likelihood of side-selling). Moreover, when we examine non-members, we find a treatment effect in the opposite direction. A one unit increase in the CRRA increases the likelihood of side-selling by 11%. The opposite signs of these treatment effects in the two subgroups may indicate different underlying preferences at work. Cooperative members would like to remain loyal to the cooperative except when they are liquidity constrained and sell to the local trader by necessity. Non-members would like to maximize their profit and sell to the local trader except when they are liquidity constrained and sell to the certain-price buyer out of necessity.

Finally, we use administrative data from the cooperative to examine the moderating effect of member loyalty on the treatment effects of additional income. We measure member loyalty as the number of years in 2013-24 that a participant who is a member has sold anything to the cooperative. Figure 2.11 shows the distribution of member loyalty. We estimate this effect using Equation 2.5, which incorporates loyalty as a moderator. Table 2.15 shows these results. At the baseline, we find that additional income is associated with a 49% decrease in side-selling by a hypothetical new member (loyalty of 0) This association decreases by 4% per year. At the mean value of loyalty (9.3 years), it is 12%. These results suggest that the larger number of marketing years a member sells to the cooperative, the less a liquidity constraint affects the decision to side-sell.

2.6.5 Limitations

This experiment is the first that we know of that examines the determinants of side-selling. It suffers from at three limitations. First, we designed the state space of the experiment to correspond to the number of rounds (60), so that all participants would face all possible scenarios over the course of the three games. New technology in adaptive experiments would allow us to expand the state space.[14] For example, we could test more than three possibilities for the mean price of the uncertain-price buyer, more than four possibilities for harvest quantity, or more than one amount of additional income. With a larger state space, we could adapt the possibilities that participants are presented with in subsequent rounds based on participant performance in the initial rounds.

Second, the framing of the certain-price buyer was done verbally, while the other randomization was performed physically: small coffee bags for the coffee harvest, a die for the price of the

[14]For example, the Bayesian adaptive choice experiment software developed by Drake et al. (2024).

uncertain-price buyer, and play money for the additional income.[15] This indigenous population may understand tactile variation better than verbal variation. In addition, the services offered by the framed buyers (microcredit and technical assistance) did not affect the results in the game. In real life, microcredit would smooth consumption and technical assistance would affect harvest quantity. Subsequent experiments could add this functionality using a mobile phone or tablet instead of tactile elements.

Finally, the allocation decisions of individual participants did not affect the outcomes of other participants. In real life, a cooperative survives or fails on the basis of the joint decision of its members. Hopfensitz and Miquel-Florensa (2017) provides an example of an experiment in which cooperative member behavior varies depending on the behavior of non-members and the presence of a punishment mechanism for side-selling. Their work provides examples of elements that we could incorporate into a future experiment as well.

2.7 Conclusion

In the past 30 years, many developing countries have shifted the way they support rural communities from a state-led approach to a market-led approach. As a result, agricultural cooperatives have emerged that offer many of the same services to their members in the present as state commodity boards in the past: a guaranteed purchase price, microcredit, and technical assistance. The big difference from state commodity boards is that agricultural cooperatives depend on the continued patronage of their producer members to finance their services. Weak institutions often prevent them from enforcing this condition. Moreover, many of the services like microcredit and technical assistance help members over the long-run, but because of liquidity constraints, members often seek to maximize profit over the short-run. Thus, side-selling threatens cooperatives' ability to offer these services, and understanding the drivers of side-selling behavior is imperative for their continued existence.

We have presented the results of a framed field experiment that examines four possible determinants of side-selling behavior for indigenous coffee farmers in Mexico. . The experiment abstracts

[15] Enumerators read from a standardized script.

the most important decision of many smallholder producers for their household economy: how and to whom they market their cash crops. In our experiment, participants can market as much as their harvest as they like to each of a certain-price and an uncertain-price buyer. Unlike many previous studies, our experiment does not employ the distinction between the delayed payment of a cooperative and the immediate payment of a local trader. We also do not restrict participants' options in subsequent rounds based on their performance in the present round.

Our results extend beyond coffee and beyond Mexico. They provide several concrete policy recommendations to cooperatives to reduce the incidence of side-selling among their members. First, we find an overall lower incidence of side-selling (18%) than in many contexts, which confirms the preference of smallholder producers for price certainty. Since eliminating delayed payments reduces the incidence of side-selling, we encourage cooperatives to find upstream financing so that they can pay their members at the moment of delivery just like local traders.

Second, the incidence of side-selling is affected slightly by harvest size. This effect is consistent with the distinction between selling at the farmgate or at the market originally proposed by Fafchamps and Hill (2005). It means that cooperatives must be attentive to the fixed costs associated with the market decisions of and reduce or eliminate these fixed costs through the use of regional collection points or even visits to the farmgate.

Third, access to credit and technical assistance does not affect producer behavior in the short term. However, in the medium term, access to microcredit can help producers weather unexpected shocks. Moreover, in the long term, technical assistance has the potential to dramatically improve producer yields. Liverpool-Tasie et al. (2020) point out that in a situation without formal contracts, cooperatives or producers may need subsidies to realize these long-term benefits.

Fourth, our additional income treatment confirmed the effectiveness of direct subsidies to producers. In the Mexican context, our subsidies are not infeasible; they are of the same magnitude as the conditional cash transfer programs of the past and present. The moderated treatment effects that we find suggest that these subsidies would be especially effective in ensuring the loyalty of cooperative members in the early years of their membership.

Finally, cooperatives need to find mechanisms to enforce sanctions on members who do not market their harvest through the cooperative. Michler and Wu (2020) provides a framework for relational contracts in contexts without formal contract enforcement. Casaburi and Macchiavello (2015) suggest that the mere threat of sanctions may be as effective as the sanctions themselves.

Governments and non-governmental organizations alike implemented market-based reforms with great enthusiasm and promise. Several decades later, they still face challenges in realizing their potential in improving the welfare of smallholder producers. The results we present here suggest a few incremental improvements to improve their effectiveness and long-term sustainability.

2.8 Exhibits

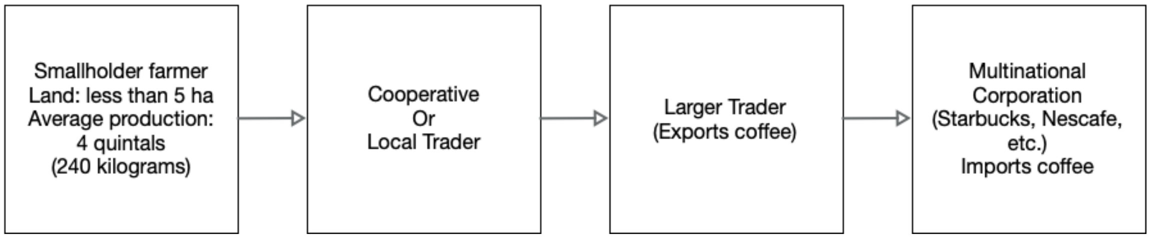

Figure 2.1: Coffee Value Chain

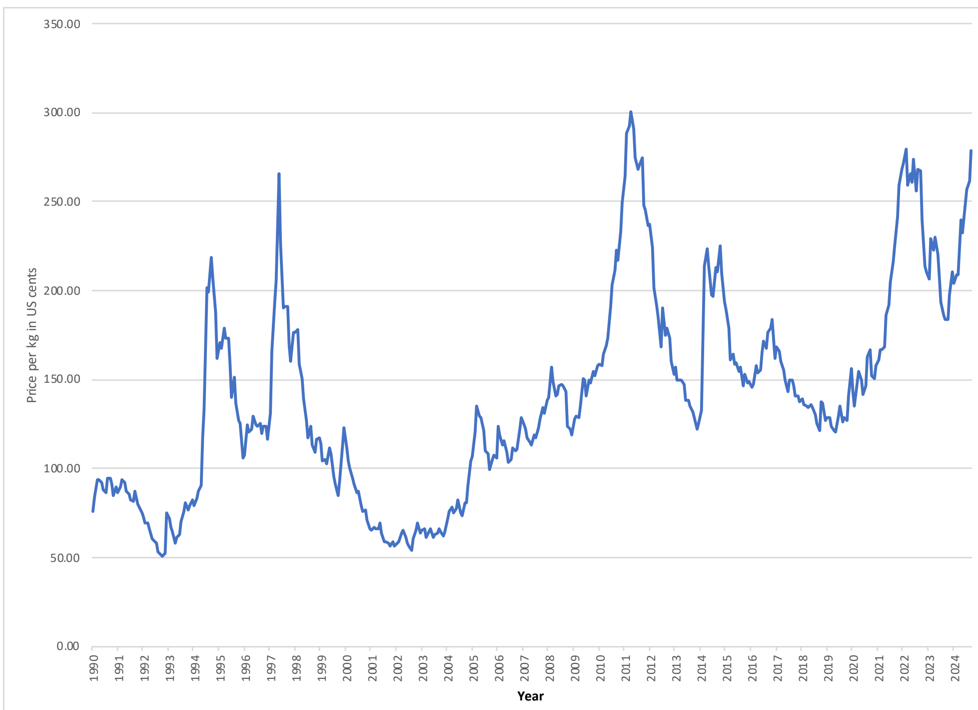

Figure 2.2: World Price of Arabica Coffee
Source: International Monetary Fund, Global price of Coffee, Other Mild Arabica [PCOFFOTMUSDM], etrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/PCOFFOTMUSDM, May 28, 2025.

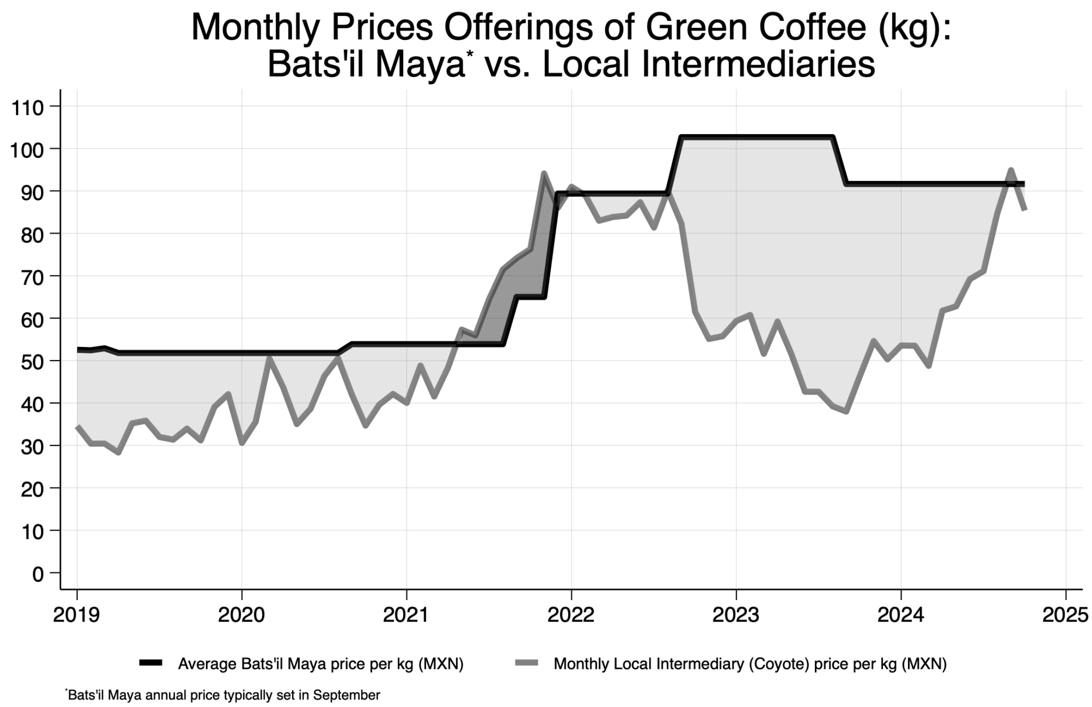

Figure 2.3: Coffee Cooperative vs Local Trader Price (2019-2024)
Source: Administrative Data from Ts'umbal Xitalha' Coffee Cooperative

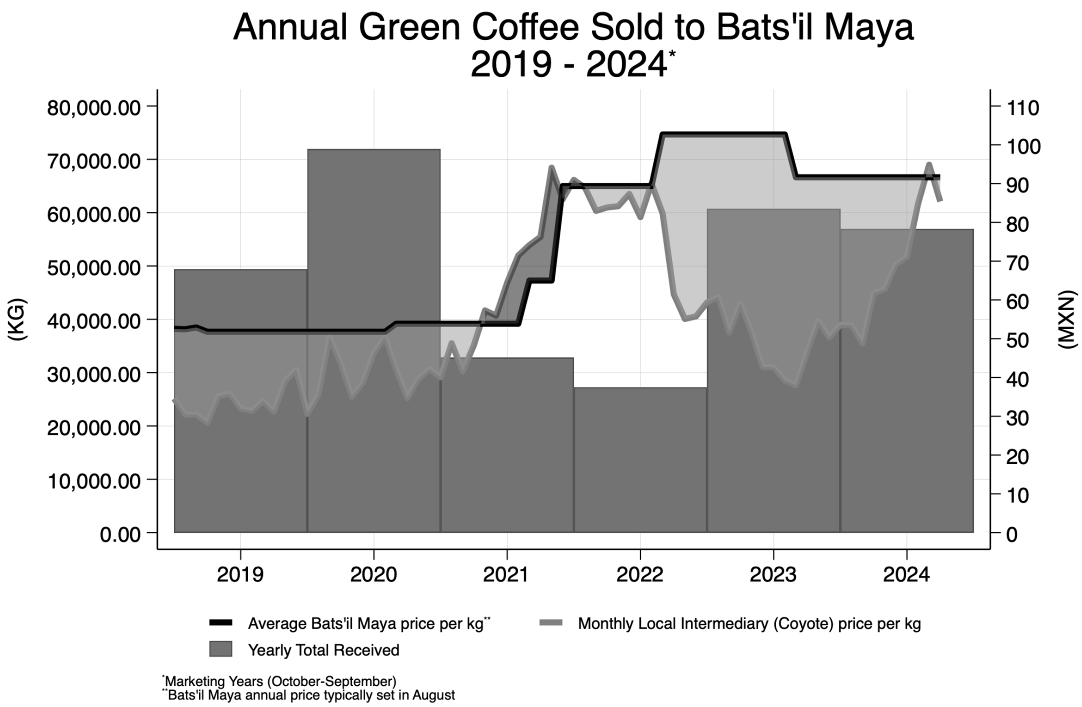

Figure 2.4: Coffee Deliveries and Market Prices (2019-2024)
Source: Administrative Data from Ts'umbal Xitalha' Coffee Cooperative

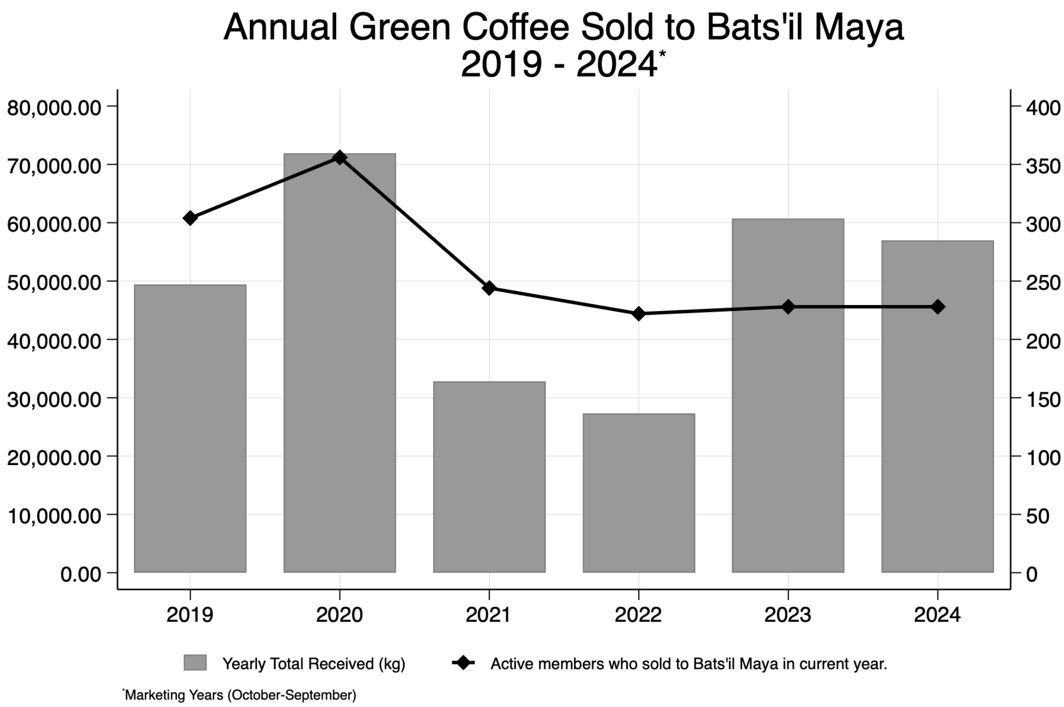

Figure 2.5: Coffee Deliveries and Members (2019-2024)
Source: Administrative Data from Ts'umbal Xitalha' Coffee Cooperative

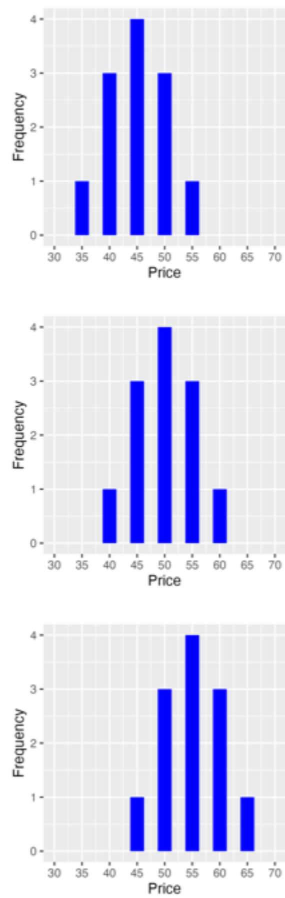

Figure 2.6: Uncertain-Price Buyer Distributions
This figure shows the three possible distributions of the price offered by the uncertain-price buyer. All three distributions are multinomial distributions with 5 support points that approximate a normal distribution. The three distributions have mean values of MXN 45, MXN 50, and MXN 55. In each distribution, the mean appears 4 times, support points MXN 5 above or below the mean appear 3 times, and support points MXN 10 above or below the mean appear once.

Quintals Sold to Certain Price Buyer (60kg)								
0	1	2	3	4	5	6	7	8
Quintals Sold to Uncertain Price Buyer (60kg)								
8	7	6	5	4	3	2	1	0

Total Revenue from Sales to Both Buyers

	Revenue from Sale to Certain Buyer: (Quantity sold to Certain Buyer x $50 MXN)								
	0	3,000	6,000	9,000	12,000	15,000	18,000	21,000	24,000
Price per kg. (MXN)	**Revenue from Sale to Uncertain Buyer: (Quantity sold to Uncertain Buyer x Dice Result)**								
40	19,200	16,800	14,400	12,000	9,600	7,200	4,800	2,400	0
45	21,600	18,900	16,200	13,500	10,800	8,100	5,400	2,700	0
50	24,000	21,000	18,000	15,000	12,000	9,000	6,000	3,000	0
55	26,400	23,100	19,800	16,500	13,200	9,900	6,600	3,300	0
60	28,800	25,200	21,600	18,000	14,400	10,800	7,200	3,600	0

Figure 2.7: Representative Payoff Table
Participants are shown one of 12 tables of this form each round to aid them in their decision of how much of their harvest to allocate to the certain-price buyer. The tables vary according to the harvest quantity and mean price offered by the uncerain-price buyer. Appendix A contains all 12 tables.

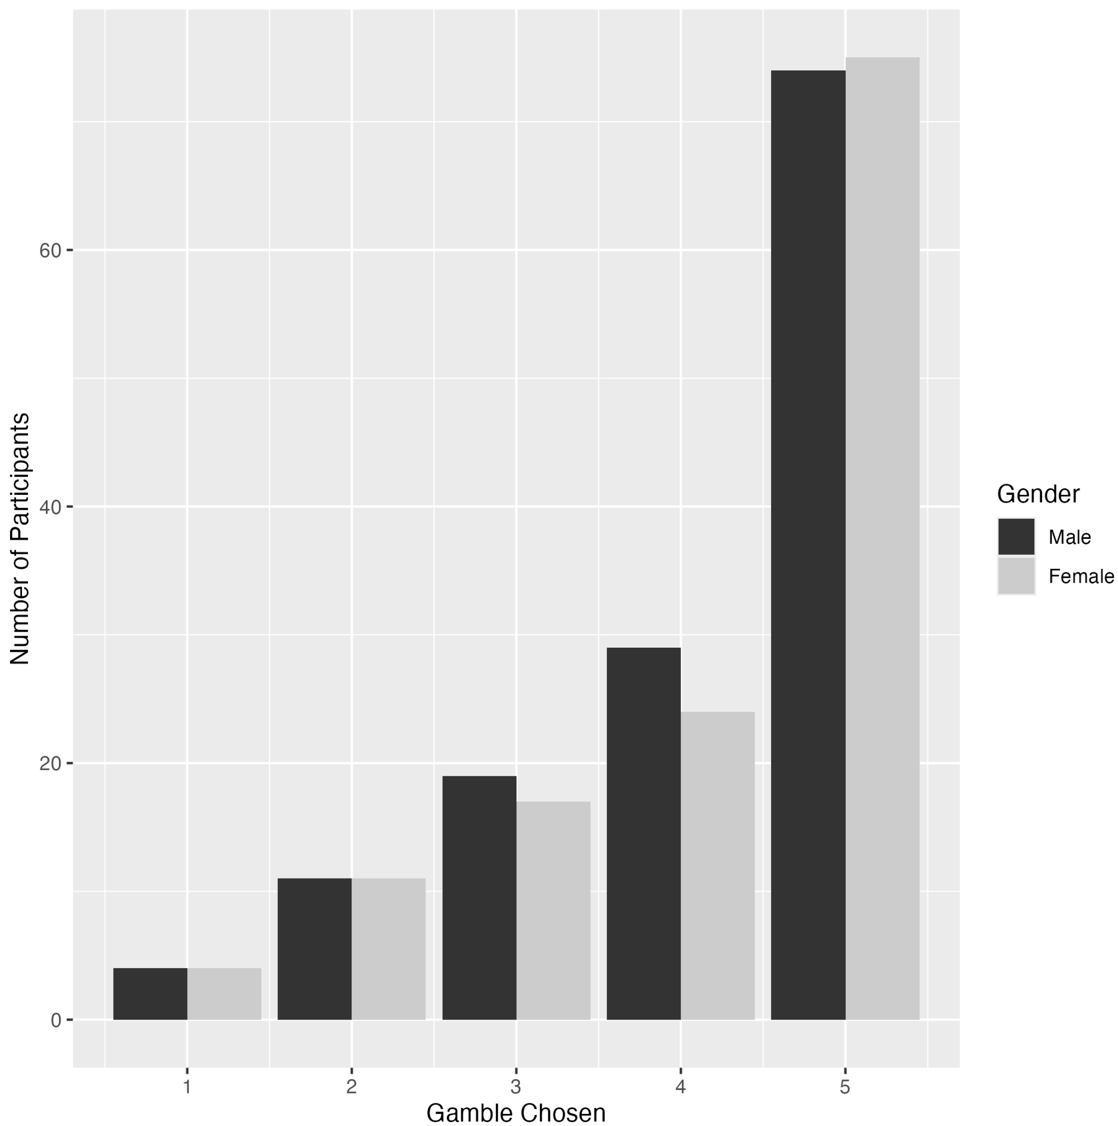

Figure 2.8: Lottery Gamble Choices by Gender
This figure displays a histogram of gamble choices from a no-loss lottery based on Eckel and Grossman (2008). Table 2.3 describes the choices. It is comparable to Figure 1 in that paper. Here we do not see differences between the gamble choices of men and women.

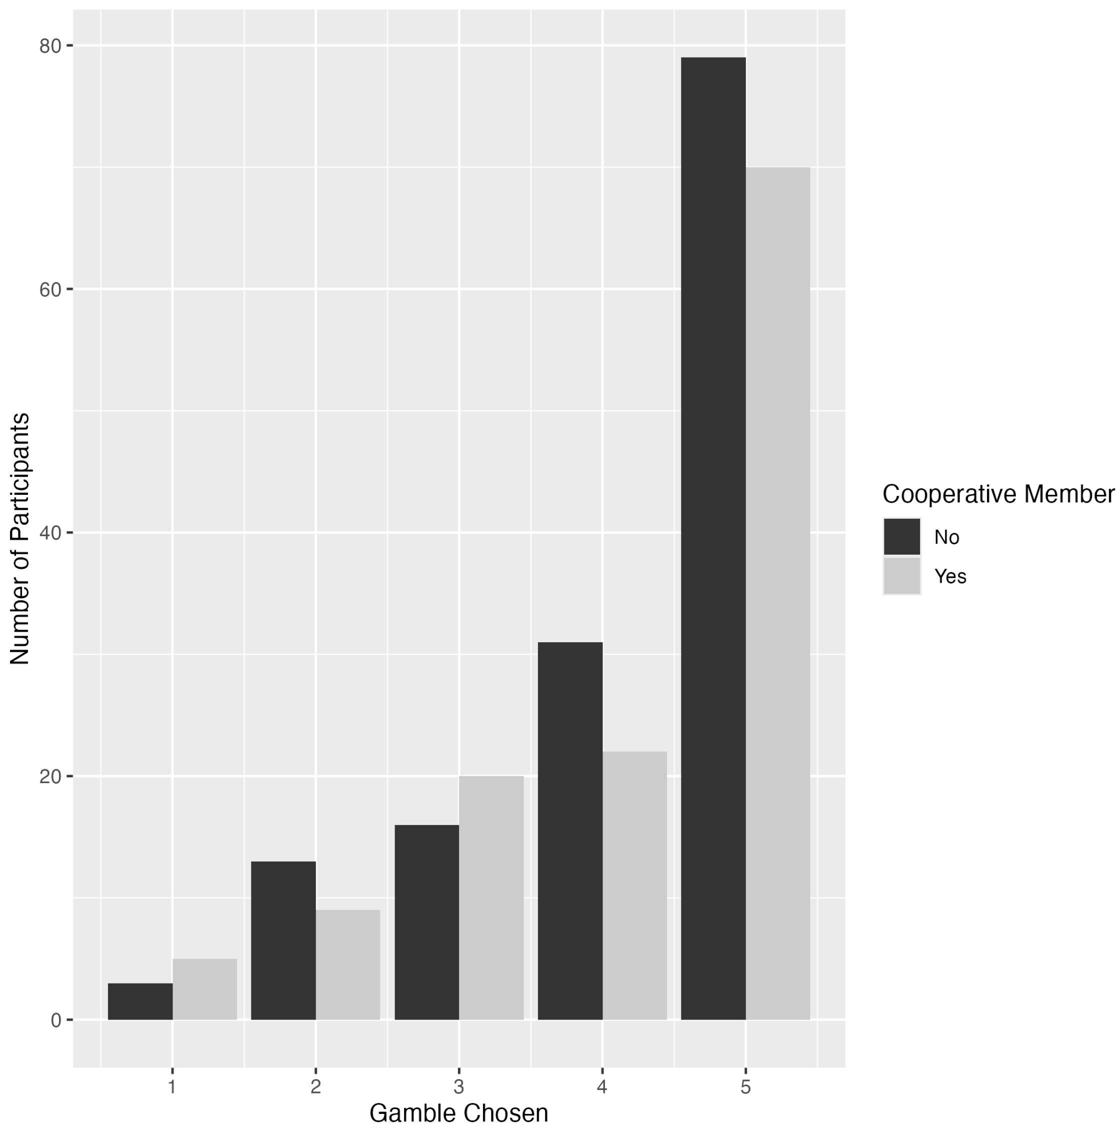

Figure 2.9: Lottery Gamble Choices by Cooperative Membership Status
This figure displays a histogram of gamble choices from a no-loss lottery based on Eckel and Grossman (2008). Table 2.3 describes the choices. It is broken down by cooperative membership status of the participants. Cooperative members are slightly more risk-averse than cooperative non-members.

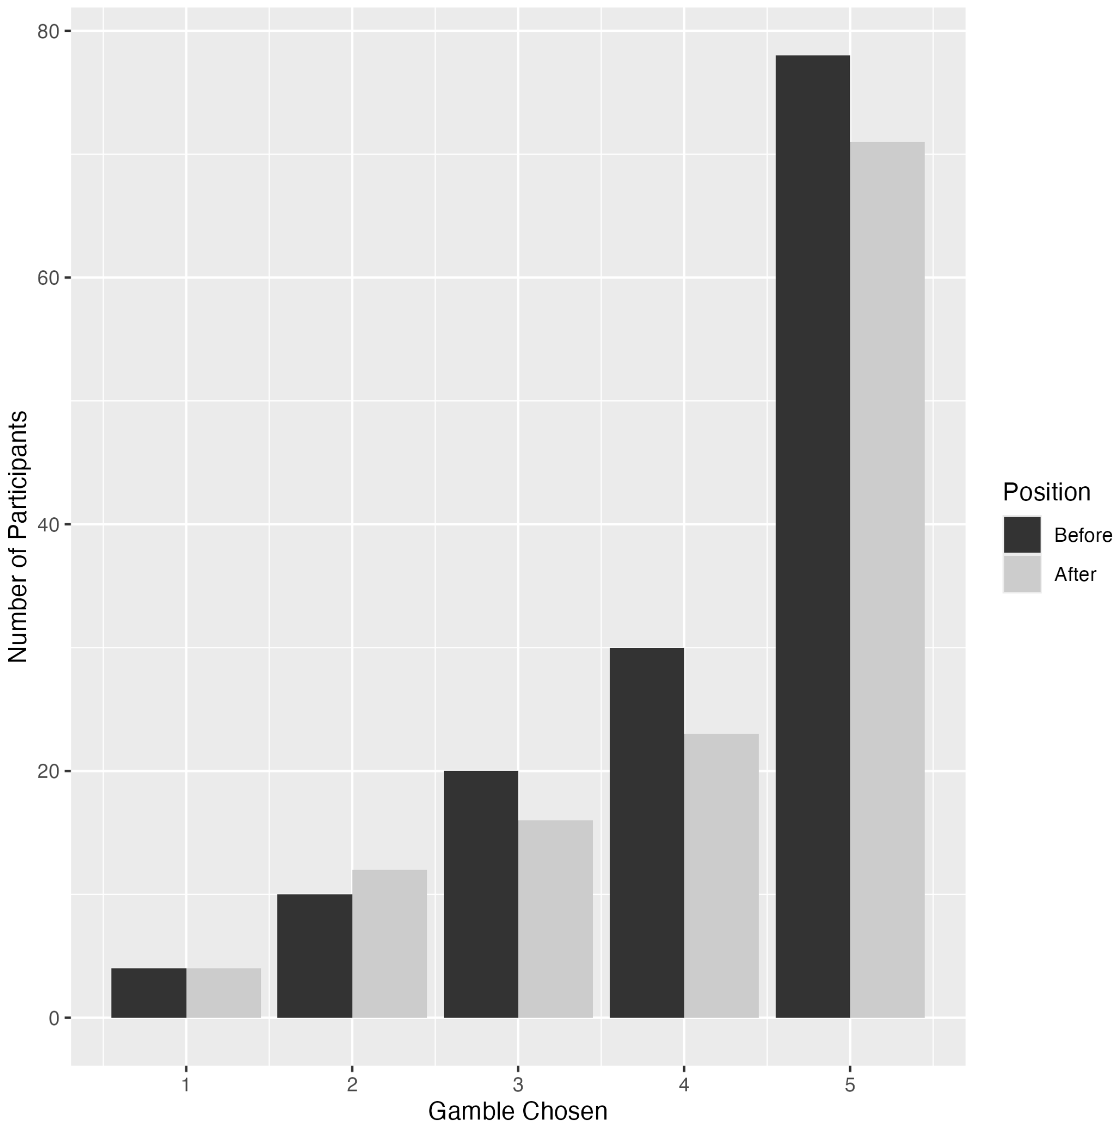

Figure 2.10: Lottery Gamble Choices by Position
This figure displays a histogram of gamble choices from a no-loss lottery based on Eckel and Grossman (2008). Table 2.3 describes the choices. It is broken down by whether participants completed the lottery before or after the game. It suggests that the distribution of responses is not associated with the lottery position.

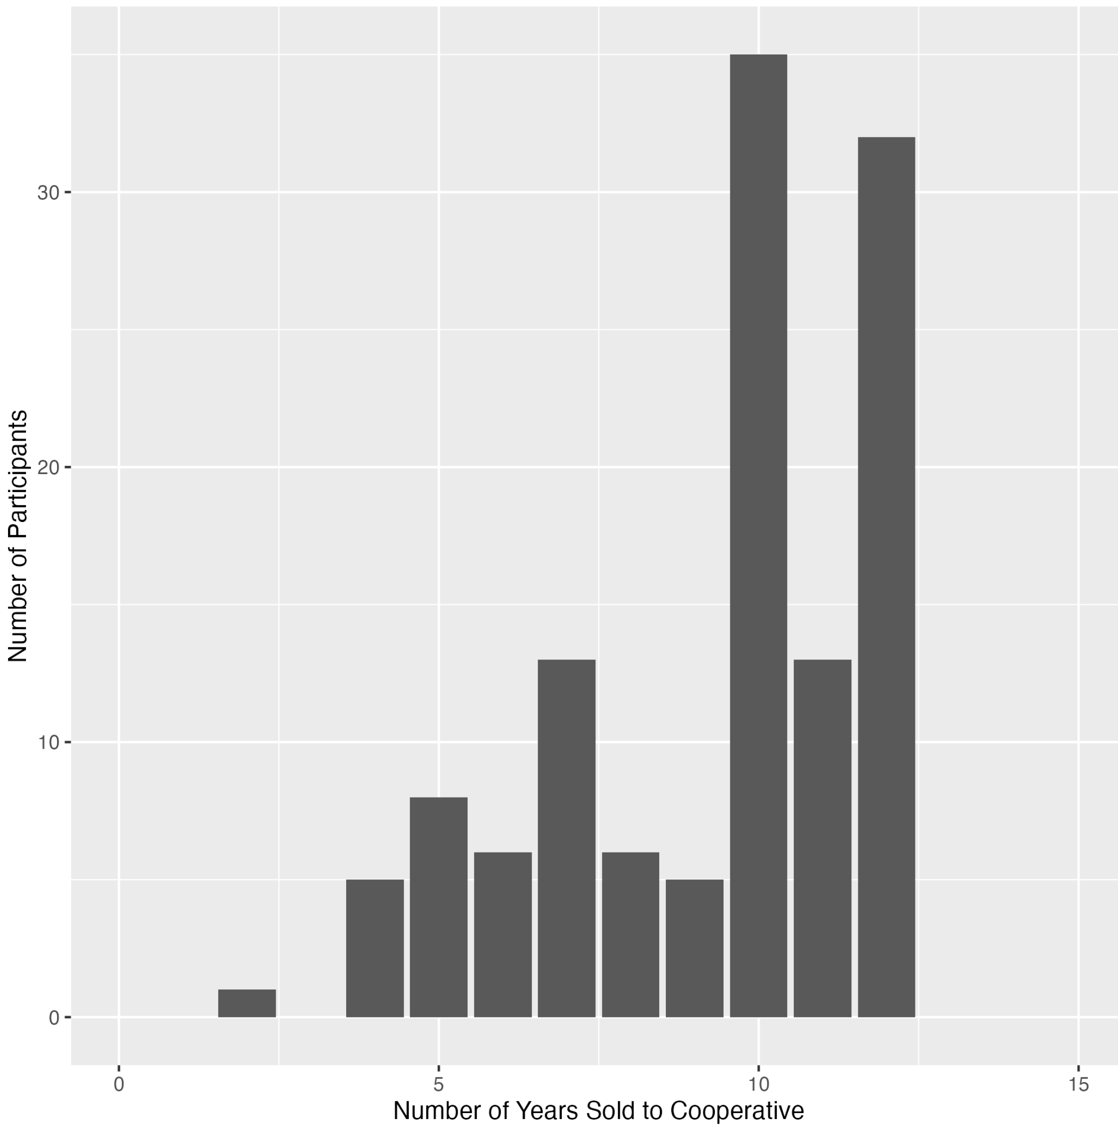

Figure 2.11: Cooperative Member Loyalty
This figure displays a histogram of the number of years in 2013-24 that participants who are cooperative members (n=124) delivered coffee to the cooperative. It is based on administrative data from the Ts'umbal Xitalha' coffee cooperative.

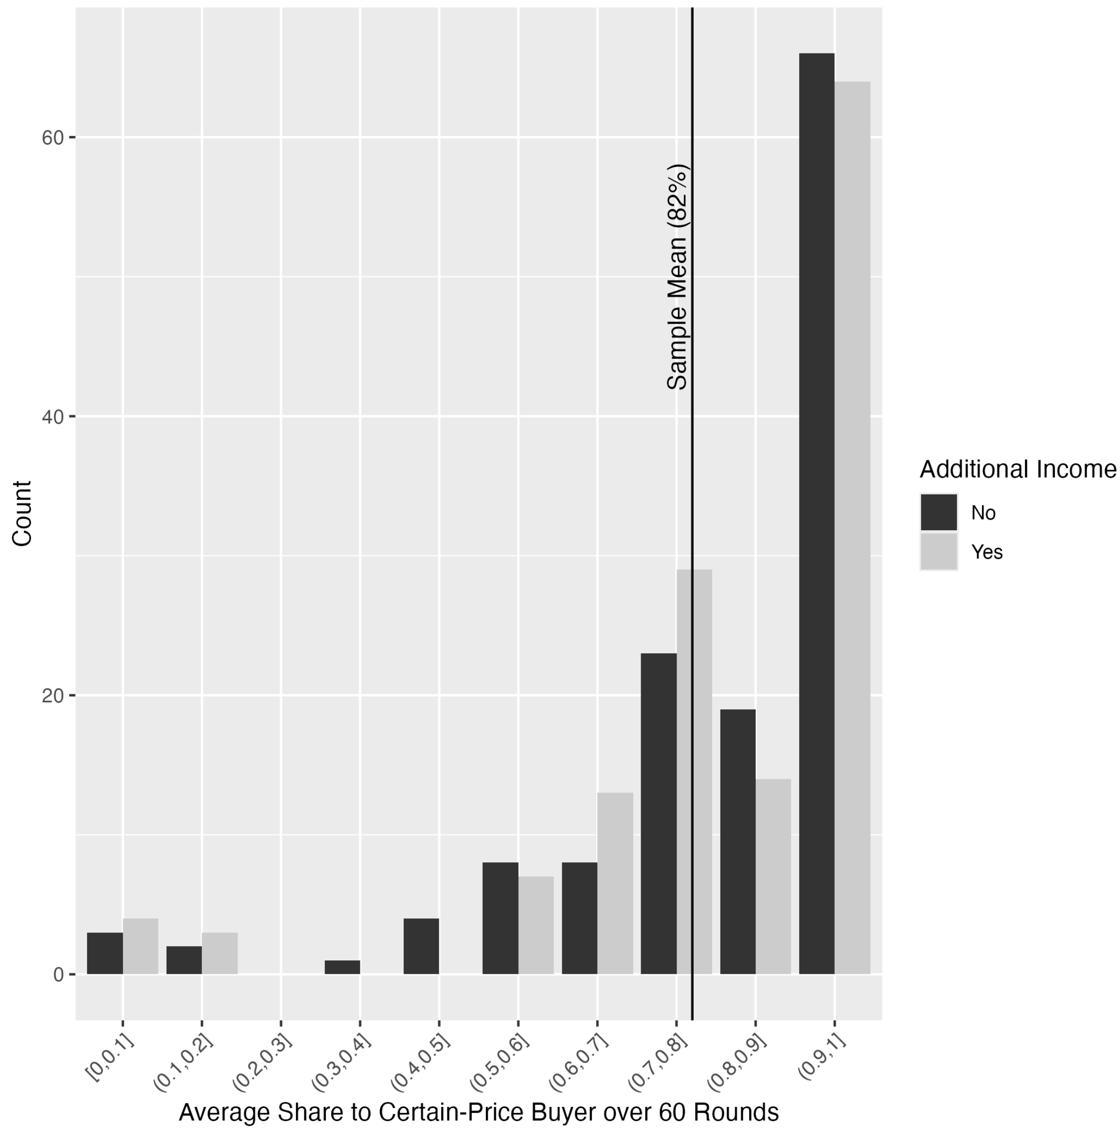

Figure 2.12: Total Margin by Treatment Status

This figure displays a histogram of the average share of harvest that participants allocate to the certain-price buyer over all 60 rounds of the experiment, broken down by treatment status. There are 268 total participants in the experiment. 134 receive the treatment, which is MXN 3,000 of additional income in every round.

Figure 2.13: Total Margin by Cooperative Membership Status

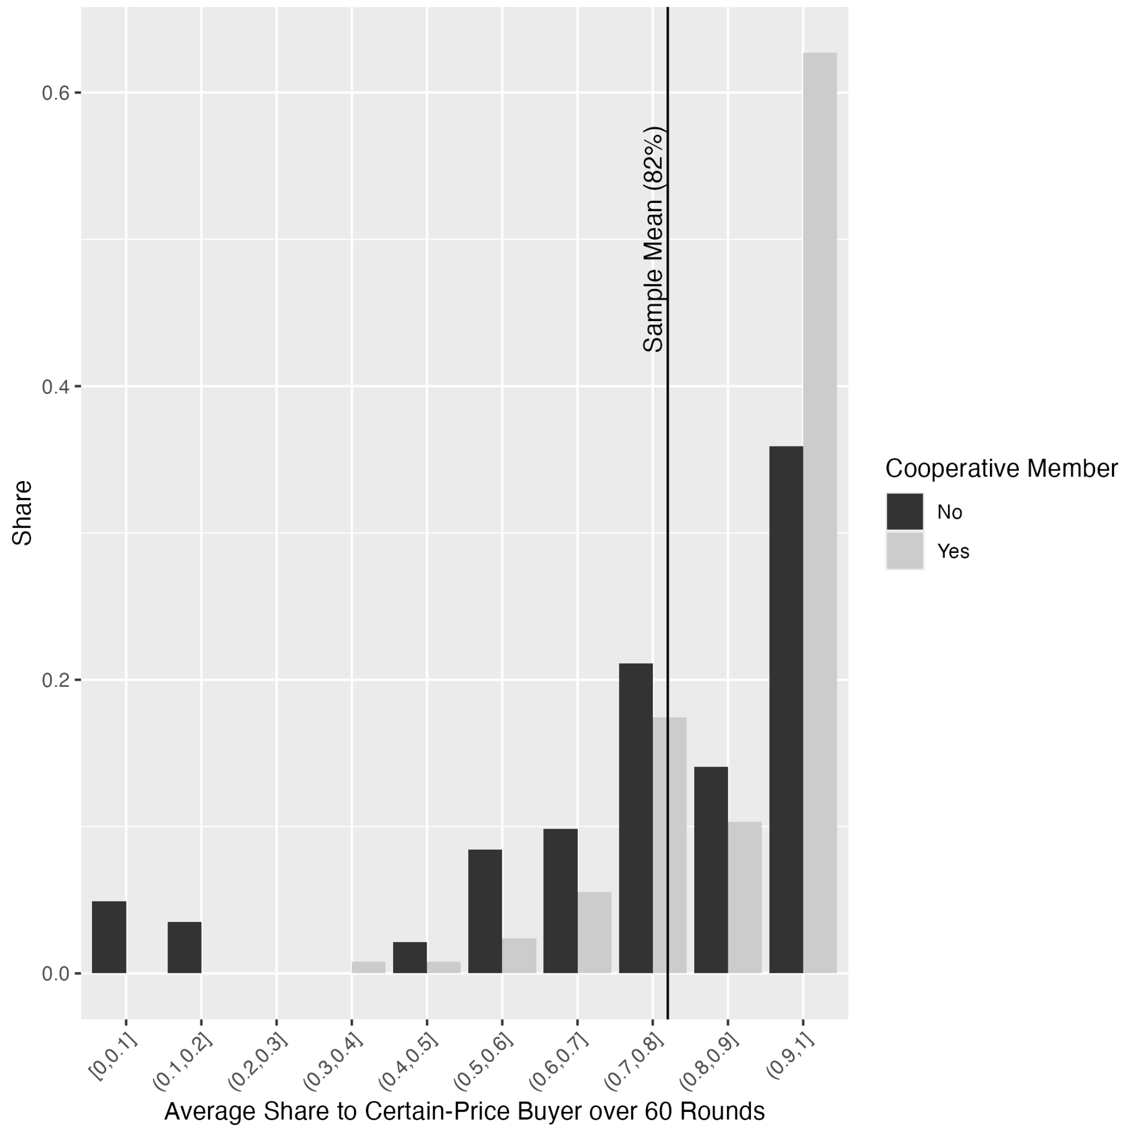

This figure displays a histogram of average share allocated to certain-price buyer over all 60 rounds by participants, broken down by cooperative membership status.

Table 2.1: Descriptive statistics at the participant level

	N	Yes	No	Mean	SD
Exit Survey					
Gender (1 = Female)	268	131	137	0.489	0.501
Age	268	—	—	43.593	15.587
Completed Only Middle School (1 = Yes)	268	37	231	0.138	0.346
Completed High School (1 = Yes)	268	29	239	0.108	0.311
Administrative Data					
Cooperative Member (1 = Yes)	268	126	142	0.470	0.500
Years Sold to Cooperative	124	—	—	9.347	2.509
Preliminary Activities					
Can read/write (1 = Yes)	268	199	69	0.743	0.438
Understands arithmetic (1 = Yes)	268	268	0	1.000	0.000
Understands percentages (1 = Yes)	268	266	2	0.993	0.086
Understands probability (1 = Yes)	268	200	68	0.746	0.436
Additional income treatment (1 = Yes)	268	134	134	0.500	0.501
CRRA (from Eckel-Grossman Lottery)	268	—	—	0.530	0.655
Practice game (1 = Yes)	268	228	40	0.851	0.357
Outcome of Interest					
Overall Margin	268	—	—	0.821	0.221
Extensive Margin	268	58	210	0.216	0.413
Intensive Margin	210	—	—	0.772	0.225

40 participants did not complete the practice game because of enumerator error.
Overall Margin is average allocation to certain-price buyer across 60 rounds.
Extensive Margin is 1 if a participant always allocates entire harvest to certain-price buyer across 60 rounds, 0 otherwise.
Intensive Margin is the average allocation for the subset of participants for whom Extensive Margin is not 1.

Table 2.2: Game Order

	Order	Count
Lottery Before		
	Lottery, Game 1, Game 2, Game 3	26
	Lottery, Game 1, Game 3, Game 2	26
	Lottery, Game 2, Game 1, Game 3	22
	Lottery, Game 2, Game 3, Game 1	24
	Lottery, Game 3, Game 1, Game 2	24
	Lottery, Game 3, Game 2, Game 1	20
Subtotal	—	142
Lottery After		
	Game 1, Game 2, Game 3, Lottery	19
	Game 1, Game 3, Game 2, Lottery	25
	Game 2, Game 1, Game 3, Lottery	15
	Game 2, Game 3, Game 1, Lottery	23
	Game 3, Game 1, Game 2, Lottery	23
	Game 3, Game 2, Game 1, Lottery	21
Subtotal	—	126
Total	—	268

All participants completed three games and an Eckel-Grossman risk preference lottery before or after the three games.
The order of the lottery and the games was determined with a roll of a 12-sided die.

Table 2.3: Gamble choices, expected payoff, and risk[1]

Choice	Event	Probability (%)	Payment (MXN)	Expected Payoff	Risk[2]	CRRA[3]
1	A	50%	10000	10000	0	$r > 2$
	B	50%	10000			
2	A	50%	15000	11250	3750	$0.67 < r < 2$
	B	50%	7500			
3	A	50%	20000	12500	7500	$0.38 < r < 0.67$
	B	50%	5000			
4	A	50%	25000	13750	11250	$0.20 < r < 0.38$
	B	50%	2500			
5	A	50%	30000	15000	15000	$r < 0.20$
	B	50%	0			

[1] Adapted from Table 1 in Eckel and Grossman (2008)
[2] Measured as standard deviation of expected payoff.
[3] Calculated as the range of r in the function $U(x) = x^{1-r}/(1-r)$ for which the subject chooses each gamble assuming constant relative risk aversion utility.

Table 2.4: Field visits to regional centers served by Ts'umbal Xitalha'

		Participants		
	Dates	Non-Members	Members	Total
Agua Dulce Tehuacan	15 July	9	12	21
Chilón	N/A	—	—	—
Coquilte'el	20 July	13	12	25
Nuevo Progreso	3 Aug; 22 Aug	45	10	55
Paraiso Chic'otanil	14 July	4	21	25
San Jose Veracruz	29 June; 2 Aug	18	29	47
Tzubute'el	19 July	6	20	26
Yaxwinic	30 June; 1 July	45	16	61
Ye'tal Ts'ahc	N/A	—	—	—
Yochibha	28 June	2	6	8
Total	—	142	126	268

Field visits were conducted in Summer 2022.
For logistical reasons, we could not visit two of the ten regional centers.
After all of the field visits were completed, we used the TX member list to determine whether experiment participants were in cooperative member families.

Table 2.5: Descriptive statistics at the participant-round level

	Game 1		Game 2		Game 3		Pooled	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Experimental Variables								
Harvest 2 quintals (1 = Yes)	0.156	(0.363)	0.154	(0.361)	0.152	(0.359)	0.154	(0.361)
Harvest 4 quintals (1 = Yes)	0.198	(0.398)	0.208	(0.406)	0.193	(0.394)	0.199	(0.399)
Harvest 6 quintals (1 = Yes)	0.305	(0.460)	0.319	(0.466)	0.327	(0.469)	0.317	(0.465)
Harvest 8 quintals (1 = Yes)	0.341	(0.474)	0.320	(0.466)	0.329	(0.470)	0.330	(0.470)
Mean of Uncertain-Price Buyer MXN 45 (1 = Yes)	0.204	(0.403)	0.211	(0.408)	0.203	(0.402)	0.206	(0.405)
Mean of Uncertain-Price Buyer MXN 50 (1 = Yes)	0.431	(0.495)	0.431	(0.495)	0.418	(0.493)	0.426	(0.495)
Mean of Uncertain-Price Buyer MXN 55 (1 = Yes)	0.365	(0.481)	0.358	(0.480)	0.379	(0.485)	0.368	(0.482)
Outcome of Interest								
Allocation to Certain-Price Buyer	0.820	(0.275)	0.826	(0.268)	0.818	(0.277)	0.821	(0.273)
Observations								
Participants	268	—	268	—	268	—	268	—
Rounds	5360	—	5360	—	5360	—	16080	—

Table 2.6: Impact on Share to Certain-Price Buyer by Game

	Dependent variable:		
	Share Sold to Certain-Price Buyer		
	Game 1	Game 2	Game 3
	(1)	(2)	(3)
Harvest 2 quintals (1 = Yes)	−0.035*** (0.009)	−0.036*** (0.009)	−0.015 (0.010)
Harvest 6 quintals (1 = Yes)	0.003 (0.006)	0.001 (0.006)	0.018*** (0.006)
Harvest 8 quintals (1 = Yes)	−0.021** (0.008)	−0.021*** (0.006)	−0.014* (0.007)
Mean of Uncertain-Price Buyer MXN 45 (1 = Yes)	−0.025*** (0.007)	−0.024*** (0.007)	−0.022*** (0.007)
Mean of Uncertain-Price Buyer MXN 55 (1 = Yes)	−0.005 (0.006)	0.002 (0.005)	−0.003 (0.006)
Linear Time Trend	0.0001 (0.001)	−0.00001 (0.0004)	0.001 (0.0004)
Participant Fixed Effects	Y	Y	Y
Participants	268	268	268
Rounds	60	60	60
Baseline Allocation	0.820	0.826	0.818
Observations	5,360	5,360	5,360

*p<0.1; **p<0.05; ***p<0.01
Standard errors are clustered at the participant level.
Reference harvest is 4 quintals.
Reference mean of price offered by uncertain-price buyer is MXN 50.
In columns (1), (2), and (3), certain-price buyer offers MXN 50.
In column (2), certain-price buyer also offered microcredit to participant last year.
In column (3), certain-price buyer is a cooperative that offered microcredit and technical assistance last year.

Table 2.7: Impact on Share to Certain-Price Buyer

	Dependent variable:
	Share Sold to Certain-Price Buyer
Harvest 2 quintals (1 = Yes)	-0.030^{***} (0.007)
Harvest 6 quintals (1 = Yes)	0.007^{*} (0.004)
Harvest 8 quintals (1 = Yes)	-0.017^{***} (0.006)
Mean of Uncertain-Price Buyer MXN 45 (1 = Yes)	-0.024^{***} (0.005)
Mean of Uncertain-Price Buyer MXN 55 (1 = Yes)	-0.003 (0.004)
Game 2 (Microcredit)	0.001 (0.007)
Game 3 (Coop with Microcredit and Technical Assistance)	-0.011 (0.011)
Linear Time Trend	0.0002 (0.0003)
Participant Fixed Effects	Y
Participants	268
Rounds	60
Baseline Allocation	0.821
Observations	16,080

*p<0.1; **p<0.05; ***p<0.01
Standard errors are clustered at the participant level.
Reference harvest is 4 quintals.
Reference mean of price offered by uncertain-price buyer is MXN 50.

Table 2.8: Participant Level Outcomes

	Average Allocation to Certain-Price Buyer Over 60 Rounds		
	Overall Margin	Extensive Margin	Intensive Margin
	(1)	(2)	(3)
MXN 3000 Additional Income	0.001 (0.025)	0.100* (0.051)	−0.019 (0.030)
Female (1=Yes)	0.007 (0.027)	0.029 (0.056)	−0.007 (0.033)
Age	−0.0003 (0.001)	−0.0005 (0.002)	−0.0004 (0.001)
Completed Only Middle School (1=Yes)	−0.099* (0.052)	−0.159** (0.068)	−0.080 (0.057)
Completed High School (1=Yes)	−0.044 (0.054)	0.020 (0.093)	−0.053 (0.071)
Played Practice Game (1=Yes)	−0.137*** (0.035)	−0.275*** (0.088)	−0.116** (0.050)
Understands Probability (1=Yes)	−0.091*** (0.031)	−0.161** (0.066)	−0.052 (0.037)
Can Read/Write (1=Yes)	0.020 (0.037)	0.036 (0.063)	0.013 (0.044)
Constant	0.971*** (0.099)	0.485*** (0.170)	0.913*** (0.115)
Game Order, Lottery Position, Lottery Outcome Controls	Y	Y	Y
Observations	268	268	210
R^2	0.134	0.110	0.123

*p<0.1; **p<0.05; ***p<0.01
In column (2), the dependent variable is a dummy which equals 1 if the participant allocates the entire harvest to the certain-price buyer in all rounds; 0 otherwise. Column (3) presents the same regression as column (1) on the subset of participants for whom the dummy variable is 0.
All three columns present heteroskedasticity-robust standard errors.

Table 2.9: Participant Level Outcomes Moderated by CRRA

	Average Allocation to Certain-Price Buyer Over 60 Rounds		
	Overall Margin	Extensive Margin	Intensive Margin
	(1)	(2)	(3)
MXN 3000 Additional Income	0.006 (0.033)	0.067 (0.066)	−0.009 (0.038)
CRRA	0.016 (0.020)	−0.041 (0.044)	0.030 (0.023)
Additional Income * CRRA	−0.001 (0.028)	0.075 (0.085)	−0.013 (0.031)
Female (1=Yes)	0.003 (0.028)	0.038 (0.055)	−0.014 (0.034)
Age	−0.0003 (0.001)	−0.0001 (0.002)	−0.0004 (0.001)
Completed Only Middle School (1=Yes)	−0.099* (0.052)	−0.148** (0.069)	−0.080 (0.058)
Completed High School (1=Yes)	−0.050 (0.053)	0.035 (0.093)	−0.059 (0.070)
Played Practice Game (1=Yes)	−0.145*** (0.034)	−0.296*** (0.087)	−0.124** (0.049)
Understands Probability (1=Yes)	−0.089*** (0.031)	−0.155** (0.065)	−0.052 (0.037)
Can Read/Write (1=Yes)	0.019 (0.037)	0.041 (0.062)	0.006 (0.045)
Constant	1.010*** (0.092)	0.551*** (0.168)	0.943*** (0.110)
Game Order, Lottery Position, Lottery Outcome Controls	Y	Y	Y
Observations	268	268	210
R^2	0.132	0.129	0.119

*p<0.1; **p<0.05; ***p<0.01
In column (2), the dependent variable is a dummy which equals 1 if the participant allocates the entire harvest to the certain-price buyer in all rounds; 0 otherwise. Column (3) presents the same regression as column (1) on the subset of participants for whom the dummy variable is 0.
All three columns present heteroskedasticity-robust standard errors.

Table 2.10: Impact on Share to Certain-Price Buyer by Cooperative Membership Status

	Dependent variable:	
	Share Sold to Certain-Price Buyer	
	Members	Non-Members
	(1)	(2)
Harvest 2 quintals (1 = Yes)	−0.025*** (0.009)	−0.036*** (0.010)
Harvest 6 quintals (1 = Yes)	0.008** (0.003)	0.005 (0.006)
Harvest 8 quintals (1 = Yes)	−0.001 (0.006)	−0.032*** (0.009)
Mean of Uncertain-Price Buyer MXN 45 (1 = Yes)	−0.007* (0.004)	−0.041*** (0.009)
Mean of Uncertain-Price Buyer MXN 55 (1 = Yes)	−0.0003 (0.004)	−0.005 (0.006)
Game 2 (Microcredit)	0.013* (0.008)	−0.009 (0.011)
Game 3 (Coop with Microcredit and Technical Assistance)	0.011 (0.011)	−0.032* (0.019)
Linear Time Trend	0.00002 (0.0002)	0.0004 (0.0004)
Participant Fixed Effects	Y	Y
Participants	126	142
Rounds	60	60
Observations	7,560	8,520

*p<0.1; **p<0.05; ***p<0.01
Standard errors are clustered at the participant level.
Reference harvest is 4 quintals.
Reference mean of price offered by uncertain-price buyer is MXN 50.

Table 2.11: Participant Level Outcomes (Cooperative Members)

	Average Allocation to Certain-Price Buyer Over 60 Rounds		
	Overall Margin	Extensive Margin	Intensive Margin
	(1)	(2)	(3)
MXN 3000 Additional Income	0.014 (0.024)	0.166** (0.083)	−0.012 (0.032)
Female (1=Yes)	0.023 (0.032)	0.0005 (0.116)	0.014 (0.045)
Age	−0.0005 (0.001)	−0.002 (0.004)	−0.0004 (0.002)
Completed Only Middle School (1=Yes)	−0.040 (0.052)	−0.213* (0.120)	−0.011 (0.064)
Completed High School (1=Yes)	0.070 (0.051)	0.172 (0.186)	0.085 (0.086)
Played Practice Game (1=Yes)	−0.080* (0.041)	−0.305** (0.128)	−0.051 (0.067)
Understands Probability (1=Yes)	−0.085*** (0.027)	−0.214** (0.102)	−0.058 (0.042)
Can Read/Write (1=Yes)	−0.052 (0.032)	−0.070 (0.105)	−0.054 (0.041)
Constant	1.091*** (0.109)	0.880** (0.388)	0.979*** (0.144)
Game Order, Lottery Position, Lottery Outcome Controls	Y	Y	Y
Observations	126	126	89
R^2	0.204	0.210	0.161

*p<0.1; **p<0.05; ***p<0.01
In column (2), the dependent variable is a dummy which equals 1 if the participant allocates the entire harvest to the certain-price buyer in all rounds; 0 otherwise. Column (3) presents the same regression as column (1) on the subset of participants for whom the dummy variable is 0.
All three columns present heteroskedasticity-robust standard errors.

Table 2.12: Participant Level Outcomes (Cooperative Non-Members)

	Average Allocation to Certain-Price Buyer Over 60 Rounds		
	Overall Margin	Extensive Margin	Intensive Margin
	(1)	(2)	(3)
MXN 3000 Additional Income	−0.035 (0.043)	0.023 (0.059)	−0.036 (0.045)
Female (1=Yes)	0.070 (0.048)	0.115* (0.062)	0.047 (0.053)
Age	−0.001 (0.002)	−0.001 (0.002)	−0.0002 (0.002)
Completed Only Middle School (1=Yes)	−0.097 (0.069)	−0.088 (0.090)	−0.101 (0.076)
Completed High School (1=Yes)	−0.049 (0.071)	0.005 (0.109)	−0.051 (0.088)
Played Practice Game (1=Yes)	−0.208*** (0.061)	−0.161 (0.131)	−0.246*** (0.077)
Understands Probability (1=Yes)	−0.106** (0.052)	−0.145* (0.082)	−0.047 (0.056)
Can Read/Write (1=Yes)	0.117* (0.062)	0.177*** (0.063)	0.088 (0.067)
Constant	0.923*** (0.152)	0.143 (0.197)	0.932*** (0.167)
Game Order, Lottery Position, Lottery Outcome Controls	Y	Y	Y
Observations	142	142	121
R^2	0.206	0.155	0.202

*p<0.1; **p<0.05; ***p<0.01
In column (2), the dependent variable is a dummy which equals 1 if the participant allocates the entire harvest to the certain-price buyer in all rounds; 0 otherwise. Column (3) presents the same regression as column (1) on the subset of participants for whom the dummy variable is 0.
All three columns present heteroskedasticity-robust standard errors.

Table 2.13: Participant Level Outcomes Moderated by CRRA (Cooperative Members)

	Average Allocation to Certain-Price Buyer Over 60 Rounds		
	Overall Margin	Extensive Margin	Intensive Margin
	(1)	(2)	(3)
MXN 3000 Additional Income	0.007 (0.031)	0.066 (0.112)	−0.003 (0.042)
CRRA	−0.021 (0.018)	−0.061 (0.080)	−0.022 (0.020)
Additional Income * CRRA	0.018 (0.031)	0.182 (0.133)	−0.017 (0.036)
Female (1=Yes)	0.018 (0.034)	0.016 (0.112)	0.007 (0.046)
Age	−0.0004 (0.001)	−0.0004 (0.003)	−0.0004 (0.002)
Completed Only Middle School (1=Yes)	−0.037 (0.053)	−0.233* (0.127)	−0.001 (0.066)
Completed High School (1=Yes)	0.060 (0.052)	0.196 (0.204)	0.044 (0.093)
Played Practice Game (1=Yes)	−0.093** (0.039)	−0.299** (0.119)	−0.070 (0.066)
Understands Probability (1=Yes)	−0.082*** (0.027)	−0.216** (0.097)	−0.054 (0.042)
Can Read/Write (1=Yes)	−0.043 (0.033)	−0.039 (0.101)	−0.052 (0.042)
Constant	1.160*** (0.084)	0.831** (0.330)	1.101*** (0.122)
Game Order, Lottery Position, Lottery Outcome Controls	Y	Y	Y
Observations	126	126	89
R^2	0.221	0.246	0.169

*p<0.1; **p<0.05; ***p<0.01

In column (2), the dependent variable is a dummy which equals 1 if the participant allocates the entire harvest to the certain-price buyer in all rounds; 0 otherwise. Column (3) presents the same regression as column (1) on the subset of participants for whom the dummy variable is 0.

All three columns present heteroskedasticity-robust standard errors.

Table 2.14: Participant Level Outcomes Moderated by CRRA (Cooperative Non-Members)

	Average Allocation to Certain-Price Buyer Over 60 Rounds		
	Overall Margin	Extensive Margin	Intensive Margin
	(1)	(2)	(3)
MXN 3000 Additional Income	−0.008 (0.057)	0.058 (0.074)	−0.013 (0.059)
CRRA	0.045 (0.037)	−0.064 (0.057)	0.062 (0.039)
Additional Income * CRRA	−0.043 (0.048)	−0.112 (0.075)	−0.029 (0.054)
Female (1=Yes)	0.060 (0.050)	0.130** (0.064)	0.033 (0.056)
Age	−0.001 (0.002)	−0.0003 (0.002)	−0.0003 (0.002)
Completed Only Middle School (1=Yes)	−0.100 (0.070)	−0.121 (0.088)	−0.093 (0.077)
Completed High School (1=Yes)	−0.064 (0.072)	0.028 (0.107)	−0.071 (0.086)
Played Practice Game (1=Yes)	−0.210*** (0.062)	−0.183 (0.126)	−0.241*** (0.079)
Understands Probability (1=Yes)	−0.105** (0.052)	−0.143* (0.083)	−0.047 (0.055)
Can Read/Write (1=Yes)	0.116* (0.064)	0.203*** (0.070)	0.074 (0.070)
Constant	0.913*** (0.154)	0.228 (0.189)	0.903*** (0.172)
Game Order, Lottery Position, Lottery Outcome Controls	Y	Y	Y
Observations	142	142	121
R^2	0.211	0.182	0.216

*p<0.1; **p<0.05; ***p<0.01
In column (2), the dependent variable is a dummy which equals 1 if the participant allocates the entire harvest to the certain-price buyer in all rounds; 0 otherwise. Column (3) presents the same regression as column (1) on the subset of participants for whom the dummy variable is 0.
All three columns present heteroskedasticity-robust standard errors.

Table 2.15: Participant Level Outcomes Moderated by Loyalty (Cooperative Members)

	Average Allocation to Certain-Price Buyer Over 60 Rounds		
	Overall Margin	Extensive Margin	Intensive Margin
	(1)	(2)	(3)
MXN 3000 Additional Income	−0.0005 (0.111)	0.416 (0.347)	−0.043 (0.154)
Female (1=Yes)	0.019 (0.035)	−0.004 (0.108)	0.012 (0.049)
Age	−0.001 (0.001)	−0.001 (0.003)	−0.001 (0.002)
Completed Only Middle School (1=Yes)	−0.046 (0.049)	−0.235 (0.153)	−0.012 (0.061)
Completed High School (1=Yes)	0.049 (0.062)	0.164 (0.194)	0.045 (0.105)
Played Practice Game (1=Yes)	−0.094** (0.038)	−0.325*** (0.119)	−0.073 (0.059)
Understands Probability (1=Yes)	−0.087*** (0.031)	−0.239** (0.097)	−0.056 (0.047)
Can Read/Write (1=Yes)	−0.053 (0.033)	−0.092 (0.104)	−0.061 (0.049)
Years Sold to Cooperative	−0.005 (0.009)	−0.0004 (0.029)	−0.004 (0.012)
Years Sold * Additional Income	0.001 (0.011)	−0.029 (0.036)	0.003 (0.016)
Constant	1.221*** (0.140)	0.952** (0.437)	1.143*** (0.186)
Game Order, Lottery Position, Lottery Outcome Controls	Y	Y	Y
Observations	124	124	88
R^2	0.222	0.266	0.153

*p<0.1; **p<0.05; ***p<0.01
In column (2), the dependent variable is a dummy which equals 1 if the participant allocates the entire harvest to the certain-price buyer in all rounds; 0 otherwise. Column (3) presents the same regression as column (1) on the subset of participants for whom the dummy variable is 0.
All three columns present heteroskedasticity-robust standard errors.

Chapter 3

Information Decay and Cooperative Entry under Risk

3.1 Introduction

Adoption rates of potentially welfare-improving production technologies remain stubbornly low in many contexts (Suri & Udry, 2022). Social networks play an important role in technology adoption by alleviating information frictions that inhibit adoption (Munshi, 2014). However, information transmission in social networks breaks down over space and time, and poorly connected firms suffer as a result. A better understanding of information decay would provide insight into how to reinforce these social networks. Strengthening them could, in turn, increase adoption and both individual and overall welfare.

This paper studies the effect of temporal and spatial lag on entry into two different cooperatives by indigenous producers who experience periods of seasonal drought. One cooperative is a coffee cooperative that offers technical training and price insurance to existing smallholder coffee producers. The other cooperative is a honey cooperative that offers these coffee producers an additional source of income during the coffee off-season. By temporal decay, we mean that it takes several years for producers who enter to experience the benefits of both cooperatives and spread information about these benefits to their neighbors. By spatial decay, we mean that the farther this information travels, the less it influences entry decisions. Both cooperatives operate in a remote area of rural southern Mexico with limited road connectivity that is isolated from outside influence and thus free of many of the usual confounders of the study of cooperative entry. We have an unusually rich data set: panel data that span 22 years with the complete set of entry decisions into both cooperatives and the locations of their coffee plots.

Our setting is ideal for studying the effect of temporal and spatial lag on entry into both cooperatives. By varying the method (first differences vs fixed effects) and the sample (11 year vs 22 year panel), we can capture the effect of temporal lag on entry. Moreover, the spatial organization of the producers exhibits a network structure: producers are organized in villages, which are then organized in regions. We consider three different levels of spatial spillovers: the direct effect of the adoption rate in a producer's village, the indirect effect of the adoption rate in neighboring villages within the same region, and the overall effect of the adoption rate of villages in other regions. Thus, we can capture the effect of spatial lag on cooperative entry as well.

These Mexican producers suffer from the effects of seasonal drought (Dobler-Morales & Bocco, 2021). We use the geolocation of producers' coffee plots and villages to augment the entry network above with periods of drought from the Standardized Precipitation-Evapotranspiration Index (Vicente-Serrano et al., 2010). The SPEI distinguishes between severe drought (between -1.5 and -2) and extreme drought (below -2). Thus, we can study the entry into both cooperatives in a context with a particular type of climate shock, seasonal drought. In particular, we are interested in whether, in the face of seasonal drought, producers with stronger networks enter the cooperatives with higher probability than producers with weaker networks. This heterogeneity would account for the direct effect of the network in mitigating the frictions, information, and otherwise that impede cooperative entry. We are also interested in how periods of seasonal drought in neighboring villages affect cooperative entry. These indirect effects could indicate how information on the effectiveness of membership in the coffee cooperative or honey production against seasonal drought affects the entry decision.

We estimate two types of models on the network graph of entry decisions. The first model allows us to study information lag over time over the 22-year period. We estimate a linear-in-means model that regresses the adoption rate within a producer's village in the previous year on his own decision to adopt (Bramoullé et al., 2009). We augment the baseline specification with the number of periods of severe and extreme drought in the previous year and interact these drought measures with the adoption rate of the village. Moreover, in line with recent work by Millimet and Bellemare (2023), we compare results from a specification with producer fixed effects and one with first differences to examine the effect of temporal lags: the differential effect of adoption rates in prior years and adoption decisions in the past year. The fixed effect specifications include the village adoption rate for all previous years. The first difference specifications only include the village adoption rate from the previous year.

The second model allows us to examine the effect of spatial lags: the differential effect of the adoption rate in the producer's own village, adoption rate in neighboring villages in the same region, and the adoption rate in villages in other regions. We estimate a spatial lag model in the style of Halleck Vega and Elhorst (2015) that uses a weighting matrix to incorporate the indirect effect of the shares of cooperative members and drought measures from neighboring villages alongside the direct effect of the adoption rate and drought measures from a producer's own village. By varying the weighting matrix, we compare three different models of information decay: one that weights neighboring villages within the same region equally, another that weights them by inverse distance, and a third that includes all villages across all regions, also weighting them by inverse distance.

Our results are as follows. Using the linear-in-means models, we find network effects in entry into both the coffee and the honey cooperative. In the model with producer fixed effects, the point estimate of the difference between living in a village with no adopters (network strength of 0) and a village with all adopters (network strength of 1) is around 50% for coffee and 40% for honey. The effect size is the same in both the short panel and the long panel. That means that a 10% increase in the village adoption rate of either cooperative in one year affects the probability that a producer in the same village will adopt the cooperative by 5% or 4% in the following year. In the model with first differences, the effect size decreases. In the short panel, it is 10% for the coffee cooperative and 12% for the honey cooperative. In the long panel, it is null for the coffee cooperative and 7% for the honey cooperative.

In general, periods of severe drought increase and periods of extreme drought decrease the probability of joining the coffee cooperative. Village network strength moves these effects in the opposite direction. For a period of severe drought, the base effect is 2% to 4% with a network effect of 3% to 4% in the opposite direction. For a period of extreme drought, the base effect is -2% to -6% with a network effect of 2% to 4% in the opposite direction. In the case of honey, the presence of periods of drought themselves does not affect the entry decision, but the interaction between the periods and network strength does, but only in one of the models: the one with producer fixed effects over the long panel.

Using the spatial lag model, for the coffee cooperative we find a direct effect that ranges from 30% to 35% in the short panel and 36% to 40% in the long panel. As in the linear-in-means results, the spatial-lag results show that periods of extreme drought decrease the likelihood of joining the cooperative. With a binary contiguity matrix at the regional level, we find an indirect network effect of 60% in the short panel which drops to 43% in the long panel. With an inverse distance matrix at the regional level, the indirect effect decreases to 45% and 43%. When we include villages in all regions, the indirect effects increase to 67% and 76%. We also find an indirect effect in the short term of periods of severe and extreme drought within the region and globally.

With the honey cooperative, we find direct network effects in the short panel of 49% and in the long panel of 39%. We find no direct effects of periods of either type of drought. In the short panel, we find indirect effects of periods of severe drought (3.6%) and extreme drought (8.1%) across the survey region. As in the linear-in-means model, we find very little effect of periods of drought on the entry decision for the honey cooperative, either in the short panel or the long panel. Thus, producers who experience drought do not look for alternative income in the form of honey production.

Our results contribute to a literature that uses network theory to analyze the effect of social networks on producer decisions. In particular, our work is closely related to the literature on technology adoption. Foster and Rosenzweig (1995) first document the role of peer learning in the adoption of high-yielding seed varieties in India. Conley and Udry (2010) use surveys to define information neighborhoods for pineapple producers and distinguish between nearby and farther-away peer effects as producers learn the correct amount of fertilizer. Next, A. Banerjee et al. (2013) expand their model so that even non-adopters can provide information as they examine the diffusion of microfinance in Indian communities. Recently, Beaman et al. (2021) explicitly model the network structure even more by considering not only the presence but the quantity of links among producers. They find that a threshold model explains the adoption of pit planting better than a simple contagion model.[1] We improve on their work by considering different temporal scales: last year with first differences versus a complete history with fixed effects.

We also contribute to a literature that studies the ability of social networks to protect against climate shocks. In the past decade, the availability of high-quality remote sensing data has opened up new research possibilities (Dell et al., 2014). The initial work of Robert Townsend (1994) shows that village networks provide insurance for unexpected consumption expenses, since villagers borrow money from each other. More recent work by Kinnan et al. (2024) shows how health shocks propagate through village networks. In our case, we are interested in how social networks provide information about the benefits of a potential technology and the working capital to adopt it. We also have a uniquely rich network. Our work contributes to a new literature that studies ex ante and ex post adaptation to climate change (Carleton et al., 2024).

Finally, we combine two different econometric techniques in a novel way to examine temporal and spatial decay in peer effects. The study of peer effects extends beyond cooperative entry to many classes of decisions (Bramoullé et al., 2020). Our study is one of the first to use panel data and the first that we know of to use two different panel lengths. [2] Moreover, we are the first to use first differences in addition to fixed effects (Millimet & Bellemare, 2023) to control for individual heterogeneity. Similarly, the study of spatial lag comes from the political science literature, for example the impact of the policy of neighboring countries on the policy of a particular country (Yesilyurt & Elhorst, 2017). To our knowledge, we are the first to compare estimation results from a linear-in-means model and from a spatial lag model to study temporal and spatial frictions.

Our paper proceeds as follows. Section 2 describes the context and the two cooperatives. Section 3 describes our entry network and drought data. Section 4 gives the empirical specifications for the linear-in-means and spatial lag models. Section 5 explains the results of both models. Section 6 concludes.

[1] "Pit planting" is an improved way to plant maize in Africa.
[2] Bramoullé et al. (2020) only gives three other examples.

3.2 Background and Context

In this section, we first describe the context of our study and the two issues facing producers. Next, we describe the two cooperatives and how they address these issues. Finally, we describe conceptually how the entry decisions of other producers in the same village and in different villages would affect a given producer's decision to enter either cooperative.

3.2.1 The Problem: Seasonal Drought and Coffee Leaf Rust

Our context is the state of Chiapas, Mexico, which is the largest coffee producing state in Mexico. Most coffee producers are smallholder producers with less than 5 hectares of land, like our population. These particular coffee plots are located on the sides of hills at altitude under a shade canopy and as part of a larger ecosystem (Soto-Pinto et al., 2000). Cooperatives have been highly operative throughout the region since the 1990s, when the Mexican government ended its subsidy programs (Martinez-Torres, 2006). They have functioned as extension programs, teaching farmers a variety of ways to respond to climate change (Soto-Pinto et al., 2012).

Smallholder agricultural producers often depend on income from one cash crop in order to finance the purchase of all items that they cannot produce themselves. Thus, they are particularly vulnerable to adverse production shocks that affect this cash crop. In our context, smallholders produce coffee, but the issues that we describe below could apply to any other cash crop, such as cacao.

We consider two vulnerabilities in particular: seasonal drought and coffee leaf rust. Seasonal drought is one of the channels through which climate change affects agricultural productivity (Ortiz-Bobea et al., 2021). To measure seasonal drought, we use the the Standardised Precipitation-Evapotranspiration Index (SPEI) compiled by Vicente-Serrano et al. (2010), which we describe in more detail in Section 3.3.1. Membership in a coffee cooperative in response to seasonal drought is an example of a producer's adaptive response to climate change (Carleton et al., 2024).

In addition to drought, which affects a variety of cash crops, coffee in particular is affected by coffee leaf rust (CLR), a fungus that affects Arabica coffee plants worldwide (Rhiney et al., 2021). Beginning in 2012, Mexico and Central America experienced an outbreak of CLR that significantly reduced production. The incidence of CLR is related to climate change. The increased heat of climate change makes coffee plants more susceptible to CLR. In addition, common agricultural practices such as monoculture and deforestation also make coffee plants more susceptible to the disease.

3.2.2 Mitigating Technologies

Both the coffee cooperative and the honey cooperative provide strategies to counter the effect of seasonal drought and coffee leaf rust on agricultural productivity and thus producer welfare.

We examine the question of membership in a coffee cooperative or in a honey cooperative by borrowing from the framework of technology adoption, in particular the notion of learning-from-others (Foster & Rosenzweig, 2010). Our approach contrasts with previous work that examines the determinants of producer entry into contract farming (Bellemare & Bloem, 2018) and fair trade arrangements (Dragusanu et al., 2014). Many producer cooperatives offer some version of these services: a guaranteed purchase price to insure production, microcredit to smooth consumption, and technical assistance to learn improved production techniques. These services are funded by upstream contracts. To our knowledge, we are the first to consider membership in a cooperative under the framework of technology adoption.

Much of the literature on technology adoption considers the adoption of improved inputs such as High Yield Variety (HYV) seeds and fertilizer (Foster & Rosenzweig, 2010). Producers will adopt the technology if the expected benefit outweighs the cost. Early work borrowed the notion of learning-by-doing from the endogenous growth literature using the target input model (Romer, 1994). In this model, producers observe the effect of a particular amount of input (usually an amount of fertilizer) and the resulting output. Over time, they learn to calibrate the amount of input to the amount of output.

One drawback to a purely learning-by-doing approach is that it may take many attempts for a producer to determine the correct input by trial and error. Thus, Foster and Rosenzweig (1995) introduce the notion of learning by observing others. In effect, every time a producer observes a neighbor's experience with a particular technology, it allows him to approximate more closely the optimal amount of the technology. The effectiveness of learning from others depends on the assumption that the experience of a neighbor is more similar to the producer's own than not, as Munshi (2014) points out. However, understanding the role of social networks in technology adoption has emerged as a key to increasing technology adoption (Beaman et al., 2021).

In the following, we describe in more detail the two cooperatives, how they mitigate the effect of seasonal drought and coffee leaf rust, and the decision problem the producer faces in deciding whether to join them.

Coffee Cooperative

Smallholder coffee producers suffer both from price risk and quantity risk. The price risk comes from output price volatility. They must sell their production to intermediaries whose prices vary depending on the international price of coffee. We consider the quantity risk that producers suffer due to the effects of climate change and coffee leaf rust. Coffee grows best at altitude in a wet, tropical climate. Thus, seasonal drought negatively impacts production. Improved agronomic techniques from technical assistance workshops offered by coffee cooperatives could mitigate these negative effects. Coffee leaf rust affects coffee plants directly by permanently reducing production. To mitigate the effects of CLR, producers must replace coffee plants with disease-resistant varieties. Coffee cooperatives both develop these plants and subsidize their planting. In addition, the technical assistance workshops teach alternatives to monoculture and deforestation that make coffee landscapes overall less susceptible to risk.

When deciding whether to join a coffee cooperative, a producer considers the expected cost and the expected benefit of technical assistance workshops. Uncertainty is involved in both of these estimates. Producers' neighbors can help them reduce the uncertainty around the expected benefit of the technical assistance workshop by sharing their own experience. In a drought situation, they can provide first-hand experience of how effective the techniques are in mitigating the effects. Similarly, replacing coffee plants due to CLR requires access to improved coffee plants, as well as labor and material costs. Cooperative membership could grant access to these plants, help in planting them, and a more certain estimate of whether the plants actually work against CLR.

Honey Cooperative

For the producers in our area of study, membership in a honey cooperative functions as a different kind of technology than membership in a coffee cooperative. Membership in a coffee cooperative offers producers the opportunity to improve their coffee production, while membership in a honey cooperative offers them the opportunity to diversify their income and insure themselves against the quantity risk that climate change and CLR pose to their coffee production.

Anderzén et al. (2020) highlight the benefits of beekeeping as a livelihood diversification strategy for a similar population of coffee growers to our own. They find that beekeeping is associated with a reduction in the incidence of food insecurity because it provides a separate source of income that comes at a different time as income from the coffee harvest. In general, coffee producers in the region are diversifying as a result of climate change (Eakin et al., 2012).

Despite the benefits of beekeeping, coffee producers have been reluctant to adopt the practice (Anderzén et al., 2024). One factor is that the technology is unfamiliar. Another factor is the initial capital investment. It takes a certain number of bees and specialized equipment to start beekeeping. Finally, they are concerned about a market for honey. Our partner honey cooperative provides training and the loan of the initial equipment. In addition, it certifies the honey as organic and provides market access to sell it in other parts of Mexico. As the number of producers' peers who adopt honey production goes up, the uncertainty around the welfare effect of adopting honey collapses. For many producers, this change in their cost-benefit analysis leads to an increase in the likelihood of entry.

3.3 Data and Descriptive Statistics

In this section, we describe the data that we use to analyze membership in the coffee cooperative and the honey cooperative. Our analysis leverages spatial variation in the location of producers and temporal variation in the timing of cooperative entry. Moreover, the spatial variation allows us to cross reference producers' locations and remote sensing drought data so that we can analyze the effect of the drought on producers' decisions to join one or both cooperatives.

3.3.1 Spatial Extent

Figure 3.1 shows a network graph of entry into the coffee cooperative. Producers are divided into regions which are, in turn, subdivided into villages. Colors indicate the year of entry. Clusters of the same color visually identify groups of producers in the same village or region that entered in the same year. This clustering reveals that producers who enter the coffee cooperative together tend to live with other producers in the same village.

Figure 3.2 shows a network graph of entry into the honey cooperative. Once again, producers are grouped into regions that branch out into villages. In contrast to the organization of coffee producers, this clustering reveals that producers who enter the honey cooperative together tend to be the only producers in their village and in many cases the only producers in their region. Moreover, even at the end of the time period, not every region has a honey producer.

Our region of interest includes substantial variation in altitude and climate. Figure 3.3 shows the variation in altitude. Of the 498 producers in our data set, we have geolocated the coffee plot of 244 of them; for the other 254, we use the coordinates of a nearby village. The combined set of elevations is normally distributed with most elevations in the range of 500 to 1500 meters. This altitude variation gives us substantial variation in rainfall and temperature for the region of interest.

We extracted Standardized Precipitation Evapotranspiration Index (SPEI) values for every coffee plot or nearby village for the years 2002-2024 using Google Earth Engine. The SPEI is a gridded measure of drought that uses variation from the mean in both precipitation and temperature over the past three months to build a rolling monthly drought index. Vicente-Serrano et al. (2010) gives more information about the calculation of the SPEI and the associated improvements over the SPI (Standardized Precipitation Index) and the self-reported Palmer scale. The SPEI has two thresholds that define the magnitude of drought conditions: **severe drought** if the drought index is between -1.5 and -2 and **extreme drought** if the drought index is below -2.

Figure 3.4 shows the average monthly value of the SPEI index for the area of interest over the 22-year interval. Horizontal lines indicate the thresholds for severe and extreme droughts. Figure 3.5 relates the average number of producers who join the coffee cooperative each year and the average months of drought each year. We see that the number of producers who join the coffee cooperative decreases substantially in years with one or more months of drought. The year 2018 is an exception to this trend. Figure 3.6 relates the average number of members who join the honey cooperative each year and the average months of drought each year. We see that entry into the honey cooperative is somewhat inversely related to the number of months of drought in a year.

3.3.2 Temporal Extent

Our analysis is based on a unique data set of 22 years of entry into the coffee cooperative and the honey cooperative. The length of this panel allows us to perform our analysis on a shorter and longer time horizon.

We break up the data into two 11-year periods based on the two sets of administrative data that we merged. From these initial periods to 2013, we have self-reported entry dates in both cooperatives. From 2013-24, we have administrative records from both cooperatives that indicate whether the producers marketed coffee through the coffee cooperative or honey through the honey cooperative. This break in the type of data provides a natural way to run our analysis over two different time horizons: a short-term time horizon and a medium-term time horizon. Thus, we will run the empirical analysis we describe in the next section on a short 11-year panel from 2013-24 and a long 22-year panel from 2002-24.

Table 3.1 summarizes the entry patterns at the village level. Table 3.2 summarizes the entry patterns at the member level. As in the network graphs in Figures 3.1 and 3.2 above, we see quite different entry patterns among the two cooperatives.

The coffee cooperative began in 2002 with four producers in one village. In 2013, 274 producers in 63 villages had joined, with a substantial portion of villages in all but one region. By 2024, 484 producers in 120 villages had joined, almost all of the sample.

The honey cooperative began in 2005 with three producers in two villages. In 2013, 25 producers in 16 villages had joined, one or two villages apiece in most of the regions. By 2024, 43 producers in 24 villages had joined, more villages in the same regions but no new regions.

These differential entry patterns motivate our use of direct and indirect effects and different specifications for the indirect effects in the models in the following section.

3.4 Empirical Framework

3.4.1 Linear-In-Means Model

Basic Model

We estimate a linear-in-means model to estimate the effect of the share of cooperative members in a producer's village on the entry decision of the producer. Bramoullé et al. (2009) gives an overview of these models, which are often used in the peer effects literature to determine the association between an outcome variable for an individual and the mean of the same outcome variable for an individual's reference group. In our case, the reference group is the individual's village, as is typical in rural settings (Munshi, 2014).

The outcome is a binary indicator y^z_{ijt} of whether producer i in village j adopted cooperative z in year t. Cooperative is indexed by $z \in c, h$ where c denotes the coffee cooperative and h the honey cooperative.

We define the **network strength** N^z_{ijt} of a producer i in village j at time t for cooperative z as the share of producers in village j that have adopted z, excluding producer i. Network strength ranges from 0 (no other members in the village) to 1 (all other producers in the village are members).

$$N^z_{ijt} = \frac{1}{n_j - 1} \sum_{k=1, k \neq i}^{n_j} y^z_{kjt} \tag{3.1}$$

We use the lagged value of a producer's network N_{ijt-1} to estimate the effect of village network on the producer's entry decision.

Next, we incorporate the yearly measure of periods of severe drought and extreme drought that we described in Section 3.3.1. The SPEI uses rolling periods of three months. The index value for a given month uses precipitation and temperature data from that month and the prior two months. We add the level effects of the number of periods of both types of drought in the previous year. The spatial resolution of our drought data allows us to compute these measures at the producer level, either using the location of a producer's coffee plot or the nearby village.

We denote the number of periods of severe drought and extreme drought experienced by producer i in community j in year t as D^s_{ijt} and D^e_{ijt}, respectively, and group them in a 2x1 vector $\mathbf{D_{ijt}}$ for notational convenience. Similarly, we denote the individual effects of these periods of drought on entry as δ_{ijt}^{zs} and δ_{ijt}^{ze} and group them in the 2x1 vector $\delta_{\mathbf{ijt}}$.

In addition, we add interaction terms to capture the differential effect of the number of periods of each type of drought on entry depending on the network strength. We denote the interaction effects as γ_{ijt}^{zs} and γ_{ijt}^{ze} and group them in a 2x1 vector $\gamma_{\mathbf{ijt}}^{\mathbf{z}}$ in Equations (3.2) and (3.3).

Finally, we present two specifications of this model that control for producer time-invariant characteristics in different ways. Equation (3.2) uses the number of periods of both types of droughts, the overall village share of cooperative members, producer fixed effects, and whether the producer joined the cooperative in the current year. Equation (3.3) uses the change in the number of periods of both types of droughts, the share of village members who adopted the cooperative in the previous year, and whether the producer joined the cooperative the current year. Both specifications incorporate time fixed effects.

$$y_{ijt}^{z} = \alpha_1^z + \beta_1^z N_{ijt-1}^z + \delta_{\mathbf{1}}^{\mathbf{z}} \mathbf{D}_{\mathbf{ijt-1}} + \gamma_{\mathbf{1}}^{\mathbf{z}} \mathbf{D}_{\mathbf{ijt-1}} N_{ijt-1}^z + \phi_{1i}^z + \xi_{1t}^z + \epsilon_{1ijt}^z \tag{3.2}$$

$$\Delta y_{ijt}^{z} = \alpha_2^z + \beta_2^z \Delta N_{ijt-1}^z + \delta_{\mathbf{2}}^{\mathbf{z}} \boldsymbol{\Delta} \mathbf{D}_{\mathbf{ijt-1}} + \gamma_{\mathbf{2}}^{\mathbf{z}} \boldsymbol{\Delta} \mathbf{D}_{\mathbf{ijt-1}} N_{ijt-1}^z + \xi_{2t}^z + \Delta \epsilon_{2ijt}^z \tag{3.3}$$

Identification

We first discuss identification of β^z, the effect of an increase in the share of producers who join cooperative z in time period $t-1$ on the probability that a given producer will join the cooperative in time period t. This coefficient of interest is present in the first specification above, Equation (3.2).

One threat to identification is time-varying shocks that affect all producers, such as large-scale drought, heat, or market shocks. In both equations, we rely on year-fixed effects ξ_t^z to control for these shocks. Another threat to identification is time-invariant unobservable producer characteristics such as ability or education. In Equation (3.2), we rely on producer fixed effects ϕ_i^z to control for these characteristics.

However, recent work by Millimet and Bellemare (2023) suggests that in long panel setups the identification assumption for fixed effects may not hold. Unobservable unit-level heterogeneity may not be constant across all of the time periods of the a given panel. Both of our panels, the short

11-period one and the longer 22-period one, are much longer than the typical three- or four-period panels used in applied research.[3] For this reason, we also estimate first-difference versions of these specifications in Equation (3.3). However, as we mention in the previous and the following sections, it would be a mistake to think of the first difference specification in Equation (3.3) as a different version of Equation (3.2). It really is a completely different model representing entry behavior on a short one-year time period instead of a multi-year year (11 or 22) time period.

We next discuss identification of δ^{zs} and δ^{ze}, the effect of a the number of periods of severe drought and extreme drought. Because we use year fixed effects above, these coefficients capture the average effect of variation in the intensity of drought in the cross section. We consider such short-term climate variation as an exogenous weather shock and use the panel specification recommended by Dell et al. (2014). Similarly to the concerns about the produce fixed effects above, however, these authors also note that the length of our two panels blurs the line between a short-term effect, which goes one way from weather to producer behavior, and a medium-term effect, which may involve an adaptive response on the part of the producer. Thus, we compare estimates of these coefficients on both panels to look for evidence of a temporal lag in adoption.

Finally, we discuss identification of $\gamma^{\mathbf{z}}$: γ^{zs} and γ^{ze}. These scalars capture the joint effect of one additional period of severe or extreme drought, respectively, and the strength of the producer's network on the probability of a producer's entry in a cooperative. If the sign of γ^{z} is the same (opposite) as the sign of β^{z}, then a stronger network increases (decreases) the effect of drought on entry or a drought increases (decreases) the effect of the network.

3.4.2 Spatial Lag Model

We also estimate an Spatially Lagged X (SLX) model to allow for the effect of spatial spillovers on the entry decisions of the members of a producer's village on the entry decision of the producer. Halleck Vega and Elhorst (2015) gives an overview of these models. The SLX model is one of a set of spatial econometric models that are used to model processes with spatial spillover effects. These models incorporated spatially lagged versions of the explanatory variables on the right-hand

[3]For example, McKenzie (2012) describes a more common scenario where researchers move from three to four or five waves of a survey.

side along with their direct counterparts to capture the indirect effect of changes in x in other spatial regions on the independent variable y. In our setting, the SLX model will incorporate spatially lagged versions of the network strength of other villages, as well as the number of periods of extreme and severe drought that these villages experience.

$$y^z_{ijt} = \alpha^z_3 + \beta^z_3 N^z_{ijt-1} + \delta^{\mathbf{z}}_{\mathbf{3}}\mathbf{D_{ijt-1}} + \mathbf{W}\mathbf{N^z_{t-1}}\theta_3 + \mathbf{W}\mathbf{D_{t-1}}\lambda_3 + \epsilon^z_{3ijt} \tag{3.4}$$

The key element of the SLX model is the weighting matrix $\mathbf{W}$, which specifies how the spatially lagged dependent variables enter the estimation. The choice of $\mathbf{W}$ depends on the underlying theory of how the spatial process works. In all cases, the diagonal elements of $\mathbf{W}$ are 0, so that the direct effect of N^z_{ijt-1} does not enter the equation a second time. In practice, researchers often estimate equations with several different specifications of weighting matrices and compare the estimated results with a Durbin-Wu-Hausman test.

We use three different weighting matrices. All three weighting matrices have zeroes down the diagonal so that the direct effect does not enter the estimating equation more than once. The scalars of the off-diagonal elements are calculated in one of three ways below.

1. **Binary Contiguity.** This weighting matrix assigns a weight of 1 to each of the villages in the same region as the village j. The matrix is row-normalized so that the weights in each row add up to 1.

$$w^1_{jk} = 1 \tag{3.5}$$

2. **Inverse Distance - Region.** This weighting matrix assigns a weight to each village k in the same region as village j according to the inverse of the distance d_{jk} between j and k. The matrix is scaled by the largest eigenvalue.

$$w^2_{jk} = \frac{1}{d_{jk}} \tag{3.6}$$

3. **Inverse Distance - All Villages.** This weighting matrix assigns a weight to each village

k i the sample according to the inverse of the distance d_{jk} between j and k. The matrix is scaled by the largest eigenvalue.

$$w_{jk}^3 = \frac{1}{d_{jk}} \tag{3.7}$$

3.4.3 Inference

Here we describe how we calculate the standard errors and perform hypothesis tests for both the linear-in-means model and the spatial-lag model. We begin with the linear-in-means model. Because we are using two-way fixed effects in Equations (3.2) and time fixed effects along with first differences in Equation (3.3), one practice would be to cluster by producer and year in both equations.

Abadie et al. (2023) argues that this practice results in standard errors that are too conservative and proposes two considerations when considering the level of clustering: **a design component** and **a treatment assignment mechanism**. In our case, we are not estimating our equations on a sample but instead analyzing a diffusion process over a whole population. Moreover, every year contains a producer, and every producer eventually contains an entry year, so every cluster is treated. For this reason, we do not cluster our standard errors by village and year. Instead, we use heteroskedasticity-robust standard errors.

Next, we turn to the spatial lag model. Here, the appropriate use of standard errors is an active area of research, so we follow the guidelines of a recent working paper by Xu and Wooldridge (2022). They use the two-part framework above that consists of a design component and a treatment assignment mechanism. Instead of clustering standard errors by the spatial unit (in our case the village), they suggest using spatial heteroskedasticity and autocorrelation consistent standard errors.

3.5 Results and Discussion

In this section, we present results from the estimation of the linear-in-means model and the spatial lag model from the previous section. We compare and contrast the estimation results of both models on the short panel and the long panel.[4] In addition to the direct effects of the village adoption rate in the presence of the two types of drought shocks, we are also interested in the interaction between the network adoption rate and the drought shocks. For the linear-in-means models, we compare estimation results from the producer fixed effects specification in Equation (3.2) and estimation results from the first-difference specification in Equation (3.3). For the spatial lag model in Equation (3.4), we compare estimation results from three different specifications of the weighting matrix.

3.5.1 Linear-In-Means Results

Here we present results from estimating Equation (3.2) and Equation (3.3) on the short and the long panel for entry into the coffee cooperation and entry into the honey cooperative.

Since Equation (3.2) uses producer fixed effects and Equation (3.3) uses first differences, we can compare the long-term impact of drought and the overall village adoption rate of both cooperatives with the short-term impact of drought and the village adoption rate of both cooperatives in the previous year.

Moreover, since we estimate both specifications on a short panel with 11 years of entry decisions and a long panel with 22 years of entry decisions, we can compare the estimation results for two time horizons. In particular, we argue that the results from the short panel capture a diffusion process in progress, and the results from the long panel capture the same diffusion process from start to finish.

Entry into Coffee Cooperative

Table 3.3 presents the results for the entry into the coffee cooperative. First, we focus on columns 1 and 2, which use producer fixed effects. A 10% increase in the previous time period of the overall membership rate of the village increases by 5% the probability that a given producer will join the coffee cooperative in the present time period.

Column 1 gives the effect of periods of severe drought and extreme drought on entry in the short panel. Here, each additional period of severe drought adds 4% to the probability that a

[4] Recall that short panel contains 11 years of entry decisions and the long panel contains 22 years of entry decisions.

producer will join the coffee cooperative. The additional network effect could eliminate this effect. On the other hand, each additional period of extreme drought decreases by 6% the probability that a producer will join. The additional network effect could be up to 4% in the opposite direction. One possibility is that the network corrects a producer's initial belief that the coffee cooperative will not help in situations of extreme drought.

Column 2 presents the same results estimated over the long panel. The point estimate of the network effect is very similar to the network effect in the short panel, as is the effect of extreme drought and the interaction between the network and extreme drought. The main difference between the estimation results on the short panel and the estimation results on the long panel is in the coefficient of the interaction effect between the network strength and the presence of one or more periods of extreme drought. In the long panel, the sign of this effect is positive, instead of negative in the short panel. One possibility is that the network reinforces a producer's initial belief that the coffee cooperative will not help in situations of extreme drought.

Next, we turn to columns 3 and 4, which estimate first difference versions of our model on the short panel. Due to the first differencing, the dependent variable Δy^c_{ijt} takes the value 1 only in the year when the producer joins the coffee cooperative. On the right-hand side, the first-differencing collapses the independent variables in the same way. The village network covariate ΔN^c_{ijt-1} is just the share of producers in the producer's reference group (village) who joined the cooperative in the previous year, not the total share of producers in the reference group who joined the cooperative up until the current year. Similarly, the drought covariate ΔD_{ijt-1} is an increase (or decrease) in the number of periods of extreme or severe drought from the previous year.

Thus, the point estimates in the estimation results for the first difference specifications capture responses to shocks and not trends. In column 3, the point estimates of Equation (3.3), are nonzero and statistically significant. If 10% of a producer's village joins the coffee cooperative in a given year, then there is a 1% chance that a producer will also join the cooperative. A period of severe drought in a given year increases the probability of joining by 2.5%. The network effect can mitigate this probability of joining by as much as 4%. In the short panel, we do not see an immediate response to extreme droughts. In column 4, which uses the long panel, we do not see a network effect at all. Only the effects of periods of drought remain. An additional period of extreme drought decreases the probability that a producer joins by 1.5%. The village network could potentially reverse that effect.

We summarize the estimation results from the linear-in-means models as follows. The strength of the village network affects a producer's decision to join the cooperative in both the short panel and the long panel. The effect comes from not only the immediate decisions of other producers in the same village to join the cooperative in the previous year, but also the cumulative decisions of other producers in the same village to join the cooperative up until the present year. Extreme drought discourages entry into the coffee cooperative and the village network reduces this effect. In the short panel, severe drought encourages entry into the coffee cooperative, and the village network also reduces this effect. One possible explanation is that the village network updates the producers' beliefs about whether cooperative membership is beneficial against severe drought and extreme drought.

Entry into Honey Cooperative

Next we turn to the estimation results of the linear-in-means model on entry into the honey cooperative. Once again, we estimate two specifications on both the short panel and the long panel. The specifications differ in that they use two different methods to control for producer-level unobservables: producer fixed effects and first differences. Table 3.4 presents the results.

First, we turn to column 1, which estimates Equation (3.2) on the short panel. Recall that this specification uses producer fixed effects. A 10% increase in the strength of the village network causes a 5% increase in the probability that a producer will join the honey cooperative. We find no effect of periods of severe drought or periods of extreme drought, either on their own or interacted with network strength.

Column 2 presents the results of estimating the same equation on the long panel. The network effect decreases slightly. A 10% increase in the strength of the village network causes a 4% increase in the probability that a producer will join the honey cooperative. Once again, periods of drought on their own do not affect entry, but in the presence of the network, a period of extreme drought increases the probability of entry by 5.1% and a period of severe drought decreases it by 3.9%. Perhaps in the first 11 years of the cooperative the network was less active in periods of severe drought and more active in periods of extreme drought.

Now we turn to estimation results from Equation (3.3) in columns 3 and 4. As we noted in the previous section, these specifications use first differences and thus capture the immediate effect of the village adoption rate of the the honey cooperative in the prior year on the probability of a producer joining the honey cooperative in a given year. In both columns, the effect of the village network is similar. A 10% increase in the village adoption rate of the honey cooperative is associated with a 1.2% increase in the short term or a 0.7% increase in the long term that a producer in the same village will join the honey cooperative. Neither point estimate is statistically significant. Wee see no effect of periods of drought, either on their own or interacted with network strength.

3.5.2 Spatial Lag Results

Entry into Coffee Cooperative

Table 3.5 presents the results of the estimation of Equation (3.4) on the short panel and the long panel with each of the three weighting matrices described in Section 3.4: a binary continguity matrix of other villages in the same region, an inverse distance matrix of other villages in the same region, and an inverse distance matrix of all other villages across all regions.

In all six columns, we see a direct effect of network strength on a producer's decision to join the coffee cooperative. The effect size ranges from 30% to 40%. Like the network effect in the linear-in-means model, we interpret this coefficient to mean that a 10% increase in the membership rate in a producer's village is associated with an increase of 3% to 4% in the probability that the producer will join the coffee cooperative. The direct effect of an additional period of severe drought is only associated with entry in column 2, which shows the estimation results for the short panel with the regional inverse distance weighting matrix. In contrast, the direct effect of an additional period of extreme drought is associated in all columns but column 2 with a 2% and a 4% decrease in the probability that a producer will join the coffee cooperative. In most columns, this result is statistically significant at the 10% level. This coefficient has the same sign and magnitude as the corresponding coefficient in the linear-in-means results in Table 3.3.

We next move to the indirect network effects. The effect sizes across the specifications are not directly comparable because the weighting matrices are normalized in different ways. The binary contiguity weighting matrix in columns 1 and 4 is row normalized, while the inverse distance matrices in columns 2, 3, 5, and 6 are normalized by the largest eigenvalue in each matrix. However, we see a substantial indirect effect of the network in the six columns, larger than the direct effect. Thus, we confirm the presence of spatial spillovers.

Finally, we examine the indirect effects of drought. Columns 1 and 4 show that an additional period of severe drought or extreme drought in another village in the same region increases the probability that a producer will join the coffee cooperative by 2% - 5%. The effect sizes increase in columns 3 and 6, which use weighting matrices that take into account all villages in the study area.

Finally, we use the log-likelihood score at the bottom of the table to compare the specifications of the three weighting matrices for the two panels. For the short panel, the specification in column 3 with all villages fits the data better. For the long panel, the specification in column 5 with only villages in the same region fits the data better. This difference may indicate that as the coffee cooperative spread, initially only peer effects in the same region mattered but then peer effects in the whole area of study became more important. This model of diffusion reflects the descriptive trends that we saw in Table 3.2.

Entry into Honey Cooperative

Table 3.6 presents the results of the estimation of Equation (3.4) on the short panel and the long panel with each of the three weighting matrices described in Section 3.4.

As in the estimation results for the entry into the coffee cooperative in the previous section, we see a strong network effect of nearly 50% in the short panel and 40% in the long panel. We see little direct effect of periods of severe or extreme drought.

The magnitude and sign of the indirect effect of the network varies depending on the specification of the weighting matrix and the length of the panel. We first consider the short panel. Column 1 shows a positive network effect with the binary contiguity matrix. For columns 2 and 3, incorporating the inverse distances at the regional and all village level causes the effect size to change sign. This paradoxical result reflects the fact that the honey producers in a given region are typically concentrated in one village and that across the area of interest the honey producers are concentrated in a few regions. Thus, the entry decision of a given producer is inversely related to the entry decision of a producer not in the same village or region.

Next, we move to the long panel. For column 4, the indirect effect of the network is positive and significant, as in column 1 in the short panel. Using the regional inverse distance weighting matrix in column 5 eliminates the indirect effect. Using the inverse distance weighting matrix with all villages in column 6 brings it back even more strongly. Thus, in the long term, the honey cooperative is spreading throughout the region of interest.

At the village level, in the short term, both periods of severe drought and periods of extreme drought are associated with the entry into the honey cooperative. These associations are not present in the long term.

Finally, we examine the log-likelihood values at the bottom of the table. For both the short panel and the long panel, neither the regional or all-village inverse distance weighting matrices improves the model fit over the binary contiguity one.

3.5.3 Limitations

Both classes of our models suffer from limitations. Since our network graph is undirected, causal identification of the linear-in-means results is threatened by time-invariant shocks that affect producers in the same village or region. For example, improvement of roads or the destruction of a key bridge could affect a producer's entry decision through the channel of market access. As Bramoullé et al. (2009) describe, one way to control for these shocks is to instrument neighbors with neighbors-of-neighbors. However, this approach works only with directed graphs.

Along these lines, we note another important assumption for causal identification: that the network structure is stable and exogenous. The length of the time periods in question raises questions about that assumption, though a unique feature of this setting is that the indigenous tend to stay in the same place that their families have inhabited for generations.

In addition, our use of producer fixed effects and first differences represents two extremes: assuming stable individual heterogeneity across 11 (or 22) periods or only using variation from the previous period. In the real world, a substantial portion of producer fixed effects probably hold stable for a "Goldilocks mean" of five periods or so.[5] This limitation affects our spatial lag model as well, since we estimate it with producer fixed effects but not with first differences.

Finally, we choose to measure drought as two discretized periods of **severe drought** and **extreme drought** instead of a continuous variable of rainfall and temperature as in other studies in the climate shock literature described by Dell et al. (2014). We use the two SPEI categories because the SPEI index combines the magnitudes of deviation from the mean of temperature and rainfall and not just level effects. At the same time, the opposite signs of the coefficients in the severe and extreme drought raise the question of whether this discretization is artificial. The producers in our population experience the weather in a continuous way.

3.6 Conclusion

In this paper, we have analyzed the effect of temporal and spatial lag on entry into two cooperatives that help smallholder producers in the face of one type of climate shock, seasonal drought. We are not the first to examine the determinants of contract farming, cooperative entry, or technology adoption. However, we bring a new approach. We apply two econometric methods—the linear-in-means model from the peer effects literature and the spatial-lag model from the regional science literature —in a novel way with a uniquely rich data set to analyze the way these technologies have diffused through an extremely isolated population over time and space.

Crucially, we do not make the unrealistic assumption of the absence of spillovers but explicitly take them into account and turn them into an object of study. We find differences in the adoption patterns of the coffee cooperative and the honey cooperative. They differences give insight into how information about the coffee cooperative and honey cooperative is transmitted across time and space in the presence of climate shocks.

Our results give insight into similar contexts in the developing world. It takes several years to learn about a new technology, whether it is learning by doing or learning from others. Moreover, it takes time for information about a new technology to be transmitted across space. Policy makers continue to lament the low uptake of many welfare-improving technologies. They and their

[5] The reference to the short story "Goldilocks and the Three Bears" by Robert Southey here refers to a quantity that is neither too short nor too long but just right.

implementing partners would do well to consider these temporal and spatial lags as they promote them.

3.7 Exhibits

Figure 3.1: Entry Network of Coffee Cooperative

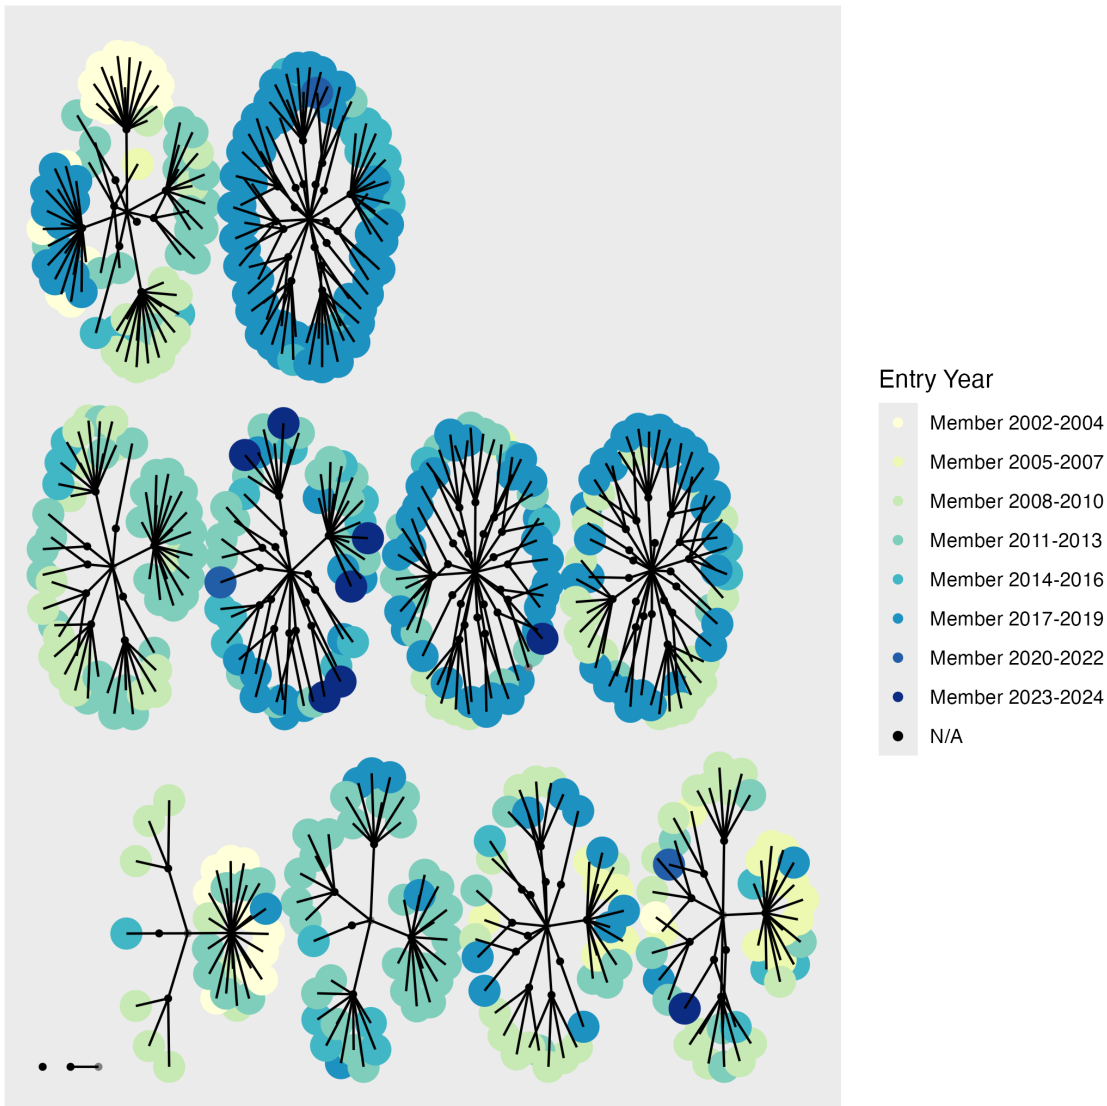

This figure displays the network of coffee cooperative members. Nodes indicate coffee producers. Colors indicate the year of entry. Nodes are grouped by village and then villages are grouped into regions. Black nodes are placeholders to position producers.

Figure 3.2: Entry Network of Honey Cooperative

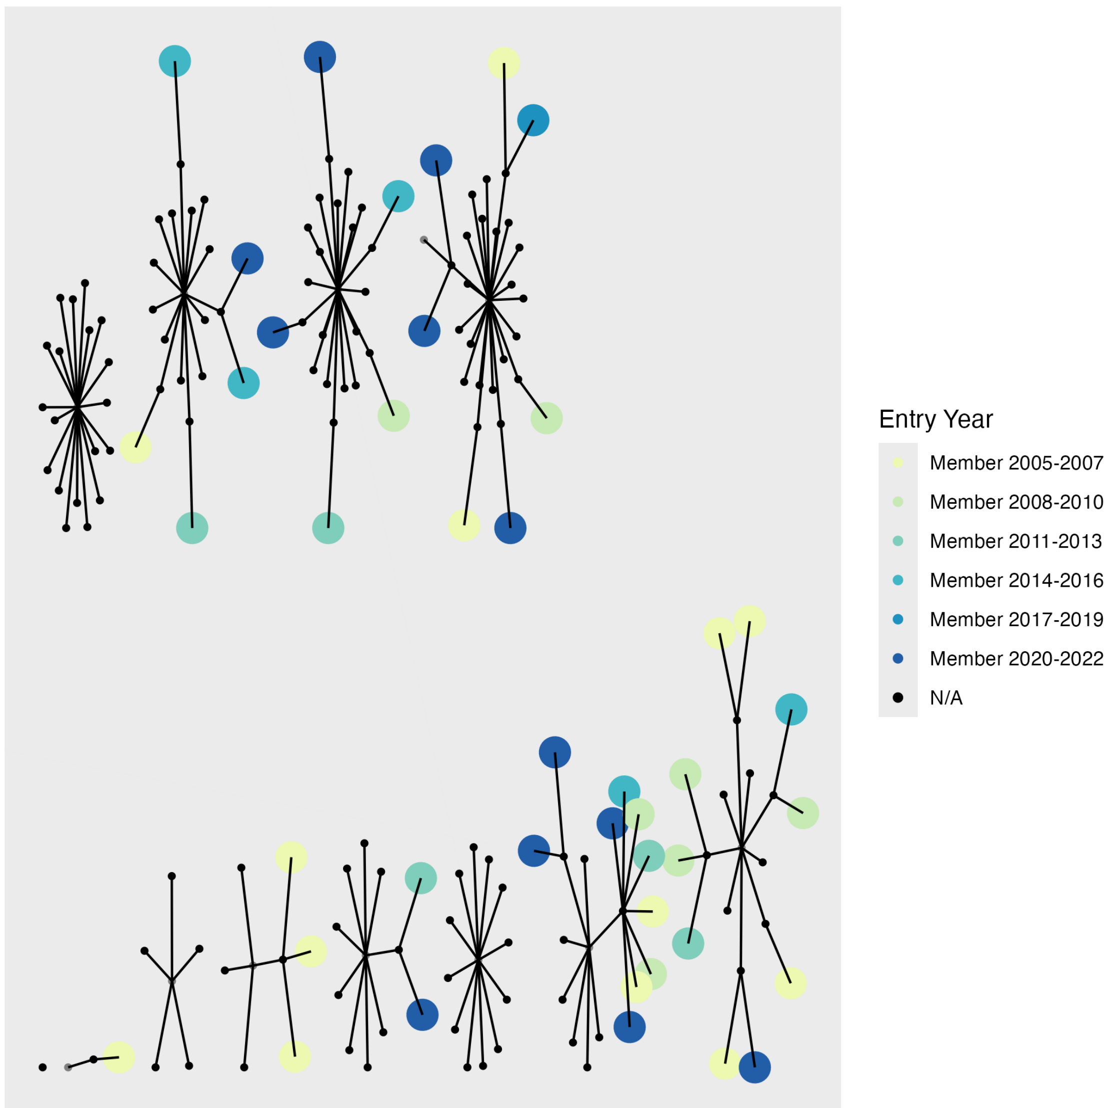

This figure displays the network of honey cooperative members. Nodes indicate honey producers. Colors indicate the year of entry. Nodes are grouped by village and then villages are grouped into regions. Black nodes are placeholders to position producers.

Figure 3.3: Elevation of Producers

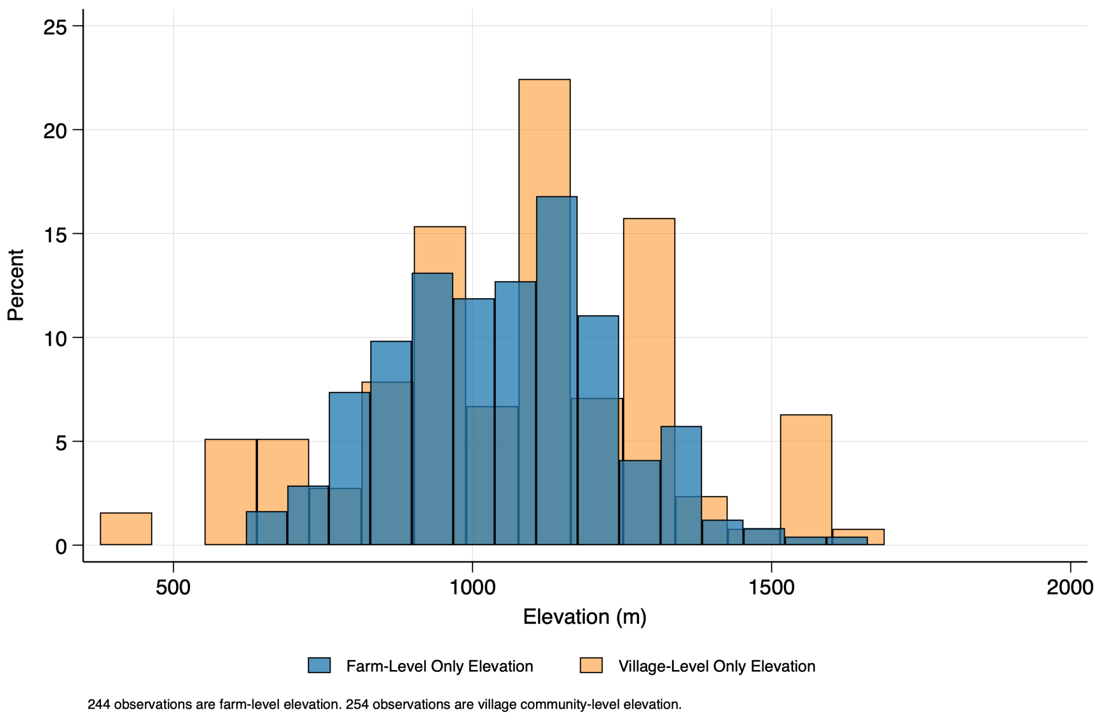

This figure displays the distribution of the elevation of the universe of producers. For the subsample of producers whose coffee plots have been geolocated, the foreground bars indicate the plot elevation. For the remaining producers, the background bars indicate the village elevation.

Figure 3.4: Monthly Variation in SPEI by Year

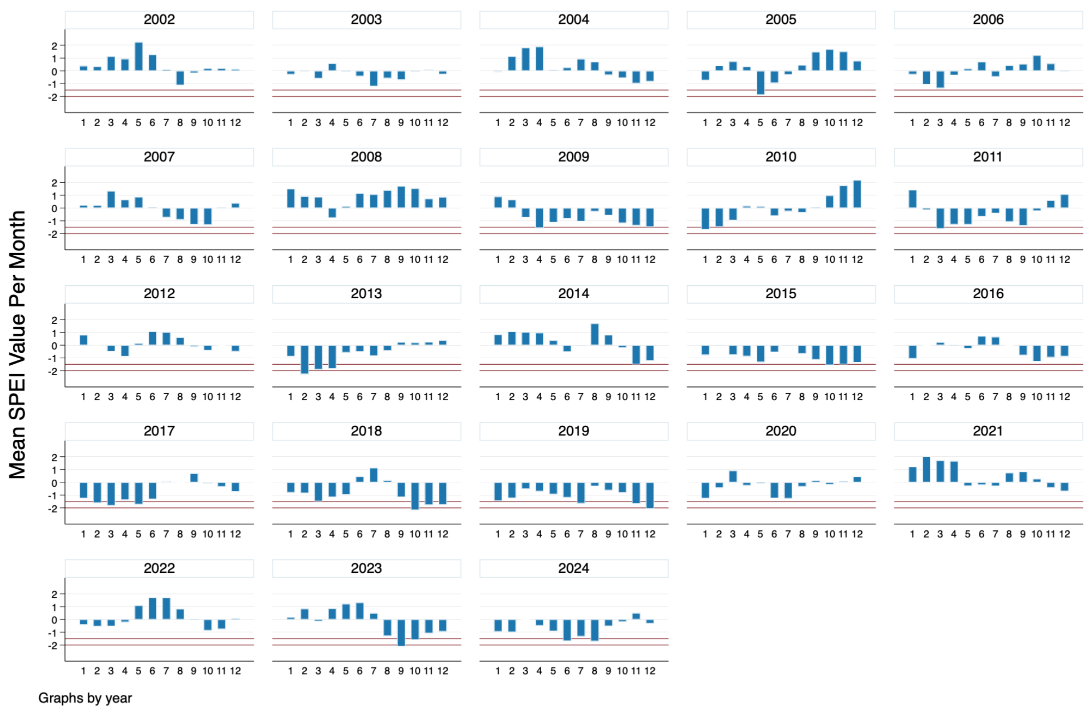

This figure displays the mean value of the SPEI (Standardised Precipitation-Evapotranspiration Index) by month from 2002-2024 for the survey region. Vicente-Serrano et al. (2010) gives more information about the data. Horizontal lines show the thresholds for severe drought (below -1.5) and extreme drought (below -2).

Figure 3.5: Coffee Cooperative Entry and Drought by Year

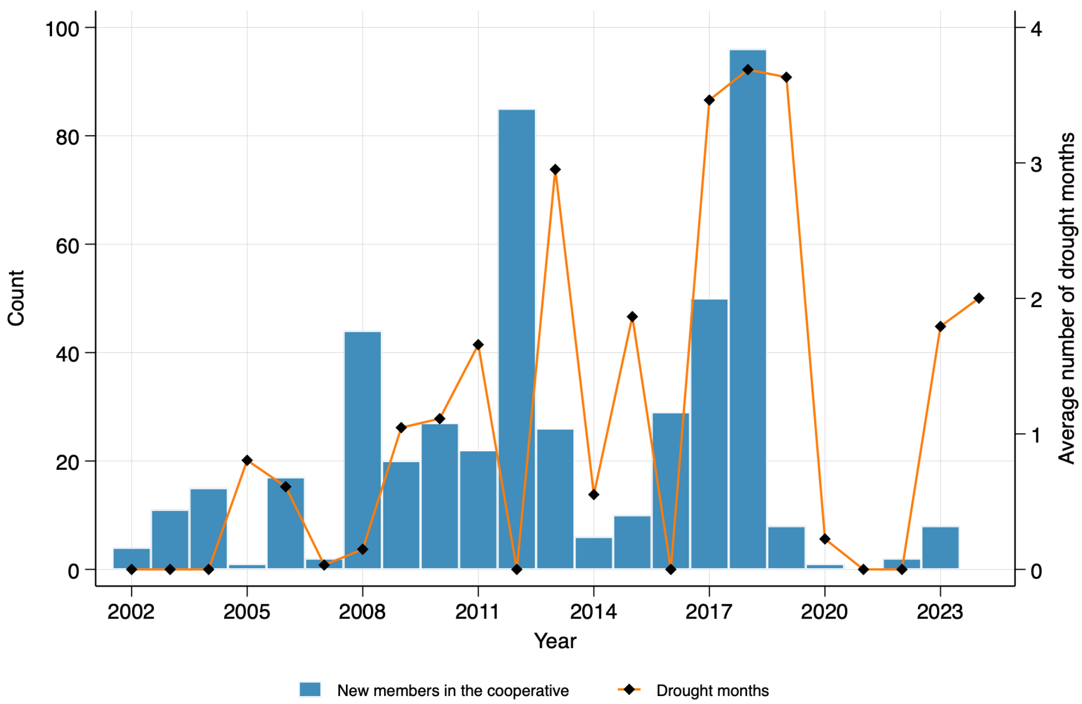

This figure displays the number of producers who entered the Batsil Maya coffee cooperative each year and the number of three-month periods of severe or extreme drought based on the SPEI.

Figure 3.6: Honey Cooperative Entry and Drought by Year

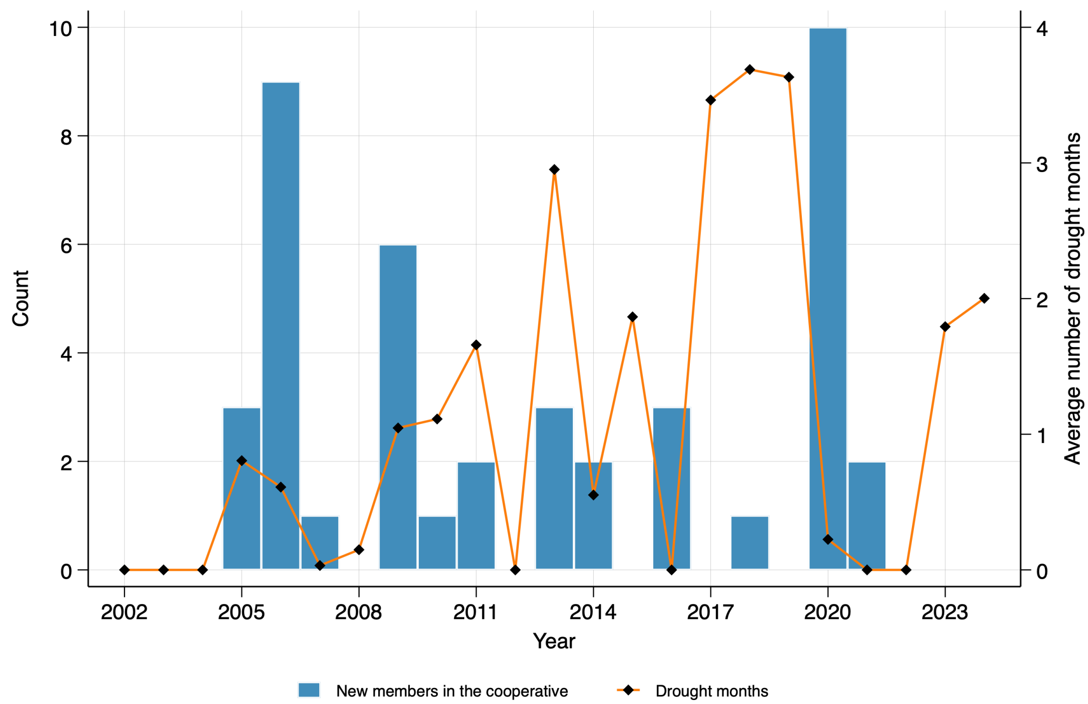

This figure displays the number of producers who entered the Chabtic honey cooperative each year and the number of three-month periods of severe or extreme drought based on the SPEI.

Table 3.1: Village-Level Entry

			Coffee			Honey		
	Region	Villages	2002	2013	2024	2005	2013	2024
	1	5	0	4	5	0	0	0
	2	15	0	7	15	1	2	4
	3	9	0	9	9	0	1	1
	4	9	1	7	8	0	5	5
	5	4	0	3	4	1	1	1
	6	22	0	9	21	1	3	5
	7	9	0	7	8	0	2	3
	8	20	0	8	19	0	2	5
	9	19	0	1	19	0	0	0
	10	12	0	8	12	0	0	0
Total	—	124	1	63	120	3	16	24

This table shows the number of villages in each region with members of the coffee and honey cooperatives at three time periods: the first year, a middle year, and the latest year in the dataset.

Table 3.2: Individual-Level Entry

			Coffee			Honey		
	Region	Individuals	2002	2013	2024	2005	2013	2024
	1	39	0	29	39	0	0	0
	2	49	0	18	46	1	2	5
	3	52	0	49	52	0	1	2
	4	72	4	53	67	0	8	10
	5	28	0	26	28	1	3	3
	6	51	0	18	50	1	3	7
	7	43	0	34	42	0	6	11
	8	58	0	21	54	0	2	5
	9	71	0	1	71	0	0	0
	10	35	0	25	35	0	0	0
Total	—	498	4	274	484	3	25	43

This table shows the number of individuals in each region with members of the coffee and honey cooperatives at three time periods: the first year, a middle year, and the latest year in the dataset.

Table 3.3: Linear-In-Means Estimates for Entry into Coffee Cooperative

	Joins Coffee Cooperative (1=Yes)			
	FE Short	FE Long	FD Short	FD Long
	(1)	(2)	(3)	(4)
Network Strength (Village)	0.533***	0.524***	0.104***	0.003
	(0.038)	(0.028)	(0.029)	(0.014)
Periods of Severe Drought	0.039***	0.019***	0.025***	0.004
	(0.008)	(0.006)	(0.007)	(0.005)
Periods of Extreme Drought	−0.063***	−0.019	−0.002	−0.015*
	(0.023)	(0.014)	(0.012)	(0.008)
Severe Drought x Network Strength	−0.040***	−0.031***	−0.038***	−0.017***
	(0.007)	(0.006)	(0.007)	(0.005)
Extreme Drought x Network Strength	0.048***	−0.036***	−0.001	0.012*
	(0.018)	(0.014)	(0.011)	(0.006)
Year Fixed Effects	Y	Y	Y	Y
Producers	498	498	498	498
Years	11	22	11	22
Observations	5,478	10,956	4,980	10,458
Adjusted R^2	0.078	0.124	0.100	0.065

The dependent variable is a dummy that indicates whether a producer joined the cooperative in a given year.
Columns 1 and 2 use producer fixed effects to control for unobserved heterogeneity.
Columns 3 and 4 use first differences to control for unobserved heterogeneity.
Columns 1 and 3 use a short 11 year panel. Columns 2 and 4 use a long 22 year panel.
Network Strength is the share of producers in the same village that joined the cooperative by the previous year.
Periods of Severe Drought (-2 <SPEI <= 1.5) and Periods of Extreme Drought (SPEI <= -2) are calculated by matching a producer's coffee plot or village and SPEI drought data from the previous year.
Standard errors are heteroskedasticity-robust.

Table 3.4: Linear-In-Means Estimates for Entry into Honey Cooperative

	Joins Honey Cooperative (1=Yes)			
	FE Short	FE Long	FD Short	FD Long
	(1)	(2)	(3)	(4)
Network Strength (Village)	0.496***	0.411***	0.124	0.068
	(0.127)	(0.101)	(0.123)	(0.068)
Periods of Severe Drought	−0.0004	−0.002	−0.001	0.0005
	(0.002)	(0.002)	(0.001)	(0.001)
Periods of Extreme Drought	0.003	0.002	0.002	0.004*
	(0.006)	(0.004)	(0.004)	(0.003)
Severe Drought x Network Strength	−0.008	−0.039***	0.007	0.002
	(0.010)	(0.014)	(0.006)	(0.005)
Extreme Drought x Network Strength	0.008	0.051***	−0.001	−0.004
	(0.015)	(0.018)	(0.005)	(0.005)
Year Fixed Effects	Y	Y	Y	Y
Producers	498	498	498	498
Years	11	22	11	22
Observations	5,478	10,956	4,980	10,458
Adjusted R^2	0.019	0.016	0.017	0.008

The dependent variable is a dummy that indicates whether a producer joined the cooperative in a given year.
Columns 1 and 2 use producer fixed effects to control for unobserved heterogeneity.
Columns 3 and 4 use first differences to control for unobserved heterogeneity.
Columns 1 and 3 use a short 11 year panel. Columns 2 and 4 use a long 22 year panel.
Network Strength is the share of producers in the same village that joined the cooperative by the previous year.
Periods of Severe Drought (-2 <SPEI <= 1.5) and Periods of Extreme Drought (SPEI <= -2) are calculated by matching a producer's coffee plot or village and SPEI drought data from the previous year.
Standard errors are heteroskedasticity-robust.

Table 3.5: Spatial Lag Estimates for Entry into Coffee Cooperative

	Joins Coffee Cooperative (1=Yes)					
	(1)	(2)	(3)	(4)	(5)	(6)
Direct Effect - Network	0.303***	0.347***	0.399***	0.363***	0.359***	0.397***
	(0.024)	(0.023)	(0.022)	(0.014)	(0.014)	(0.013)
Direct Effect - Severe Drought	0.002	0.019**	0.001	−0.005	0.006	−0.010
	(0.008)	(0.009)	(0.008)	(0.007)	(0.008)	(0.007)
Direct Effect - Extreme Drought	−0.027	−0.004	−0.037*	−0.029*	−0.021	−0.032*
	(0.020)	(0.022)	(0.021)	(0.017)	(0.018)	(0.017)
Indirect Effect - Network	0.594***	0.454***	0.673***	0.433***	0.425***	0.757***
	(0.034)	(0.033)	(0.060)	(0.022)	(0.019)	(0.039)
Indirect Effect - Severe Drought	0.023**	−0.002	0.064***	0.024***	0.006	0.056***
	(0.010)	(0.010)	(0.018)	(0.009)	(0.008)	(0.015)
Indirect Effect - Extreme Drought	0.054**	0.001	0.166***	0.016	−0.002	0.038
	(0.025)	(0.027)	(0.049)	(0.020)	(0.020)	(0.038)
Weighting Formula	Binary	Inv Dist	Inv Dist	Binary	Inv Dist	Inv Dist
Other Villages	Region Only	Region Only	All	Region Only	Region Only	All
Producer and Year Fixed Effects	Y	Y	Y	Y	Y	Y
Producers	498	498	498	498	498	498
Years	11	11	11	22	22	22
Log-Likelihood	733	661	602	-154	-88	-180
Observations	5,478	5,478	5,478	10,956	10,956	10,956
Adjusted R^2	0.656	0.647	0.639	0.746	0.749	0.745

The dependent variable is a dummy that indicates whether a producer joined the cooperative in a given year.
Columns 1, 2, and 3 use a short 11 year panel. Columns 4, 5, and 6 use a long 22 year panel.
Network Strength is the share of producers in the same village that joined the cooperative by the previous year.

Table 3.6: Spatial Lag Estimates for Entry into Honey Cooperative

	Joins Honey Cooperative (1=Yes)					
	(1)	(2)	(3)	(4)	(5)	(6)
Direct Effect - Network	0.494***	0.498***	0.492***	0.390***	0.389***	0.386***
	(0.062)	(0.062)	(0.062)	(0.037)	(0.037)	(0.037)
Direct Effect - Severe Drought	−0.002	−0.008*	−0.011*	−0.003	0.0002	0.002
	(0.003)	(0.004)	(0.006)	(0.003)	(0.004)	(0.004)
Direct Effect - Extreme Drought	−0.0003	−0.007	−0.020	−0.005	0.011	0.006
	(0.006)	(0.012)	(0.015)	(0.005)	(0.009)	(0.010)
Indirect Effect - Network	0.149**	−0.062*	−0.265***	0.089**	0.002	0.234***
	(0.065)	(0.032)	(0.062)	(0.044)	(0.038)	(0.089)
Indirect Effect - Severe Drought	0.003	0.012**	0.036**	0.002	−0.004	−0.010
	(0.003)	(0.006)	(0.016)	(0.003)	(0.005)	(0.011)
Indirect Effect - Extreme Drought	0.009	0.017	0.081*	0.014**	−0.013	−0.007
	(0.007)	(0.017)	(0.049)	(0.007)	(0.012)	(0.032)
Weighting Formula	Binary	Inv Dist	Inv Dist	Binary	Inv Dist	Inv Dist
Other Villages	Region Only	Region Only	All	Region Only	Region Only	All
Producer and Year Fixed Effects	Y	Y	Y	Y	Y	Y
Producers	498	498	498	498	498	498
Years	11	11	11	22	22	22
Log-Likelihood	6106	6116	6137	8004	8003	8015
Observations	5,478	5,478	5,478	10,956	10,956	10,956
Adjusted R^2	0.895	0.895	0.896	0.695	0.694	0.695

The dependent variable is a dummy that indicates whether a producer joined the cooperative in a given year.
Columns 1, 2, and 3 use a short 11 year panel. Columns 4, 5, and 6 use a long 22 year panel.
Network Strength is the share of producers in the same village that joined the cooperative by the previous year.

Chapter 4

Where You Go Depends on Who You Know: Social Networks as Determinants of Mexican Internal Migration

4.1 Introduction

Internal migration has overtaken birth and death as the primary source of demographic change at the regional level in many countries. In 2019, an estimated 763 million people worldwide lived outside the region in which they were born (UNESCO, 2018). Often internal migration plays an important role not only in population change but also in structural change, especially in the form of rural to urban migration. Yet despite the size and significance of this phenomenon, two important puzzles remain: who migrates and where?

In practice, internal migrants do not simply maximize the present value of two competing income streams, as Roy (1951) predicted. Nor does the introduction of self-selection by Borjas (1987) offer much clarity. In fact, breaking down migration trends by age and education only adds to the puzzle. In some situations, internal migration exhibits positive selection: the people who leave have the most to gain (Chiquiar & Hanson, 2005). In other situations, internal migration exhibits negative selection: the people who leave have the most to lose (Ibarraran & Lubotsky, 2007). These apparently contradictory results suggest that factors other than potential income gain play an important role in internal migration decisions.

Recent work has proposed that social networks at the destination may explain these seemingly contradictory results. McKenzie and Rapoport (2010) show that migrants are negatively selected in communities with a high proportion of outmigration and positively selected in communities with a low proportion of outmigration. Their work complements two other well-known studies of the effects of social networks on migration. Carrington et al. (1996) develops a model of endogenous moving costs to explain the internal migration destination choices of African-Americans in the 20th century of the United States. Munshi (2003) documents network effects in US-Mexico migration. In both cases, migrants in one time period tend to follow migrants from the same origin in previous time periods instead of simply seeking out the highest wages.

The present essay finds that social networks are positively associated with internal migration flows in Mexico from origin states to destination municipalities for men aged 25 to 55 from the Mexican population census in three recent five year intervals (1995-2000, 2005-2010, and 2010-2015). I examine this group of working age men as a sample that would be prone to migrate for work instead of education or family reunification. As a proxy for these men's social networks, I use the total number of people from the same origin state who migrate to the same destination municipality over the previous five year interval. The absence of flows from approximately 75% of the possible origin-destination pairs creates left censoring on the dependent variable. For this reason, I model the extensive and intensive margin of the men's internal migration as two distinct processes. The extensive margin refers to the opening, remaining open, or closing of potential migration corridors across the three time periods of interest. The intensive margin refers to the magnitude of the migrant that pass through these corridors at a particular time period.

At both margins, I estimate three models on these internal migration flows: a model with wage differences alone, a structural gravity model with wage differences, and a structural gravity model with wage differences and the social network proxy above. Adding additional covariates greatly increases the predictive power of the model. For a representative time period (1995-2000) at the extensive margin, differences in base wage and return to skill explain 2.5% of the variation at the extensive margin; a structural gravity model 26%; and a structural gravity model with social networks 39%. For the same time period at the intensive margin, differences in base wage and return to skill explain 5% of the variation; a structural gravity model 39%; and a structural gravity model with social networks 58%.

In addition, adding these additional factors first reduces and then eliminates the effect of the wage differences along a corridor on internal migrant flows. In the first model, wage differences are significantly associated with migration. In the second model, the association with wage differences decreases substantially but remains significant. In the third model, it decreases even more and loses significance. Moreover, the magnitude of the association increases across the three time periods. At the extensive margin, a 1% increase in the size of the social network is associated with a 4.6%, 11.9%, and 12.6% higher probability of the presence of a migration corridor, respectively, in the three time periods. At the intensive margin, the social network elasticities are 19%, 30.4%, and 31.9%. These empirical results show the importance of social networks in driving internal migration both at the extensive and intensive margins for the time periods in question.

Four aspects of this essay warrant further explanation because of their novelty. First, the novelty and the scale of my data source stand out as the first use of multiple waves of nationally-representative census data to examine internal migration with a structural gravity model augmented by social networks. The original article on the role of diasporas in international migration estimated a structural gravity model using cross-sectional data from 195 origin countries to 30 destination countries in the Organization for Economic Cooperation and Development (OECD) from 1990-2000 (Beine et al., 2011). Most previous work in Mexico and elsewhere has used small-scale panel data household surveys like the Mexican Migration Project, the Mexican Family Life Survey, or the National Survey of Rural Households (Cuecuecha & Pederzini, 2014; Durand & Massey, 2019). This work suggests that social factors have come to dominate economic factors over time (Asad & Garip, 2019). Even previous studies that estimated gravity models on Mexican census data did not include social networks in their models (Ochoa et al., 2018; Soloaga, Isidro et al., 2010). Thus this essay applies a frontier model to a large-scale data set: three waves of nationally representative Mexican census data.

Second, I follow Beine et al. (2011) in integrating the canonical models of migration from microeconomics and macroeconomics into a unified framework. In the applied microeconomics models of Roy (1951) and Borjas (1987), potential migrants consider expected wage gain at the destination. In the present case and others, these models do not often fit the data well. In contrast, applied macroeconomics models the movement of people from one region to another in terms of the relative populations of the two regions and the distance between them using the same gravity model that it uses to model the movement of goods (Anderson, 2011). Gravity models often fit the data well but lack theoretical underpinning. The integrated model in the present essay begins with an individual considering two destinations; uses a multinomial logit to extend the choice to an arbitrary number of destinations; and then considers the probability of migration to each destination at the population level to estimate the migration flow from the origin to each destination. I include destination population, distance, indigenous share at origin and destination, and urban share at origin and destination as factors that affect migration flows. Because this structural gravity model is still relatively new in the literature, I offer a simplified derivation in section 2 based on Beine et al. (2016).

Third, I use auxiliary Mincer regressions at the state and municipality level to estimate the usual labor market parameters: base salary, return to skill, and return to experience. I drop return to experience because of the small magnitude. Instead of considering a representative worker (i.e. a 25 year old internal migrant with an elementary school education), and generating an average wage at origin and destination like Falaris (1987), I difference the base salary and return to skill parameters at the origin and destination to account for heterogeneity across geographical regions in the population of potential migrants. Intuitively, less educated rural to urban migrants may be influenced more by base salary differences, whereas more educated urban to urban migrants may be influenced more by differences in return to skill. These differences go into the model as proxies for wage differences between the origin and destination. To my knowledge, I am the first person to model counterfactual wage differences in this way.

Fourth, I use a new approach to the econometric challenge of identification of network effects.

Here I use past migration flows as a proxy for social networks and examine their association with present flows. One concern is the possibility of serial correlation across time periods: unobserved time-invariant factors like individual preferences or cultural proximity that would affect both past and present migration flows along the same corridor. Previous literature has used rainfall shocks at the origin (Munshi, 2003) or railroads (Woodruff & Zenteno, 2007) as instruments to reduce the bias from this potential serial correlation. Instead, I use the presence of a migration flow in a given corridor in 1960 to control for "the taste for migration," non-economic factors that could influence the migration flow along a corridor. I justify the use of this control for both practical as well as theoretical reasons. The earliest Mexican census that asks about migration is from 1960. At this time, the structure of the Mexican economy was very different from the present day in several important ways. 1960 predates the end of the Bracero program of US agricultural visas (1964), the beginning of the maquiladora export manufacturing program (1964), the entry of Mexico into the General Agreement on Tariffs and Trade (1986), and the entry of Mexico into NAFTA (1994). Thus any migration corridors present in 1960 are due to either time-invariant cultural factors or economic factors proper to that period but not the time periods in question from 1995-2015. I argue that controlling for this taste for migration allows me to model only time-varying migration trends and estimate the effect of social networks on this migration. The inclusion of the migration taste factor decreases the associations and elasticities that I find above by at most 10%. This result suggests that the association of the social networks and migration flows does not come from unobserved time-invariant factors.

The results here matter not only in academic circles but also to policymakers who seek to accurately understand present and future migration trends. In the short term, receiving communities integrate new migrants into existing housing, jobs, schooling, and other programs. In many cases, NGOs assist with this integration. In the long term, destinations plan to adjust their infrastructure to account for future internal migration. Private enterprises as well benefit from information about future labor supply, since the population I study in this essay primarily moves for work reasons.

The essay proceeds as follows. Section 2 develops a theoretical model of migration flows from micro foundations to a structural gravity model with social networks. Section 3 gives additional background on internal migration and describes the data from the Mexican census. Section 4 describes the empirical method, including the Mincer models I use to estimate differences in base wage and return to skill. Section 5 presents the results: determinants of migration and social network elasticities at the extensive and intensive margin. Section 6 concludes.

4.2 Theoretical Framework

This essay uses a structural gravity model that unifies the canonical models of migration from microeconomics and macroeconomics. In this section, first I offer a conceptual overview of the simplifying assumptions required to consider migration in terms of expected wage gain, the monetary and non-monetary factors included in the cost of migration, and the proposed effect of social networks on reducing these costs. Next, I review the microeconomic migration model developed by Borjas (1987) and his discussion of positive and negative selection. Finally, I summarize the structural gravity model from Beine et al. (2011), which provides a bridge from a microeconomic model like that of Borjas to the gravity model that applied macroeconomics uses to analyze trade and migration flows. A key element of this essay's structural gravity model is the inclusion of social networks that reduce the cost of migration.

4.2.1 Conceptual Framework

Lucas (2021) describes the literature on rural-urban migration with a particular focus on the factors that affect migration flows and the effects of migration on origin and destination. His taxonomy of migration allows me to clarify the type of migration I will examine in this essay: potential migrants migrate when the income at the destination is higher than the income at the origin, taking into account the cost of migration. I make the following assumptions about potential migrants and their migration decisions.

1. Potential migrants decide to migrate purely based on economic reasons. This excludes other forms of migration, such as family reunification.

2. Potential migrants migrate permanently. This excludes seasonal migration or circular migration.

3. Potential migrants enjoy certain wages at the origin. This excludes the impact of uncertainty around agricultural production or potential risk aversion.

4. Potential migrants aspire to formal employment at the destination. This excludes informal employment at the destination.

5. Potential migrants always find jobs at the destination. This excludes the job search process or the possibility of unemployment or informal employment at the destination.

6. Potential migrants migrate based on a comparison between the income at the origin and the income at the destination.

A rich literature following Harris and Todaro (1970) has modeled rural-urban migration and the probability of obtaining a job in the formal sector at the destination. With the simplifying assumptions above, I sidestep this literature. Rather, the situations I consider align more with those considered by B. Banerjee (1991), who examines the case of migrants who migrate with a pre-arranged job. This approach matches that of Falaris (1987), who uses cross-sectional samples to pool movers and stayers across the possibility of multiple destinations within Venezuela, accounting for selection into migration but ignoring job search frictions and the existence of the informal sector.

Monetary gains not only come in terms of improved wages for the same occupation but also in terms of wage gains from occupational sorting. Previously Roy (1951) had proposed that individuals will choose the occupation that matches their endowment. He uses the example of occupational sorting into hunting and fishing. The endogeneity of this decision biases any attempt to estimate the effect of the occupation on the individual's income. Individual unobservables like ability could effect both the choice of occupation as well as the realized income.

Moreover, since a potential migrant considers lifetime earnings, different aged workers will approach the migration investment decision in different ways; for younger potential migrants, the potential return is larger than for older potential migrants, for example.

Within this framework of income comparison, potential migrants also consider the cost of migration. Sjaastad (1962) first introduced this notion of the cost of migration in his model, which treats migration as an investment in the migrant's own human capital. He considers two types of costs: monetary and non-monetary. Monetary costs include the cost of moving and the increase in cost of living at the destination. Non-monetary costs include the opportunity cost of lost wages while searching for a job and learning a new job; they also include the psychic cost of living away from family and friends.

Recent qualitative work has examined the role of social networks in migration, especially relative to these three costs. Garip and Asad (2016) outlines three channels by which social networks reduce the cost of migration: social facilitation, normative influence, and network externalities. Social facilitation refers to the way that past migrants reduce information frictions at the destination, reducing costs and increasing benefits, a phenomenon first identified by Yap (1977). Normative influence refers to the way that past migrants change social norms and make migration more attractive, also reducing the cost at the origin. Network externalities refer to how past migrants create a pool of common resources for future migrants.

Sociologists distinguish between the strong ties of family and friends and the "weak ties" of individuals from the same village or state (Granovetter, 1973). Davis et al. (2002) examines the effect of both types of ties on international and internal migration. In contrast, this essay only considers the effect of weak ties on internal migration.

4.2.2 Microeconomic Framework

Borjas (1987) offers a formal model of migration in terms of expected wage gain. Equations (4.1) and (4.2) decompose the expected wage for a resident of the origin or destination into an observed group mean and an unobserved disturbance. The indices 0 and 1 below indicate origin and destination. Intuitively, a given individual migrates when the gains, both observed or unobserved, outweigh the cost of migration: when the sign of the index function I below is positive.

$$\ln w_0 = \mu_0 + \epsilon_0 \tag{4.1}$$

$$\ln w_1 = \mu_1 + \epsilon_1 \tag{4.2}$$

$$I = \ln(w_1) - \ln(w_0) \tag{4.3}$$

Equation (4.5) introduces a cost of migration C. The ratio $\frac{C}{w_0}$ is the same for all individuals at the origin. The destination wages must exceed the origin wages plus the cost of migrating.

$$I > 0 \tag{4.4}$$

$$\ln(w_1) - \ln(w_0 + C) > 0 \tag{4.5}$$

$$\ln(w_1) - \ln(w_0) - \frac{C}{w_0} > 0 \tag{4.6}$$

$$(\mu_1 - \mu_0) + (\epsilon_1 - \epsilon_0) > \frac{C}{w_0} \tag{4.7}$$

Borjas' contribution is that not only the mean but also the variance of ϵ_i varies from the origin to the destination. He offers as one possibility a compressed distribution of $\epsilon_0 \sim N(0, \sigma_0^2)$ in an origin country with a low return to skill that expands to $\epsilon_1 \sim N(0, \sigma_1^2)$. Thus two potential migrants whose unobservable characteristics are nearly equivalent at the origin could see a larger difference at the destination, leading one to migrate and the other to stay. In fact, the location of an individual's unobserved characteristics in the distributions at the origin and the destination plays a key role in satisfying the migration condition above.

Borjas uses this model to consider the implications of different distributional assumptions of origin and destination unobservables.

1. Under positive selection, the best candidates migrate and outperform locals.

2. Under negative selection, below-average candidates can migrate, do worse than locals, but still earn more than at the origin, because their native country has a more unequal income distribution.

3. Finally, under refugee selection, below-average individuals can migrate and outperform locals, because the income distribution is wider in the destination.

To account for these unobservable factors, a further model incorporates schooling into the wage equation at the origin and destination as follows.

$$\ln w_0 = \mu_0 + \delta_0 s + \epsilon_0 \tag{4.8}$$

$$\ln w_1 = \mu_1 + \delta_1 s + \epsilon_1 \tag{4.9}$$

Here δ_0 and δ_1 represent the return to schooling (or skill premium) at the origin and the destination respectively. Thus under a revised version of equation (4.7) above, individuals will migrate when

$$\ln w_1 - \ln(w_0 + C) > 0 \tag{4.10}$$

$$(\mu_1 - \mu_0) + (\delta_1 - \delta_0)s + (\epsilon_1 - \epsilon_0) > \frac{C}{w_0} \tag{4.11}$$

4.2.3 Structural Gravity Model

The microeconomic model above considers the decision to migrate at the individual level. This essay will consider migration flows between origin states and destination municipalities. To bridge the gap between the individual and the aggregate, I use the the structural gravity model. Beine et al. (2016) provides a recent presentation of a structural gravity model developed from micro foundations and applied to migration.

The classical gravity model originated with Ravenstein's work studying migrant flows in the 19th century (Ramos, 2016). It has been applied successfully since then in many different contexts to the flow of goods and factors between countries (Anderson, 2011). Its empirical robustness owes to its parsimonious specification: the size of the origin, the size of the destination, and the inverse square of the distance between them. Recently available bilateral international migrant flow data has led to a renewed interest in the gravity model in studying migration. Nevertheless, until recently it was an "unconnected orphan" in the economics literature because of its lack of theoretical foundations. The structural gravity model addresses this deficiency by deriving an aggregate gravity model from an individual's decision to migrate in a framework in the microeconomic model above.

I begin with similar equations to (4.8) and (4.9) from the previous section. They consider the utility of an individual of type h staying in country i as $u_{ii}(h)$ and the same individual moving to country j as $u_{ij}(h)$. $C_{ij}(.)$ below is the cost of moving from country i to country j. I assume that it is constant for all individuals along the same migration corridor. This includes both the cost of moving as well as adapting to the destination.

$$u_{ii}(h) = w_i(h) + A_i + \epsilon_i \tag{4.12}$$

$$u_{ij}(h) = w_j(h) + A_j - C_{ij}(.) + \epsilon_j \tag{4.13}$$

The A_i and A_j terms are origin and destination characteristics that affect the desirability of living there. The error terms ϵ_i and ϵ_j are iid and follow an extreme value distribution.

Following a Mincer framework, an individual's schooling and experience will define his type h. I suppress the experience term to write the origin and destination wage equations as follows:

$$w_i(h) = \delta_i h + \mu_i + \epsilon_i \tag{4.14}$$

$$w_j(h) = \delta_j h + \mu_j + \epsilon_j \tag{4.15}$$

Now I use a multinomial logit model to extend this model from a single origin and destination to multiple destinations (indexed by k) and write the probability that a type h resident of country i will move to country j.

$$Pr\left[U_{ij}(h) = \max_k U_{ik}(h)\right] = \frac{N_{ij}}{N_i} = \frac{\exp[\delta_j h + A_j + \mu_j - C_{ij}(.)]}{\sum_k \exp[\delta_k h + A_k + \mu_k - C_{ik}(.)]} \tag{4.16}$$

In the same way, I can write the ratio of emigrants from country i to country j (movers) to residents of country i (stayers) as

$$\frac{N_{ij}}{N_{ii}} = \frac{exp[\delta_j h + A_j + \mu_j - C_{ij}(.)]}{exp[\delta_i h + A_i + \mu_i]} \tag{4.17}$$

Using logs, I next solve for the migration flow N_{ij} from location i to j by individuals of type h to obtain:

$$\ln N_{ij}(h) = (\delta_j - \delta_i)h + (A_j - A_i) + (\mu_j - \mu_i) - C_{ij}(.) + \ln N_{ii}(h) \tag{4.18}$$

Next I will model the cost function $C_{ij}(.)$. Recalling the monetary and non-monetary costs outlined in the previous section, I include the distance between location i and j, both as a proxy for the initial monetary cost of travel and the psychic non-monetary cost of being away from one's friends and family.

An important element of the model I employ in this essay is is the inclusion of the social network in the cost function.

For this reason, I also include the stock of existing migrants M_{ij} from location i presently residing in j. Carrington et al. (1996) provides an early example of the this approach, which the authors term taking into account endogenous moving costs. They incorporate the stock of existing migrants in a given Northern destination state from a given Southern origin state into the cost function of a dynamic model of migration. In the international context, Beine et al. (2011) follows a similar approach and names the this stock of existing migrants in a given destination country from a given origin country the diaspora.

Both social networks and labor market conditions change over time, so I add a time dimension to the model. I assume that individuals who migrate in time period $[t-1, t]$ make a decision based on conditions at the beginning of the time period at moment $t-1$. With this addition, I arrive at the reduced form I will empirically estimate.

$$\begin{aligned} \ln N_{ijt} = (\delta_{jt-1} - \delta_{it-1})h + (\mu_{jt-1} - \mu_{it-1}) + M_{ijt-1} - \text{distance}_{ij} \\ + (A_{jt-1} - A_{it-1}) + \ln N_{it-1} + \epsilon_{ijt} \end{aligned} \tag{4.19}$$

Munshi (2020) also adds destination networks to a Roy-Borjas migration model and proposes that this addition generates two testable predictions for the augmented model.

1. Because the diaspora adds to the wage differential and subtracts from the cost of migration, potential migrants will reject higher wage differentials to follow the diaspora.

2. As the diaspora size increases over time, individuals from farther down the ability distribution choose to migrate.

In this essay, I will test both of these predictions.

4.3 Data

The principal data source for this essay is the Mexican population census conducted by the INEGI (National Institute of Statistics and Geography) and harmonized by IPUMS. In this section, I give background on Mexican internal migration, describe the three five year periods of interest, the subsample of interest (Working Age Men), the outcome of interest (migrant flow), an important additional variable (stock of internal migrants in 1960), and additional origin and destination characteristics that I use in the empirical model in (4.19).

4.3.1 Mexican Internal Migration

Despite a long history, internal migration within Mexico has received much less attention than US-Mexico migration. Mexican internal migration began to increase in the second half of the 20th century. The first wave in the 1960s coincided with the return of many agricultural workers from the US after the ending of the Bracero program. Martin (2020) gives more background on this program, which gave nearly 5 million temporary visas in a lottery to Mexican farm workers from 1942-64. During this period, many more Mexicans moved closer to the border in the hope of receiving visas in the lottery. Moreover, the international migration networks that this program established affected subsequent migration even after its termination. To create jobs for these return migrants, the Mexican government gave tax benefits to export manufacturing plants (maquiladoras) on the Mexican side of the border. Since the creation of this program, factories in ten border cities have emerged as prominent destinations for internal migrants (Hanson, 2001). The export manufacturing sector continued to grow as Mexico opened to foreign trade: its entry into the General Agreement on Tariffs and Trade in 1986 and the passage of the North American Free Trade Agreement in 1994. Chiquiar (2005) and Arends-Kuenning et al. (2019) show that Mexico's entry into GATT caused different regions to grow more quickly or slowly owing to preexisting physical and human capital endowments; the passage of NAFTA continued these trends. As a result of both events, export manufacturing benefited and internal migration increased toward more quickly growing areas, especially those with export manufacturing.

4.3.2 Periods of Interest

IPUMS harmonizes the decennial Mexican population census from 1960 to 2010, with the exception of the 1980 census, the records of which were destroyed by the 1985 Mexico City earthquake. Beginning in 1995, the INEGI began to conduct a quinquennial census as well, and IPUMS also has harmonized this census for the years 1995, 2005, and 2015. Thus I have five candidate five year intervals: 1990-1995, 1995-2000, 2000-2005, 2005-2010, and 2010-2015. I exclude two of these intervals.

1990-1995 As the subsequent tables and maps indicate, the 1995 census had a much smaller sample than subsequent decennial or quinqennial censuses so I cannot use it to construct $flow_{1995-2000} = stock_{2000} - stock_{1995}$. In addition, the devaluation of the peso, the Zapatista uprising in Chiapas, and the passage of the North American Free Trade Agreement creates an idiosyncratic shock in this period.

2005-2010 I omit the 2005-2010 interval because the 2005 census did not include the employment module that I use to generate counterfactual wage predictions that I describe in the subsequent section.

Thus I conduct the analysis over three intervals: 1995-2000, 2000-2005, 2010-2015. Like many other censuses, the Mexican census uses stratified sampling along several demographic characteristics. Individual entries have individual population weights and household population weights. In accordance with the guidelines in Solon et al. (2015), I use the individual population weights for the descriptive statistics below and for the Mincer regressions I estimate in section 4.4.1.

4.3.3 Sample of Working Age Men

This essay estimates the effects of social networks on a particular sample of internal migrants: men from age 25 to 55, a group we call Working Age Men. In this subsection, I explain the choice of this sample using the 2000 census.

As I discuss in section 4.2, this essay's theoretical model considers migration as an investment decision in a human capital framework in which the potential migrant seeks to maximize lifetime earnings. Not all migrants move to maximize earnings, however. Some move to attend school or to follow family members.

In order to estimate the additional explanatory power of social networks over wage differences in determining migration destinations, I would like to choose a sample that would be particularly prone to migrate for work and thus sensitive to wage differences. Using this subset would put a upper bound on the effect of wage differences and a lower bound on the effect of social networks. I will use the 2000 census to justify the sample because of its completeness and a unique question was only asked in this year: the reason for migration. Table 4.1 summarizes their results by share of internal migrants.

Migration Cause	Women	Men	Working Age Men	Total
Unknown	14.632	12.725	7.779	27.357
Seeking work	8.251	13.028	7.864	21.279
Family move	9.159	6.232	3.572	15.392
Other reason, not elsewhere classified	6.225	5.576	3.850	11.801
Job relocation	3.494	6.932	5.667	10.426
Marriage or union	5.036	1.857	1.425	6.893
Study	1.802	2.019	0.536	3.821
Violence or insecurity	0.802	0.727	0.506	1.529
Health	0.805	0.698	0.483	1.503
Total	50.205	49.795	31.681	100.000

Table 4.1: Reason for Migration for Individuals Aged 16-65 (2000 Census)

First, though an almost identical number of men and women migrate, men are almost twice as likely to migrate for work than women (the reasons "seeking work" and "job relocation"). Thus I will restrict the sample to men.

Again using the 2000 census, I examine the share of male internal migrants by age in figure 4.1. This graph gives an empirical estimate of the probability of migration conditional on age. At each end of this interval, the decision to migrate for work is one of a set of options working locally (through the interval), additional education (at the lower bound), and retiring (at the upper bound). I would like to choose a interval that minimizes the probability of these two other decisions.

To choose the the lower bound of this interval, I consider the probability of education conditional on age in Figure 4.2. At age 20, 25% of men are still in school; by age 25, this share has dropped to 9%.

To choose the upper bound of the interval, I consider the probability of retirement conditional on age in Figure 4.3. At age 65, 13% of men have decided to retired; at age 55, this share has dropped to 3%.

Thus I restrict the sample to men to minimize the impact of the decision to migrate for non-work reasons and from 25 to 55 years old to minimize the impact of the decision to pursue additional education or retire.

I consider one final selection issue in the sample: the decision to select into labor. Figure 4.4 shows the share of unemployed men by age. This share is at most 2% for a given age cohort. These results suggest that most men who want can find employment, in either the formal or informal sector.

The rightmost column in Table 4.1 shows the migration reasons of our Working Age Men sample. 44% report migration for work reasons.

I conclude my description of the Working Age Men sample with two final comments. First, the migration cause question has many limitations. Only the 2000 census asks the question and almost 40% of the responses are "Other reason" or "Unknown". It serves at best as only a rough guide to choosing the appropriate sample for this analysis, even when we supplement it with employment, retirement, and schooling profiles by age.

On the other hand, the analysis of this essay does not depend entirely on finding a sample

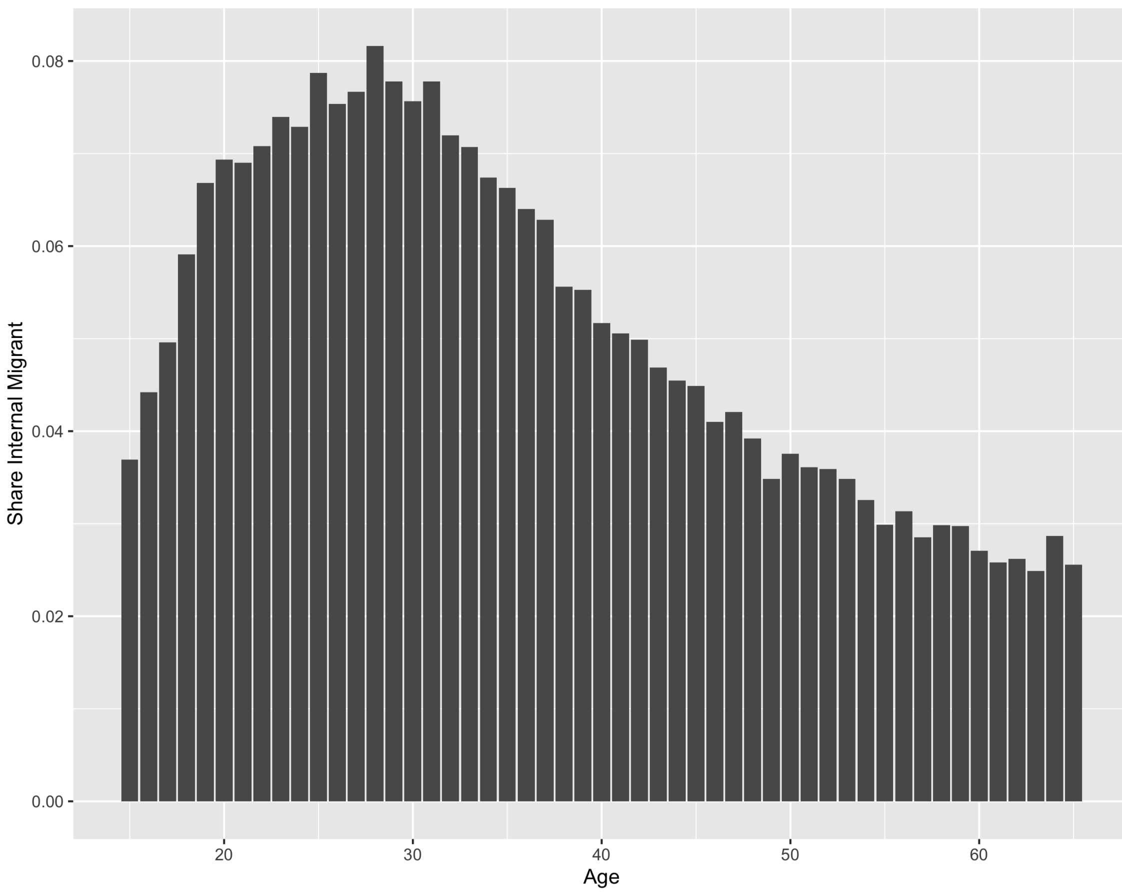

Figure 4.1: Share of Male Internal Migrants by Age (2000)

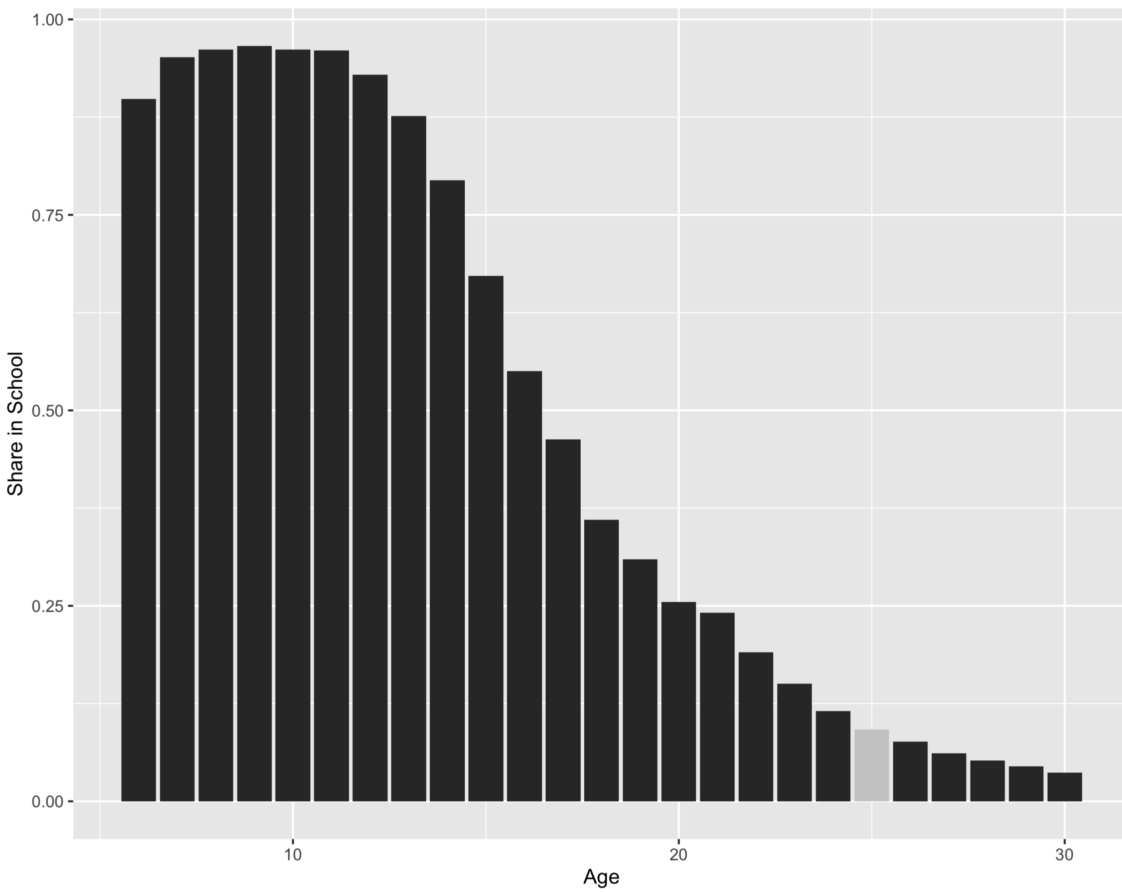

Figure 4.2: Share of Males in School by Age (2000)

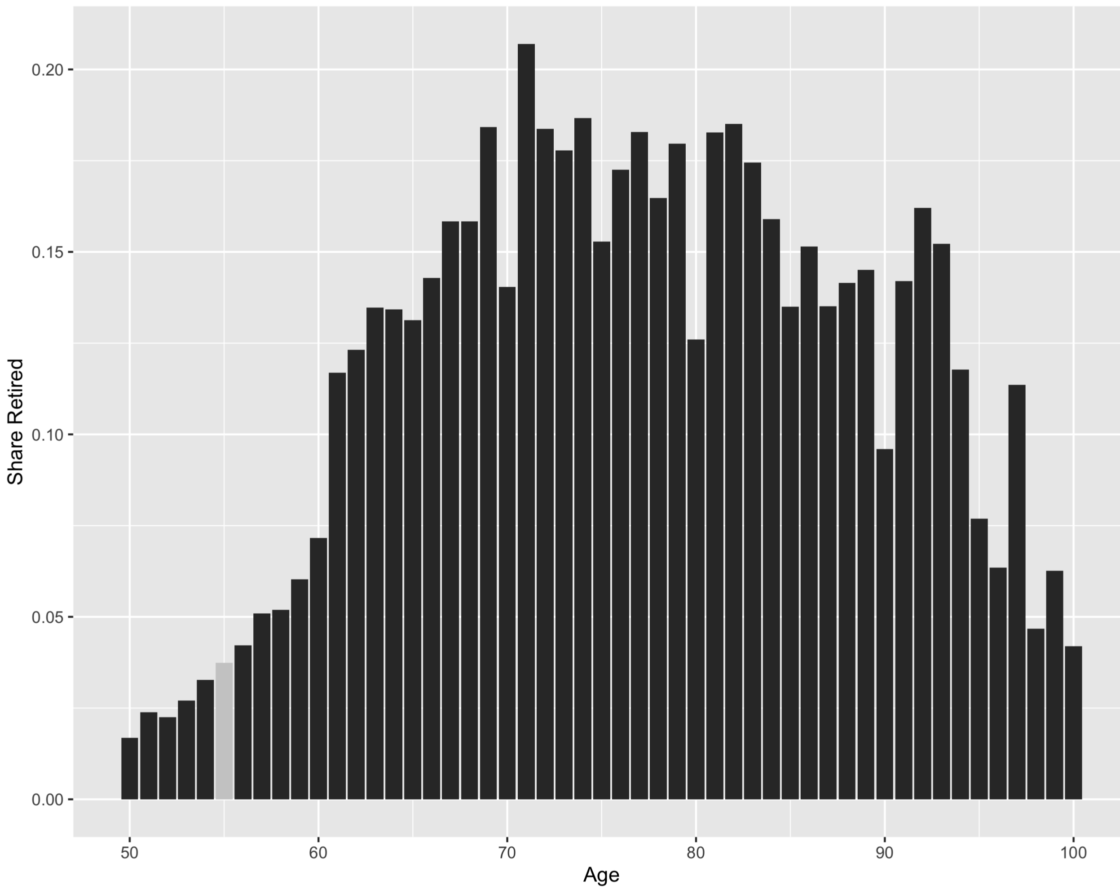

Figure 4.3: Share of Males Retired by Age (2000)

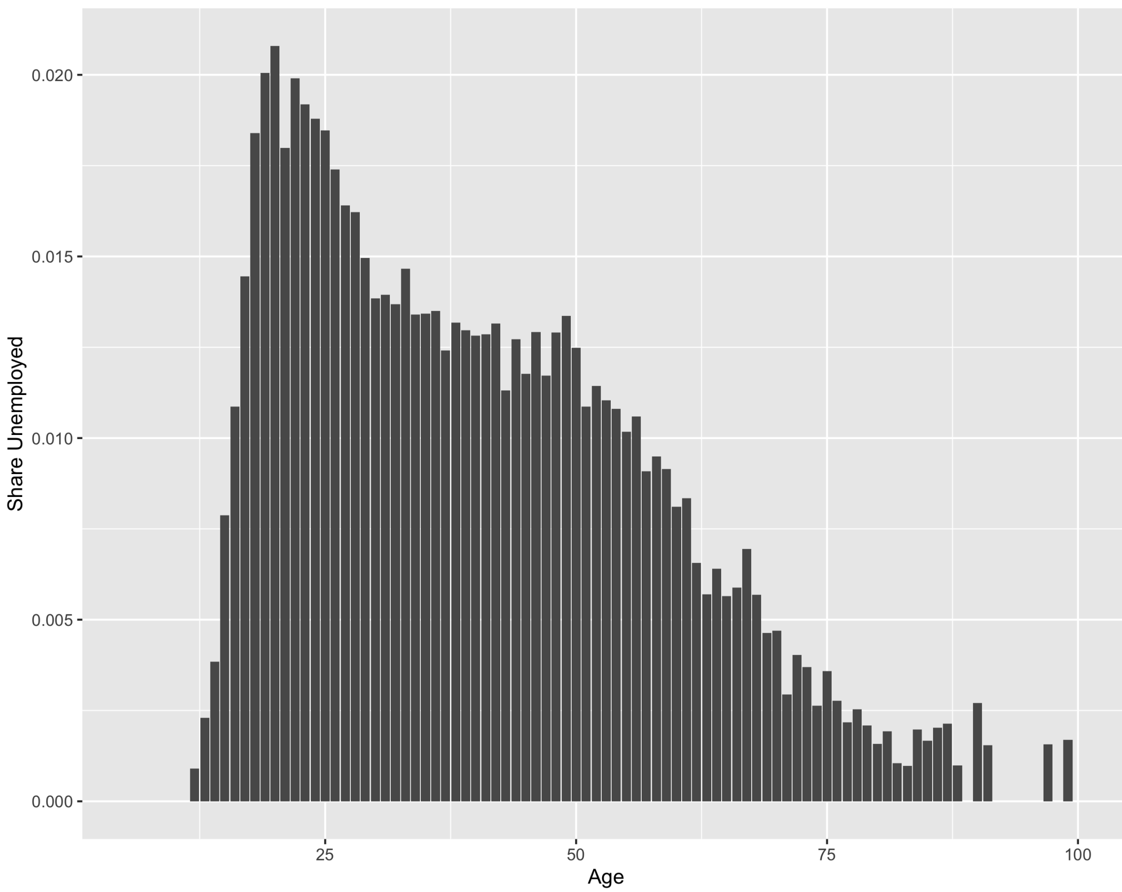

Figure 4.4: Share of Males Unemployed by Age (2000)

of individuals whose choice set consists only of working locally or migrating for work reasons. It merely examines the extent to which social networks add additional explanatory power to other reasons why individuals migrate, especially wage differences. For this reason, I use a subsample particularly sensitive to wage differences.

4.3.4 Summary Statistics of Working Age Men

Table 4.2 shows summary statistics for the population of Working Age Men in each census: income, age and schooling. I compute experience in the typical way (age - schooling - 6) to use in the Mincer regressions we will describe in section 4.4.1. Labor Force Participation (LFP) is above 90% in all intervals. The key variables of schooling and income are present for the vast majority of the sample.

Unlike other censuses, the Mexican census asks several questions about income: earned income, income from pensions, income from government support programs. I use "earned income", which is defined as monthly income in pesos. In addition, the table shows the percentage who are internal migrants.

Through the five censuses, mean income, age, and schooling increase. Income and schooling increase as a result of Mexico's economic development. The age increase shows the demographic changes of an aging population.

Table 4.3 shows a decomposition by cohort of the Working Age Population with the population shares of each cohort as well as the LFP and the internal migration shares. The age distribution shifts slightly older from 1995 to 2015. LFP peaks in the 35-39 cohort but is above 90% in almost all cells except in the final five year interval. Internal migration is highest for the youngest cohort and steadily decreases. This empirical fact is typical of rural-urban migration in developing countries (Lucas, 2021).

		Shares				Income		Age		Schooling		Experience	
Year	Total	LFP	School	Income	Migrant	Mean	SD	Mean	SD	Mean	SD	Mean	SD
1995	14875665	0.949	0.994	0.951	0.066	1602.962	23.858	36.624	0.058	7.845	0.037	22.779	0.075
2000	16497274	0.917	0.971	0.964	0.061	3597.730	12.014	36.908	0.008	8.306	0.005	22.602	0.010
2005	18472940	NA	0.973	NA	0.042	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
2010	20775439	0.931	0.989	0.929	0.068	6035.489	20.827	37.909	0.014	9.333	0.009	22.576	0.017
2015	22747904	0.911	0.997	0.918	0.053	7240.150	19.833	38.198	0.012	10.280	0.009	21.918	0.015

Table 4.2: Working Age Men Overall by Year

	Share				Working				Int Migrant			
Cohort	1995	2000	2010	2015	1995	2000	2010	2015	1995	2000	2010	2015
25-29	0.243	0.235	0.199	0.194	0.946	0.909	0.912	0.891	0.082	0.078	0.087	0.067
30-34	0.211	0.207	0.192	0.186	0.961	0.933	0.944	0.929	0.084	0.073	0.087	0.066
35-39	0.187	0.185	0.189	0.179	0.963	0.934	0.949	0.930	0.063	0.061	0.072	0.057
40-44	0.146	0.153	0.162	0.172	0.952	0.927	0.943	0.924	0.057	0.049	0.059	0.048
45-50	0.118	0.120	0.139	0.142	0.945	0.908	0.932	0.911	0.041	0.041	0.048	0.038
50-54	0.096	0.101	0.119	0.127	0.906	0.866	0.893	0.874	0.040	0.036	0.038	0.033

Table 4.3: Working Age Men by Cohort

4.3.5 Outcome of Interest: Migrant Flow

Mexico is divided into 32 federal entities (31 states and the federal district of Mexico City). For simplicity I refer to the entities below as states. Each state is divided further into municipalities.

The census asks about migration in two different ways. In 1960, it asks if individuals presently residing in a given municipality have moved from another state during their lifetime. In the 1990 and subsequent Mexican censuses, individuals presently residing in a particular municipality report their state of residence five years ago.

By totaling the number of residents of a municipality who were born in another state (in the case of 1960) or who lived in a different state five years ago (in the case of the other time periods), I can construct a measure of a possible internal migration corridor originating in one state and terminating in a municipality in another state. This method omits both temporary migrants as well as migrants who lived in a third location, either within Mexico or abroad, in the intervening five years.

With this definition, the three intervals of interest, I compute the migrant flow in the period $[t-1, t]$ from origin state s to destination municipality m and denote it as $flow_{smt}$. As a point of comparison, I also compute the internal migrant stock in 1960, $flow_{1960}$.

Table 4.4 provides information about the internal migrant flows in each of the three periods of interest as well as the internal migrant stock in 1960. The Destinations column indicates the number of municipalities surveyed in the census. The Possible Flows column multiplies Destinations by 31 to indicate the number of potential migration corridors that could be captured. The Active Flows indicates the number of migration corridors that were actually captured.

The number of possible flows is increasing across time periods for two reasons. First, the number of municipalities is increasing. As of 2021, there are 2471 municipalities. Second, the coverage of the Mexican census is improving.

Computing the flows this way allows me to detect $32 \cdot 2471 = 79072$ origin-destination combinations. Since I do not consider migrations within the same federal entity, I can remove 2471 of these combinations, placing an upper bound of 76601 detectable migrant flows. The number of flows in the last time period approaches this upper bound.

A central empirical question of this essay is whether migration from a particular origin state to a destination is associated with further migration along the same corridor in subsequent time periods. Thus the first set of columns compares migration corridors from the three periods of interest to to the corridors open in 1960 and the two subsequent sets of columns compare the second and third period of interest to the first and the third to the second, respectively. In all cases, the Stay and Closed columns in these comparison columns adds up to the Active Flows column for the reference period; the Stay and Opened columns adds up to the Active Flows column from the current period.

		Flows		Compared to 1960			Compared to 1995-2000			Compared to 2000-2005		
Period	Dest	Possible	Active	Stay	Closed	Opened	Stay	Closed	Opened	Stay	Closed	Opened
1960	252	7812	4195	NA	NA	NA	NA	NA	NA	NA	NA	NA
1995-2000	627	19437	6654	2263	1932	4391	NA	NA	NA	NA	NA	NA
2000-2005	1978	61318	10052	2707	1488	7345	4286	2368	5766	NA	NA	NA
2010-2015	2309	71579	17307	2873	1322	14434	4715	1939	12592	7223	2829	10084

Table 4.4: Summary of Migrant Flows by Period

Figures 4.5, 4.6, and 4.7 show the logged flow of internal migrants by destination municipality for each of the three intervals in question.

4.3.6 Origin and Destination Characteristics

The theoretical model that I developed in the previous section also includes origin and destination characteristics that could affect the desirability of migration. For each of the three time periods of interest, I use the value of the characteristic from the start of the time period.

Population I compute the population of a state or municipality by using the total number of people in the most recent decennial census. Figure 4.8 shows the logged municipal population at year 2000.

Indigenous Share I compute the indigenous share for a state or municipality using the number of people in the most recent decennial census who report being indigenous. I divide this number by the population. The 1990 census does not ask this question, so I do not include this characteristic for that time period. Figure 4.9 shows this share.

Urban Share I compute the urban share for a state or municipality using the number of people in the most recent decennial census who report living in an urban area. I divide this number by the total population. Figure 4.11 shows this share for the year 2000.

Border Potential migrants could migrate to the border as a destination as a first step toward migration to the US. In a similar way, migrants recently deported from the US could originate at the border. I assign a dummy variable to all origin and destinations on the US/Mexico border. Figure 4.10 shows border municipalities.

Distance. IPUMS provides shapefiles for all Mexican states and municipalities. I compute the $distance_{sm}$ as the distance between the centroid of s and the centroid of m.

4.4 Empirical Framework

Here I describe in detail the empirical method I use to estimate the structural gravity model that I develop in the previous section, concluding with equation (4.19). First, I describe the auxiliary Mincer regressions that I use to estimate the labor market difference parameters. Second, I describe

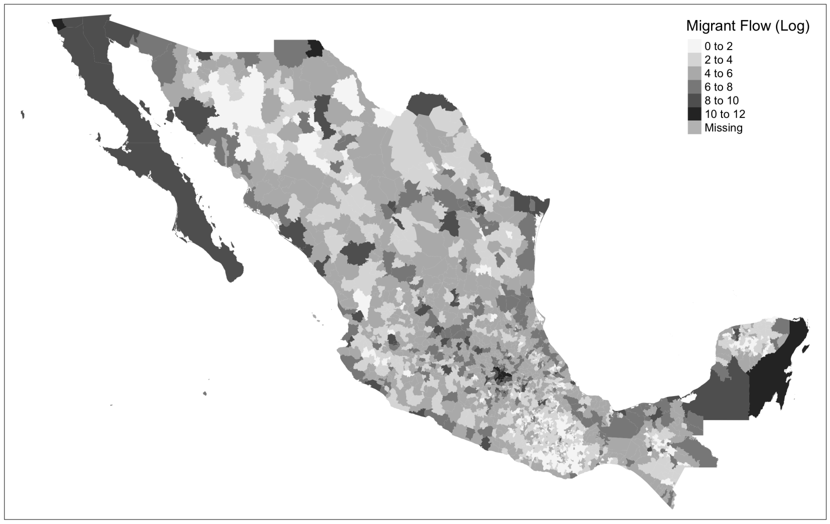

Figure 4.5: Internal Migration Flow from 1995-2000

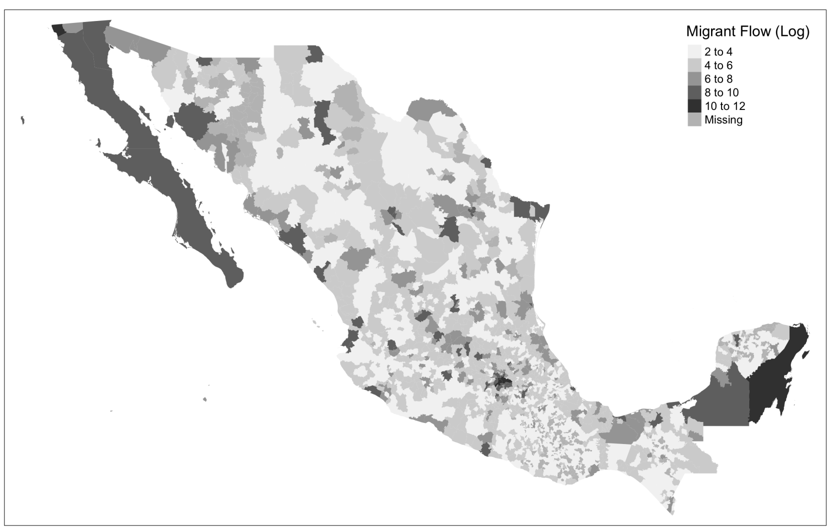

Figure 4.6: Internal Migration Flow from 2000-2005

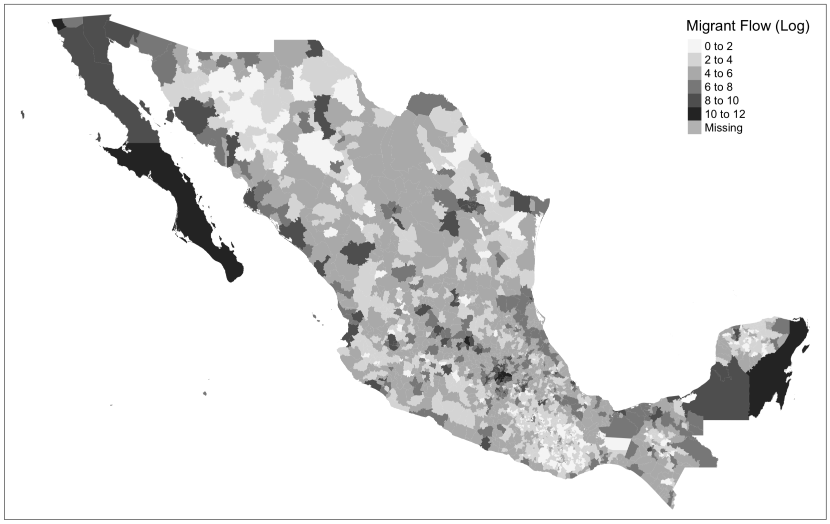

Figure 4.7: Internal Migration Flow from 2010-2015

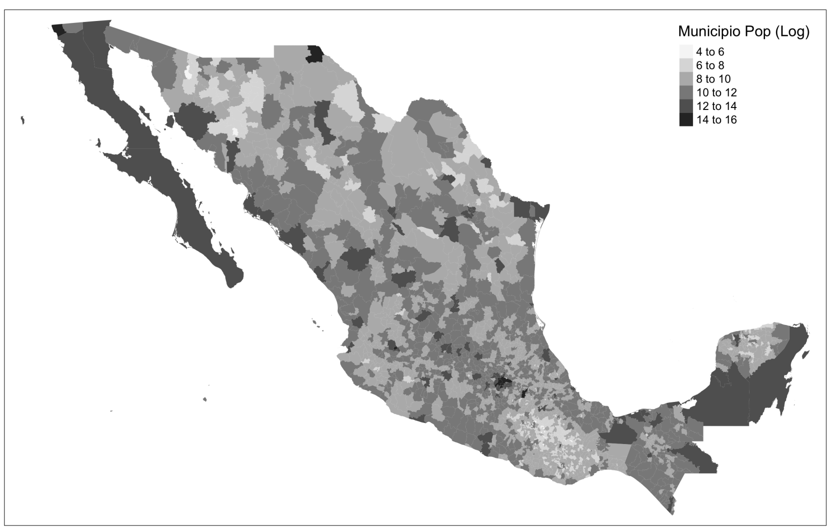

Figure 4.8: Municipality Population (2000)

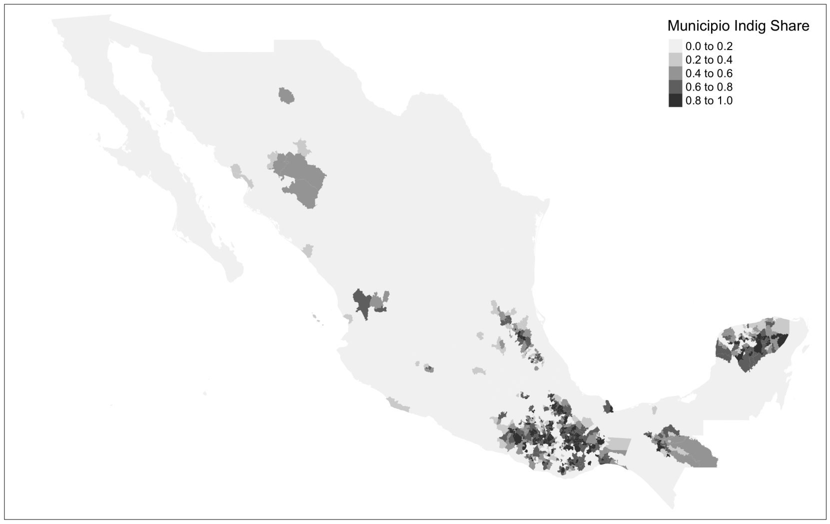

Figure 4.9: Municipality Indigenous Household Share (2000)

Figure 4.10: Border Municipalities

how I estimate the main model, the effect of social networks on the migrant flow $flow_{smt}$ from a state to a municipality at a given time period. I use OLS to estimate separately the extensive and intensive margin to account for the possibility of left censoring in $flow_{smt}$. Third, I describe a potential threat to identification, the presence of time-invariant factors that could affect migrant flows across multiple time periods. I propose the use of the presence of a migrant flow along the same corridor in the year 1960 to control for this possibility, which I call a "taste of migration." Finally, I discuss inference issues and the use of clustering at the destination state level.

4.4.1 Labor Market Differences

I do not use population weights for the main regressions here. and for the Mincer regressions I describe in section 4.4.1.

Both Falaris (1987) and Beine et al. (2011) use average wage at the destination to capture the effect of wage differential on migration. Using only the mean wage, however, omits the heterogeneous effects of varying levels of schooling and experience on an individual's decision to migrate. As table 4.2 indicates, average education level changes across the time periods of interest. Moreover, the effect of the difference in return to skill could vary depending on the presence of positive or negative selection, as I describe in section 4.2.1.

In order to capture the expected wage differential in a flexible way that accounts for this heterogeneity, I suppress the h from the theoretical model and instead incorporate $\mu_j - \mu_i$ and $\delta_j - \delta_i$ directly into the empirical model. I denote these differences as $\tilde{\alpha}$ and $\tilde{\beta}$ respectively. Note that in the case of internal migration origin country i becomes origin state s and destination country j becomes destination municipality m.

$$\mu_{jt} - \mu_{it} = \alpha_{mt} - \alpha_{st} = \tilde{\alpha}_{mst} \tag{4.20}$$

$$\delta_{it} - \delta_{it} = \beta_{mt} - \beta_{st} = \tilde{\beta}_{mst} \tag{4.21}$$

Now I must estimate $\tilde{\alpha}_{mst}$ and $\tilde{\beta}_{mst}$. To do this, I use Mincer equations for the Working Age Men in each state s and municipality m at the beginning of each time period $t \in \{1995, 2000, 2010\}$.

As I describe in the previous section, I use "earned income in the past month" as the income variable and exclude elements of the sample with unknown education levels and income.

I estimate the parameters of a Mincer regression separately on each state and municipality in Mexico for the three time periods of interest and store the α and β coefficients.

$$\log(\text{income}_{tmi}) = \alpha_{tm} + \beta_{tm}\text{educ}_{tmi} + \gamma_{tm}\text{exp}_{tmi} + \delta_{tm}\text{exp}^2_{tmi} + \epsilon_{tmi} \tag{4.22}$$

$$\log(\text{income}_{tsi}) = \alpha_{ts} + \beta_{ts}\text{educ}_{tsi} + \gamma_{ts}\text{exp}_{tsi} + \delta_{ts}\text{exp}^2_{tsi} + \epsilon_{tsi} \tag{4.23}$$

For these parameters, m indexes the municipality, s indexes the state, and i indexes the individual within the municipality or state. The α_{tm} and α_{ts} parameters represent the base salary of the origin state and destination municipality. The β_{tm} and β_{ts} parameters represent the return to skill at origin state and destination municipality.

Figures 4.12, 4.13, and 4.14 show the logged base salary at the beginning of the three time periods in question. Figures 4.15, 4.16, and 4.17 show the skill premium at the beginning of the three time periods in question.

I find little cross-sectional variation in the γ and δ parameters related to the return to experience, so I do not include them in the model.

As table 4.2 indicates, the small sample size of the 1995 census results in the inability estimate labor market parameters for some municipalities, which appear in grey in the associated maps.

Using these estimates, I compute the parameter differences for each combination of origin state and destination municipality in each of the three time periods. I use these values in the estimation of the main model below.

$$\tilde{\alpha}_{smt} = \alpha_{tm} - \alpha_{ts} \tag{4.24}$$

$$\tilde{\beta}_{smt} = \beta_{tm} - \beta_{ts} \tag{4.25}$$

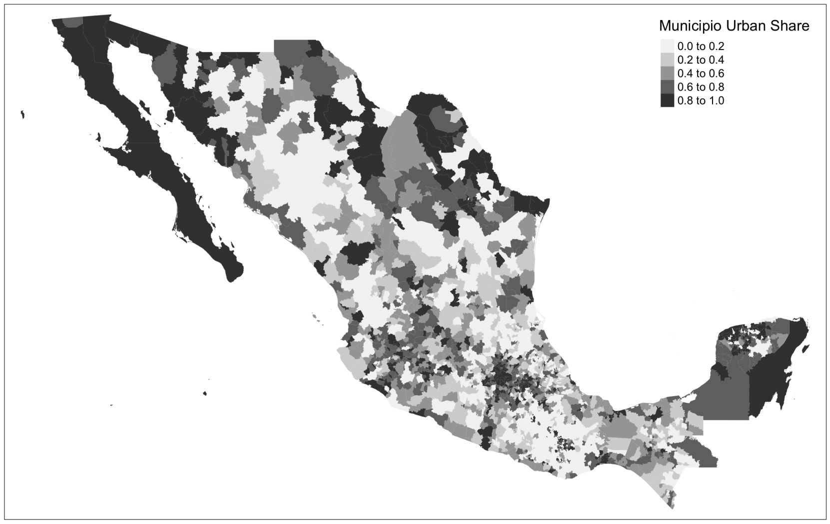

Figure 4.11: Municipality Urban Household Share (2000)

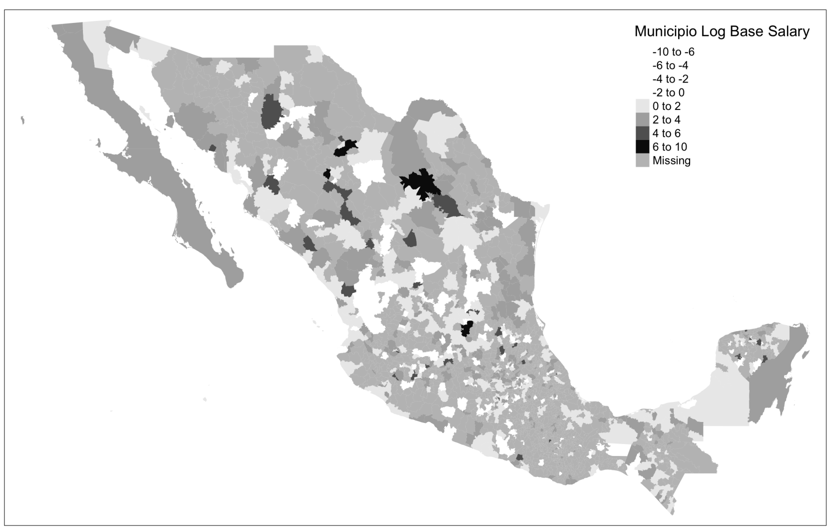

Figure 4.12: Logged Base Salary by Municipality (1995)

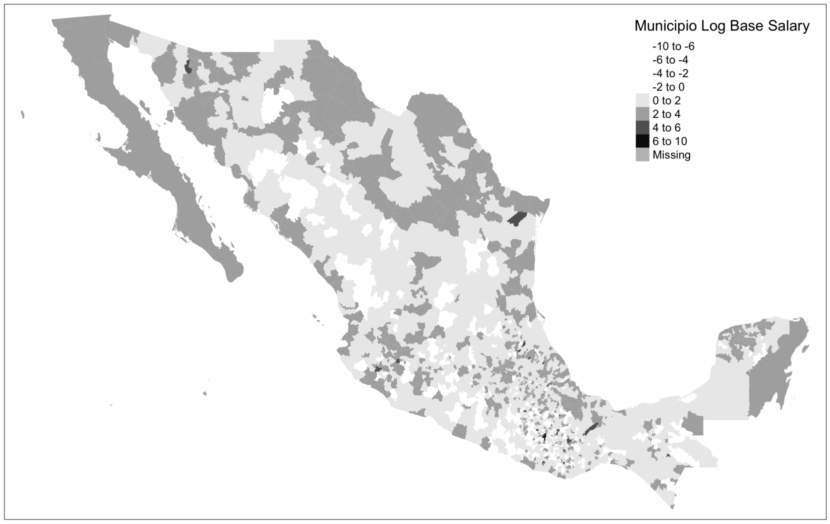

Figure 4.13: Logged Base Salary by Municipality (2000)

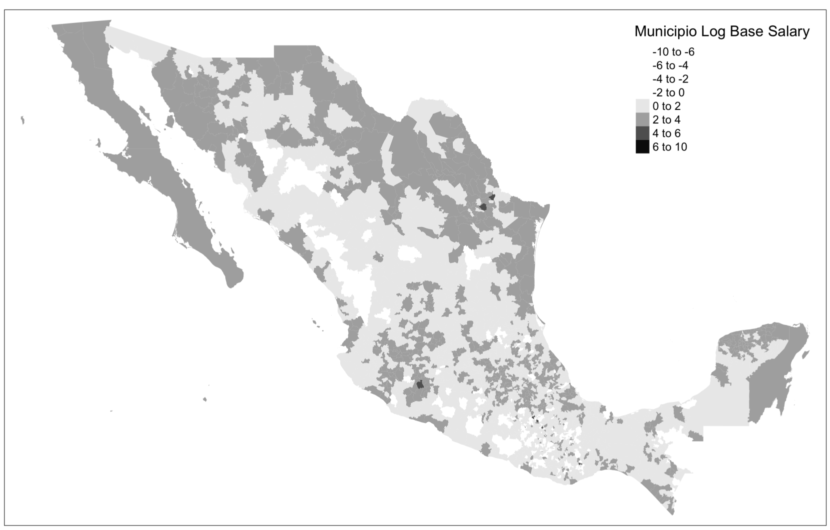

Figure 4.14: Logged Base Salary by Municipality (2010)

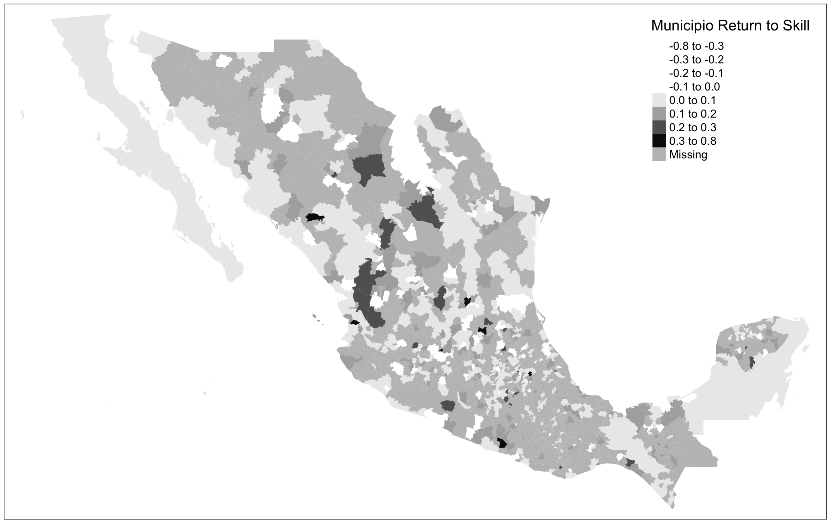

Figure 4.15: Return to Skill by Municipality (1995)

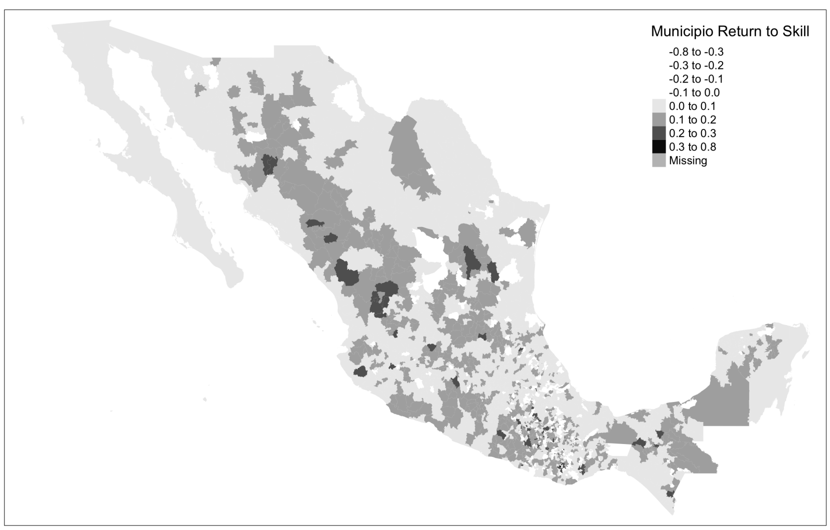

Figure 4.16: Return to Skill by Municipality (2000)

4.4.2 Estimation

Recall that I consider aggregate migrant flows $flow_{smt}$ from a Mexican state s to a Mexican municipality m in five year intervals $[t-1,t]$. Section 4.2.1 gives more detail on this definition of internal migration and section 4.3.5 gives more information about these particular migrant flows.

I begin with equation (4.19) as follows:

$$\begin{aligned}\log(\text{flow}_{smt}) = \delta_0 + \delta_1\tilde{\alpha}_{smt-1} + \delta_2\tilde{\beta}_{smt-1} \\ + \delta_3\text{flow}_{smt-1} + \delta_4\text{distance}_{sm} + \gamma A_s + \phi A_m + \epsilon_{smt}\end{aligned} \tag{4.26}$$

Here I use $flow_{smt}$ as N_{ijt}, the number who migrate from country i to country j at time period t. I use $flow_{smt-1}$ as M_{ijt}, the size of the diaspora from country i already present in country j at time t.

Recall that A_s and A_m are vectors of other origin state and destination municipality characteristics that could influence migrant flow. I populate them with three characteristics: urban share at $t-1$, indigenous share at $t-1$, and presence on the US/Mexico border. In addition, to match the $log(N_{it-1})$ term in the model, I include the population of the origin state in A_s.

Extensive and Intensive Margins

As table 4.4 indicates, an econometric challenge to estimating equation (4.26) is the presence of zero flow values: 66%, 84%, and 76% respectively in the three time periods of interest.

This issue is not unfamiliar in the gravity model literature. Beine et al. (2011) uses two-stage Heckman estimators as a robustness check on OLS estimates of a model very similar to ours. The intuition behind the Heckman estimator is that two separate processes are operating: the first selects into or activates a migration corridor and the second determines the magnitude of the flow through it. Treating these processes as one risks biasing the estimated effect of the treatment, in this case the social network.

The Heckman estimator corrects for selection bias: the same factors that affect the presence of a migration corridor could also affect the flow through that some corridor. Here I am interested in examining the presence of a migration corridor as a process in its own right. For this reason, I estimate separately the extensive and intensive margin of the effect of social network on migration flows.

The extensive margin refers to the effect of the social network on whether a corridor opens or remains open during a time period. I define a new dependent variable $flowpresent_{smt} = 1[flow_{smt} > 0]$ as an indicator that takes the value of 1 if there are internal migrants from state s residing in municipality m at time t and 0 otherwise. I will estimate a version of equation (4.26) above with this new indicator variable.

To perform this estimation, I use an Linear Probability Model for two reasons. First, very few of the predicted values are out of the $[0, 1]$ range, so I see no advantage to a logit or probit model. Second, I can directly interpret the coefficients of the LPM in the subsequent results section. The coefficient of interest is δ_3, the effect of the social network on the presence of a flow.

$$\begin{aligned} \text{flowpresent}_{smt} = \delta_0 + \delta_1 \tilde{\alpha}_{smt-1} + \delta_2 \tilde{\beta}_{smt-1} \\ + \delta_3 \text{flow}_{smt-1} + \delta_4 \text{distance}_{sm} + \gamma A_s + \phi A_m + \epsilon_{smt} \end{aligned} \tag{4.27}$$

I estimate equation (4.27) on the entire sample. For the subsample for which $flowpresent_{smt} = 1$ at each time period, I estimate equation (4.26) to obtain the intensive margin.

The estimation technique in this essay differs from other literature that uses gravity models to estimate migration flows. First, many authors use the Pseudo-Poisson Maximum Likelihood Model developed by Silva and Tenreyro (2006) to account for potential bias in the effect of determinants of migration because of the censoring on the dependent variable. Because I estimate these margins separately, we do not use this estimator. Neither does Beine et al. (2011).

Second, even though I could construct a panel data set of our state-municipality migrant flows across the time periods of interest, I do not, because we expect the effects of the determinants of migration to vary over time. As section 4.3.1 describes, I would like to see the effect of the structural changes of the Mexican economy and the changing value of the outside option of migrating to the US on the estimated effects of the various determinants of migration in our model in our time periods of interest. In particular, I would like to see if the effect of social networks changes over time.

4.4.3 Identification

Identifying network effects poses statistical challenges. In this section, I group these challenges in two categories: across space (the way that the relative desirability of one destination influences another in the same time period) and time (the way that the patterns of migration in one time period affect migration in a subsequent time period through the channel of social networks). Both types of challenges relate to SUTVA (the Stable Unit Value Treatment Assumption) described in Morgan and Winship (n.d.) and elsewhere.

First, I consider the challenges across space. In the multinomial logit model that I develop in section 2, potential migrants do not simply choose to stay or leave; instead, they choose among a variety of destinations. In the aggregate setup here, each unit is a potential migration corridor from an origin state to a destination municipality. The treatment is the magnitude of the migration flow along the same corridor from the previous time period. For SUTVA to hold, the treatment received by one unit must have no relationship to the treatment received by another unit across space or time In other words, the migrant flows along a given corridor in the previous time period $flow_{smt-1}$ must be unaffected by the migrant flow along other corridors $flow_{-s-m-1}$ in the same previous time period.

This statement is not true. In a given time period, the migration decisions of the population of Working Age Men in a given state satisfies a population balancing equation: the number of men who do not migrate plus the number of men who migrate to each destination must sum to the total population of the state. Intuitively, the sums are related in this way: increasing the social network in one potential destination municipality decreases the size of the social network in another potential destination municipality. Thus the strong form of SUTVA does not hold in this case.

In fact, changing the distribution of the destination municipalities of migrants from the same origin state in a previous period affects the relative desirability of those destinations in the current time period. The relevant question is the relative magnitude of these effects. I argue that these effects are so small as to not warrant consideration because they are so diffuse. From a given state, a potential migrant considers nearly 2000 destinations. In a given time period, 10% to 20% of these destination corridors are open: 200 to 400 destinations. The decision of one migrant would seem not to affect the decision of another migrant very much.

Recent literature in biostatistics has developed techniques to address this particular relaxation of SUTVA, a situation of allocation of a common resource where the treatment status of each unit affects the treatment status of every other unit (Miles et al., 2019). Further analysis could apply these techniques to the present situation to empirically verify this intuitive argument.

Next, I consider challenges across time. As Manski (1993) and Munshi (2020) point out, identifying network effects poses statistical challenges. In this case, since I am estimating the effect of $flow_{smt-1}$ on $flow_{smt}$, a serially correlated shock across two time periods could generate a spurious correlation between the flows.

One approach would use an instrumental variable to estimate the effect of $flow_{smt-1}$. In the Mexico-US context, two authors use exogenous shocks at the origin: rainfall (Munshi, 2003) and the presence of railroad networks (Woodruff & Zenteno, 2007). In the European context, Beine et al. (2011) use three different instrumental variables: diplomatic representation of one country in another, the presence of a guest worker program, and conflicts in the origin country. Munshi (2020) also discusses the possibility of exploiting variation in network quality or conditions at the destination. By using an instrument correlated with origin conditions but not destination conditions, all of these approaches hope to correct for any serially correlated shock.

This essay proposes a different approach: a simple model of a taste for migration. A taste for migration between an origin and a destination is a factor such as a common climate or cultural connection that the model here does not account for. The simplest version of such a taste would be a time-invariant dummy variable between an origin-destination pair. A more sophisticated version could (1) vary continuously depending on the origin or destination and (2) vary depending on the combination of time period.

As a proxy for this taste for migration, I use the presence of a stock of migrants from the same origin state at the destination municipality in 1960. Section 4.3.5 gives more details about this variable. Note that I use the presence of a stock of all internal migrants, not simply working age men. The reason is that I want to measure the effect of time-invariant factors that would drive migration along a particular corridor and affect all potential migrants equally.

The reason I use the stock of migrants in 1960 is practical as well as historical. The earliest available Mexican census is in 1960. In addition, 1960 comes before the end of the Bracero program and the beginning of the export manufacturing program that we described in section 4.3.1. If these social networks represent long-run processes, using 1960 stocks as a control allows us to account for the unobservable initial conditions that started these processes and separate their effect from the effect of interest, the ongoing role of social networks in keeping migrant corridors open and inducing migration flow through these corridors. In the next section, we will present estimates of our model with and without this control. To my knowledge, this essay is the first one to model a taste for migration in this way.

4.4.4 Inference

In addition to identification issues that could bias the estimation of the effect of social networks, inference issues could interfere with the estimation of the significance of these effects. In particular, within-state correlation of unobservable factors related to migration destinations could affect the estimation of the significance of the effects of social networks on migration to these destinations. These factors include state-level policy decisions or industry-specific economic factors that could affect the labor demand and thus the migration flow across a particular state, for example. To address these inference issues, I cluster the standard errors by destination state using the standard cluster-robust variance-covariance matrix estimator.

4.5 Results and Discussion

In this section I will present the three main results of the empirical analysis, which estimates the extensive and intensive margin of internal migration separately. First, a structural gravity model with social networks provides additional explanatory power over a standard structural gravity model and a Roy-Borjas model with only differences in base wage and skill premium in estimating both margins across three time periods. Second, the effect of social networks monotonically increases at both margins across all three time periods; the effect of the other factors varies according to the migration climate. Third, the results hold up under a robustness check: a time-invariant migration taste factor.

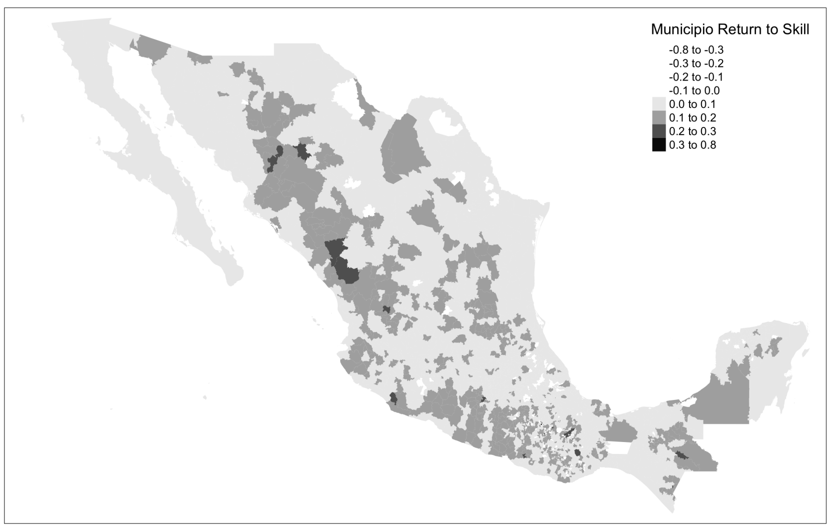

Figure 4.17: Return to Skill by Municipality (2010)

Table 4.5: Internal Migration in Period 2000-2005

	Dependent variable:		
	Internal Migrant Flow Present		
	(1)	(2)	(3)
Base Salary Diff	0.063^{***} (0.011)	0.013^{**} (0.006)	0.006^{**} (0.003)
Return to Skill Diff	0.859^{***} (0.155)	0.081 (0.061)	0.039 (0.035)
Distance (log)		-0.124^{***} (0.009)	-0.055^{***} (0.006)
Dest Population (log)		0.106^{***} (0.016)	0.055^{***} (0.010)
Origin on Border		0.052^{***} (0.015)	0.036^{***} (0.011)
Dest on Border		0.231^{***} (0.047)	0.085^{***} (0.029)
Origin Urban Share		0.237^{***} (0.026)	0.101^{***} (0.013)
Dest Urban Share		0.078^{***} (0.019)	0.045^{***} (0.012)
Origin Indig Share		0.073^{*} (0.039)	0.024 (0.022)
Dest Indig Share		0.058^{***} (0.019)	0.029^{***} (0.010)
Social Network (log)			0.119^{***} (0.005)
Constant	0.196^{***} (0.021)	0.570^{***} (0.129)	0.202^{***} (0.070)
Observations	61,318	61,318	61,318
R^2	0.025	0.262	0.393
Adjusted R^2	0.025	0.262	0.393
Residual Std. Error	0.366 (df = 61315)	0.318 (df = 61307)	0.288 (df = 61306)

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Standard errors clustered by destination state

Table 4.6: Internal Migration in Period 2000-2005

	Dependent variable:		
	Internal Migrant Flow (log)		
	(1)	(2)	(3)
Base Salary Diff	0.374*** (0.052)	0.079** (0.035)	0.022 (0.018)
Return to Skill Diff	4.355*** (0.903)	0.173 (0.449)	−0.359 (0.221)
Distance (log)		−0.396*** (0.038)	−0.103*** (0.026)
Dest Population (log)		0.463*** (0.039)	0.229*** (0.021)
Origin on Border		0.143** (0.066)	0.110*** (0.040)
Dest on Border		0.785*** (0.186)	0.324*** (0.107)
Origin Urban Share		0.927*** (0.155)	0.196** (0.092)
Dest Urban Share		0.137 (0.096)	0.027 (0.074)
Origin Indig Share		0.631*** (0.210)	0.175 (0.145)
Dest Indig Share		0.462*** (0.141)	0.295*** (0.076)
Social Network (log)			0.304*** (0.017)
Constant	3.278*** (0.060)	2.391*** (0.389)	1.068*** (0.296)
Observations	10,052	10,052	10,052
R^2	0.051	0.392	0.582
Adjusted R^2	0.051	0.391	0.581
Residual Std. Error	1.128 (df = 10049)	0.903 (df = 10041)	0.749 (df = 10040)

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors clustered by destination state

4.5.1 Internal Migration at the Extensive Margin

Recall that table 4.4 provides a summary of the extensive margin across time: the presence of migration corridors. The models I present here estimate the factors that cause a corridor to open from one time period to the next or the factors that cause a corridor to stay open. For simplicity I will examine the extensive margin in one of the three time periods, but the analysis applies as well to the other two time periods. I choose the 2000-2005 time period because of better data availability at the start of the period in 2000.

Tables 4.5 and 4.5 shows three specifications of the extensive and intensive margins of migration to a destination municipality in the time period (2000-2005). Since I use a linear probability model for the extensive margin, I can directly interpret the coefficients as percentage increases in the probability of a migrant flow.

Specification 1: Wage Differentials

In specification (1) with only wage differentials, a one point increase in the base salary gap is associated with an 6.3% increase in the probability of a flow. An 0.1 increase in the return to skill difference is associated with a 8.6% increase in the probability of a flow. Because the dependent variables are calculated as coefficient differences from Mincer regressions, it is difficult to interpret these magnitudes directly. The sign and relative increase match the predictions of models that rely on wage differentials, however. Figures 4.13 and 4.16 show the base salary and return to skill by municipality for 2000. The relatively low R^2 value of 0.025 indicates the poor predictive power of this model.

Specification 2: Structural Gravity Factors

Specification (2) adds other structural gravity factors: distance, destination population, presence on the border for origin and destination; urban share for origin and destination; and indigenous share for origin and destination. I discuss the impact of these factors in turn.

First, large cities drive migration; a one percent increase in destination population increases the probability of a migrant flow by 10%. Distance works against migration by increasing migration cost. Increasing the distance by a factor of 2.7 decreases the probability of a migrant flow by 12%.

An origin on the border could indicate an individual who has already migrated once or who has been deported. These individuals are predisposed to migration. It increases the probability of migration by 8%. A destination on the border increases the probability of migration by 23%. These municipalities attract migrants, either because of better jobs or the possibility of subsequent migration to the United States.

In this time period, the presence of a higher share of indigenous at both the origin and destination is associated with migration At the origin, the indigenous are likely to live in communities where there are little other economic opportunities than subsistence agriculture. At the destination: the indigenous, more than other groups, migrate where other indigenous are present. Figure 4.9 seems to confirm this trend. Visually it suggests that indigenous are migrating from their traditional communities in southern Mexico to work in the tourism industry around Cancun on the Yucatan Peninsula. An increase in 10% of the share at the origin increases the likelihood of a migrant corridor by 0.7%; an increase in 10% of the share at the destination increases the likelihood of a migrant corridor by 0.6%.

In this model, the wage difference factors do not matter as much. The impact of a one point base salary difference drops to 1.3% and remains significant. The impact of the return to skill difference loses significance. I can hypothesize that large, urban cities, especially on the border, offer the sort of labor markets that would have higher base salaries and reward education. Thus the structural gravity factors absorb the impact of the differences in labor models.

Overall, specification 2 has much more predictive power than specification 1, with an R^2 of 0.262.

Specification 3: Social Networks

The third specification reveals the core result of this essay. Here I augment the previous specification with the logged size of the social network, the migrant stock from the given origin state in the destination municipality at the beginning of the time period.

An increase of one log point in the size of the social network increases the probability of a migrant flow by 12%. Overall, specification 3 has even more predictive power than specification 2, with an $R^2 of 0.393$.

When I compare specification 3 to specification 2, the impact of the other factors in the model decreases by half: a one point impact in base salary, the presence of the origin or destination on the border, or the impact of the urban share at the origin or destination. These factors remain significant.

In the case of indigenous share, the impact at the origin loses significance, while the impact at the destination decreases by half. This result suggests that social networks especially play a role in the migration of the indigenous.

4.5.2 Internal Migration at the Intensive Margin

Next I use the same three specifications to analyze migration at the intensive margin: the magnitude of the internal migrant flow for the subset of possible migration corridors which are activated in a given time period.

Specifications (1) and (2) function in the same way their counterparts in specifications (1) and (2) in the extensive margin case. Initially, differences in the base salary and return to skill seem to drive migration but the introduction of structural gravity factors dramatically reduces the explanatory power of these factors. The same structural factors that drove the presence of a migrant flow also drive the magnitude of the flow.

Examining the role of social networks in specification (3) confirms the core result of this essay. In this log-log model, I can interpret this coefficient as a social network elasticity. An increase in 1% of the size of the social network in one time period causes an increase in 0.3% of the magnitude of the migrant flow in the subsequent time period. The addition of social networks to the model decreases the magnitude of the effects of other factors by half or more. In particular, it eliminates the effect of indigenous share at the origin or urban share at the destination, suggesting the size of the social network accounts for the variation previously explained by these factors. The predictive power of this model is quite high at 0.582.

4.5.3 Migration Climate

For the potential migrant, internal migration and international migration are related. In practice, the decision to internally migrate takes into consideration both the expected value of staying as well as the expected value of international migration. An extension of this model would incorporate all three options—staying, internal migration, and internal migration—into one integrated model.

For the three periods in question, I note certain contextual factors through which the outside option of US-Mexico migration varies in the three periods that we study. Villarreal (2014) and Durand and Massey (2019) provide more background on these trends.

1. From 1995-2000, migration from Mexico to the US was increasing, owing to the recent passage of NAFTA and a relatively porous border.

2. From 2000-2005, border security increased as a result of the September 11 attacks. The Mexican economy continued to recover from the 1994 ”tequila crisis” and associated devaluation of the peso.

3. From 2010-2015, the global economic recession of 2008 caused the return of an estimated 500,000 temporary migrants from the US to Mexico. US-Mexico migration peaked and begun to decline.

Examining the determinants of internal migration in these three time periods will allow us to indirectly examine the overall ”migration climate” of individuals and the relative ease or difficulty of these outside options.

In addition to the value of the outside option, changing conditions in the Mexican economy have also affected internal migration. As part of an overall shift in Mexico's economy from rural agriculture to urban industry, the labor market of the rural agricultural sector has changed. Residents of rural areas have sorted into productive farmers, who continue to make a profit despite challenging market environments, and non-productive farmers, who have abandoned subsistence farming in search of other opportunities. These other opportunities include local non-agricultural work and migration within Mexico in addition to international migration (Charlton & Taylor, 2016). Our use of migration flows from Mexican population census data complements the household-level migration histories from the National Rural Mexican Household Survey panel that these authors use. We do not examine international migration flows from the origin states directly. Nevertheless, the changes we observe in the determinants of migration over the three time periods of interest match these authors' conclusions about the increasing availability of high-skilled non-farm job opportunities in Mexico.

Table 4.7: Extensive Margin of Internal Migration Across Time

	Dependent variable:		
	Internal Migrant Flow Present		
	1995-2000	2000-2005	2010-2015
	(1)	(2)	(3)
Base Salary Diff	0.006^{*} (0.004)	0.006^{**} (0.003)	0.012^{***} (0.003)
Return to Skill Diff	-0.084 (0.088)	0.039 (0.035)	0.153^{**} (0.065)
Distance (log)	-0.170^{***} (0.008)	-0.055^{***} (0.006)	-0.110^{***} (0.009)
Dest Population (log)	0.114^{***} (0.009)	0.055^{***} (0.010)	0.034^{***} (0.005)
Origin on Border	0.082^{***} (0.015)	0.036^{***} (0.011)	0.057^{***} (0.016)
Dest on Border	0.215^{***} (0.044)	0.085^{***} (0.029)	0.052 (0.036)
Origin Urban Share	0.126^{***} (0.041)	0.101^{***} (0.013)	0.151^{***} (0.044)
Dest Urban Share	0.162^{***} (0.025)	0.045^{***} (0.012)	0.057^{***} (0.009)
Origin Indig Share		0.024 (0.022)	-0.003 (0.028)
Dest Indig Share		0.029^{***} (0.010)	-0.005 (0.008)
Social Network (log)	0.046^{***} (0.005)	0.119^{***} (0.005)	0.126^{***} (0.009)
Constant	1.197^{***} (0.095)	0.202^{***} (0.070)	1.180^{***} (0.101)
Observations	19,437	61,318	71,579
R^2	0.323	0.393	0.324
Adjusted R^2	0.323	0.393	0.324
Residual Std. Error	0.391 (df = 19427)	0.288 (df = 61306)	0.352 (df = 71567)

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Standard errors clustered by destination state

Table 4.8: Intensive Margin of Internal Migration Across Time

	Dependent variable:		
	Internal Migrant Flow (log)		
	1995-2000	2000-2005	2010-2015
	(1)	(2)	(3)
Base Salary Diff	0.039^{***} (0.004)	0.022^{***} (0.003)	0.087^{***} (0.003)
Return to Skill Diff	−0.059 (0.088)	-0.359^{***} (0.035)	0.619^{***} (0.065)
Distance (log)	-0.481^{***} (0.008)	-0.103^{***} (0.006)	-0.193^{***} (0.009)
Dest Population (log)	0.609^{***} (0.009)	0.229^{***} (0.010)	0.437^{***} (0.005)
Origin on Border	0.017 (0.015)	0.110^{***} (0.011)	0.056^{***} (0.016)
Dest on Border	0.832^{***} (0.044)	0.324^{***} (0.029)	0.142^{***} (0.036)
Origin Urban Share	0.707^{***} (0.041)	0.196^{***} (0.013)	0.338^{***} (0.044)
Dest Urban Share	0.476^{***} (0.025)	0.027^{**} (0.012)	0.138^{***} (0.009)
Origin Indig Share		0.175^{***} (0.022)	0.052^{*} (0.028)
Dest Indig Share		0.295^{***} (0.010)	-0.091^{***} (0.008)
Social Network (log)	0.190^{***} (0.005)	0.304^{***} (0.005)	0.319^{***} (0.009)
Constant	1.605^{***} (0.095)	1.068^{***} (0.070)	-0.375^{***} (0.101)
Observations	6,654	10,052	17,307
R^2	0.536	0.582	0.675
Adjusted R^2	0.535	0.581	0.675
Residual Std. Error	1.119 (df = 6644)	0.749 (df = 10040)	0.835 (df = 17295)

Note: *p<0.1; **p<0.05; ***p<0.01

Standard errors clustered by destination state

4.5.4 Social Networks Across Time

Next I use our structural gravity and social networks model across the three time periods of interest at the extensive and intensive margin. I present the results in Table 4.5.3 and 4.5.3.

The effect of the social network is monotonically increasing across time periods. A one log point increase in the size of the social network causes a 4.1%, 11.6%, or 12.7% increase in the probability of a migrant flow, respectively. Conditional on the presence of a migrant flow, a 1% increase in the size of the network increases the flow by 17.3%, 28.6%, or 30.6%.

Changes in the effect and significance of the other factors in the model reflect changing economic conditions in Mexico.

Base salary difference matters slightly in all three time periods, while return to skill only matters in the third time period. This result matches other literature like Charlton and Taylor (2016) which suggests that Mexico is in the late stages of a structural transition where education and non-farm opportunities have increased together.

Distance matters less over time. I imagine the increasing presence of communications technology like cell phones and the Internet as well as the ease of travel within Mexico reducing the effect of physical and psychic costs that distance models.

I see urban-urban migration across all three time periods at both margins. The presence of this migration instead of purely rural-urban migration contributes to the evidence for a structural transformation in Mexico.

In contrast, the effect of a destination on the border drops across three time periods. This result matches US-Mexico migration trends, which peaked in 2007. It suggests decreasing transit migration from the first to the second time period and an absence of this type of migration in the third time period.

In the two time periods for which I have data, indigenous migration above and beyond general trends plays a role in the first but not the second. In 2005, the indigenous tended to migrate where other indigenous were present, but not in 2015 ten years later.

Table 4.9: Extensive Margin of Internal Migration Across Time

	Dependent variable:		
	Internal Migrant Flow Present		
	1995-2000	2000-2005	2010-2015
	(1)	(2)	(3)
Base Salary Diff	0.006 (0.015)	0.006 (0.018)	0.012 (0.041)
Return to Skill Diff	−0.084 (0.299)	0.044 (0.221)	0.152 (0.478)
Distance (log)	-0.163^{***} (0.029)	-0.055^{**} (0.026)	-0.110^{***} (0.024)
Dest Population (log)	0.105^{***} (0.026)	0.053^{***} (0.021)	0.034 (0.029)
Origin on Border	0.077 (0.062)	0.035 (0.040)	0.057 (0.039)
Dest on Border	0.190 (0.158)	0.077 (0.107)	0.053 (0.101)
Origin Urban Share	0.134 (0.198)	0.105 (0.092)	0.151 (0.114)
Dest Urban Share	0.159 (0.119)	0.045 (0.074)	0.057 (0.040)
Origin Indig Share		0.027 (0.145)	−0.003 (0.095)
Dest Indig Share		0.028 (0.076)	−0.005 (0.066)
Social Network (log)	0.041^{***} (0.012)	0.116^{***} (0.017)	0.127^{***} (0.013)
Stock Present in 1960	0.141	0.073	−0.008
Constant	1.183^{***} (0.388)	0.207 (0.296)	1.180^{***} (0.227)
Observations	19,437	61,318	71,579
R^2	0.328	0.395	0.324
Adjusted R^2	0.328	0.394	0.324
Residual Std. Error	0.389 (df = 19426)	0.288 (df = 61305)	0.352 (df = 71566)

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors clustered by destination state

Table 4.10: Intensive Margin of Internal Migration Across Time

	Dependent variable:		
	Internal Migrant Flow (log)		
	1995-2000	2000-2005	2010-2015
	(1)	(2)	(3)
Base Salary Diff	0.036^{**} (0.015)	0.027 (0.018)	0.094^{**} (0.041)
Return to Skill Diff	−0.087 (0.299)	−0.197 (0.221)	0.766 (0.478)
Distance (log)	-0.411^{***} (0.029)	-0.082^{***} (0.026)	-0.182^{***} (0.024)
Dest Population (log)	0.535^{***} (0.026)	0.200^{***} (0.021)	0.421^{***} (0.029)
Origin on Border	−0.033 (0.062)	0.086^{**} (0.040)	0.052 (0.039)
Dest on Border	0.690^{***} (0.158)	0.255^{**} (0.107)	0.072 (0.101)
Origin Urban Share	0.768^{***} (0.198)	0.265^{***} (0.092)	0.384^{***} (0.114)
Dest Urban Share	0.463^{***} (0.119)	0.030 (0.074)	0.139^{***} (0.040)
Origin Indig Share		0.260^{*} (0.145)	0.102 (0.095)
Dest Indig Share		0.285^{***} (0.076)	−0.085 (0.066)
Social Network (log)	0.175^{***} (0.012)	0.286^{***} (0.017)	0.306^{***} (0.013)
Stock Present in 1960	0.613	0.388	0.324
Constant	1.391^{***} (0.388)	1.052^{***} (0.296)	-0.405^{*} (0.227)
Observations	6,654	10,052	17,307
R^2	0.554	0.594	0.679
Adjusted R^2	0.553	0.593	0.678
Residual Std. Error	1.097 (df = 6643)	0.738 (df = 10039)	0.830 (df = 17294)

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors clustered by destination state

4.5.5 Migration Taste Factor

As I mention in the section 4.4.3 one threat to the identification of the effect of social networks is the presence of serial correlation: some unobserved factor that affects a migration flow in one period and in the next.

I use a simple model of such a factor that we call a taste for a particular migration corridor. Tables 4.5.4 and 4.5.4 show revised estimations with the inclusion of this taste factor. It is significant across 5 of the 6 estimated equations, which indicates the role of taste in migration. On the other hand, it does not change appreciably the magnitude, sign, or significance of our results at the extensive margin. At the intensive margin, it brings down the magnitude of the social network coefficient by approximately 10%. These results suggest that the effects of social networks that I have estimated are not caused by a mere taste for certain migration corridors. Further work could examine the role of this taste for migration in the initial conditions that jumpstarted these social networks.

4.6 Conclusion

This essay has used a structural gravity model to estimate the effect of social networks on the extensive margins and intensive margins of migrant flows from origin states to destination municipalities in Mexico over the time periods 1995-2000, 2000-2005, and 2010-2015. The extensive margin refers to the presence of a migration corridor; the intensive margin refers to the magnitude of migrant flow through this corridor. It uses a sample of working age men from 25-55 who would most likely migrate for economic instead of non-economic reasons to compare the explanatory power of a Roy model, a structural gravity model without social networks, and a structural gravity model with social networks. In the third model, it aims to identify the effect of the social networks.

I find two main results.

First, in all three time periods, the model reveals a rich set of factors other than wage differences and return to skill that are associated with the presence of migration corridors and the magnitude of migrant flow through them. These factors include urbanization share, indigenous share, and presence on the US/Mexico border. I am interested in particular in identifying the effect of the presence and size of a social network of migrants from the same origin state who migrated in the previous time period. To identify this effect, I use the presence of a migrant flow from 1960 along the same corridor to control for serially-correlated unobservables that could influence the size of the social network across time periods. I call this control a "taste for migration" along a particular corridor.

In a representative time period (2000-2005), a model that includes only differences in base salary and return to skill explains 2.5% of the variation in the presence of migration corridors; additional structural gravity factors 26%; and the social network 39% The corresponding models of the magnitude of the migrant flow explain 5%, 39%, and 58% of the variation, respectively. Thus social networks add explanatory power to structural gravity models at the extensive and intensive margins. Both models vastly outperform standard Roy migration models that focus only on individual utility maximization.

Moreover, across all three time periods, the effect of social networks is monotonically increasing. At the extensive margin, a 1% increase in the size of the social network increases by 5%, 12%, and 13% the likelihood of a migration corridor. At the intensive margin, the equivalent social network elasticities are 19%, 30%, and 32%.

Estimating these models separately for each time period reveals changes in the Mexican economy that affect internal migration: including increased educational opportunities and the decreasing appeal of migration to the US.

These results contribute to a strand of literature in the economics of migration that argues not only for the importance of considering social networks or diaspora effects but that these effects dominate economic effects as drivers of migration. Moreover, they do so in a novel context of internal migration instead of international migration.

One novelty of our approach, the use of Mexican census data, is also its limitation. Since we use the nationally representative Mexican census collected every five years instead of smaller scale panel data surveys that ask for more detailed migration histories, we can only account for long-term permanent migration that occurs at most one time per time period. Moreover, we consider here the weak ties of migrants from the same state instead of using a more granular measure of migrants from the same municipality.

Moreover, though we restrict the sample to men aged 25-55 who would tend to migrate for work, we still observe that many of the individuals in this sample migrate for reasons other than employment. Other data sources, such as the quarterly Mexican Survey of Occupation and Employment (ENOE), would provide a targeted look at internal migrants who obtain formal employment at a much higher temporal frequency. These data sources would provide the ability to understand more deeply the mechanism by which new migrants help existing migrants find jobs.

Finally, further research in the origin of migrant networks in specific contexts such as Mexico is needed to provide a better understanding for what causes corridors to develop in the first place. Ideally, this research would reveal an instrument that could be used to credibly identify the effect of these social networks.

Nevertheless, the results here provide useful tools for several sets of actors as they seek to predict and respond to internal migration trends: local and state governments in receiving communities that must accommodate new populations; export manufacturing factories and other sources of employment for new migrants; and NGOs that facilitate their integration into receiving communities.

Bibliography

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2023). When Should You Adjust Standard Errors for Clustering?*. *The Quarterly Journal of Economics*, *138*(1), 1–35. https://doi.org/10.1093/qje/qjac038

Alemu, D., Guinan, A., & Hermanson, J. (2021). Contract farming, cooperatives and challenges of side selling: Malt barley value-chain development in Ethiopia. *Development in Practice*, *31*(4), 496–510. https://doi.org/10.1080/09614524.2020.1860194

Anderson, J. E. (2011). The gravity model. *Annual Review of Economics*, *3*(1), 133–160. https://doi.org/10.1146/annurev-economics-111809-125114

Anderzén, J., Guzmán Luna, A., Luna-González, D. V., Merrill, S. C., Caswell, M., Méndez, V. E., Hernández Jonapá, R., & Mier y Terán Giménez Cacho, M. (2020). Effects of on-farm diversification strategies on smallholder coffee farmer food security and income sufficiency in Chiapas, Mexico. *Journal of Rural Studies*, *77*, 33–46. https://doi.org/10.1016/j.jrurstud.2020.04.001

Anderzén, J., Vandame, R., Ocampo, B., Merrill, S. C., Jonapá, R. H., Nájera, O. A., Anderson, C. R., & Méndez, V. E. (2024). Multiple values of beekeeping (with A. mellifera) as an element of diversified, agroecological coffee farms in Chiapas, Mexico. *Renewable Agriculture and Food Systems*, *39*, e29. https://doi.org/10.1017/S1742170524000164

Arana-Coronado, J. J., Trejo-Pech, C. O., Velandia, M., & Peralta-Jimenez, J. (2019). Factors Influencing Organic and Fair Trade Coffee Growers Level of Engagement with Cooperatives: The Case of Coffee Farmers in Mexico. *Journal of International Food & Agribusiness Marketing*, *31*(1), 22–51. https://doi.org/10.1080/08974438.2018.1471637

Arends-Kuenning, M., Baylis, K., & Garduño-Rivera, R. (2019). The effect of NAFTA on internal migration in mexico: A regional economic analysis. *Applied Economics*, *51*(10), 1052–1068. https://doi.org/10.1080/00036846.2018.1524976

Asad, A. L., & Garip, F. (2019). Mexico-u.s. migration in time: From economic to social mechanisms. *The ANNALS of the American Academy of Political and Social Science*, *684*(1), 60–84. https://doi.org/10.1177/0002716219847148

Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2013). The Diffusion of Microfinance. *Science*, *341*(6144), 1236498. https://doi.org/10.1126/science.1236498

Banerjee, B. (1991). The determinants of migrating with a pre-arranged job and of the initial duration of urban unemployment: An analysis based on indian data on rural-to-urban migrants. *Journal of Development Economics*, *36*(2), 337–351. https://doi.org/10.1016/0304-3878(91)90040-3

Barrett, C. B. (1996). On price risk and the inverse farm size-productivity relationship. *Journal of Development Economics*, *51*(2), 193–215. https://doi.org/10.1016/S0304-3878(96)00412-9

Beaman, L., BenYishay, A., Magruder, J., & Mobarak, A. M. (2021). Can Network Theory-Based Targeting Increase Technology Adoption? *American Economic Review*, *111*(6), 1918–1943. https://doi.org/10.1257/aer.20200295

Beine, M., Bertoli, S., & Fernández-Huertas Moraga, J. (2016). A practitioners' guide to gravity models of international migration. *The World Economy*, *39*(4), 496–512. https://doi.org/10.1111/twec.12265

Beine, M., Docquier, F., & Özden, Ç. (2011). Diasporas. *Journal of Development Economics*, *95*(1), 30–41. https://doi.org/10.1016/j.jdeveco.2009.11.004

Bellemare, M. F., Barrett, C. B., & Just, D. R. (2013). The Welfare Impacts of Commodity Price Volatility: Evidence from Rural Ethiopia. *American Journal of Agricultural Economics*, *95*(4), 877–899. https://doi.org/10.1093/ajae/aat018

Bellemare, M. F., & Bloem, J. R. (2018). Does contract farming improve welfare? A review. *World Development*, *112*, 259–271. https://doi.org/10.1016/j.worlddev.2018.08.018

Bellemare, M. F., Lee, Y. N., & Just, D. R. (2020). Producer Attitudes Toward Output Price Risk: Experimental Evidence from the Lab and from the Field. *American Journal of Agricultural Economics*, *102*(3), 806–825. https://doi.org/10.1002/ajae.12004

Bellemare, M. F., Lee, Y. N., & Novak, L. (2021). Contract farming as partial insurance. *World Development*, *140*, 105274. https://doi.org/10.1016/j.worlddev.2020.105274

Bernard, T., & Spielman, D. J. (2009). Reaching the rural poor through rural producer organizations? A study of agricultural marketing cooperatives in Ethiopia. *Food Policy*, *34*(1), 60–69. https://doi.org/10.1016/j.foodpol.2008.08.001

Bhuyan, S. (2007). The "People" Factor in Cooperatives: An Analysis of Members' Attitudes and Behavior. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie*, *55*(3), 275–298. https://doi.org/10.1111/j.1744-7976.2007.00092.x

Binswanger, H. P. (1980). Attitudes Toward Risk: Experimental Measurement in Rural India. *American Journal of Agricultural Economics*, *62*(3), 395–407. https://doi.org/10.2307/1240194

Bobrow-Strain, A. (2007). *Intimate enemies: Landowners, power, and violence in Chiapas*. Duke University Press. Retrieved July 14, 2023, from http://site.ebrary.com/id/10217154

Borjas, G. J. (1987). Self-selection and the earnings of immigrants. *The American Economic Review*, *77*(4), 531–553.

Boyd, C. M., & Bellemare, M. F. (2020). The Microeconomics of Agricultural Price Risk. *Annual Review of Resource Economics*, *12*(1), 149–169. https://doi.org/10.1146/annurev-resource-100518-093807

Boyd, C. M., & Bellemare, M. F. (2022). Why not insure prices? Experimental evidence from Peru. *Journal of Economic Behavior & Organization*, *202*, 580–631. https://doi.org/10.1016/j.jebo.2022.08.004

Bramoullé, Y., Djebbari, H., & Fortin, B. (2009). Identification of peer effects through social networks. *Journal of Econometrics*, *150*(1), 41–55. https://doi.org/10.1016/j.jeconom.2008.12.021

Bramoullé, Y., Djebbari, H., & Fortin, B. (2020). Peer Effects in Networks: A Survey. *Annual Review of Economics*, *12*(1), 603–629. https://doi.org/10.1146/annurev-economics-020320-033926

Caputo, V., & Just, D. R. (2022). The economics of food related policies: Considering public health and malnutrition. In *Handbook of Agricultural Economics* (pp. 5117–5200). Elsevier. https://doi.org/10.1016/bs.hesagr.2022.03.008

Cardenas, J. C., & Carpenter, J. (2013). Risk attitudes and economic well-being in Latin America. *Journal of Development Economics*, *103*, 52–61. https://doi.org/10.1016/j.jdeveco.2013.01.008

Carleton, T., Duflo, E., Jack, B. K., & Zappalà, G. (2024). Chapter 4 - Adaptation to climate change. In L. Barrage & S. Hsiang (Eds.), *Handbook of the Economics of Climate Change* (pp. 143–248). North-Holland. https://doi.org/10.1016/bs.hesecc.2024.10.001

Carrington, W. J., Detragiache, E., & Vishwanath, T. (1996). Migration with endogenous moving costs. *The American Economic Review*, *86*(4), 909–930.

Casaburi, L., & Macchiavello, R. (2015). Loyalty, Exit, and Enforcement: Evidence from a Kenya Dairy Cooperative. *American Economic Review*, *105*(5), 286–290. https://doi.org/10.1257/aer.p20151076

Casaburi, L., & Reed, T. (2022). Using Individual-Level Randomized Treatment to Learn about Market Structure. *American Economic Journal: Applied Economics*, *14*(4), 58–90. https://doi.org/10.1257/app.20200306

Charlton, D., & Taylor, J. E. (2016). A declining farm workforce: Analysis of panel data from rural mexico. *American Journal of Agricultural Economics*, *98*(4), 1158–1180. https://doi.org/10.1093/ajae/aaw018

Charness, G., Gneezy, U., & Imas, A. (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization*, *87*, 43–51. https://doi.org/10.1016/j.jebo.2012.12.023

Chiquiar, D. (2005). Why mexico's regional income convergence broke down. *Journal of Development Economics*, *77*(1), 257–275. https://doi.org/10.1016/j.jdeveco.2004.03.009

Chiquiar, D., & Hanson, G. H. (2005). International migration, self-selection, and the distribution of wages: Evidence from mexico and the united states. *Journal of Political Economy*, *113*(2), 239–281. https://doi.org/10.1086/427464

Conley, T. G., & Udry, C. R. (2010). Learning about a New Technology: Pineapple in Ghana. *American Economic Review*, *100*(1), 35–69. https://doi.org/10.1257/aer.100.1.35

Cuecuecha, A., & Pederzini, C. (2014). *Migration and remittances from mexico: Trends, impacts, and new challenges*.

Davis, B., Stecklov, G., & Winters, P. (2002). Domestic and international migration from rural mexico: Disaggregating the effects of network structure and composition. *Population Studies*, *56*(3), 291–309.

Dell, M., Jones, B. F., & Olken, B. A. (2014). What Do We Learn from the Weather? The New Climate-Economy Literature. *Journal of Economic Literature*, *52*(3), 740–798. https://doi.org/10.1257/jel.52.3.740

Dobler-Morales, C., & Bocco, G. (2021). Social and environmental dimensions of drought in Mexico: An integrative review. *International Journal of Disaster Risk Reduction*, *55*, 102067. https://doi.org/10.1016/j.ijdrr.2021.102067

Dragusanu, R., Giovannucci, D., & Nunn, N. (2014). The Economics of Fair Trade. *Journal of Economic Perspectives*, *28*(3), 217–236. https://doi.org/10.1257/jep.28.3.217

Drake, M., Payró, F., Thakral, N., & Tô, L. T. (2024). Bayesian Adaptive Choice Experiments.

Durand, J., & Massey, D. S. (2019). Evolution of the mexico-u.s. migration system: Insights from the mexican migration project. *The ANNALS of the American Academy of Political and Social Science*, *684*(1), 21–42. https://doi.org/10.1177/0002716219857667

Eakin, H., Benessaiah, K., Barrera, J. F., Cruz-Bello, G. M., & Morales, H. (2012). Livelihoods and landscapes at the threshold of change: Disaster and resilience in a Chiapas coffee community. *Regional Environmental Change*, *12*(3), 475–488. https://doi.org/10.1007/s10113-011-0263-4

Eckel, C. C., & Grossman, P. J. (2008). Chapter 113 Men, Women and Risk Aversion: Experimental Evidence. In C. R. Plott & V. L. Smith (Eds.), *Handbook of Experimental Economics Results* (pp. 1061–1073). Elsevier. https://doi.org/10.1016/S1574-0722(07)00113-8

Eckel, C. C., & Londono, N. C. (2021). How to Tame Lab-in-the-Field Experiments. In D. P. Green & J. N. Druckman (Eds.), *Advances in Experimental Political Science* (pp. 79–102). Cambridge University Press. https://doi.org/10.1017/9781108777919.007

Ewusi Koomson, J., Donkor, E., & Owusu, V. (2022). Contract farming scheme for rubber production in Western region of Ghana: Why do farmers side sell? *Forests, Trees and Livelihoods*, *31*(3), 139–152. https://doi.org/10.1080/14728028.2022.2079007

Fafchamps, M., & Hill, R. V. (2005). Selling at the Farmgate or Traveling to Market. *American Journal of Agricultural Economics*, *87*(3), 717–734. https://doi.org/10.1111/j.1467-8276.2005.00758.x

Falaris, E. M. (1987). A nested logit migration model with selectivity. *International Economic Review*, *28*(2), 429–443. https://doi.org/10.2307/2526735

Finkelshtain, I., & Chalfant, J. A. (1991). Marketed Surplus under Risk: Do Peasants Agree with Sandmo? *American Journal of Agricultural Economics*, *73*(3), 557–567. https://doi.org/10.2307/1242809

Fischer, E., & Qaim, M. (2014). Smallholder Farmers and Collective Action: What Determines the Intensity of Participation? *Journal of Agricultural Economics*, *65*(3), 683–702. https://doi.org/10.1111/1477-9552.12060

Folch, A., & Planas, J. (2019). Cooperation, Fair Trade, and the Development of Organic Coffee Growing in Chiapas (1980–2015). *Sustainability*, *11*(2), 357. https://doi.org/10.3390/su11020357

Foster, A. D., & Rosenzweig, M. R. (1995). Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy*, *103*(6), 1176–1209. https://doi.org/10.1086/601447

Foster, A. D., & Rosenzweig, M. R. (2010). Microeconomics of Technology Adoption. *Annual Review of Economics*, *2*(1), 395–424. https://doi.org/10.1146/annurev.economics.102308.124433

Frijters, P., Kong, T. S., & Liu, E. M. (2015). Who is coming to the artefactual field experiment? Participation bias among Chinese rural migrants. *Journal of Economic Behavior & Organization*, *114*, 62–74. https://doi.org/10.1016/j.jebo.2015.03.006

Garip, F., & Asad, A. L. (2016). Network effects in mexico–u.s. migration: Disentangling the underlying social mechanisms. *American Behavioral Scientist*, *60*(10), 1168–1193. https://doi.org/10.1177/0002764216643131

Geng, X., Janssens, W., & Kramer, B. (2023). Liquid milk: Savings, insurance and side-selling in cooperatives. *Journal of Development Economics*, *165*, 103–142. https://doi.org/10.1016/j.jdeveco.2023.103142

Gerard, A., Lopez, M. C., Clay, D. C., & Ortega, D. L. (2021). Farmer cooperatives, gender and side-selling behavior in Burundi's coffee sector. *Journal of Agribusiness in Developing and Emerging Economies*, *11*(5), 490–505. https://doi.org/10.1108/JADEE-05-2020-0081

Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, *78*(6), 1360–1380.

Halleck Vega, S., & Elhorst, J. P. (2015). THE SLX MODEL. *Journal of Regional Science*, *55*(3), 339–363. https://doi.org/10.1111/jors.12188

Hanson, G. H. (2001). U.s.–mexico integration and regional economies: Evidence from border-city pairs. *Journal of Urban Economics*, *50*(2), 259–287. https://doi.org/10.1006/juec.2001.2217

Harris, J. R., & Todaro, M. P. (1970). Migration, unemployment and development: A two-sector analysis. *The American Economic Review*, *60*(1), 126–142.

Harrison, G. W., & List, J. A. (2004). Field Experiments. *Journal of Economic Literature*, *42*(4), 1009–1055. https://doi.org/10.1257/0022051043004577

Holt, C. A., & Laury, S. K. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, *92*(5), 1644–1655. https://doi.org/10.1257/000282802762024700

Hopfensitz, A., & Miquel-Florensa, J. (2017). Mill ownership and farmer's cooperative behavior: The case of Costa Rica coffee farmers. *Journal of Institutional Economics*, *13*(3), 623–648. https://doi.org/10.1017/S1744137416000527

Ibarraran, P., & Lubotsky, D. (2007). Mexican immigration and self-selection: New evidence from the 2000 mexican census. *Mexican immigration to the United States*, 159–192.

Keenan, M., Fort, R., & Vargas, R. (2024). Shocked into side-selling? Production shocks and organic coffee farmers' marketing decisions. *Food Policy*, *125*. https://doi.org/10.1016/j.foodpol.2024.102631

Kinnan, C., Samphantharak, K., Townsend, R., & Vera-Cossio, D. (2024). Propagation and Insurance in Village Networks. *American Economic Review*, *114*(1), 252–284. https://doi.org/10.1257/aer.20220892

Klein, K. K., Richards, T. J., & Walburger, A. (1997). Determinants of Co-operative Patronage in Alberta. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie*, *45*(2), 93–110. https://doi.org/10.1111/j.1744-7976.1997.tb00195.x

Liverpool-Tasie, L. S. O., Wineman, A., Young, S., Tambo, J., Vargas, C., Reardon, T., Adjognon, G. S., Porciello, J., Gathoni, N., Bizikova, L., Galiè, A., & Celestin, A. (2020). A scoping review of market links between value chain actors and small-scale producers in developing regions. *Nature Sustainability*, *3*(10), 799–808. https://doi.org/10.1038/s41893-020-00621-2

Lucas, R. E. B. (2021, November 18). *Crossing the divide: Rural to urban migration in developing countries*. https://doi.org/10.1093/oso/9780197602157.001.0001

Macchiavello, R., & Morjaria, A. (2015). The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports. *American Economic Review*, *105*(9), 2911–2945. https://doi.org/10.1257/aer.20120141

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, *60*(3), 531–542. https://doi.org/10.2307/2298123

Markelova, H., Meinzen-Dick, R., Hellin, J., & Dohrn, S. (2009). Collective action for smallholder market access. *Food Policy*, *34*(1), 1–7. https://doi.org/10.1016/j.foodpol.2008.10.001

Martin, P. (2020, July 10). *Wilson center: Mexican braceros and US farm workers* [Wilson center]. Retrieved May 8, 2022, from https://www.wilsoncenter.org/article/mexican-braceros-and-us-farm-workers

Martinez-Torres, M. E. (2006). *Organic Coffee: Sustainable Development by Mayan Farmers* (Paper edition). Ohio University Press.

Mattos, F. L., & Zinn, J. (2016). Formation and adaptation of reference prices in grain marketing: An experimental study. *Agricultural Economics*, *47*(6), 621–632. https://doi.org/10.1111/agec.12260

McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics*, *99*(2), 210–221. https://doi.org/10.1016/j.jdeveco.2012.01.002

McKenzie, D., & Rapoport, H. (2010). Self-selection patterns in mexico-u.s. migration: The role of migration networks. *The Review of Economics and Statistics*, *92*(4), 811–821.

Michler, J. D., & Wu, S. Y. (2020). Relational Contracts in Agriculture: Theory and Evidence. *Annual Review of Resource Economics*, *12*(1), 111–127. https://doi.org/10.1146/annurev-resource-101719-034514

Miles, C. H., Petersen, M., & van der Laan, M. J. (2019). Causal inference when counterfactuals depend on the proportion of all subjects exposed. *Biometrics*, *75*(3), 768–777. https://doi.org/10.1111/biom.13034

Millimet, D., & Bellemare, M. F. (2023). Fixed Effects and Causal Inference. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4467963

Mojo, D., Fischer, C., & Degefa, T. (2017). The determinants and economic impacts of membership in coffee farmer cooperatives: Recent evidence from rural Ethiopia. *Journal of Rural Studies*, *50*, 84–94. https://doi.org/10.1016/j.jrurstud.2016.12.010

Morgan, S. L., & Winship, C. (n.d.). *Counterfactuals and causal inference: Methods and principles for social research* (2nd edition). Cambridge University Press.

Moya, A. (2018). Violence, psychological trauma, and risk attitudes: Evidence from victims of violence in Colombia. *Journal of Development Economics*, *131*, 15–27. https://doi.org/10.1016/j.jdeveco.2017.11.001

Mujawamariya, G., D'Haese, M., & Speelman, S. (2013). Exploring double side-selling in cooperatives, case study of four coffee cooperatives in Rwanda. *Food Policy*, *39*, 72–83. https://doi.org/10.1016/j.foodpol.2012.12.008

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the u. s. labor market*. *The Quarterly Journal of Economics*, *118*(2), 549–599. https://doi.org/10.1162/003355303321675455

Munshi, K. (2014). Community Networks and the Process of Development. *Journal of Economic Perspectives*, *28*(4), 49–76. https://doi.org/10.1257/jep.28.4.49

Munshi, K. (2020). Social networks and migration. *Annual Review of Economics*, *12*(1), 503–524. https://doi.org/10.1146/annurev-economics-082019-031419

Newbery, D. M. G., & Stiglitz, J. E. (1981). *The theory of commodity price stabilization: A study in the economics of risk*. Clarendon Press.

Nyamamba, J. S., Ayuya, O. I., & Sibiko, K. W. (2022). Determinants of side selling behaviour in emerging sorghum supply chains in Kisumu County, Kenya. *Cogent Economics & Finance*, *10*(1), 1986932. https://doi.org/10.1080/23322039.2021.1986932

Ochoa, L. C. M., Ponce, R. A. C., & Espinosa, M. d. L. R. (2018). Migración interna en México. *Paradigma económico. Revista de economía regional y sectorial*, *10*(2), 39–59.

Ortiz-Bobea, A., Ault, T. R., Carrillo, C. M., Chambers, R. G., & Lobell, D. B. (2021). Anthropogenic climate change has slowed global agricultural productivity growth. *Nature Climate Change*, *11*(4), 306–312. https://doi.org/10.1038/s41558-021-01000-1

Palm-Forster, L. H., & Messer, K. D. (2021, January 1). Chapter 80 - Experimental and behavioral economics to inform agri-environmental programs and policies. In C. B. Barrett & D. R. Just (Eds.), *Handbook of Agricultural Economics* (pp. 4331–4406). Elsevier. https://doi.org/10.1016/bs.hesagr.2021.10.006

Papke, L. E., & Wooldridge, J. M. (1996). Econometric Methods for Fractional Response Variables With an Application to 401 (K) Plan Participation Rates. *Journal of Applied Econometrics*, *11*(6), 619–632. http://www.jstor.org/stable/2285155

Pfeiffer, L., López-Feldman, A., & Taylor, J. E. (2009). Is off-farm income reforming the farm? Evidence from Mexico. *Agricultural Economics*, *40*(2), 125–138. https://doi.org/10.1111/j.1574-0862.2009.00365.x

Pitts, S. (2019). Value Chain Integration as an Alternative to Fair Trade for Chiapas Coffee Farmers. In B. S. Sergi & C. C. Scanlon (Eds.), *Entrepreneurship and Development in the 21st Century* (pp. 103–138). Emerald Publishing Limited. https://doi.org/10.1108/978-1-78973-233-720191007

Pitts, S. (2023). The role of the farmer and their cooperative in supply chain governance: A Latin American small producer perspective. In *Handbook of Research on Cooperatives and Mutuals* (pp. 172–184). Edward Elgar Publishing. https://doi.org/10.4337/9781802202618.00019

Ramos, R. (2016). Gravity models: A tool for migration analysis. *IZA World of Labor*. https://doi.org/10.15185/izawol.239

Reglas de Operación del Programa Sembrando Vida. (2022, December 30). Secretaria de Bienestar. Retrieved October 9, 2023, from https://www.gob.mx/cms/uploads/attachment/file/788907/SV2023_DOF.pdf

Renard, M.-C., & Breña, M. O. (2010). The Mexican Coffee Crisis. *Latin American Perspectives*, *37*(2), 21–33. https://doi.org/10.1177/0094582X09356956

Rhiney, K., Guido, Z., Knudson, C., Avelino, J., Bacon, C. M., Leclerc, G., Aime, M. C., & Bebber, D. P. (2021). Epidemics and the future of coffee production. *Proceedings of the National Academy of Sciences*, *118*(27). https://doi.org/10.1073/pnas.2023212118

Robert Townsend. (1994). Risk and Insurance in Village India. *Econometrica*, *62*, 539–591.

Romer, P. M. (1994). The Origins of Endogenous Growth. *Journal of Economic Perspectives*, *8*(1), 3–22. https://doi.org/10.1257/jep.8.1.3

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, *3*(2), 135–146.

Sandmo, A. (1971). On the Theory of the Competitive Firm Under Price Uncertainty. *The American Economic Review*, *61*(1), 65–73.

Shumeta, Z., D'Haese, M., & Verbeke, W. (2018). A Two-Step Econometric Estimation of Covariates of Side Selling: The Case of Coffee Cooperatives in Southwest Ethiopia. *The Journal of Development Studies*, *54*(10), 1775–1791. https://doi.org/10.1080/00220388.2017.1324146

Silva, J. M. C. S., & Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, *88*(4), 641–658. https://doi.org/10.1162/rest.88.4.641

Sjaastad, L. A. (1962). The costs and returns of human migration. *Journal of Political Economy*, *70*(5), 80–93.

Soloaga, Isidro, Lara Ibarra, Gabriel, & Wendelspiess, Florian. (2010). Determinantes de la migración interestatal: 1995-2000 y 2000-2005. In N. Lustig, Yunez, Antonio, & Sabido, Alfonso Castañeda (Eds.), *Economía rural* (pp. 172–194). El Colegio de México.

Solon, G., Haider, S. J., & Wooldridge, J. M. (2015). What are we weighting for? *Journal of Human Resources*, *50*(2), 301–316. https://doi.org/10.3368/jhr.50.2.301

Soto-Pinto, L., Castillo-Santiago, M., & Jimenez-Ferrer, G. (2012). Agroforestry Systems and Local Institutional Development for Preventing Deforestation in Chiapas, Mexico. In *Deforestation Around the World*. IntechOpen. https://doi.org/10.5772/35172

Soto-Pinto, L., Perfecto, I., Castillo-Hernandez, J., & Caballero-Nieto, J. (2000). Shade effect on coffee production at the northern Tzeltal zone of the state of Chiapas, Mexico. *Agriculture, Ecosystems & Environment*, *80*(1), 61–69. https://doi.org/10.1016/S0167-8809(00)00134-1

Suri, T., & Udry, C. (2022). Agricultural Technology in Africa. *Journal of Economic Perspectives*, *36*(1), 33–56. https://doi.org/10.1257/jep.36.1.33

UNESCO. (2018, October 29). *Internal migration — 2019 GEM report.*

Vicente-Serrano, S. M., Beguería, S., & López-Moreno, J. I. (2010). A Multiscalar Drought Index Sensitive to Global Warming: The Standardized Precipitation Evapotranspiration Index. *Journal of Climate*, *23*(7), 1696–1718. https://doi.org/10.1175/2009JCLI2909.1

Villarreal, A. (2014). Explaining the decline in mexico-u.s. migration: The effect of the great recession. *Demography*, *51*(6), 2203–2228. https://doi.org/10.1007/s13524-014-0351-4

Woldeyohanes, T., Heckelei, T., & Surry, Y. (2017). Effect of off-farm income on smallholder commercialization: Panel evidence from rural households in Ethiopia. *Agricultural Economics*, *48*(2), 207–218. https://doi.org/10.1111/agec.12327

Woldie, G. A. (2010). Optimal Farmer Choice of Marketing Channels in the Ethiopian Banana Market. *Journal of Agricultural & Food Industrial Organization*, *8*(1). https://doi.org/10.2202/1542-0485.1298

Wollni, M., & Fischer, E. (2015). Member deliveries in collective marketing relationships: Evidence from coffee cooperatives in Costa Rica. *European Review of Agricultural Economics*, *42*(2), 287–314. https://doi.org/10.1093/erae/jbu023

Woodruff, C., & Zenteno, R. (2007). Migration networks and microenterprises in mexico. *Journal of Development Economics*, *82*(2), 509–528.

Wright, D. R., Bekessy, S. A., Lentini, P. E., Garrard, G. E., Gordon, A., Rodewald, A. D., Bennett, R. E., & Selinske, M. J. (2024). Sustainable coffee: A review of the diverse initiatives and governance dimensions of global coffee supply chains. *Ambio*, *53*(7), 984–1001. https://doi.org/10.1007/s13280-024-02003-w

Wuepper, D., Bukchin-Peles, S., Just, D., & Zilberman, D. (2023). Behavioral agricultural economics. *Applied Economic Perspectives and Policy*, 1–12. https://doi.org/10.1002/aepp.13343

Xu, R., & Wooldridge, J. M. (2022). *A Design-Based Approach to Spatial Correlation.* arXiv: 2211.14354 [econ]. https://doi.org/10.48550/arXiv.2211.14354

Yap, L. Y. L. (1977). The attraction of cities: A review of the migration literature. *Journal of Development Economics*, *4*(3), 239–264. https://doi.org/10.1016/0304-3878(77)90030-X

Yesilyurt, M. E., & Elhorst, J. P. (2017). Impacts of neighboring countries on military expenditures: A dynamic spatial panel approach. *Journal of Peace Research*, *54*(6), 777–790. https://doi.org/10.1177/0022343317707569

Appendices

Appendix A

Appendix Materials for Essay 1

A.1 General Instructions for Participants

This is an experiment about individual decision making under price uncertainty. We are trying to understand how people allocate their sales to different buyers when they are unsure of the sales price. We have designed simple decision-making games in which we will ask you to make choices in a series of situations. In this experiment you have to imagine you are producing a certain amount of coffee and that, like in real life, the sale price is uncertain.

You will spend about two hours in this study playing games, for which you will be compensated with at least one food voucher coupon for your participation. In addition, you may earn between one and six additional coupons based on luck and how you play the game. Finally, you may receive an additional coupon in a lottery game. The amount will be paid to you in money vouchers that can be redeemed for food items at the end of the experiment.

1. You will play three sets of games and a lottery. Each one has its own instructions.
2. You should make your own decisions. Do not discuss your decisions with other participants or other members of the family.
3. Please turn off your cell phone, radio, or television.

4. You need to have a good understanding of how your decisions affect your game payoff. Please ask at any time during the session if you have any questions.

A.2 Instructions for Enumerators

1. Ask the following screening questions and enter in tablet (regardless of their answers):

   (a) Have you ever sold coffee you or your family has produced? Yes/No

   (b) Do you know how to read and write? Yes/No

2. Next say "Now, I will ask you some math questions to make sure you will understand the games we will play". Ask the following questions:

   (a) What is 40% of MXN 100?

   (b) If you produced 17 bags of coffee and sold 9, how many do you have left to sell?

   (c) Imagine there is bag with 3 blue balls and 7 red balls. You draw one ball. Is it more likely that it is a red or a blue ball?

3. The participant rolls the die and the tablet automatically assigns the order in which each participant will play the three games and the lottery. The lottery is played either before or after the three games. The three games are played in a random order. Make sure to read the instructions of the corresponding section.

4. Give the farmer the information sheet about the experiment, explain the experiment, and answer any questions he/she would have. Remind the farmer that he/she must complete the whole experiment to receive compensation. Make sure all the fields are filled and that the form is signed. Give a copy of the consent to the farmer.

5. The tablet will automatically determine if the participants get all of their income from coffee or if they receive another source of income. If the participants receive another source of income, give them the MXN 3,000 coupon.

Lottery	Choice	Probability (%)	Payout (MXN)
1	A	50%	10,000
	B	50%	10,000
2	A	50%	15,000
	B	50%	7,500
3	A	50%	20,000
	B	50%	5,000
4	A	50%	25,000
	B	50%	2,500
5	A	50%	30,000
	B	50%	0

Figure A1: Lottery Table shown to Participants

6. The tablet will automatically assign the order in which each participant will play the three games and the lottery. The lottery will be played either before or after all three games of sales allocation. The three games of sales allocation will be played in random order. Make sure you read the instructions for each corresponding section.

7. Once you finish with the three games (1, 2, and 3) and the lottery, you will ask the questions in the final questionnaire. Finally, determine the farmer's total compensation.

A.3 Instructions for Participants: Lottery

In the following table, you have five possible lotteries. Each one has two possibilities: A and B. Each possibility has a 50% chance of occurring and has its own payment. Choose which lottery you prefer (only one).

A.4 Instructions for Enumerators: Lottery

Record the lottery chosen by the participant. Roll a die. If the number is from 1 to 6, add compensation A to the participant's compensation. If the number is from 7 to 12, add compensation B to the participant's compensation.

A.5 Instructions for Participants: Game 1

A.5.1 Tasks

1. In this game you have to imagine that you are producing coffee and that you can sell it to two different buyers.

2. Extra Income.

    (a) **If the participant has extra income,** give them the extra income coupon. In addition to the income from the coffee sale, you received MXN 3,000 from another source this month.

    (b) **If the participant does not have extra income.** The income from this coffee sale is the only income you will receive this year.

3. Buyer 1 offers a fixed price always equal to MXN 50 per quintal, and Buyer 2 offers a variable price that goes from MXN 35 to MXN 65, depending on the case, that follows a given distribution, and which will be realized at the end of the harvest.

4. In each round, that resembles an agricultural season, we will let you know what is the quantity that you will harvest: 2 quintals, 4 quintals, 6 quintals, or 8 quintals. In each round, you will have to allocate all of your harvest between Buyer 1 and Buyer 2. You will allocate your harvest between both buyers in 1 quintal increments, and there is no possibility of storage. Your goal is to allocate your harvest in a way that maximizes your profit.

5. Buyer 2 offers a variable price that follows a given distribution. You will roll a 12-sided die to know under which price scenario (of 3 scenarios) you will be playing.

If the dice gives 1,2,3,4: Scenario 1. Buyer 1 (fixed price=50MXN) Buyer 2 (offers variable price in MXN)

If the die gives 5, 6, 7, or 8: Scenario 2 Buyer 1 (fixed price=50MXN) Buyer 2 (offers variable price in MXN)

If the die gives 9, 10, 11, or 12: Scenario 3 Buyer 1 (fixed price=50MXN) Buyer 2 (offers variable price in MXN)

6. After you allocate all your harvest between Buyer 1 (that offers a fixed price equal to MXN 50 per quintal of coffee) and Buyer 2 (that offers a variable price), the price of Buyer 2 will be realized, and we will let you know the profit (quantity times price) you earned in each round.

7. In all rounds, the resulting harvest is a result of chance and not your effort. In every round, your profit from selling coffee will be between MXN 4,200 and MXN 31,200 MXN. The more the harvest, the more your profit will be.

8. You will get a minimum profit of MXN 4,200 if the harvest is 2 quintals, you allocate 2 quintals to Buyer 2, and Buyer 2's realized price (in Scenario 3) is MXN 35 per quintal. You will get a maximum profit of MXN 31,200 if the harvest is 8 quintals, you allocate all 8 quintals to Buyer 2, and the Buyer 2's realized price is MXN 65 (in Scenario 2) (See Profit Tables 1-12).

9. You will first play ten rounds of practice games. After the practice games, you will play twenty rounds of the real game. In the real games, your profits will increase your game payoff, but not your compensation for participating in the experiment.

A.5.2 Keep in mind

1. The coffee you produce is all of the same quality and of an average quality, so you can only sell your coffee at the one price offered by each buyer.

2. The quantity of coffee harvested depends on chance and not on your own effort.

3. You cannot store the commodity produced or profits between rounds. Each round of the game has its own profit.

4. (If the participant has extra income) In addition to the income from the coffee sale, you received MXN 3,000 from another source this year.

5. (If the participant has no extra income) The income from this coffee sale is the only income you will receive this year.

A.5.3 Payoffs

1. Your payoff from the game will be based on your performance on the real (not the practice) rounds of the game.

2. At the end of the experiment, we will total your profit from all of the real rounds of the game and divide it by 250000.

3. Then we will pay you the total quantity in our experimental bills, which you will be able to exchange for non-perishable goods.

A.6 Instructions for Participants: Game 2

The instructions are the same as for Game 1 in Section A.5. The only difference is in point 3.

Buyer 1 offers a fixed price always equal to MXN 50 per quintal, and Buyer 2 offers a variable price that goes from MXN 35 to MXN 65, depending on the case, that follows a given distribution, and which will be realized at the end of the harvest. **Buyer 1 offered you a microcredit last year**.

A.7 Instructions for Participants: Game 3

The instructions are the same as for Game 1 in Section A.5. The only difference is in point 3.

Buyer 1 offers a fixed price always equal to MXN 50 per quintal, and Buyer 2 offers a variable price that goes from MXN 35 to MXN 65, depending on the case, that follows a given distribution, and which will be realized at the end of the harvest. **Buyer 1 is a cooperative that offered you a microcredit and technical assistance last year.**

A.8 Exit Survey

A.8.1 The Producer and His/Her Family

1. Village
2. Before coming here today, were you hungry?
3. On a scale of 1 to 10, how much do you like the weather today?
4. Gender
5. Age
6. Educational Level
7. How many members do your families have?
8. How many are above 65 years old?
9. How many are below 12 years old?

A.8.2 Income

1. How much money did you get from the sale of your crops last year?
2. Did you or another family member worked for money on the farm of another family last year?
3. How much did you get paid?
4. Did you or another member of your family work for pay in another city within Mexico last year?
5. How much did you get paid?
6. Do you get Sembrando Vida?[1]
7. How much do you get every two months?

[1]A conditional cash transfer program initiated in 2018. More details can be found in *Reglas de Operación Del Programa Sembrando Vida* (2022)

8. During the past year, in how many months did you not have enough food to feed yourself or your family? You may select more than one.
   (a) January
   (b) February
   (c) March
   (d) April
   (e) May
   (f) June
   (g) July
   (h) August
   (i) September
   (j) October
   (k) November
   (l) December

A.8.3 Farm

1. Which of the following animals do you have?
   (a) Poultry
   (b) Horses
   (c) Mules
   (d) Donkeys
   (e) Sheep or Goats
   (f) Cattle
   (g) Pigs
2. Which of the following crops do you have?

(a) Corn

(b) Beans

(c) Coffee

(d) Zucchini

(e) Chayote Squash

(f) Chile

(g) Banana

(h) Sugar cane

(i) Oranges or Mandarines

(j) Yuca

(k) Sweet potato

(l) Papaya

(m) Mango

3. How many plots of land do you have?

4. How many total hectares do you have?

5. What percentage of your farm is for your own consumption?

6. If you had enough money, what would you prefer to do?

   (a) Plant more coffee

   (b) Plant more staples

   (c) Produce honey (or produce more honey)

   (d) Plant fruit trees

   (e) Buy livestock (or more livestock)

   (f) None of the above.

A.8.4 Coffee

1. How long have you been growing coffee?
2. What varieties of coffee do you presently have in your parcel?
   (a) Typica (o criolla)
   (b) Bourbón
   (c) Maragogype
   (d) Geisha
   (e) Tabi
   (f) Caturra
   (g) Mundo Novo
   (h) Garnica
   (i) Catimor
   (j) Pacamara
   (k) Oro Azteca
   (l) Robusta
3. How much coffee did you grow last year?
4. How much coffee did you sell to an intermediary in the past year?
5. What is the highest price that you received from an intermediary last year?
6. How do you sell your coffee to the intermediary?
   (a) The intermediary comes to my parcel.
   (b) The intermediary comes to my village
   (c) The intermediary comes to the nearest population center.
7. How much coffee did you sell to a cooperative last year?

8. What is the highest price that you received from a cooperative last year?

9. How do you sell your coffee to the cooperative?

    (a) The cooperative comes to my parcel.

    (b) The cooperative comes to my village.

    (c) The cooperative comes to the nearest population center.

10. What's the best price you've received for your coffee?

11. What's the worst price you've received for your coffee?

12. Why do you sell more coffee to the cooperative than the intermediary?

13. Why do you sell more coffee to the intermediary than the cooperative?

A.8.5 Honey

1. Do you produce honey?

2. How much honey did you produce last year?

3. What's the best price that you have received for your honey?

4. On a scale of 1 (not interested) to 10 (very interested) how interested are you in producing honey?

5. Do you know someone who produces honey?

Figure A2: Scenario 1

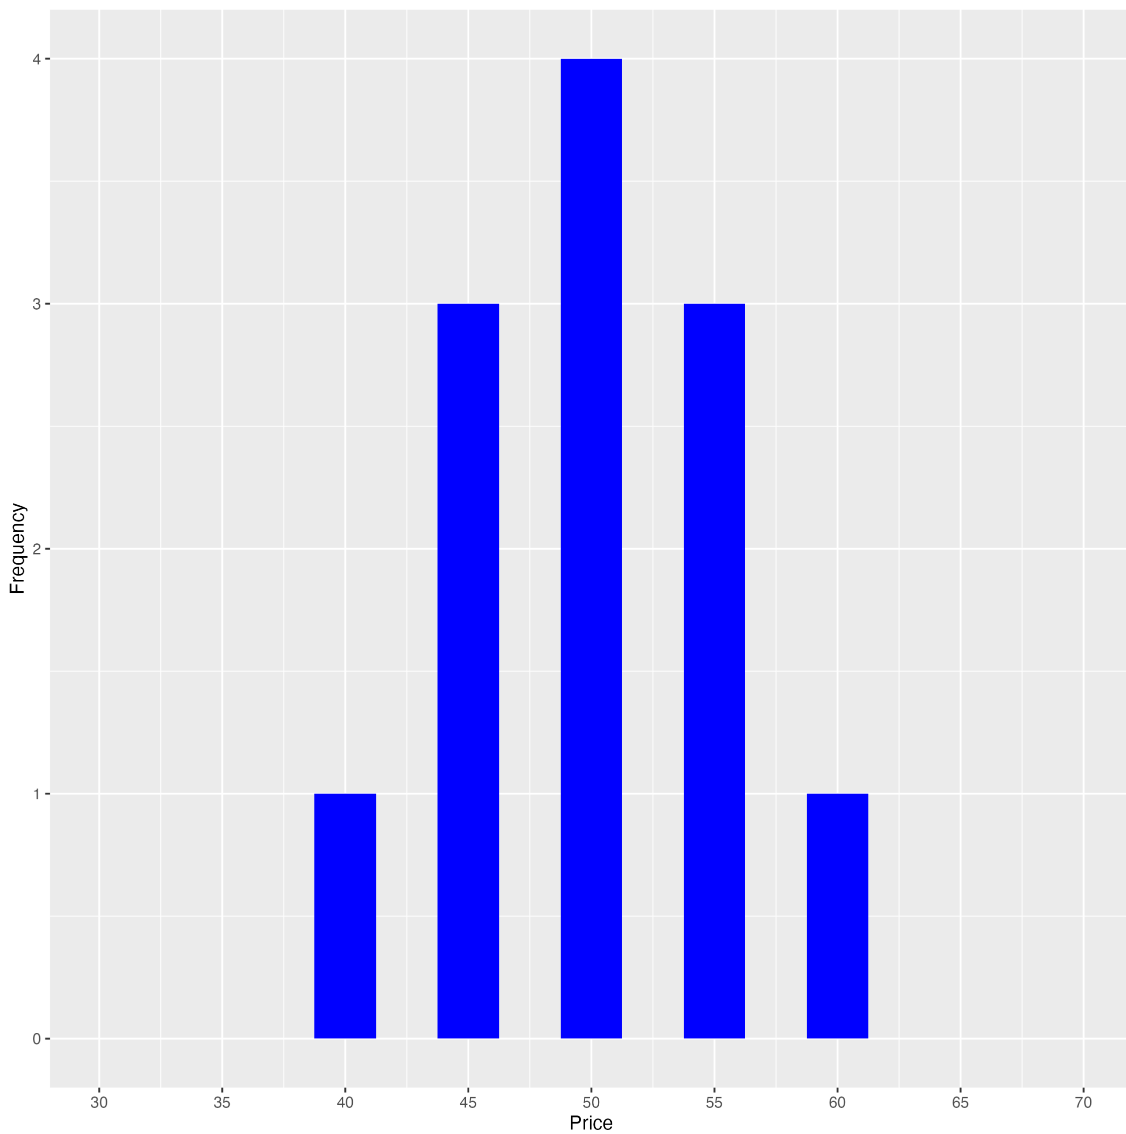

Figure A3: Scenario 2

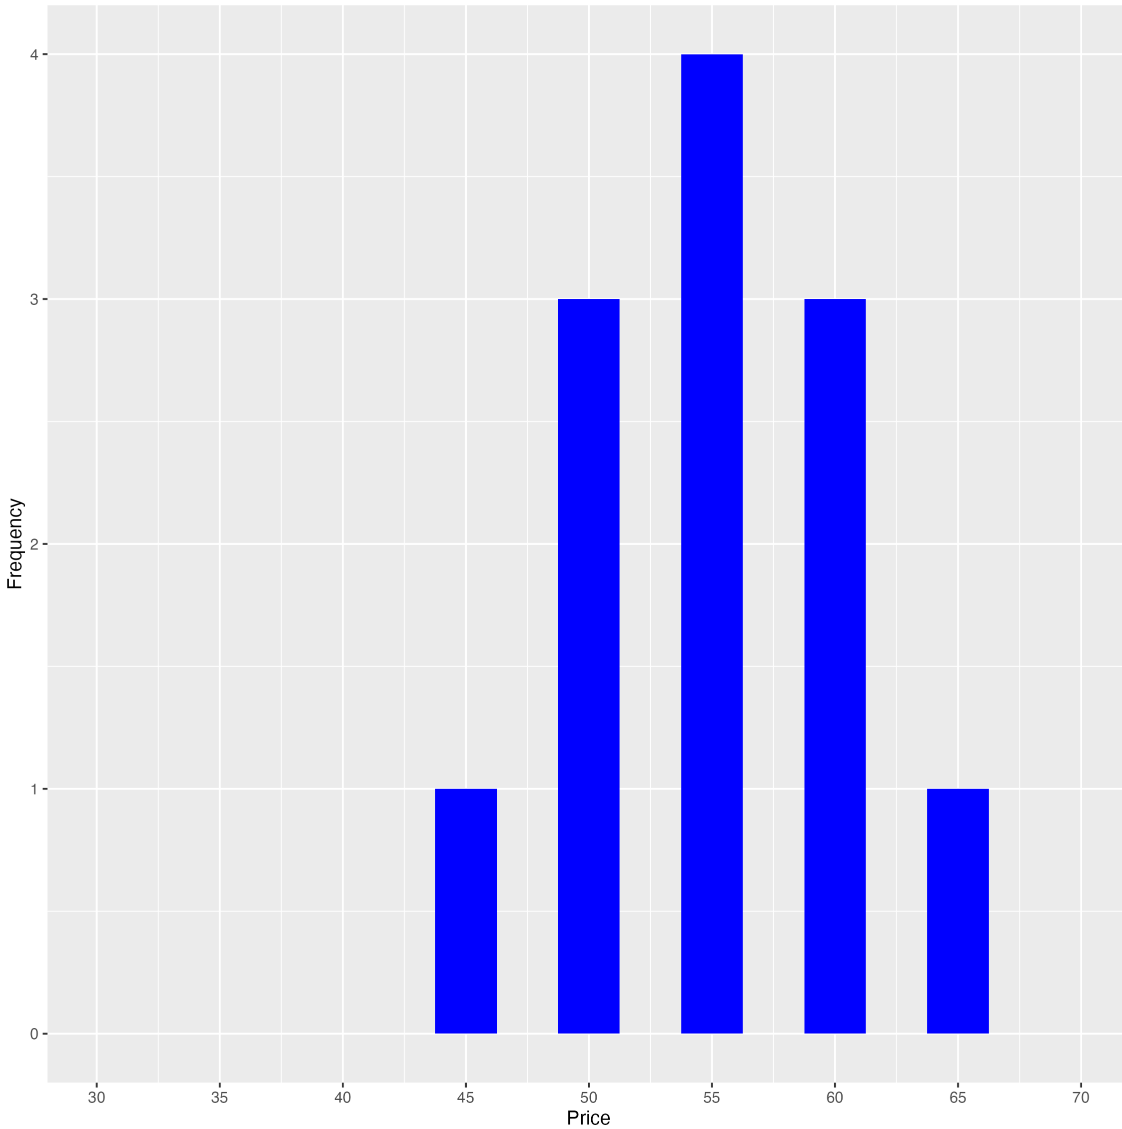

Figure A4: Scenario 3

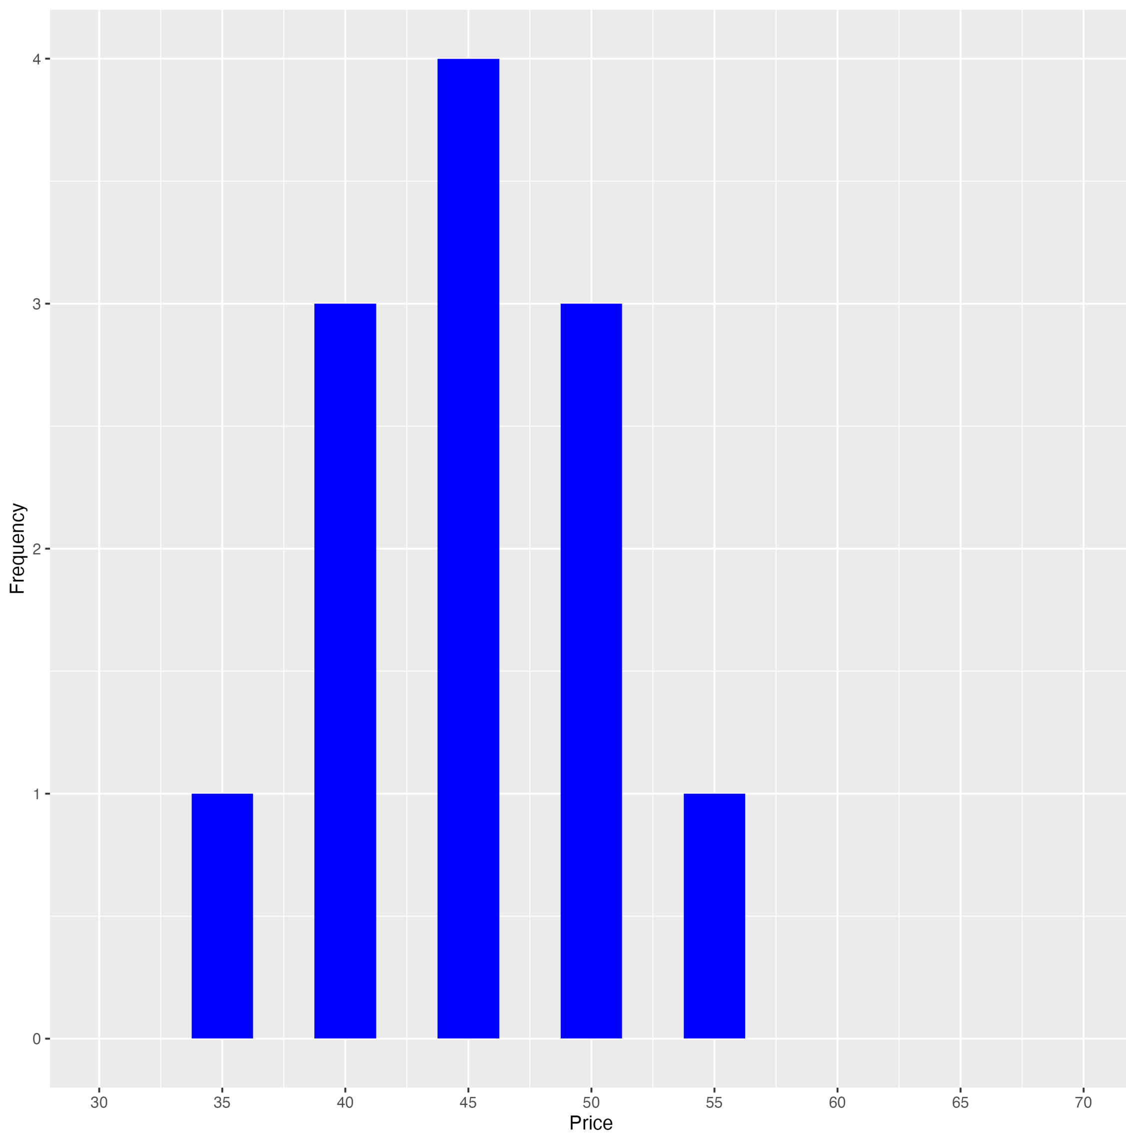

Figure A5: Profit Table for 2 quintal harvest, Buyer 2 Scenario 1

	Quintals Sold to Buyer 1 (60kg)		
	0	**1**	**2**
	Quintals Sold to Buyer 2 (60kg)		
	2	**1**	**0**
	Total Revenue from Sales to Both Buyers		
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)		
	0	3000	6000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)		
35	4200	2100	0
40	4800	2400	0
45	5400	2700	0
50	6000	3000	0
55	6600	3300	0

Figure A6: Profit Table for 2 quintal harvest, Buyer 2 Scenario 2

	Quintals Sold to Buyer 1 (60kg)		
	0	**1**	**2**
	Quintals Sold to Buyer 2 (60kg)		
	2	**1**	**0**
	Total Revenue from Sales to Both Buyers		
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)		
	0	3000	6000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)		
40	4800	2400	0
45	5400	2700	0
50	6000	3000	0
55	6600	3300	0
60	7200	3600	0

Figure A7: Profit Table for 2 quintal harvest, Buyer 2 Scenario 3

	Quintals Sold to Buyer 1 (60kg)		
	0	1	2
	Quintals Sold to Buyer 2 (60kg)		
	2	1	0
	Total Revenue from Sales to Both Buyers		
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)		
	0	3000	6000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)		
45	5400	2700	0
50	6000	3000	0
55	6600	3300	0
60	7200	3600	0
65	7800	3900	0

Figure A8: Profit Table for 4 quintal harvest, Buyer 2 Scenario 1

	Quintals Sold to Buyer 1 (60kg)				
	0	1	2	3	4
	Quintals Sold to Buyer 2 (60kg)				
	4	3	2	1	0
	Total Revenue from Sales to Both Buyers				
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)				
	0	3000	6000	9000	12000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)				
35	8400	6300	4200	2100	0
40	9600	7200	4800	2400	0
45	10800	8100	5400	2700	0
50	12000	9000	6000	3000	0
55	13200	9900	6600	3300	0

Figure A9: Profit Table for 4 quintal harvest, Buyer 2 Scenario 2

	Quintals Sold to Buyer 1 (60kg)				
	0	1	2	3	4
	Quintals Sold to Buyer 2 (60kg)				
	4	3	2	1	0
	Total Revenue from Sales to Both Buyers				
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)				
	0	3000	6000	9000	12000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)				
40	9600	7200	4800	2400	0
45	10800	8100	5400	2700	0
50	12000	9000	6000	3000	0
55	13200	9900	6600	3300	0
60	14400	10800	7200	3600	0

Figure A10: Profit Table for 4 quintal harvest, Buyer 2 Scenario 3

	Quintals Sold to Buyer 1 (60kg)				
	0	1	2	3	4
	Quintals Sold to Buyer 2 (60kg)				
	4	3	2	1	0
	Total Revenue from Sales to Both Buyers				
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)				
	0	3000	6000	9000	12000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)				
45	10800	8100	5400	2700	0
50	12000	9000	6000	3000	0
55	13200	9900	6600	3300	0
60	14400	10800	7200	3600	0
65	15600	11700	7800	3900	0

Figure A11: Profit Table for 6 quintal harvest, Buyer 2 Scenario 1

	Quintals Sold to Buyer 1 (60kg)						
	0	1	2	3	4	5	6
	Quintals Sold to Buyer 2 (60kg)						
	6	5	4	3	2	1	0
	Total Revenue from Sales to Both Buyers						
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)						
	0	3000	6000	9000	12000	15000	18000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)						
35	12600	10500	8400	6300	4200	2100	0
40	14400	12000	9600	7200	4800	2400	0
45	16200	13500	10800	8100	5400	2700	0
50	18000	15000	12000	9000	6000	3000	0
55	19800	16500	13200	9900	6600	3300	0

Figure A12: Profit Table for 6 quintal harvest, Buyer 2 Scenario 2

	Quintals Sold to Buyer 1 (60kg)						
	0	1	2	3	4	5	6
	Quintals Sold to Buyer 2 (60kg)						
	6	5	4	3	2	1	0
	Total Revenue from Sales to Both Buyers						
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)						
	0	3000	6000	9000	12000	15000	18000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)						
40	14400	12000	9600	7200	4800	2400	0
45	16200	13500	10800	8100	5400	2700	0
50	18000	15000	12000	9000	6000	3000	0
55	19800	16500	13200	9900	6600	3300	0
60	21600	18000	14400	10800	7200	3600	0

Figure A13: Profit Table for 6 quintal harvest, Buyer 2 Scenario 3

	Quintals Sold to Buyer 1 (60kg)						
	0	1	2	3	4	5	6
	Quintals Sold to Buyer 2 (60kg)						
	6	5	4	3	2	1	0
	Total Revenue from Sales to Both Buyers						
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)						
	0	3000	6000	9000	12000	15000	18000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)						
45	16200	13500	10800	8100	5400	2700	0
50	18000	15000	12000	9000	6000	3000	0
55	19800	16500	13200	9900	6600	3300	0
60	21600	18000	14400	10800	7200	3600	0
65	23400	19500	15600	11700	7800	3900	0

Figure A14: Profit Table for 8 quintal harvest, Buyer 2 Scenario 1

	Quintals Sold to Buyer 1 (60kg)								
	0	1	2	3	4	5	6	7	8
	Quintals Sold to Buyer 2 (60kg)								
	8	7	6	5	4	3	2	1	0
	Total Revenue from Sales to Both Buyers								
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)								
	0	3000	6000	9000	12000	15000	18000	21000	24000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)								
35	16800	14700	12600	10500	8400	6300	4200	2100	0
40	19200	16800	14400	12000	9600	7200	4800	2400	0
45	21600	18900	16200	13500	10800	8100	5400	2700	0
50	24000	21000	18000	15000	12000	9000	6000	3000	0
55	26400	23100	19800	16500	13200	9900	6600	3300	0

Figure A15: Profit Table for 8 quintal harvest, Buyer 2 Scenario 2

	Quintals Sold to Buyer 1 (60kg)								
	0	1	2	3	4	5	6	7	8
	Quintals Sold to Buyer 2 (60kg)								
	8	7	6	5	4	3	2	1	0
	Total Revenue from Sales to Both Buyers								
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)								
	0	3000	6000	9000	12000	15000	18000	21000	24000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)								
40	19200	16800	14400	12000	9600	7200	4800	2400	0
45	21600	18900	16200	13500	10800	8100	5400	2700	0
50	24000	21000	18000	15000	12000	9000	6000	3000	0
55	26400	23100	19800	16500	13200	9900	6600	3300	0
60	28800	25200	21600	18000	14400	10800	7200	3600	0

Figure A16: Profit Table for 8 quintal harvest, Buyer 2 Scenario 3

	Quintals Sold to Buyer 1 (60kg)								
	0	1	2	3	4	5	6	7	8
	Quintals Sold to Buyer 2 (60kg)								
	8	7	6	5	4	3	2	1	0
	Total Revenue from Sales to Both Buyers								
	Revenue from Sale to Buyer 1 (Quantity Sold to Buyer 1 x $50 MXN)								
	0	3000	6000	9000	12000	15000	18000	21000	24000
Price per kilogram (MXN)	Revenue from Sale to Buyer 2 (Quantity Sold x Dice Result)								
45	21600	18900	16200	13500	10800	8100	5400	2700	0
50	24000	21000	18000	15000	12000	9000	6000	3000	0
55	26400	23100	19800	16500	13200	9900	6600	3300	0
60	28800	25200	21600	18000	14400	10800	7200	3600	0
65	31200	27300	23400	19500	15600	11700	7800	3900	0

Appendix B

Appendix Materials for Essay 3

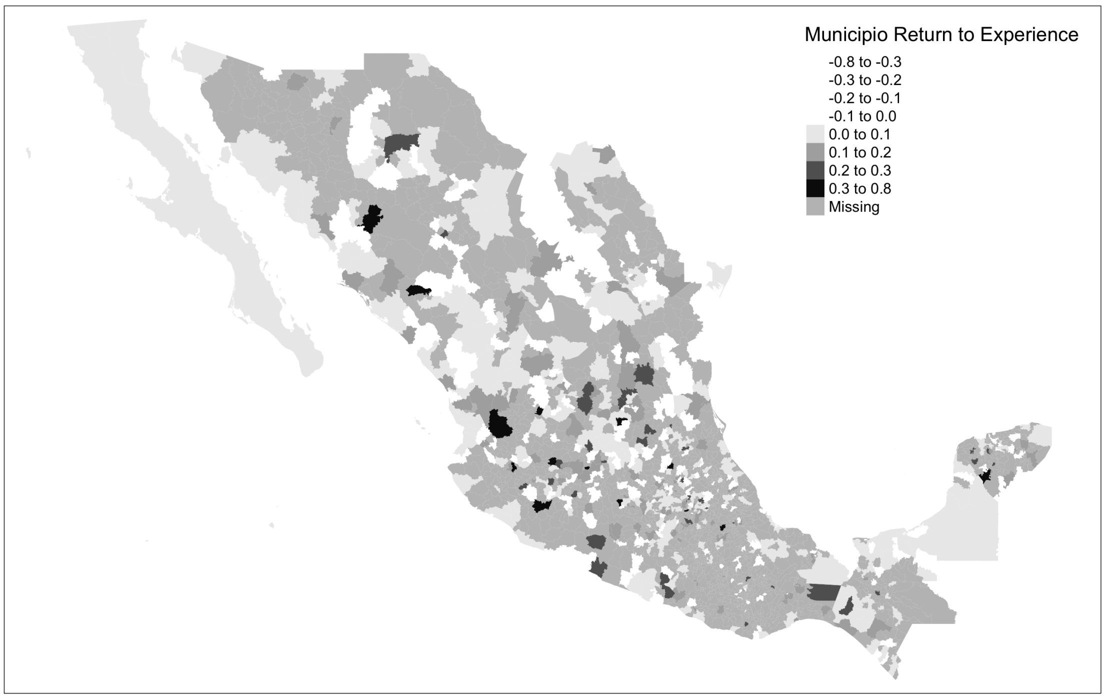

Figure B1: Return to Experience by Municipality (1995)

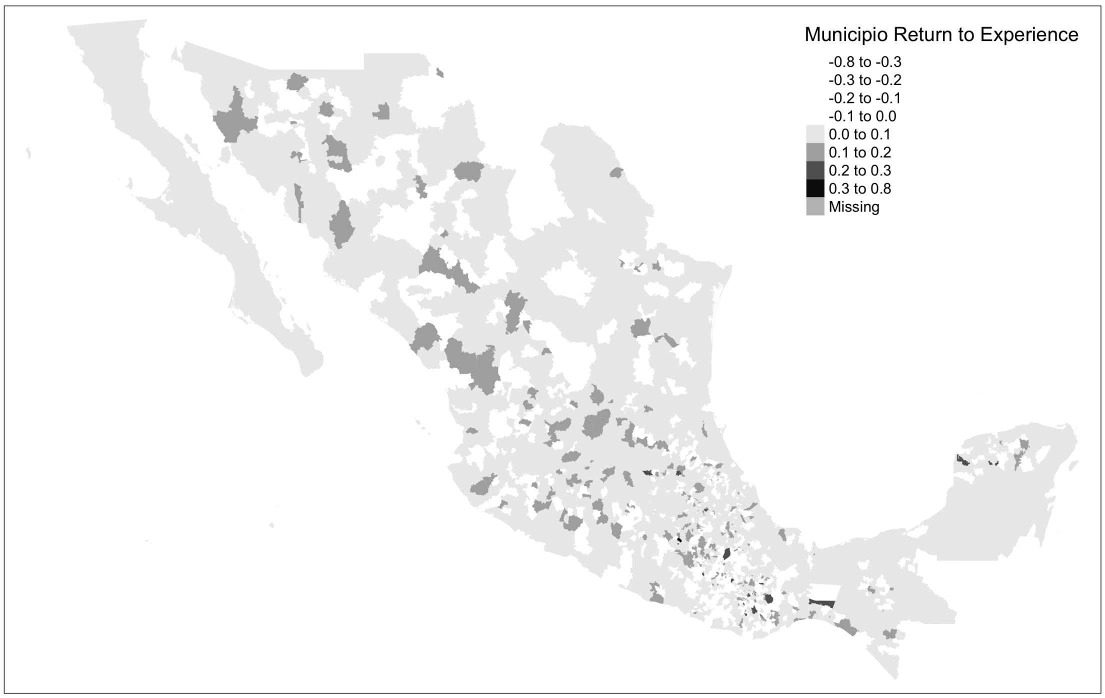

Figure B2: Return to Experience by Municipality (2000)

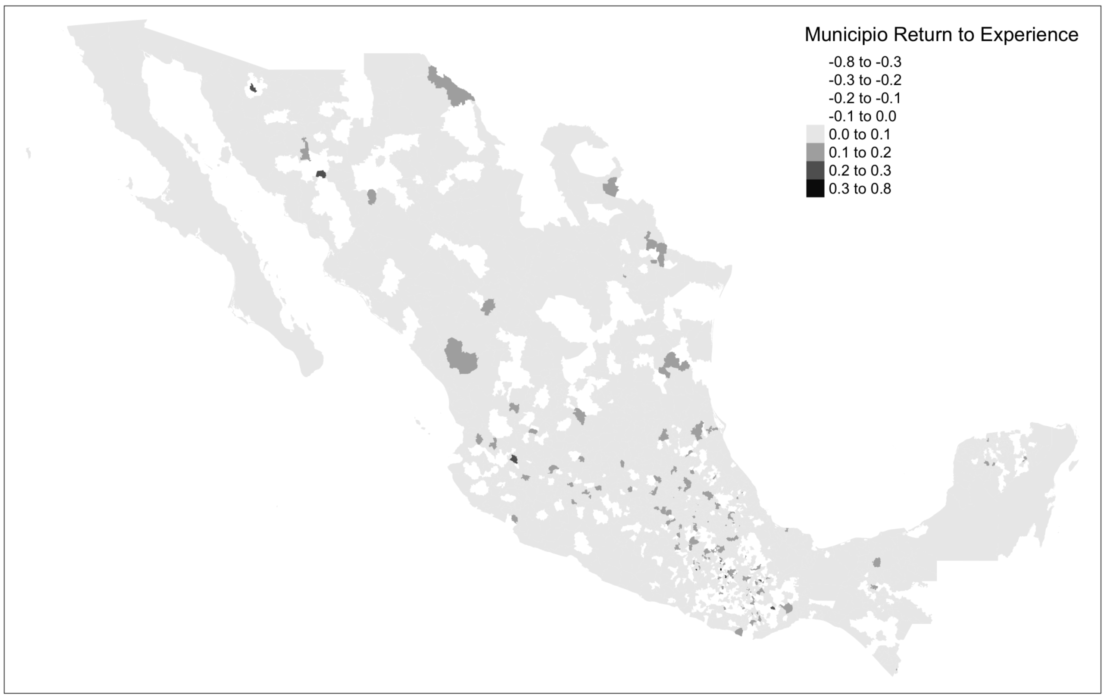

Figure B3: Return to Experience by Municipality (2010)