

ENTREPRENEURSHIP, JOB SEARCH AND OCCUPATIONAL CHOICE IN DEVELOPING COUNTRIES

A Dissertation Submitted for the Degree of
Doctor of Philosophy in Economics



MICHAEL KOELLE
Lady Margaret Hall
University of Oxford

Trinity 2017

Abstract

This thesis is about entrepreneurship within the context of occupational choice in a frictional labour market. It unites two perspectives on entrepreneurship: the perspective of the entrepreneur as a firm, and the perspective of the entrepreneur as a self-employed worker. My emphasis is on developing countries, where a large share of the workforce are self-employed entrepreneurs. I combine theory with empirical applications using labour-force panel data from Mexico, Colombia and Ghana. The thesis consists of four chapters.

The model in Chapter 1 provides a theoretical framework to capture the nexus of occupational and firm size choices of entrepreneurs. In the model, these two choices are non-separable. Entrepreneurs face a trade-off: they could either make the most of their entrepreneurial potential, or run a smaller firm while waiting to take the next attractive job opportunity. Better opportunities in wage employment increase incentives to keep firms small.

Chapter 2 tests this model. To obtain variation in alternative employment opportunities for entrepreneurs, I exploit trade-induced changes to wage employment across local labour markets in Mexico. Using a large dataset built from several sources, I find evidence that is consistent with the theoretical predictions. I use guidance from theory, supplementary evidence, and an

event study design to distinguish the occupational choice mechanism from competing channels.

Chapter 3 develops an empirical model where the observed allocation of individuals to occupations is a product of both the attractiveness and the availability of job opportunities. This type of selection mechanism is central to each of the other chapters in this thesis. The chapter is concerned with the identification of such a model that combines comparative advantage and search frictions. I estimate the model on panel data from urban Colombia and find evidence for both search frictions and selection on comparative advantage.

Finally, Chapter 4 builds and estimates a structural life-cycle model of sectoral choice between self-employment and wage-employment, costly job search, and capital accumulation in Ghana. We incorporate two central types of market frictions: credit constraints for entrepreneurs, and search frictions in the wage labour market. This provides a unified framework to think about career paths in developing countries. We link the estimated structural model to recent experimental literature in development economics, and use simulated experiments to interpret the mechanisms behind behavioural responses to labour market shocks.

Approximate length of this thesis:

65,000 words

Acknowledgements

This thesis marks the capstone of a long education, my personal *Bildungsroman*. I am thankful for all those teachers who have taught me along these years; who inspired me, challenged me, who stoked my curiosity, and who sometimes left me in awe. I am truly grateful to my parents, who knew how to instil in me a love of learning and knowledge for their own sake.

My most sincere thanks go to my supervisors Simon Quinn and Margaret Stevens – two of the greatest of all the teachers I ever had – for their advice, support, and encouragement. I am particularly indebted to Simon for the opportunities he has opened up for me, the generosity of his time, and his unrelenting optimism in collaborating on what forms the last chapter of this dissertation.

I would like to acknowledge and thank all the individuals whose conversations, discussions and comments have had an influence on my work on this dissertation: Orazio Attanasio, Martin Browning, Ian Crawford, Marcel Fafchamps, Paolo Falco, James Fenske, Manuel Fernández, Markus Goldstein, Doug Gollin, Juan Carlos Hallak, Leander Heldring, Clément Imbert, Michael Keane, Kalina Manova, Peter Neary, Carlo Pizzinelli, Marlon Seror, Raúl Sánchez de la Sierra, Kenneth Wolpin, and Christopher Woodruff. A special round of thanks to the Oxford Firms and Develop-

ment Student Research Group who have accompanied me during all or most of this journey, and whose countless influences are impossible to point-identify: Rachel Cassidy, Elwyn Davies, Jonathan Lain, Pavel Luengas Sierra, Muhammad Meki, Chris Roth, Daniela Scur, Anisha Sharma, Hannah Uckat, and Marc Witte. I am grateful to Pavel for assistance with some of the Mexican data, and to Paolo and especially Jono for their help with and patience in explaining me the GUHPS.

Funding for this research, during my graduate studies for the MPhil and DPhil, has been provided by a number of organisations: Economic and Social Research Council (ESRC); *Deutscher Akademischer Austauschdienst* (DAAD) and *Heinrich Boell Stiftung* of Germany; the Department of Economics and Lady Margaret Hall at Oxford. I appreciate the generous support of these organisations.

While it is true that this thesis is solely the product of my own work unless where otherwise indicated, it is equally true that from start to finish, for every single day of the journey, this whole process has been one piece of incredible teamwork with the most amazing teammate I could have imagined, my wife Samira. Words cannot describe how obliged I am for the unconditional support, loving patience, cheerful encouragement, and faithful yet sincere advice. *Gracias, mi equipo.*

Contents

1	Microenterprises and the Lure of Wage Work: Theory	13
1.1	Introduction	13
1.1.1	Chapter Overview	18
1.1.2	Related literature	19
1.2	Model	26
1.2.1	Conceptual framework	26
1.2.2	A model of investment under on-the-job-search . . .	28
1.3	Solution and characterisation	32
1.3.1	Optimal firm size	32
1.3.2	Existence and uniqueness	33
1.3.3	Local comparative statics	34
1.3.4	Equilibrium selection and global comparative statics	36
1.4	Discussion	40
1.4.1	Intuition	40
1.4.2	Heterogeneous entrepreneurial ability	43
1.4.3	Consistency with stylised facts	44
1.5	Extension: Intertemporal investment	47
1.6	Testable implications	52
1.7	Conclusion	53

Appendix A	57
A.1 Parametric assumptions in model simulations	57
2 Microenterprises and the Lure of Wage Work: Evidence from Mexican Export Manufacturing	59
2.1 Introduction	59
2.1.1 Related literature	63
2.2 Empirical strategy	66
2.2.1 Empirical specifications	69
2.3 Data	70
2.3.1 Mexican export manufacturing	73
2.3.2 Self-employment in the Mexican labour market . . .	76
2.3.3 Transitions out of self-employment	78
2.3.4 How well does the model match the setting?	80
2.4 Empirical results	83
2.4.1 The predictive power of exports for employment . .	84
2.4.2 Export shocks and local labour markets	85
2.4.3 Export shocks and occupational choice	88
2.4.4 Export shocks and entrepreneurial choice	91
2.5 Mechanisms and Heterogeneity	93
2.5.1 Testing for heterogeneous effects	93
2.5.2 Alternative explanations	95
2.5.3 The role of expectations	100
2.5.4 Robustness	107
2.6 Conclusion	108
Appendix B	111
B.1 Map of Mexican export manufacturing in 2005	111

B.2	Additional results on sector shares	112
B.3	Additional results on heterogeneity	114
B.4	Additional results on wages	121
B.5	Robustness checks	124
3	Occupational Choice under Search Frictions: A Roy Search Model	135
3.1	Introduction	135
3.1.1	Related literature	137
3.2	Empirical Framework	139
3.2.1	Expressing the attractiveness of wage-employment .	140
3.2.2	The selection problem	141
3.2.3	A Roy Search Model	143
3.2.4	Identification	146
3.2.5	Relation to literature on identification of selection models	158
3.2.6	Maximum likelihood estimation of the RSM	161
3.2.7	The Roy Search Model and the Lure of Wage Work .	162
3.3	Data	164
3.3.1	Background	164
3.3.2	Descriptive statistics	165
3.4	Results	169
3.4.1	First step: An estimated Roy Search Model	170
3.4.2	Second step: Measures of alternative employment opportunities	172
3.4.3	Third step: Outcome regressions	173
3.5	Robustness	176
3.5.1	An alternative formulation of the RSM	176

3.5.2	Sensitivity of outcome regressions to alternative RSM formulations	176
3.6	Conclusion	179
Appendix C		181
C.1	Derivation of Likelihood Function	181
C.2	Data definitions and summary statistics	184
C.3	Likelihood Function of the alternative model	185
C.4	Additional tables and figures	186
4	Self-Employment, Capital and Job Search: Evidence from a Structural Model	189
4.1	Introduction	189
4.1.1	Related literature	192
4.2	Data and descriptive evidence	195
4.3	Model	199
4.3.1	Wages and incomes	203
4.3.2	Permanent heterogeneity	205
4.3.3	Consumption	206
4.3.4	Search frictions	207
4.4	Identification	210
4.4.1	Identification of structural parameters	211
4.4.2	Inference	214
4.5	Estimation results	215
4.5.1	Goodness of fit	215
4.5.2	Parameter estimates	218
4.5.3	Sensitivity of estimates to moments	222
4.6	Interpretation	224

4.6.1	Structural replication of field experiments	224
4.6.2	Entrepreneurship: Necessity or opportunity?	228
4.6.3	Wage work: Stepping stone to self-employment? . .	231
4.7	Conclusion	234
Appendix D		237
D.1	Data and variable construction	237
D.2	Computational details	239
D.3	Goodness of fit	240
D.4	Replication of the Flypaper Effect	247

Chapter 1

Microenterprises and the Lure of Wage Work: Theory


1.1 Introduction

Small firms, owned and operated by an entrepreneur, unite several economic elements. They are firms that generate profits by turning inputs into output, and as employers provide for an important share of jobs. Small firms are also a vehicle of self-employment for entrepreneurs. This is particularly true in developing countries, where a large share of the urban workforce are self-employed entrepreneurs with small businesses, as Figure 1 documents. Despite their individual sizes, as a group, small-scale entrepreneurs create a significant number of wage jobs. In Mexico, for instance, 28 of percent of the workforce are wage-employed in firms of ten workers or fewer. About the same number are entrepreneurs running such firms.

Given this central role of entrepreneurs in labour markets of developing countries, researchers and policymakers have experimented with many dif-

S

FIGURE 1
INTERNATIONAL COMPARISON OF NON-RURAL NON-AGRICULTURAL
SELF-EMPLOYMENT



Notes: Own calculation. Data restricted to capital cities or metropolitan areas with at least 1M population for the following sources: 2007 ENOE (Mexico), 2007 ESLF (Colombia) and WIEGO (2009) unified dataset compiled from: 2003 1-2-3 survey (Mali, Burkina Faso, Senegal); 2006 ENAHO (Peru), 2007 LFS (Vietnam), and 2007 LFS tabulations (UK). Data restricted to non-agricultural self-employment for 2004 Microcensus tabulations (Germany) and 2007 CPS tabulations (US).

ferent policies to alleviate potential constraints on firms and entrepreneurs.¹ Yet, an important puzzle prevails: microenterprises on average tend to stay small, in spite of apparently very high returns to capital.² While previous work has paid careful attention either to the entrepreneur and their firm *or* to the self-employed worker, it has largely ignored the dual role of the en-

¹ For example, the World Bank Group invested USD18.7bn into projects related to entrepreneurship between 2000 and 2013. I review the literature on policy evaluations aimed at promoting entrepreneurship in developing countries in Section 1.1.2

² Experimental evidence suggests monthly returns to capital may be as high as 20 percent in Mexico (McKenzie and Woodruff (2008)), 15 percent in Ghana (Fafchamps, McKenzie, Quinn, and Woodruff (2014)) or 5 percent in Sri Lanka (De Mel, McKenzie, and Woodruff (2008)).

trepreneur as a firm *and* as a self-employed worker.³

In this thesis, I unite these two perspectives on entrepreneurship. In the present chapter, I theoretically address the question of whether entrepreneurs' decisions about their firm, and their choices about their own employment, are separable or not. I propose a particular mechanism through which such a non-separability matters: the anticipation of a better outside opportunity for entrepreneurs in the future creates incentives for under-investment in firms. Many entrepreneurs in developing countries lack attachment to their firm, and most of those who leave their firm do so to take up a wage job. Entrepreneurs in Mexico, Colombia and Ghana – the countries where the empirical analysis of this thesis takes place – all follow a remarkably similar pattern, as Table 1 attests: about 20 percent of self-employed men are found in wage employment a year later.

In my model, non-separability arises from four related market failures that constrain entrepreneurs' information and choice sets: (i) the choice between entrepreneurship and wage employment is discrete; (ii) management and ownership of a firm are tied together; (iii) capital investments are firm-specific; and (iv) job opportunities and wages are uncertain. Forward-looking entrepreneurs discount returns to investing into firm-specific capital stock,

³ This duality is evident in the use of language and the structure of most existing data. Economists' language tends to juxtapose "enterprises" and "firms" against "self-employment" and "subsistence entrepreneurship", often framing these two as disjoint groups. For instance, Lerner and Schoar (2010) (p. 2) synthesize the state of thinking on entrepreneurship as following:

One of the main findings that emerge [...] is that there are two extremely distinct "flavors" of entrepreneurship – subsistence entrepreneurs and transformational entrepreneurs. The former manage microbusinesses that do not grow beyond providing employment for the owner, or at most, for his or her family. The latter type of entrepreneurship is the real engine of growth for the economy, since these entrepreneurs start firms that grow rapidly and create jobs and innovation in the economy.

Administrative and survey data generally only cover one of (a) large firms or establishments, (b) microenterprises, and (c) wage- and self-employed workers.

TABLE 1
ANNUAL TRANSITION RATES OUT OF SELF-EMPLOYMENT, MEN [IN %]

Country	NE _{t+1}	WE _{t+1}	SE _{t+1}	Total
Colombia	9.32	22.61	68.07	100.00
Ghana	9.20	20.25	70.55	100.00
Mexico	13.11	19.74	67.15	100.00

Notes: Working-age (15-64) men self-employed in period t and observed in two adjacent periods. WE=Wage employed, SE=Self-employed (includes entrepreneurs), NE=Not gainfully employed (includes unpaid workers). *Sources:* 2005-2014 ENOE (Mexico), 2007-2010 ESLF (Colombia), 2005-2013 GHUPS (Ghana).

with the discount factor given by the probability that they eventually take up a wage offer. They under-invest into their business relative to a world where entrepreneurial choices are made independently and only depend on the returns to capital.

In Chapter 2, I test this model for small firms in Mexico with a well-identified research design. I exploit exogenous determinants of job creation in large export manufacturing establishments across local labour markets to identify the effects of wage employment opportunities on entrepreneurs. I document that the critical assumptions of my model closely match this setting, and that my design isolates market-level shocks to job postings. I find evidence that is consistent with the mechanism where alternative employment opportunities for entrepreneurs reduce growth of small firms. With a species analysis, a case study of large plant expansion announcements in the Mexican automotive industry, and an analysis of differential timing of effects, I distinguish between competing mechanisms behind these findings. The supplementary evidence suggests that it is indeed the expectation of job opportunities, and not other mechanisms, which drive the results.

In Chapter 3, I consider the empirical content of a joint model of occupa-

tional choice and earnings in the presence of search frictions. I focus on the occupational choice between wage-employment and self-employment. Search frictions pose a constraint to the availability of jobs, and they drive a wedge between the attractiveness of wage-employment (defined by comparative advantage) and the observed allocation of individuals across occupations. I develop an empirical framework that I call the Roy Search Model (RSM), which embodies the key features of this setting. I show how the RSM can be identified with the kind of labour force data that is widely available even in developing countries. I estimate the model using panel data from urban Colombia and find evidence for both search frictions and selection on comparative advantage. Finally, I use the estimated attractiveness of wage-employment to devise an alternative test for the model in Chapter 1, from a different setting and with a different empirical methodology.

In the final chapter, I leverage theory and empirics into a structural estimation exercise. In Chapter 4, co-authored with Simon Quinn, we build a structural life-cycle model of occupational choice in urban Ghana. Our model allows for endogenous career choices, costly job search, physical and human capital accumulation, and a rich structure of unobserved heterogeneity. We incorporate two central types of market frictions: credit constraints for entrepreneurs, and search frictions in the wage labour market. With this, we provide a unified framework to a body of experimental literature on the role for such frictions in limiting career opportunities in developing countries. We validate our model by replicating two recent field experiments. Through the lens of the model, we shed light on the mechanisms underlying behavioural responses to labour market shocks in developing countries.

1.1.1 Chapter Overview

In this chapter, I formulate a two-period model that captures how discrete occupational choice can affect investment decisions of entrepreneurs. I nest a standard model of optimal investment under uncertainty (with convex preferences and technology) within a two-sector labour market search model (building on Stigler (1962), Mortensen (1970) and subsequent literature). Entrepreneurs can engage in on-the-job search for a wage job. When a sufficiently attractive job offer arrives, entrepreneurs will take it and shut down their firm. Effectively, the anticipated shutdown probability drives a wedge between the cost of capital and the return to capital. I show that, viewed from the perspective of the microenterprise, it serves as an additional discount factor in the intertemporal asset allocation problem. This is the transmission channel between occupational and entrepreneurial choices. Four critical assumptions let this non-separability arise. First, the choice between being a self-employed owner-manager of a small firm (an entrepreneur) and holding a wage job is discrete. Second, management cannot be passed on to an external manager. Third, invested capital is specific to the firm and cannot be (fully) liquidated. Fourth, there is ex-ante uncertainty about the availability and attractiveness of wage jobs at the time of investment.

In the limit – when attractive wage employment opportunities are absent – the shutdown probability goes to zero, and the model collapses to a neo-classical investment model. In such a case, the problem of the firm is separable from occupational choices of the owner. Under non-separability of occupational and entrepreneurial choices, the schedule of optimal investment can be convex, implying the existence of multiple local solutions to

optimal firm size. In a dynamic setting, initial asset endowment determines which equilibrium is reached, and a ‘poverty trap’ may arise (see Kraay and McKenzie (2014)). Comparative statics show that entrepreneurs reduce investment into their firm when they anticipate wage employment to be more attractive, either because of an increase in the offer probability, or due to a stochastic improvement in the wage offer distribution. The intuition is that better outside opportunities increase the likelihood of leaving entrepreneurship, which in probability reduces the states of the world in which the fruits from investment into the firm can be reaped. Consequently, entrepreneurs choose to lower their exposure.

The model is able to replicate several stylised facts from recent literature on firms in developing countries. By formalising how choices regarding firm size can depend on occupational choices of entrepreneurs, it rationalises the ‘puzzle of high returns’. The model can further replicate the findings that access to finance has the biggest benefits for ex-ante more successful microenterprises, and that the size of start-up grants can determine firm size of new entrant entrepreneurs.

1.1.2 Related literature

This thesis contributes to several strands of literature. Here I review the literature related to the motivation, the model, and to the overall conceptual contribution of this thesis. I will discuss literature related to particular empirical applications and methodologies in each individual chapter.

Occupational choice, entrepreneurship and firm performance. First and foremost, I contribute to the literature on entrepreneurship, occupational

choice, and firm performance. This literature – which goes back at least to the seminal contribution of Lucas (1978) – has from its start been concerned with how occupational choices of individuals who are endowed with heterogeneous entrepreneurial ability can rationalise the cross-sectional firm-size distribution. More recent contributions directly building on Lucas (1978) investigate how the firm size distribution is shaped by the institutional environment (e.g. Gollin (2008) and Garicano, Lelarge, and van Reenen (2016)). Theory has identified a number of further determinants of both selection into entrepreneurship and firm performance: capital constraints (Evans and Jovanovic (1989)), human capital formation, and risk aversion (Kihlstrom and Laffont (1979)). Learning about comparative advantage (Jovanovic (1982)) or productivity shocks (Hopenhayn (1992)) have been established as the leading theoretical explanations for firm exit.

I add to this by showing how occupational choice can influence firm size decisions, on top of the usual role for entrepreneurial ability, financial constraints, or productivity shocks. Holding ability in entrepreneurship fixed, entrepreneurs with better opportunities in wage employment will choose a smaller firm size. This is because they expect a shorter duration of their spell in entrepreneurship. Therefore they expect to have fewer opportunities for earning a return on their capital investment. One strand of literature considers how alternative employment opportunities influence firm outcomes by changing entrepreneurs' incentives for risky investments into the firm (Polkovnichenko (2003), Vereshchagina and Hopenhayn (2009)). They assume a deterministic outside option given by a fixed wage or benefit. By contrast, I focus on the role of uncertainty that entrepreneurs face in the wage labour market, not business risk per se. To the best of my knowledge, this paper is the first contribution that considers this mechanism, together

with a paper that was developed in parallel by Humphries (2017). I differ from Humphries in that I focus on incentives for incumbent entrepreneurs instead of firm size decisions by new entrants. This enables me to derive testable comparable statics which allow for separating the role of opportunities in wage employment from the one of entrepreneurial ability. I also derive a number of novel implications that arise in a model which considers occupational choice and firm size jointly. The most important one is the possibility of multiple local firm size optima.

Labour markets in developing countries. The model that I propose is particularly relevant for developing countries, where the employment share of self-employment is large compared to rich countries. The literature on the nature of occupational choice in labour markets of developing countries dates back to Lewis (1954), Harris and Todaro (1970), and Fields (1975). A recurring theme is the role of self-employment and the nature of entrepreneurship, empirically assessed by Magnac (1991) and more recently Bosch and Maloney (2010), Günther and Launov (2012), and Haywood and Falco (2016). My model nests both the ‘competitive’ and ‘segmented’ view of self-employment in developing country labour markets. It takes agents as utility maximisers who optimally choose from the set of employment opportunities available to them. At the same time it acknowledges the presence of frictions which imply that not all potentially attractive opportunities are available at all times; and that there is ex ante uncertainty about them. Further, my model points out that firm size and thereby entrepreneurial profits are not determined independently. They constitute an endogenous margin of choice which interacts with occupational choice and with the economic environment.

Microenterprises in developing countries. A very active and growing empirical literature examines small firms ('microenterprises') in developing countries, frequently in the context of randomised control trials and other field research. Many empirical contributions amount to tests of the earlier theoretical literature on firm performance surveyed above, exploiting the induced random variation in endowments and constraints. Blattman, Fiala, and Martinez (2014), Fafchamps and Quinn (2016), and McKenzie (2017) report on capital grants experiments to foster high-growth entrepreneurship; and thereby implicitly on credit constraints for aspiring entrepreneurs. McKenzie and Woodruff (2008), De Mel, McKenzie, and Woodruff (2008) and Fafchamps, McKenzie, Quinn, and Woodruff (2014) do the same for existing microenterprises. Related to this, Banerjee, Karlan, and Zinman (2015) summarise the results from six randomised evaluations of microfinance around the world. Karlan and Valdivia (2011) and Karlan, Knight, and Udry (2015) document field experiments on business training and consulting.⁴ These papers amount to a joint test of whether entrepreneurial ability constrains microenterprise growth, and whether it is malleable through training or consulting. Management practices (McKenzie and Woodruff (2016)) are one form of 'ability' that have started to receive attention. A very recent strand of literature examines whether constraints in the form of labour market frictions prevent microenterprises from hiring more workers (Hardy and McCasland (2015), Cohen (2016), de Mel, McKenzie, and Woodruff (2016)).

I contribute to this discussion by considering a margin of choice by microentrepreneurs that has not received attention by this literature: the decision to leave the firm for a different occupation, for example to take up a wage job.

⁴ McKenzie and Woodruff (2014) review this broader literature and its findings.

I show that this can directly be a constraint to the size and growth of small firms. It is a potentially powerful mechanism, given the large rates at which microentrepreneurs give up their firms (as Table 1 showed). My model can rationalise the existence of high returns in microenterprises found by some of the literature (McKenzie and Woodruff (2008), De Mel, McKenzie, and Woodruff (2008) and Fafchamps, McKenzie, Quinn, and Woodruff (2014)). I also discuss how occupational choice interacts with other mechanisms, in particular ability and credit constraints. This sheds light on some of the patterns found in this literature: ex-ante more successful enterprises tend to be the biggest beneficiaries from easing financial constraints, even in settings where the average firm benefits only little.

Entrepreneurs and their firms in rich countries. An early literature brought the theoretical ideas on entrepreneurship surveyed above to US data, in a descriptive and correlational way (Evans (1987), Evans and Jovanovic (1989), Evans and Leighton (1989), Blanchflower and Oswald (1998)). More recently, Hombert, Schoar, Sraer, and Thesmar (2014) and Gottlieb, Townsend, and Xu (2016) test whether increased welfare benefits ease entry into entrepreneurship by providing a safety net.⁵ Hamilton (2000) and Levine and Rubinstein (2017)) ask whether entrepreneurs in the US earn more than wage workers. In the last couple of years, entrepreneurship in rich countries has again received increased attention, due to the rise in self-employment and similar work arrangements, notably in the gig economy (documented by Katz and Krueger (2016) for the US and Adam, Miller, and Pope (2017) for the UK). To some degree this constitutes a convergence of developing and developed country labour markets.

⁵ Bianchi and Bobba (2013) also test this idea for Mexico, and find that a cash transfer welfare programme increased entry into entrepreneurship by enhancing willingness to bear risk. I review literature on the Mexican labour market in Chapter 2.

Although the motivation and the empirical applications of this thesis focus on developing countries, the model in the present chapter is more general than that. Its mechanism relies on a small number of institutional features and market failures that can be expected to apply in rich countries just as well as in developing countries. As a case in point, the assumption that occupational choice is discrete and convex combinations of hours in different occupations are not possible is, if anything at all, more salient in the institutional frameworks of richer countries. This thesis will, however, not test that conjecture and leaves it as an interesting avenue for future research.⁶

Marginal product dispersion and ‘misallocation’. A regular empirical finding in the literature on firms and productivity is that marginal products are not equalised across firms, even within narrowly defined industries (for evidence from developing countries see e.g. Söderbom and Teal (2004) for Ghana, Bartelsman, Haltiwanger, and Scarpetta (2013) for Central and Eastern Europe and Hsieh and Klenow (2009) for India and China.) Hsieh and Klenow (2009) interpret this as evidence for misallocation, in the sense that aggregate productivity and output should be higher if factors were reallocated across firms. David and Venkateswaran (2017) question whether such misallocations are a result of distortions (understood as resulting from policy or institutions) or frictions inherent in the structure of markets. David, Hopenhayn, and Venkateswaran (2016) show that ex-post losses in productivity and output can be substantial when firms make input decisions under imperfect information.

My model replicates the empirical finding of dispersed marginal products

⁶ Humphries (2017) documents that a model very similar to mine is consistent with a number of stylised facts on entrepreneurship in rich countries. He further shows that, inserted into a life-cycle framework, such a model can replicate the rich patterns of self-employment found in Swedish data. He does not attempt at a proper test of the model, though.

of capital across firm. This arises because entrepreneurs with different option values in wage employment apply different effective discount factors on the return to capital, even if their firms are identical. Surviving firms exhibit different marginal products of capital, suggesting allocative inefficiency in an ex-post sense. With this I show that search frictions in the market for entrepreneurs' labour can be a source of information frictions for entrepreneur-run firms.

Non-separability in production and labour. On a more general level, I contribute to the strands of literature that uncover non-separability of choices across markets. Such results generally rely on missing markets because of market failures. In the case of my model, this would be the missing market for managerial labour in small firms. Prominent examples of non-separability results appear in development economics for production and consumption of household farms (De Janvry, Fafchamps, and Sadoulet (1991), Benjamin (1992); LaFave and Thomas (2016))⁷ or in labour economics for fertility and female labour force participation (Killingsworth and Heckman (1986), Eckstein and Wolpin (1989), Blundell, Costa Dias, Meghir, and Shaw (2016)). Poverty traps are a particular kind of non-separability between initial wealth and the stable equilibrium that is eventually reached in a dynamic model. Poverty traps in a production economy usually require a deviation from neoclassical assumptions about the production technology (Banerjee and Newman (1993), Piketty (1997), Aghion and Bolton (1997)) or about preferences (Banerjee and Mullainathan (2010)), in addition to missing credit markets. In this paper, I show how poverty traps can arise in a model with neoclassical (convex) technology and preferences. In-

⁷ LaFave and Thomas (2016) refer to non-separability of farm households' production and consumption, where there is a clear order to the timing of decisions, as 'non-recursivity'.

stead, the poverty trap is created by a discrete-continuous choice model with irreversible investments.

1.2 Model

1.2.1 Conceptual framework

The framework starts with the idea that owners of small businesses in urban areas have access to an active market for wage jobs. Most of their daily decisions about the business may be on the intensive margin – decisions about buying, selling, and investments. Entrepreneurs will periodically consider the extensive margin too – that is, whether to continue or to shut down their business. For large firms, exit decisions are based on alternative uses of capital, and may follow a learning story (Evans (1987), Jovanovic (1982)): firms exit the industry when they have observed their performance sufficiently to conclude that capital is not used in the firm more efficiently than in the market. For small owner-managed firms and self-employed entrepreneurs, however, the opportunity cost of their own labour is perhaps really the binding constraint to firm survival.

There is a crucial difference in how information is revealed in markets for capital and labour, which determine the respective value of alternative uses of production factors. Financial markets make *anonymous* offers of contracts to the general public, and investors are able to accept such contracts at any point in time. Employers, on the other hand, make job offers only to *specific* applicants. At the pre-contracting stage, workers only know the *distribution* of wage contracts in the market, and learn about the value of a specific contract only when an *personal offer* has been made to them. This is

formalised in the canonical partial equilibrium search model (Stigler (1962), Mortensen (1970), McCall (1970)) where workers go around searching for the optimal contract, never knowing precisely at which moment in time an acceptable offer will be made to them. This uncertainty about the opportunity cost of their own labour can have crucial implications for how entrepreneurs conduct their business.

I convey this idea in a simple but generalisable model of an individual entrepreneur. I nest an otherwise standard model of investment behaviour under uncertainty within the canonical two-sector labour market search model.⁸ The self-employed entrepreneur engages in on-the-job search for a wage job, meaning that with a certain probability, a wage offer arrives.⁹ The agent faces a distribution of wages and only learns about the value of a particular possible wage offer if and when it arrives.¹⁰ Entrepreneurs have to decide about how much to invest into their business before an eventual job offer is made. The invested capital is specific to the firm and cannot be liquidated. The reservation wage summarises entrepreneurs' optimal own future occupational choice. A static two-stage model of an individual entrepreneur embedded in a perfect capital market is sufficient to express all the key insights into the mechanism. It gives a closed-form solution for the reservation wage, which allows me to keep the algebra simple and fo-

⁸ Both of these basic models are widely used textbook models, but to my knowledge have not yet been combined in this form in the literature. The model of investment behaviour is identical to the Ramsey (1928) neoclassical growth model with a representative agent. The search model follows the standard textbook version of the Mortensen (1970) model.

⁹ Here I model search as an exogenous process. In Chapter 4 I will include both a costly search effort choice and capital accumulation into an occupational choice model between self-employment, wage employment, and non-employment. In such a model, agents have to trade off two potential investments: investments into capital to increase future earning capacity in self-employment, and investment into costly search to increase the probability of receiving a job offer.

¹⁰ For simplicity, I do not model uncertainty in business profits. Instead, incomes from self-employment are assumed to be deterministic. This implies that agents know about their entrepreneurial ability, and that there is no market risk to running a business. Since I model entrepreneurs as risk-neutral, this assumption is not restrictive. It allows me to focus on the mechanism of occupational choices.

cus on the intuition of how occupational choice interacts with investment behaviour. I provide some extensions to the model after the discussion.

The model is able to generate two key insights. First, it can rationalise why marginal products of capital differ in otherwise identical firms, and why they are higher than the market interest rate. Individuals with better prospects in wage employment will tend towards a smaller capital stock, in anticipation of the possibility that they may shut down the firm – and lose their investment – when an attractive job offer arrives. The model’s comparative statics are directly useful as an empirical test of this prediction. An improvement in either the offer probability or the wage offer distribution shifts incentives towards a smaller firm size. Second, the model can generate multiple local solutions to the conditions for optimal firm size. Changes in the economic environment may then cause agents into select to a different, distant equilibrium.

1.2.2 A model of investment under on-the-job-search

There are two occupations: self-employment and wage-employment. The agent is initially self-employed as the owner of a business (an entrepreneur) with access to a production technology $\pi(k)$. The timing is as following. She rents capital k on a perfectly competitive market at rental rate r . Capital is contracted and then, before production is realised, a wage offer is drawn from a compound lottery. There is a probability λ that the agent receives a wage offer. This allows for the presence of search frictions. If an offer is made, it is itself drawn from a distribution $F(w)$. The stochastic distributions are known to the agent. The agent decides to accept or to reject the wage offer. Production, wage payments, and consumption are realised.

Four assumptions, which I maintain throughout, are critical for the occupational choice mechanism that I model:

Assumption 1.1 *Market environment*

- a) DISCRETE OCCUPATIONAL CHOICE: *An individual can either be self-employed or wage-employed at a given point in time.*
- b) MISSING MARKET FOR MANAGERIAL LABOUR: *Production in a firm only takes place when the firm owner is self-employed.*
- c) FIRM-SPECIFIC CAPITAL: *Capital invested in the firm cannot be liquidated.*
- d) SEARCH FRICTIONS AND WAGE DISPERSION: *A wage offer arrives with probability $0 \leq \lambda < 1$. $F(w; a)$ is the continuous distribution function of wage offers w parameterised by a that admits first order stochastic dominance in a :*

$$F(w, a) = \int_{-\infty}^w f(x, a) dx \quad \text{with} \quad \int_{-\infty}^{+\infty} f(x, a) dx = 1;$$

$$F_a(w, a) = \frac{\partial F(w, a)}{\partial a} \leq 0.$$

Together, the assumption of a missing market for managerial labour and discreteness of occupational choices embody the principle that ownership and control of small firms cannot be separated.¹¹ Firm-specificity of capital im-

¹¹ Fama and Jensen (1983) in their essay “Separation of ownership and control” justify this as following (p. 307):

[C]ombining decision and risk-bearing functions is efficient in small noncomplex organizations because the benefits of unrestricted risk sharing and specialization of decision functions are less than the costs that would be incurred to control the resulting agency problems. The proprietorships, partnerships, and closed corporations observed in small scale production and service activities are the best examples of classical entrepreneurial firms in which the major decision makers are also the major residual risk bearers.

plies that occupational choices are not costless.¹² The assumption of search frictions and wage dispersion introduces uncertainty about the arrival and value of wage-employment opportunities, which is crucial for occupational choices to have any effect on input choices. The parameterisation in Assumption (1.1d) lets me discuss an improvement in the wage distribution by means of comparative statics calculus on parameter a . My notation omits this parameter where not required.

I further assume a neoclassical production function which fulfils the Inada conditions:

Assumption 1.2 *Neoclassical production function:*

The production function $\pi(k)$ is twice continuously differentiable and obeys the Inada conditions:

$$\begin{aligned}\pi(0) &= 0; & \lim_{k \rightarrow 0} \pi'(k) &= \infty; \\ \pi'(k) &> 0; & \lim_{k \rightarrow \infty} \pi'(k) &= 0; \\ \pi''(k) &< 0.\end{aligned}$$

For the static model, I make the following assumption on preferences:

Assumption 1.3 *Preferences*

RISK NEUTRALITY: *The agent's utility function over consumption is given by $u(c) = c$.*

The assumption of risk neutrality isolates the profit motives for capital in-

¹² This cost is borne upon liquidation, at the end of life of a firm. Alternatively, one could assume one-time entry costs that are proportional to the capital stock of a firm, and obtain a very similar solution. What is important is that this cost is independent of the lifetime of the firm.

vestment from other motives such as consumption smoothing.¹³ Given risk neutrality, expected utility equals expected income. At the stage when capital is chosen, occupational choices can only be known in expectations. Expected utility at this stage can be written as

$$EU(k) = \left\{ (1 - \lambda) \cdot [\pi(k) - rk] + \lambda \cdot (E \max(w, \pi(k)) - rk) \right\}, \quad (1.1)$$

where the expectation is taken over the wage distribution. Given that w is monotonic as we integrate along its distribution, there will be a unique reservation wage $z = \pi(k)$ above which all wage offers are accepted. We can rewrite the expression (1.1) as:

$$EU(k) = \left\{ \pi(k) \cdot [1 - \lambda(1 - F(\pi(k)))] + \lambda \int_{\pi(k)} w \cdot f(w) dw - rk \right\}. \quad (1.2)$$

In this representation, there is a single choice variable, capital, which determines both firm size and occupational choice, given technology and the economic environment. The cost of capital has to be borne in all states of the world, whereas returns are only realised when the agent is self-employed.

I define firm size to be synonymous with the capital stock of a firm. Given the assumptions, there is a deterministic one-to-one mapping between capital stock and output.

¹³ Since accepted wage offers will always be higher than profits, and debt repayments are independent of the state of the world, a risk-averse agent might use capital to make a transfer in expectation from the higher-income (wage-employment) to the lower-income (self-employment) state. These incentives are ruled out by assuming risk neutrality.

1.3 Solution and characterisation

I start by characterising the solution to optimal firm size in the absence of changes in the economic environment. I show that a solution to optimal firm size is guaranteed to exist, but that local optima need not be unique. Next, I show how firm size in the neighbourhood of an optimum responds to changes in the economic environment. I derive local comparative statics with respect to a first-order stochastic improvement in the wage offer distribution, and to an increase in the probability that a wage offer is made. I then discuss equilibrium selection and show how that even under multiple local equilibria, the model gives clear global comparative statics prediction.

1.3.1 Optimal firm size

The agent chooses capital to maximise her expected utility. Occupational choice is discrete, and depends on incomes in the two sectors (wage employment, self-employment) as well as on the degree of search frictions. The necessary first-order condition for optimal capital choice is given by:

$$EU_k = \pi'(k) \cdot [1 - \lambda(1 - F(\pi(k)))] - r = 0. \quad (1.3)$$

The returns to capital are multiplied by the probability of staying in self-employment, which I denote $\delta(k)$:

$$\delta(k) = 1 - \lambda(1 - F(\pi(k))). \quad (1.4)$$

If this probability $\delta(k)$ is less than one – i.e. if entrepreneurs take up a wage job with a positive probability – then the firm chooses an optimal size where returns to capital can be substantially higher than the rental rate of capital r .

In other words, the returns from capital are discounted by the endogenous probability of leaving the firm. This is the key insight of this model.

The slope of the marginal utility (i.e. the second derivative of the objective function) is given by

$$EU_{kk} = \pi''(k) \cdot \delta(k) + \pi'(k) \cdot \delta'(k). \quad (1.5)$$

The first-order condition is necessary for a solution to the entrepreneur's problem. The sufficient second-order condition for a local maximum requires (1.5) to be negative at a solution to the first-order condition (1.3). The next section deals with existence and uniqueness of local optima.

1.3.2 Existence and uniqueness

The following proposition summarises the results on existence and uniqueness of solutions to the entrepreneur's problem.

Proposition 1.1 *(Existence and uniqueness) There exists at least one solution to the first-order condition (1.3). There may exist more than one solution. If there is exactly one solution, it is a local maximum. If there is more than one solution, there may be multiple local maxima.*

To prove existence, first note the limits of the marginal product of capital and of the probability of staying in entrepreneurship, which are given by:

$$\begin{aligned} \lim_{k \rightarrow 0} \pi' &= \infty; & \lim_{k \rightarrow 0} \delta(k) &= 1 - \lambda; \\ \lim_{k \rightarrow \infty} \pi' &= 0; & \lim_{k \rightarrow \infty} \delta(k) &= 1. \end{aligned}$$

Together, this implies that EU_k approaches infinity in the limit as capital goes to zero, and that it approaches a negative bound $(-r)$ for very large values of capital. By continuity, there must be a crossing point where $EU_k = 0$. This establishes existence of a solution to (1.3). Given that expected utility increases in the lower limit and decreases in the upper limit, there further must be a solution which is a maximum.

To consider uniqueness, note that although the marginal utility in (1.3) needs to eventually turn negative somewhere in order to satisfy the limit conditions, it does not need to decrease for every level of capital. This follows from the sign of (1.5), which is in general ambiguous. If it is possible that marginal utility increases locally, expected utility is not necessarily globally concave. Therefore, multiple solutions to the first-order condition (1.3) may exist. Together with what we established about the slope of expected utility in the limit, this means that if there are multiple solutions to the first-order condition, there must be multiple local maxima (unless all but one of the solutions are saddle points).

If there are multiple local maxima, only one of these will be a global optimum which satisfies utility maximisation. I will return to the problem of selecting among multiple solutions after discussing how changes to the economic environment affect the choice of capital in the neighbourhood of local solutions.

1.3.3 Local comparative statics

I use comparative statics calculus to describe how local optima change with changes in the economic environment. In particular, I evaluate the effects of changes to the attractiveness and availability of wage-employment. Ex-

ploiting Assumption (1.1d), I model an improvement in the wage distribution as a first order stochastically-dominant shift in the wage distribution by increasing the parameter a . I obtain the partial derivative

$$\frac{dk}{da} = -\frac{EU_{ka}}{EU_{kk}} = -\frac{\lambda \cdot \frac{\partial F(\pi(k);a)}{\partial a} \cdot \pi'(k)}{EU_{kk}} \leq 0, \quad (1.6)$$

where the inequality follows from $EU_{kk} < 0$. A stochastic improvement in the wage offer distribution reduces optimal firm size. The inequality is strict whenever the shock to offers changes the probability of staying in self-employment $\delta(k)$. In (1.6), this is the case where $\frac{\partial F(\pi(k);a)}{\partial a} < 0$; so at levels of k where $\pi(k)$ overlaps with the range of the wage distribution $F(w; a)$. Then, an increase in a increases the mass of wages above the reservation wage $\pi(k)$. For a given reservation wage, it is now more likely that a wage offer is accepted. This brings down the optimal capital investments and in turn the reservation wage, until a new optimum that satisfies (1.3) is reached.

The same qualitative result obtains when the offer probability λ is hit with a shock, instead of the wage distribution, since:

$$\frac{dk}{d\lambda} = \frac{(1 - F(\pi(k); a)) \cdot \pi'(k)}{EU_{kk}} \leq 0. \quad (1.7)$$

Again, the inequality is strict whenever a shock to the wage distribution transmits into a shock to $\delta(k)$. This is the case when $F(\pi(k), a) < 1$. More wage offers only change the probability of staying if there is at least some wage offer that could be acceptable.

I summarise these local comparative statics in Lemma 1.1:

Lemma 1.1 (*Local comparative statics*) (i) *An improvement in the wage offer distribution or* (ii) *an increase in the offer probability reduces firm size*

in the neighbourhood of a local maximum, if it changes the probability of staying in self-employment $\delta(k)$.

1.3.4 Equilibrium selection and global comparative statics

The possibility of multiple local optima requires us to examine not only local, but also global, comparative statics. With multiplicity, a small change in the economic environment can change the order in which local optima are ranked. This would then imply large changes in chosen capital as the agent jumps between local optima.

Ranks are a function of the expected utility associated with each of the local optima. There is no analytical way to express a ranking of solutions in a general; we would need to know all the primitives of the model in order to numerically evaluate expected utility. However, we can use comparative statics on the expected utility function to analytically establish the direction and relative magnitude in which expected utility, and therefore the rank evaluation of local maxima, are being shifted.

First, we can show that both an increase in a and an increase in λ increase expected utility as a function of capital. Using integration by parts, rewrite (1.2) as

$$EU(k) = \pi(k) - rk - \lambda \int_{\pi(k)}^{\infty} [F(w; a) - 1] dw. \quad (1.8)$$

Then,

$$EU_a(k) = -\lambda \cdot \int_{\pi(k)} \frac{\partial F(w; a)}{\partial a} dw \geq 0; \quad (1.9)$$

$$\begin{aligned} EU_\lambda(k) &= -\pi(k) \cdot [1 - F(\pi(k))] + \int_{\pi(k)} w \cdot f(w) dw \\ &= \int_{\pi(k)} (1 - F(w)) dw \geq 0, \end{aligned} \quad (1.10)$$

where the last step again involves performing an integration by parts. We can once more differentiate the partial derivatives (1.9) and (1.10) to show that the magnitude of the change in expected utilities decreases with k :

$$EU_{ak} = \lambda \cdot \frac{\partial F(\pi(k); a)}{\partial a} \cdot \pi'(k) \leq 0; \quad (1.11)$$

$$EU_{\lambda k} = -(1 - F(\pi(k))) \cdot \pi'(k) \leq 0. \quad (1.12)$$

The inequalities are strict under the same conditions that were discussed in the previous section. I summarise the results in Lemma 1.2.

Lemma 1.2 (*Comparative statics of expected utility*) (i) *An improvement in the wage offer distribution, or (ii) an increase in the offer probability increases utility at all levels of capital, if it changes the probability of staying in self-employment $\delta(k)$. The increase is stronger the lower the level of capital.*

These results allow us to make clear predictions about comparative statics without having to evaluate the expected utility, nor having to know which of potentially multiple solutions was selected before the change in the economic environment. Our results all point in the same direction: improvements in wage employment translate into a lower firm size. This can occur either through a local marginal change to the new optimum (Lemma 1.1) or

through a jump from a higher to a lower optimum (a direct implication of Lemma 1.2). We cannot analytically distinguish between which of these two mechanisms is at play. The global comparative statics that I summarise in Proposition 1.2, however, do not require us to make this distinction:

Proposition 1.2 (*Global comparative statics*) (i) *An improvement in the wage offer distribution (an increase in a) weakly decreases optimal firm size.* (ii) *An increase in the offer probability λ weakly decreases optimal firm size. Firm size strictly decreases in cases (i) and (ii) if the increase in parameter changes the probability of staying in self-employment $\delta(k)$.*


In Figure 2, I illustrate the characterisations of the model which we obtained so far, using a parametric form for the functions π and f . I use a Cobb-Douglas production function, and a lognormal wage distribution.¹⁴ The left-hand figure depicts the marginal utility EU_k and the first-order condition (the crossing point of EU_k with the vertical axis). I illustrate two scenarios: a baseline scenario; and an alternative scenario where a small upward shift in the wage offer distribution has taken place.

I first focus on the baseline scenario, depicted in blue dash-dotted curves. The left-hand figure graphs marginal utility. There are three crossing points with the x-axis, marked with thinly dotted vertical lines. These are the three solutions to the first-order conditions. The right-hand side depicts expected utility. The level of capital associated with each of the crossing points is marked in the same way as before. The odd crossing points (A and C) are local maxima, the crossing point in the middle B is a local minimum. If there are multiple local maxima, only one of these will be a global optimum

¹⁴ Appendix A.1 lists the details of parametric assumption for each figure.

FIGURE 2

MULTIPLE SOLUTIONS, LOCAL AND GLOBAL COMPARATIVE STATICS



which satisfies utility maximisation. In this example, C is the global maximum. The thin, dotted horizontal line marks the maximum expected utility in the baseline case, for reference.

Consider what happens after a wage offer improvement (graphed in red solid curves). Marginal utility shifts by a small amount – while this is only visible in parts of the distribution, it can be verified that this is true everywhere. As a result, both local firm size optima shift downward (although only the downward shift of C is visible, with the new solution marked by an unlabeled grey dashed line). This illustrates the local comparative statics result in Lemma 1.1. The shift in marginal utility is related to the shift in $\delta(k)$. At point A , most wage offers are already accepted at baseline. The wage improvement therefore does little to alter $\delta(k)$, the probability of staying. The next section provides a more extensive discussion on the intuition.

The shift in expected utility, shown in the right-hand side of Figure 2, is much more dramatic than the shifts in marginal utility and in local optima. Expected utility improves visibly as a result of better wages. Moreover, these improvements are larger for smaller values of capital. This illustrates Lemma 1.2. As a result, the new solution close to C no longer is a global optimum. The change in the environment leads the entrepreneur to choose a much smaller firm size at A .

1.4 Discussion


In this section, I offer additional discussion on mechanisms, implications and predictions of the model. I start by providing intuition for why a either a low or a high-capital equilibrium can be potentially optimal, and in what sense they constitute underinvestment. Second, I discuss the implications of introducing cross-sectional differences in entrepreneurial ability. Finally, I show how the model is consistent with a number of stylised facts on the heterogeneity of microenterprises.

1.4.1 Intuition

How does non-separability between entrepreneurial and occupational choices arise in this model? Inspecting the first-order condition (1.3), we see that the difference to the neoclassical model arises from introducing the probability of staying, $\delta(k)$, into the firm's problem. The effective discounting of profits drives a wedge between the marginal product of capital and the interest rate. This is how the model expresses that when entrepreneurs consider their own occupational choices, they may decide to run a smaller firm than their potential would allow them to. The fact that effective discounting varies with

firm size also leads to the possibility that both A and C are self-sustaining optima. I discuss the intuition behind these observations with Figure 3.

FIGURE 3
INTUITION: EFFECTIVE DISCOUNT FACTOR AS A TRANSMISSION CHANNEL



In the right-hand panel, I plot the model in the rates-of-return space. Any equilibrium requires the marginal product of capital π' to equal the effective cost of capital r/δ . This is clear from the first order condition. Potential equilibria are given by the local optima A and C from above. At the low-capital solution A , the level of profits that the entrepreneur makes from the small capital stock k_A is low. This directly results in a low reservation wage $z_A = \pi(k_A)$. Most wage offers will be above z_A , and the entrepreneur will accept them. Ex ante, the probability of leaving the firm is high, reaching λ if all potential offers would be accepted. With a high probability of leaving the firm, a forward-looking entrepreneur applies a high effective discount factor, and will choose a small capital stock. Hence, the choices of capital

stock and reservation wage at A are mutually consistent. Similarly, at point C , a high capital stock implies a high reservation wage z_C , which means that few wage offers will be attractive to the entrepreneur. She stays in self-employment with a high probability – so it makes sense for her to choose a large capital stock.

The right-hand panel of Figure 3 also shows the neoclassical benchmark case, in which the firm's decision are separable from the entrepreneur's occupational choices. My model reduces to the neoclassical model if $\lambda = 0$ – if there are no offers of wage-employment – or if profits strictly dominate wages. In the neoclassical model, $\delta(k)$ is always equal to one: entrepreneurs compare the return to capital in the firm only to the opportunity cost of capital r . Optimal firm size under that model is reached at C_R .

Whenever the effective discount factor $\delta(k)$ is smaller than one, optimal firm size is smaller than the potentially achievable maximum C_R . We can verify this from the first-order condition (1.3). This means that entrepreneurs invest less into their firms – they choose a smaller capital stock – than they would in a world where occupational and entrepreneurial choices are separable. From an ex-post perspective (after uncertainty about wage employment has been revealed and entrepreneurs have made their occupational choices) this can be seen as a form of underinvestment: the surviving firm does not equate its marginal productivity with its marginal cost.

The ex-post wedge between marginal productivity and marginal cost is particularly strong when the low-capital solution A is selected as an equilibrium. As Lemma 1.2 has shown, this is more likely when there are many attractive wage employment opportunities for entrepreneurs. Consequently, the firm size that the entrepreneur chooses is particularly far from the bench-

mark. By contrast, when the high-capital equilibrium C is selected, the model solution is closer to the benchmark. A sufficient condition for a unique solution to the first order condition is $\pi'(k)\delta(k)$ decreasing in k ; which is equivalent to $\delta^{-1}(k)/\pi'(k)$ satisfying the monotone likelihood ratio property. Intuitively, if the rise in the probability of staying in entrepreneurship is less steep than the fall in the marginal product of capital, then discounted (expected) returns fall monotonically. If expected returns are monotonic, then only one solution to the return to capital satisfies the optimality conditions.

1.4.2 Heterogeneous entrepreneurial ability

I explore the effects of heterogeneous entrepreneurial ability with panels A and B of Figure 4. I first vary entrepreneurial ability while keeping all other parameters constant.¹⁵ This changes the relative location of entrepreneurial profits and the wage offer distribution in the cross-section. In particular, higher ability shifts the curve of the marginal productivity π' upward, and simultaneously shifts the r/δ curve downward. I then contrast comparative statics of a given wage increase at different levels of ability.

The middle panels take the previous case as a baseline, where C is the optimal firm size, as before. Lower ability entrepreneurs (left panels) choose a low-capital solution A_1 . Higher ability entrepreneurs (right panels) choose a high-capital solution C_2 . In the illustrated examples, the optima for both low and high ability entrepreneurs are unique; whereas baseline ability entrepreneurs select among two local maxima.

¹⁵ Here, higher entrepreneurial ability implies a Hicks-neutral upward productivity shift in the production function.

Now consider again an increase in the wage offer distribution through an increase in a . Note that the discount factor depends on both ability (through the production function) and on the state of the labour market (through f and λ). Therefore, a different δ applies in each case. For ease of readability, my notation continues to simply use δ ; understand as being parameter-specific. In each of the upper panels, the desired rate of return (the effective cost of capital) changes to a new level; graphed as solid red curves. The equilibrium adjustment occurs through a mix in changes to local optima and in changes to equilibrium selection, as can be seen from the lower panels. For lower ability entrepreneurs, optimal firm size changes locally to A'_1 . For higher ability entrepreneurs, a low-capital local maximum begins to exist, but it is still optimal to maintain a high-capital firm; although a slightly lower one than before (at $C'_2 < C_2$). Entrepreneurs with baseline ability switch between local optima (from C to A').


Importantly, in each of the three cases, optimal firm size decreases as a result of the wage shock: Proposition 1.2 holds for any admissible production function. While entrepreneurial ability in itself is an important determinant of firm size, it does not alter the qualitative conclusions we draw from our comparative statics.

1.4.3 Consistency with stylised facts


Before I proceed to finding an empirical setting that allows me to carry out an empirical test of these comparative statics, I will discuss a number of stylised facts from recent literature on firms in developing countries which the model is able to replicate.

FIGURE 4
HETEROGENEOUS ENTREPRENURIAL ABILITY


A. Marginal products and discounted returns



B. Expected utility and global comparative statics



C. The heterogeneous effect of microcredit



High returns to capital. The model is able to rationalise the observation that despite returns to capital that are much higher than the prevailing inter-

est rate for loans, microenterprises do not grow (McKenzie and Woodruff (2008), De Mel, McKenzie, and Woodruff (2008) and Fafchamps, McKenzie, Quinn, and Woodruff (2014)). This is particularly salient in a low-capital equilibrium (point A in Figure 2). The figure illustrates how the return to capital (π') is much higher than the interest rate r at this point. Even when a high-capital equilibrium is selected, a wedge exists between the returns to capital in the firm and the interest rate as long as there is a positive probability of receiving a wage offer higher than the reservation wage.

The introduction of microfinance has transformative impacts only for the ex-ante largest microenterprises. Randomised evaluations of microfinance have found that, while transformative effects for the average microentrepreneur with improved access to finance are absent, there can be large positive effects on and business profits for the ex ante most successful entrepreneurs. There is evidence on this from Mexico (Angelucci, Karlan, and Zinman (2015)) and India (Banerjee, Duflo, Glennerster, and Kinnan (2015), Banerjee, Breza, Duflo, and Kinnan (2015)). Related to this, Fafchamps, McKenzie, Quinn, and Woodruff (2014) find that effects on profits and capital stock from cash and in-kind grants for female microentrepreneurs in Ghana are larger at the top of the distribution.

My model is able to replicate this pattern. I illustrate in panels C of Figure 4. I take the situation from the top panels of the Figure and assume that because of existing credit constraints, only low equilibria (at A_1 , A and A_2) are selected. Now I introduce microcredit, which allows each entrepreneur to borrow an additional amount Δk . As the figure illustrates, this is not enough to bring the lower-ability entrepreneurs to a firm size

which would lead to a higher expected utility. For example, the mid-ability entrepreneur would be able to achieve a firm size B with credit, which is actually a utility-minimising choice. Therefore, these entrepreneurs will not borrow for productive purposes. The lower-ability entrepreneur would not prefer to run a larger firm even if she faced no credit constraints at all. The high-ability entrepreneur could build a larger firm by borrowing Δk which would increase her expected utility; she therefore will choose to take out a loan. In a repeated version of the model, the entrepreneur could grow her microenterprise into a small firm.

Large capital grants can help establish large and lasting firms. This is a corollary of the microfinance case. Potential entrepreneurs that are given a large start-up grant can immediately jump to a high equilibrium (such as C or C_2), provided they have sufficient entrepreneurial ability. Entrepreneurs with a small start-up grant will start small and remain small, absent any further large shock to access to finance. This matches evidence from start-up grants in several sub-Saharan countries (Fafchamps and Quinn (2016), McKenzie (2017)).

1.5 Extension: Intertemporal investment

In this section, I discuss an extension to the basic model. The basic model is a static model with two stages that are divided by the revelation of information about the state of the labour market. Importantly, complete credit markets imply that the agent can reach any level of capital she finds optimal – which sometimes can result in a very small firm. An important question which remains is whether the non-separability of occupational and en-

entrepreneurial choices can leave agents stuck with a low level of capital. This application of the model to microfinance already exhibited a scenario where this is the case: in the bottom-middle panel of Figure 4, the entrepreneur is unable to reach the utility-maximising capital stock at C and is stuck at A' . Here, I introduce a two-period model of intertemporal investment under borrowing constraints which addresses this question further.

Timing of the model: Consider the following extension to the model in section 1.2.2. There are two periods. In period one, the self-employed agent owns capital k_1 and earns income $\pi(k_1)$. Still in period one, she has to decide how much to invest into future firm-specific capital k_2 . The agent consumes what she does not invest: $c_1 = \pi(k_1) + k_1 - k_2$ and receives utility $u(c_1)$ from consumption. At the start of period two, wage offers are drawn from the compound lottery, and the agent decides whether to accept or reject if presented with an offer. Production, wage payments, and consumption are realised.

I introduce two additional assumptions. I specify an assumption on credit constraints, and replace Assumption 1.3 on risk neutrality with an alternative assumption on risk aversion and discounting. The main motivation behind this change is to guarantee an interior solution to the capital accumulation problem. If the credit constraint never binds at an optimal choice, then we can analyse the problem using unconstrained optimisation techniques.

Assumption 1.4 *Credit constraints*

The agent cannot borrow, but she can invest up to the total of her wealth into capital: $k_2 \leq k_1 + \pi(k_1)$.

Assumption 1.5 *Preferences*

- a) **CONCAVE UTILITY:** *The utility function $u(c)$ is twice continuously differentiable, with $u'(c) > 0$ and $u''(c) < 0$.*
- b) **DISCOUNTING:** *The discount factor is $0 < \beta \leq 1$.*

The entrepreneur maximises her expected discounted utility V which depends on initial capital k_1 and the capital choice in the second period k_2 :

$$V(k_1, k_2) = u(\pi(k_1) + k_1 - k_2) + \beta \cdot \left[\lambda \cdot \int_{\pi(k_2)} \left(u(w) - u(\pi(k_2)) \right) \cdot f(w) dw + u(\pi(k_2)) \right] \quad (1.13)$$

where I express the problem in analogous form to (1.2), noting that there is again a unique reservation wage $z = \pi(k_2)$. The first order condition for optimal firm size is given by

$$V_2 = -u'(\pi(k_1) + k_1 - k_2) + \beta \cdot \delta(k_2) \cdot u'(\pi(k_2)) \cdot \pi'(k_2) = 0. \quad (1.14)$$

Condition (1.14) expresses the Euler equation that characterises optimal asset allocation across periods. It implicitly defines the optimal capital choice given initial capital, $k_2(k_1)$. A particular interest lies in steady state solutions where $k_2(k_1) = k_1$.¹⁶ In a steady state solution, the Euler equation simplifies to

$$\beta \cdot \delta(k_2) \cdot \pi'(k_2) = 1. \quad (1.15)$$

We can characterise the solution to the extended model as following. First, equation (1.15) equals the first-order condition from the static baseline model

¹⁶ In analogy to models with many periods, I refer to an equilibrium that fulfils this condition as a ‘steady state equilibrium’.

whenever $\beta = 1/r$. That is, the local optima of the static model with complete capital markets are steady-state solutions in a dynamic model under capital constraints. This mirrors well-known equivalences in macroeconomics between closed and open economy models. Second, whenever $\delta(k_2) = 1$, the model reduces again to a standard neoclassical benchmark model; in this case the Ramsey (1928) model of endogenous capital accumulation and growth. Third, the equivalence with local optima of the static model means that the steady-state equilibria inherit the properties of static solutions. For example, the equivalence characterises these steady-state solutions as local utility maxima. To see this, write the second derivative of V as:

$$\begin{aligned} V_{22} &= u''(\pi(k_1) + k_1 - k_2) + u''(\pi(k_2)) \cdot \beta \delta(k_2) [\pi'(k_2)]^2 \\ &+ \beta u'(\pi(k_2)) \cdot [\delta(k_2)\pi'(k_2) + \delta'(k_2)\pi''(k_2)]. \end{aligned} \quad (1.16)$$

Recognising that in a static optimum, the term in brackets equals (1.5) and is negative, we can establish that (1.16) is negative at $k_2(k_1) = k_1$, which gives the result. Fourth, we can characterise the shape of the policy function $k_2(k_1)$ more generally. Loosely speaking, the marginal expected discounted utility V_2 consists of the discounted marginal utility of the chosen capital stock k_2 – which is given by (1.3) – and a term which penalises large differences between initial and chosen capital stock. The convexity of this penalty term increases with the curvature of the utility function. Fifth and finally, this implies that even when multiple static optima exist, there can be a unique optimal k_2 for every initial capital stock; in this case $k_2(k_1)$ will be a well-defined function.

FIGURE 5
DYNAMIC EXTENSION: SOLUTION AND LAW OF MOTION




Figure 5 illustrates this characterisation. I maintain the parameterisation from the baseline scenario in Figure 2, let $\beta = 1/r = 0.95$, and assume CRRA utility with a coefficient of relative risk aversion of 1.5. The left-hand panel of Figure 5 illustrates expected discounted utility as a function of the capital choice k_2 , for two different initial values k_1 . These initial values correspond to the two steady state equilibria A and C . For both initial values, V is a globally concave function; and it can be verified that this is true for *any* k_1 in the given scenario. This implies that the law of motion $k_2(k_1)$ is a well-defined function. The law of motion is illustrated in the right-hand panel of Figure 5. It follows an S-shape.

This example illustrates the potential of the model to generate dynamic inefficiency; in a way that is sometimes referred to as a ‘poverty trap’. Entrepreneurs with an initial capital stock that is low – such as A – can remain stuck in a low-capital, low profit equilibrium. This equilibrium is not deter-

ministic – after all, entrepreneurs at A leave self-employment at relatively high rates to take up better-paying wage jobs. However, it can be persistent. Those entrepreneurs who start in A and are not lucky enough to receive an attractive wage are worse off than those who start in C . Importantly, they are worse off in the initial period as well as in future periods in which they remain self-employed.

1.6 Testable implications

Testing the model requires variation in the availability or attractiveness of wage employment opportunities. We then expect to see two effects if the model has empirical relevance. First, we would expect more entrepreneurs to switch to wage-employment. Second, we would expect entrepreneurs to reduce their firm size. Furthermore, we would expect this change in firm size to occur in *anticipation* of entrepreneurs' occupational choices.

The model has been set up as a model of an individual entrepreneur. We can think of all of the parameters as individual-specific. In particular the location of the production function and the wage distribution depend on individual ability in either occupation. In the data, however, we observe a cross-section of individuals. We would expect that there exists a distribution of abilities in wage-employment and entrepreneurship, in a Roy-model sense. That is, we should not expect to directly observe the probability of leaving to wage-employment of each individual entrepreneur. This poses a challenge for the test of the model. For instance, the entrepreneurs who can expect to receive more attractive wage offers might also be the more capable entrepreneurs, able run to a larger firm. They might be less likely to be bound by credit constraints. Even their preferences might be different –

they might be the more entrepreneurial ‘type’, keen to take on risk.

Fortunately, the comparative statics derived in section 1.3.3 are robust. They hold for entrepreneurs of every level of ability, as illustrated in Figure 4. They hold locally as well as globally. They hold for risk aversion and for risk neutrality. And they hold if credit constraints are present, as well as when they are not. We can therefore exploit *changes* in wage employment opportunities, instead of cross-sectional variation. With such a test in first differences, we can be agnostic about unobserved heterogeneity in ability, preferences, or credit constraints. While all of these may exist, they will not constitute mechanisms that confound the test. A test in first differences will hold these factors fixed.

1.7 Conclusion

Our traditional understanding of small-scale entrepreneurship makes a stark distinction between the ‘transformative’ entrepreneur as a firm, and the ‘subsistence’ entrepreneur as a self-employed worker. In this chapter, I unite these two perspectives. I propose a new mechanism to explain size differences among entrepreneur-run firms. In my framework, occupational and entrepreneurial choices of small firm owners are not separable. There is a tension between investment into future profitable business opportunities in the firm, and investment into search for a future wage job. By providing a simple but general model, I work out the assumptions that are crucial for this tension to arise: discrete occupational choice, a missing market for managerial labour, firm-specific investments, and uncertainty about wage employment opportunities at the investment stage. Under these assumptions, entrepreneurs discount returns with the probability of leaving entrepreneur-

ship. I show that a static model with complete credit markets, as well as a dynamic model with credit constraints, produce the same qualitative results. I derive local and global comparative statics, and discuss how the model can be tested empirically.

The model is deliberately kept simple. This allows me to isolate the key assumptions that the mechanism requires. Some of these might be relaxed from their strongest form. First, occupational choice is discrete. Most of the signs of comparative statics can be expected to go through if the hours choice was continuous, but the result of multiple equilibria would likely not hold up. Second, small firms can only operate if the owner herself runs and manages the firm. If we allow for external managers instead, the main results would still hold qualitatively if the manager is less productive than the owner-operator entrepreneur. For instance, their incentives might not be aligned with the incentives of the entrepreneur. Third, the firm requires some form of sunk investment that cannot costlessly be reversed. In my model, this assumption takes a particular form: investments are sunk for the duration of the loan contract for capital. We could alternatively think of other types of sunk costs. Fourth, there is uncertainty about future occupational choices at the time of the investment decision. I choose to model this uncertainty with a standard job search framework. One could think of alternative models of uncertainty.

Some of these assumptions can be restrictive. Whether the model applies will depend on the context and on the type of firm. For instance, multiple occupations – with at least one of them in an own business – are the norm rather than the exception in some labour markets, especially in very poor countries. Many self-employed activities do not require large initial

investments, such as petty trading or hawking. Finally, anecdotal evidence suggests that entrepreneurs might be able to install trusted family members to cover spells of absence from their business. Any empirical application of the model will therefore require a careful examination of the context.

Appendix A

A.1 Parametric assumptions in model simulations

I assume very standard functional forms. The production function is assumed to be Cobb-Douglas, with only capital as an input. The wage distribution is log-normal. Search frictions are binary variables that necessarily follow a Bernoulli distribution with a unique parameter, the mean λ . I maintain these functional forms throughout:

$$\pi(k) = A \cdot k^\alpha; \quad (\text{A.1})$$

$$\ln w \sim N(\mu_w, \sigma_w^2). \quad (\text{A.2})$$

The following parameter values were chosen in each simulation:

TABLE A1
PARAMETRIC ASSUMPTIONS IN MODEL SIMULATIONS

Simulation	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
r	1.05	1.05	1.05	1.05	1.05	1.05	1.05	1.05
A	1.6	1.6	1.6	1.6	1.55	1.55	1.65	1.65
α	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
μ_w	–	1.95	2.01	2.05	1.95	2.05	1.95	2.05
σ_w	–	0.3	0.3	0.3	0.3	0.3	0.3	0.3
λ	0	0.7	0.7	0.7	0.7	0.7	0.7	0.7

A.1. PARAMETRIC ASSUMPTIONS IN MODEL SIMULATIONS

These simulations appear in the figures of this chapter as following:

TABLE A2
SIMULATIONS AND MODEL ILLUSTRATION

Figure 2	Baseline scenario: (2);	Alternative scenario: (3)
Figure 3	Baseline scenario: (2);	Neoclassical benchmark: (1)
Figure 4, Baseline ability	Starting scenario: (2)	Wage increase scenario: (4)
Figure 4, Lower ability	Starting scenario: (5)	Wage increase scenario: (6)
Figure 4, Higher ability	Starting scenario: (7)	Wage increase scenario: (8)

In the dynamic extension, I assume that the utility function is CRRA:

$$u(c) = \frac{c^{1-\theta}}{1-\theta}, \quad (\text{A.3})$$

with a coefficient or relative risk aversion $\theta = 1.5$. The discount factor $\beta = 0.95$ is the inverse of the interest rate r .

Chapter 2

Microenterprises and the Lure of Wage Work: Evidence from Mexican Export Manufacturing

2.1 Introduction

In this chapter, I carry out an empirical test of the theory in Chapter 1. In particular, I test the comparative statics that I derive in section 1.3.3 of that chapter: the effect on firm size of having more attractive wage employment opportunities. This requires variation in wage-employment opportunities available to entrepreneurs that is exogenous to other determinants of the returns to capital. I exploit the differential impact on wage employment opportunities of a trade shock: the boom of Mexican manufacturing exports from 2005-2014, and the change in the composition of Mexican exports to the US as more and more high-end manufacturing moved to the country. The boom was interrupted by a bust during the US financial crisis of 2008/9, which temporarily destroyed close to half a million manufacturing jobs in Mexico. This bust and the subsequent recovery introduce additional sources

of variation.

Self-employed entrepreneurs in Mexico and their exposure to export manufacturing jobs provide an ideal testing ground for my theory. I show descriptively that the setting matches all of the key model assumptions: occupational choice is discrete, the vast majority of exiting entrepreneurs cannot hand over the management of their firm to someone else, and firm assets can only be sold at a large markdown. Furthermore, jobs in export manufacturing are financially attractive to almost all entrepreneurs, yet the markets for small firms and export establishments are mostly non-overlapping. I build a dataset that includes almost one million small firms embedded into more than 40,000 unique local labour market observations that span the decade from 2005-2014. Using the universe of social security records, I am able to build a quarterly panel of export manufacturing jobs in every municipality in Mexico, which I define as a local labour market.

The key challenge to identification is that export manufacturing employment might be correlated with factors that directly impact the profitability (and hence the returns to capital) of small firms. I isolate variation in employment that is exogenous to conditions of the local economy by interacting national-level variation in exports at the product level with a local labour market's pre-determined exposure to different export industries. The challenge which remains is that even purely exogenous labour market shocks can affect margins other than occupational choice, for instance through equilibrium adjustments of wages and prices. I overcome these concerns in two ways. First, I directly look for evidence on local price adjustments, and drop the segments of the market for which I find such evidence from the analysis. Second, closely guided by my theoretical model, I explore differences in my

estimated effects over time. In particular, my theory predicts that only anticipated future changes in wage employment opportunities should have any effect. By contrast, equilibrium adjustments will play out in the aftermath of shocks to the wage sector; and so will their effects on outcomes.

I structure my main empirical findings into four steps. First, I show that predicted changes in export manufacturing employment at the level of a municipality are a good predictor of actual changes in this variable. This confirms that the variation that I isolate constitutes a significant shock to local labour markets. I call the predicted local export manufacturing changes ‘export manufacturing shocks’. I also show that it is really the interaction between initial exposure to export industries and the national aggregate volume of exports per industry that constitutes this shock, not the individual components. Second, I document adjustments in local labour markets to export manufacturing shocks. I report how the shock affects transitions into export manufacturing, but not wages in the sector. This is consistent with a model of search frictions where shocks increase the probability of obtaining offers, but not the value of offers. Third, I find that export shocks in the local municipality increase the propensity entrepreneurs to leave self-employment. In terms of timing, this finding extends to quarters after an export manufacturing shock, but not before. Fourth, I find that export shocks negatively affect the growth of small firms: entrepreneurs exposed to local labour markets where more new wage employment opportunities are created are less likely to create net employment, compared to firms in labour markets without these opportunities.¹ Importantly, these effects are present for up to a

¹ In this empirical chapter, I measure firm size (and growth) in terms of labour inputs, not capital stock. This is due to data limitations in the labour force survey. In any well-behaved production function, labour and capital inputs move in the same direction so long as there is some degree of complementarity between the two inputs. In future versions, I hope to complement my analysis with results from the Mexican Economic Census which does measure firms’ capital.

year *before* export manufacturing employment shocks, but not after export manufacturing has arrived.

These findings are consistent with the predictions from my theoretical model. I provide three further pieces of complementary analysis to determine whether they arise through the model mechanism, or through alternative mechanisms with similar observable implications. First, I explore the heterogeneity of effects along the main dimension that is relevant for the model: whether an entrepreneur is really an ‘entrepreneurial type’ or rather a ‘wage worker type’. I find that both effects on occupational choices and on entrepreneurial choices are twice as large for the entrepreneurs who look like wage workers. Second, I directly consider alternative mechanisms. I focus on the two leading alternative explanations for my findings: local demand adjustments, and local wage adjustments. I document that for manufacturing entrepreneurs – who are more likely to operate in similar market segments as export manufacturing plants – there is indeed evidence that alternative mechanisms are at play. I do not find such evidence for non-manufacturing entrepreneurs, for whom the main results uphold. Third, I provide direct evidence which suggests that the *expectations* of employment opportunities are what drives reductions in the growth of small firms. I collect a dataset on announcements of large plant expansions in the auto manufacturing industry in Mexico, the largest among the export manufacturing industries. Using an event study approach, I document that small firm growth only responds after expansions and the job creation associated with them have been made public.

Taken together, I interpret my findings as evidence that the mechanism I postulate is indeed relevant. This has at least two important implications for policy. First, policies aimed at fostering growth of small firms need to take

into account the incentives created for entrepreneurs by the occupational choice channel, which can be an important mediating factor for the success of such policies. Second, from a macro perspective, the non-separability is a channel which causes formal employment in large firms and informal employment to move in opposite directions. Whether this is a boon or a burden depends on the direction of the shock. For instance, a negative trade shock to formal firms – such as during the U.S. crisis of 2008/09 – can be mitigated by the fact that fewer wage employment opportunities increase the incentives for local entrepreneurs to expand their firms and hire workers, who are largely informal.

2.1.1 Related literature

This chapter relates to most of the literature reviewed in section 1.1.2 of Chapter 1. I provide empirical evidence from Mexico, a large upper middle-income developing country, on a new explanation for size and performance of entrepreneur-run growth. I am motivated by recent, mostly experimental, literature on the performance of microenterprises in developing countries, including in Mexico. In addition, my empirical setting and methodology speak to two further strands of empirical literature.

Occupational choice in the Mexican labour market Much of the existing literature views occupational choice in Mexico (and in Latin America more generally) through the lens of ‘informality’, against the backdrop of legislation and social protection (e.g. Maloney (1999), Levy (2010), Levy and Schady (2013)). Here I review some examples. The early 2000s saw a number of labour market reforms, some of which impacted occupational choice. Bruhn (2013) finds that after a business registration reform in 2002, infor-

mal business owners of the ‘wage worker type’ were 20 percent more likely to take up the formal wage jobs created by the reform. A nationwide introduction of consumer credit into 700 branches of the consumer electronics chain Elektra in 2002 created a large shock to access to finance. Bruhn and Love (2014) document that this increased both entry and survival of small informal firms. The 2002-2007 rollout of a new non-contributory health insurance scheme, *Seguro Popular*, changed occupational choice. Bosch and Campos-Vazquez (2014) find it tilted the composition of employment in small and medium-sized firms towards informal jobs. Del Valle (2014) argues that keeping women with caregiving responsibilities in the labour market is a mechanism behind the increase in informality from *Seguro Popular*. The conditional cash transfer programme *Progresa/Oportunidades* is perhaps the most widely known social protection programme in Mexico (see Schultz (2004), Attanasio, Meghir, and Santiago (2011)). Bianchi and Bobba (2013) find that *Progresa* increased entry into entrepreneurship by changing the outside option of potential entrepreneurs and thereby enhancing their willingness to bear risk.

In contrast to much of the standard view on occupational choice in Mexico, I divide workers into entrepreneurs and employees, instead of into formal and informal workers. Entrepreneurs own the capital and earn the residual profit (which may contain an implicit wage) of the firm they work in and manage. Employees receive a set wage from their employer, who can be either an entrepreneur or an incorporated firm. This perspective lets me focus on the two basic *economic* functions that entrepreneurs unite: that of a self-employed worker, and that of a firm. Nevertheless, the question of formality lurks in the background. Whereas both the firms that entrepreneurs run and the workers they employ would mostly be considered as informal,

jobs in large, incorporated export manufacturing firms are formal jobs. They adhere to labour legislation, in particular to mandatory social security affiliation. My results then imply a negative correlation between formal jobs and informal work (both in self-employment and in informal wage work). In particular, the occupational choice mechanism provides disincentives for informal wage employment creation by small firms at times when formal jobs are abundant.

Local labour market effects of trade exposure With my empirical analysis that exploits a shock to the export manufacturing industry in Mexico, I contribute to the literature on the effects of trade shocks on local labour markets. Being mainly concerned with the effects of import competition from emerging economies on labour markets of advanced countries such as the US (e.g., Autor, Dorn, and Hanson (2013)), Acemoglu, Autor, Dorn, Hanson, and Price (2016)) or Germany (Dauth, Findeisen, and Suedekum (2014)) this literature has paid less attention to the flipside of the coin: job creation in developing countries as a result of offshoring and export-led growth. Notable exceptions are Atkin (2016) and Heath and Mobarak (2015) who show how local factory openings alter the opportunity cost of education in Mexico and Bangladesh. Furthermore, Dix-Carneiro (2014) and Costa, Garred, and Pessoa (2016) analyse labour market adjustment to trade shocks in Brazil, and Blattman and Dercon (2016) study occupational choice responses to random allocations of factory jobs in Ethiopia.² McCaig and Pavcnik (2014)’s analysis of trade-induced labor reallocation in Vietnam, unlike most previous work, pays explicit attention to the role of self-employment in developing country labour markets. Methodologically,

² A few working papers have also looked at adjustments of Mexican local labour markets during the U.S. financial crisis (e.g. Kaplan, Lederman, and Robertson (2011), Mendez (2014)).

my paper joins Adão (2015) in exploiting local trade shocks as identifying variation in an occupational choice model.

2.2 Empirical strategy

Testing for separability between entrepreneurial and occupational decisions choices requires variation in wage employment opportunities available to entrepreneurs. For an empirical counterpart to the comparative static (1.7), we need to find a situation where wage employment opportunities influence entrepreneurial decisions only through their effect on the probability entrepreneurs place on leaving self-employment in future periods, and not directly through their effect on profits made by the firm.

The approach I take focusses on variation in wage employment opportunities across local labour markets. I consider wage employment opportunities that are in principle available to any worker because they do not require a narrow set of skills. In particular, I exploit changes in employment created by a set of manufacturing industries dominated by blue-collar jobs in large firms that produce for export to international markets.

There are three main concerns for identification if one were to use observed variation in local export manufacturing jobs. First and foremost, omitted factors at the level of the local labour market could drive both the growth in export manufacturers, and the growth of small firms. Second, there could be a link between export manufacturers and small firms that operates through the direct profit channel, instead of through the occupational choice channel. For instance, linkages in the product markets or competition in the demand for labour could directly influence net profits of small firms. Third,

if small firms themselves are export manufacturers, a mechanical reverse causal effect could be at play.

My starting point to circumvent these identification problems is to isolate the variation in export manufacturing employment that occurs because of shocks to foreign demand for goods produced by Mexican export manufacturers. I construct a measure of job creation in the export manufacturing sector in the local labour market that is purely due to the variation in foreign demand for the goods that the local labour markets specialised in during an initial period. In particular, I express predicted job creation in export manufacturing (XME), per thousand people of working-age, in local labour market m and period t , as:

$$\left(\frac{\Delta \widehat{XME}}{WAP} \right)_{m,t} = \sum_{g \in \{XM\}} \frac{E_{m,0,g}}{WAP_{m,t-1}} \frac{\Delta x_{t,g}}{x_{0,g}}, \quad (2.1)$$

whereby I create the predicted number of jobs per industry g as the interaction of initial local employment in this industry $E_{m,0,g}$ with the change of national Mexican exports of goods from this industry between the current period t and the previous period $\frac{\Delta x_{t,g}}{x_{0,g}}$, expressed as an index normalised to initial period exports. Recent empirical work has used similar measures to allocate trade shocks to local labour markets (e.g. Adão (2015), Autor, Dorn, and Hanson (2013), Topalova (2010)). In particular, I use the variation in bilateral exports to the US, by far the biggest export market, which I allocate to local labour markets defined as municipalities. This allow me to exploit the negative shocks to exports following the US financial crisis of 2008/09, and the recovery period in subsequent years. I normalise everything by the lagged working-age population (in thousand) and sum up over all export manufacturing industries to create a single job creation figure for

a local labour market and period. I usually take a period to be a trimester.

Using predicted instead of actual local job creation provides for variation that is independent of economic forces operating at the local level other than initial exposure to export manufacturing, which can be controlled for. Focussing on a sector dominated by large, internationally integrated firms, I effectively separate the spectrum of firms in two groups: small firms run by entrepreneurs who may have an outside option in local wage employment – the group for whom I expect my model to apply – and large, incorporated export manufacturers whose choices are determined largely by international demand. However, this strategy is still exposed to the potential problem that effects on small firm growth operate through the profit channel, and not through occupational choice.

To address this remaining issue, I begin by noting that I exploit variation at fairly high frequency – a quarter – which should mitigate the influence of price adjustment in local markets, which take time to play out.³ Further, we expect some direct effects on profits to influence small firm growth in the opposite direction to wage employment opportunities: while increased job creation should reduce small firm growth through the occupational choice channel, local demand spillovers should increase it. I overcome any remaining concerns in two ways. First, spillovers to local goods or labour demand are testable as long as they are passed through to observed equilibrium prices. This allows me to detect such spillovers, and to identify market segment for which spillovers are less of a concern. Second, any future ex-

³ Shocks to manufacturing industries need not necessarily propagate through equilibrium adjustments at the local market level, but can rather spread through industrial backward and forward linkages. For the US, Acemoglu, Autor, Dorn, Hanson, and Price (2016) find that only 10-15 percent of job losses due to Chinese import competition can be explained by local adjustments. The most important channel are backward (upstream) industrial linkages with suppliers of directly exposed industries.

port manufacturing shocks should not influence current local price adjustments, thereby shutting down the direct channel. But they should influence entrepreneurial decisions through the anticipated occupational choice channel, according to what I posit in the model. I therefore look at the timing of entrepreneurial decisions and export manufacturing shocks.

I structure my main empirical analysis into four steps. The first two steps are set at the level of the local labour market, in order to understand the aggregate effects of my XME shock measure. I then move to an analysis at the individual level. I look at two outcome variables which jointly embody the model predictions that I wish to test. In step three, I investigate whether XME shocks change the probability of leaving entrepreneurship. In the fourth and final step, I test for whether entrepreneurs adjust their firm size in anticipation of XME shocks.

2.2.1 Empirical specifications

To understand the effects on export job growth on local labour markets, I regress the first-differenced outcome variable $\Delta y_{m,t}$ on the predicted quarterly change of predicted export manufacturing employment (XME) as defined in equation (2.1):

$$\Delta y_{m,t} = \alpha + \beta \left(\frac{\Delta \widehat{XME}}{WAPop} \right)_{m,t} + \mathbf{x}'_{m,t} \beta_1 + \Delta \epsilon_{m,t}. \quad (2.2)$$

By specifying the regression model in the first difference, I remove any time-invariant factors at the level of the municipality, which could influence both outcomes and the *level* of export manufacturing employment. I always control for the sources of variation that are used as inputs in the construction

of predicted XME, and which could directly shift trends in outcomes. In the vector $\mathbf{x}_{m,t}$ I control for trimester fixed effects to pick up any aggregate time series variation contained in the time series of exports; the initial XME share to capture differential trends by initial condition of the municipality; and state fixed effects to avoid spurious correlation due to different trends in regions where export manufacturing happens to be located. Per default, I always include these three sets of controls; and then may add additional controls discussed later.

I estimate the effects of wage employment opportunities in export manufacturing on individual decisions of entrepreneur i using the following regression model:

$$\Delta y_{i,m,t} = \alpha + \beta \left(\frac{\Delta \widehat{XME}}{WAP} \right)_{m,t-s} + \mathbf{x}'_{m,t} \beta_1 + \mathbf{x}'_{i,m,t} \beta_2 + \Delta \epsilon_{i,m,t}, \quad (2.3)$$

where the coefficient of interest is the effect of predicted XME on individual changes in outcomes between period t and $t - 1$. I include the same municipality-level controls as in (2.2). In $\mathbf{x}_{i,m,t}$ I additionally control for individual demographics, firm characteristics, and size of the locality of residence. I furthermore allow for the possibility of shifting the change in predicted XME by s quarters into the future or the past. This allows for picking up anticipation of, and adjustment to, XME shocks.

2.3 Data

I build a dataset that links individual survey data on the Mexican working-age population to administrative records covering the universe of formal employment in Mexican local labour markets. In order for my test with cross-

market variation to have power, I require a large sample covering many different local labour market constellations. Individual data come from 40 trimestral waves of the Mexican labour force survey, ENOE (for *Encuesta Nacional de Ocupación y Empleo*) that span the years 2005 to 2014. The ENOE is collected by the national statistical institute, INEGI, and is targeted to cover around 120,000 households every trimester. The survey has a rotating panel structure, and individuals remain in the survey for 15 months (5 waves) if they remain resident in the same dwelling.

In all of my analysis, I restrict the sample to the working-age population aged 15 to 64. I pay particular attention to the trajectories of individuals who initially own and operate a non-agricultural business with a maximum of 10 workers as their main occupation. I retain around 918,000 spells that run over two quarters, and where individuals start in own-account work in the first of the two periods. Information in the ENOE encompasses the kind of variables typically found in labour force surveys, including sociodemographics, employment information, and income. A smaller subset of the entrepreneurs surveyed for the ENOE underwent an additional, cross-sectional enterprise module called ENAMIN, carried out in the fourth quarter of 2008, 2010 and 2012.⁴

The local labour markets I consider in this analysis are defined by municipalities. At the time of the 2005 Census, Mexico was divided into 2,454 municipalities, and I use these municipality boundaries throughout. I identify municipalities by the unique identifier assigned by INEGI, and merge differ-

⁴ From the ENAMIN I obtain information on capital and investment for around 54,500 small firms with a maximum of 10 workers. However, I refrain from using this data for my analysis. The measuring tool for capital and investment changes between 2008 and 2010, so I cannot pool the three waves. This limits the power of my tests. More crucially, patterns in the data suggest that there might be misclassification in the entrepreneur status for the ENAMIN waves. There is a large spike in ‘transitions’ out of self-employment into wage-employment for the ENAMIN waves, and back into self-employment in the following ENOE wave.

ent data sources by municipality. With this, I am able to capture variation in 41,484 unique local labour market constellations across space (municipalities) and time (quarters) that each overlap with at least one observation from the ENOE. Of these, 25,875 have non-zero predicted export manufacturing employment changes.

Social security records come from IMSS, a government institution tasked both with administering the public social insurance scheme, and with providing services, in particular healthcare and pensions. Enrolment with IMSS is mandatory for all workers, thus by definition the set of all IMSS-affiliated workers is equivalent to the universe of formal employment in Mexico. From disaggregated end-of-month records provided by IMSS, I construct a dataset containing the average number of IMSS affiliates in each trimester, for each municipality and industry, identified by 3-digit NAICS codes.⁵ I exclude the Federal District of Mexico City (D.F.) from my analysis because I am not able to map IMSS records to municipality-level sub-units of the D.F.

For the construction of predicted export manufacturing employment per thousand working age population, I obtain monthly time series of Mexican exports to the US at the level of 3-digit NAICS industries from the US Census Bureau. To construct the working-age population for each municipality and quarter, I combine population counts from the 2005 and 2010 Census, and official population predictions 2010-2050 from CONAPO, the National Population Council. My measure of working-age population is hence based on a model of long-term population trends, and not exposed to short-term

⁵ IMSS maintain their own system of geographical and industry classification. I match IMSS sectors to 3-digit NAICS sectors by hand, and perform the crosswalk from IMSS geographical units to INEGI municipality boundaries constructed by IMSS.

fluctuations.

2.3.1 Mexican export manufacturing

Mexico has seen a boom in export-oriented industrialisation in recent decades, in particular since the creation of NAFTA, the free-trade zone between Mexico, the US, and Canada, in 1994. Sharing a 2000 mile land border with the largest economy in the world unsurprisingly implies that the lion's share of Mexican exports (more than 80 percent) goes to the US. While the initial phases of export manufacturing were characterised by a prevalence of low-skilled jobs, many of whom were held by young women (Atkin (2016)), by 2005 employment in this sector was roughly equally distributed between industries in the garment trade with tasks on the lower end of the skill spectrum, and high-skill intensive industries in electrical and machinery manufacturing, including automotive.⁶

I define the export manufacturing sector to consist of the ten 3-digit NAICS industries where more than 50 percent of output was exported to the US for all months in 2005-2007. This is the period immediately prior to the drop in Mexican manufacturing exports that came with the U.S. financial crisis.^{7 8} The ten industries that fall within this definition of export manufacturing are: Textile products mills; Apparel manufacturing; Leather and allied products manufacturing; Fabricated metal products manufacturing; Machinery and equipment manufacturing; Computer and electronic prod-

⁶ See Verhoogen (2008) on how in the mid-1990s Mexican manufacturing plants responded to negative trade shocks with quality upgrading; and Iacovone, Rauch, and Winters (2013) on how competition from China lead to a reallocation in Mexican manufacturing.

⁷ This definition is in spirit similar to the way Atkin (2016) defines export manufacturing industries. However, I work with the 3-digit NAICS classification system that is applied in both Mexican and US national statistics, whereas Atkin harmonises to the international ISIC classification. I also focus on export industries after 2005, while Atkin looks at the 1985-2000 period.

⁸ I obtain information on manufacturing industry production from the INEGI Monthly Manufacturing Industries Survey (EMIM).

uct manufacturing; Electrical equipment, appliance and component manufacturing; Transportation equipment manufacturing; Furniture and related products manufacturing; Miscellaneous manufacturing.

Taken together, the export manufacturing sector defined this way accounts for two thirds of all bilateral exports from Mexico to the US; around USD 13 billion of a total of USD 19 billion per month.⁹ It provided for an average of 2 million jobs during 2005-2014. Here and for the rest of the chapter, I define export manufacturing employment as formal (IMSS-affiliated) jobs in any of the ten industries. Export manufacturing employment amounts to close to 13 percent of all formal jobs, and 5.7 percent of all paid employment in Mexico. Within export manufacturing, the largest industry by employment and export volume is transportation equipment manufacturing: the automotive sector accounts for 25 percent of all export manufacturing jobs, and 37 percent of exports.

The vast majority of export manufacturing jobs - 85 percent - are blue-collar jobs and require full-time employment.^{10 11} This is mirrored by the characteristics of workers in export manufacturing: around 70 percent hold either an upper or lower secondary degree (which corresponds to 9-12 years of schooling), and they are on average younger than the general working-age population. Almost all (92 percent) export manufacturing firms are incorporated enterprises, and they are much larger than the average Mexican employer: three quarters of the jobs are in large firms with more than 50 em-

⁹ The largest exporting industry that falls not under my definition is oil and gas, which exported USD 2.4 bn per month, on average.

¹⁰ Data on job and worker characteristics come from the ENOE and IMSS data, and data on small firm characteristics come from the matched ENOE-ENAMIN. Sampling weights are used. Table 2 compares these characteristics to those of other wage workers, and those of entrepreneurs.

¹¹ Full-time employment means working at least 35 hours per week during the survey reference period (the last 7 days).

employees, 90 percent in firms with more than 20 employees, and 95 percent in firms with more than 10 employees. As a point of comparison, half of all wage jobs in Mexico are in firms with 10 workers or fewer.

FIGURE 1
EMPLOYMENT AND EXPORTS IN MEXICAN EXPORT MANUFACTURING



Source: IMSS end-of-month social security records on formal employment; US Census Bureau monthly volume of imports from Mexico in nominal USD, deflated with US CPI. Each time series normalised to their respective average annual level in 2005.

By visual inspection of the respective time series displayed in Figure 1, there is a close correlation between the real volume of exports and employment in the export manufacturing sector. Exports and employment were on a slightly increasing path until the end of 2007, then started decreasing in 2008 and fell dramatically until the middle of 2009, when they again recovered. From peak to bottom during the US financial crisis, seasonally adjusted exports fell by about 25 percent, and export manufacturing employment by about 20 percent. This collapse of manufacturing exports, and the subsequent recovery, amplify the variation in wage employment opportunities across

local labour markets over time and space that I will exploit.

2.3.2 Self-employment in the Mexican labour market

Around 67 percent of the Mexican working-age population – individuals between 15 and 64 percent of age – are gainfully employed. Of those in paid employment, a quarter are self-employed, 40 percent are in formal wage employment, and a third are in informal wage-employment. Around 11 million Mexicans were self-employed during 2005-2014, on average. The remainder of this section puts self-employment into the context of the Mexican labour market, with a particular emphasis on the firm size distribution.

Most firms owned by self-employed entrepreneurs are very small, as Table 1 attests. In fact, 80 percent of the self-employed do not have a single paid worker in their firm.¹² 15 percent of small firms have between one and four paid employees. Only 1 percent of small firms have 10 or more paid employees. Notwithstanding that most have few or no workers, small firms as a group are very important for job creation. About 40 percent of all wage jobs are created in firms with 9 employees and less. Together with the self-employment of entrepreneurs, this implies that 55 percent of all jobs in Mexico are created in small firms with at most ten workers.

Three quarters of all workers in small firms subsist in informal employment relationships. Conversely, 80 percent of all informal wage jobs are created in firms with 10 workers or less. This contrasts with export manufacturing employment, which only consists of formal jobs by definition; and which almost exclusively is created by large firms. Table 2 summarises some ad-

¹² Of these 80 percent, at least 20 percent have an unpaid worker, often a family member who helps out in the business.

TABLE 1
EMPLOYMENT BY FIRM SIZE IN MEXICO

Firm Size (incl. entrepreneur)	# En- trepreneurs	# Em- ployees	% firms	% wage employment
1	8.97M	1.16M	82.5	4.1
2-5	1.62M	8.47M	14.9	30.1
6-10	0.16M	2.84M	1.5	10.1
11+	0.12M	15.71M	1.1	55.7
Total	10.88M	28.19M	100.0	100.0

Source: ENOE, 2005-2014. Figures are obtained using sampling weights. Reported are averages across all 10 years. All employees in firms of size 1 are domestic workers.

ditional characteristics of entrepreneurs, workers in export manufacturing, and other wage workers. Entrepreneurs tend to be older and less educated than wage workers. They are less likely to work full time. The commerce sector in particular is dominated by independent businesses. There are no big differences in gender or migration status across the different occupations.

An important assumption of the model is that workers are attracted to export manufacturing by attractive job opportunities, yet that there is uncertainty about these opportunities. Figure 2 plots the distribution of entrepreneurial profits and wages in export manufacturing. Export manufacturing pay on average higher wages; these should be attractive especially for the sizable group of very low-paid entrepreneurs.¹³ Nevertheless, the distributions overlap right to the top, suggesting that there exist export manufacturing jobs that are financially attractive for almost all entrepreneurs.

¹³ Frías, Kaplan, and Verhoogen (2012) document the existence of significant wage premia paid by Mexican exporters.

TABLE 2
INDIVIDUAL AND JOB CHARACTERISTICS BY OCCUPATION


	Entrepreneur	XM worker	Other worker
Employment share	0.25	0.06	0.69
<i>Individual characteristics</i>			
Male	0.57	0.63	0.61
Inter-state migrant	0.25	0.28	0.25
<i>Education</i>			
Primary	0.41	0.19	0.24
Secondary	0.45	0.67	0.54
Tertiary	0.14	0.14	0.23
<i>Age</i>			
15-29	0.15	0.43	0.38
30-39	0.26	0.30	0.27
40-49	0.29	0.19	0.21
50-64	0.28	0.08	0.13
<i>Job characteristics</i>			
Formal (IMSS-affiliated)	0.00	1.00	0.54
Full time (min 35 hrs/week)	0.65	0.96	0.81
<i>Industry</i>			
Manufacturing	0.15	1.00	0.15
Construction & Transport	0.15	0.00	0.17
Commerce	0.37	0.00	0.17
Services	0.32	0.00	0.50
<i>Firm size (# workers)</i>			
1-5	0.98	0.02	0.35
6-10	0.02	0.03	0.11
11-50	0.00	0.13	0.23
51-250	0.00	0.25	0.14
> 250	0.00	0.58	0.17
Observations	1,272,771	285,370	3,825,407

Source: ENOE, 2005-2014. Figures are obtained using sampling weights. Entrepreneurs are self-employed (including employers); export manufacturing (XM) workers are formal wage employees in the export manufacturing industries; other workers are all other employees.

2.3.3 Transitions out of self-employment

There is a high degree of labour churning and employment transitions in the Mexican labour market. Table 3 tabulates quarterly transition rates. Of all workers who were found self-employed in a given period, only two thirds remain self-employed when interviewed in the following trimester. 15 per cent will have taken up a wage job, and a similar share will have left em-

FIGURE 2
THE DISTRIBUTION OF ENTREPRENEURIAL PROFITS AND EXPORT
MANUFACTURING WAGES



Source: 2005-2014 ENOE. The graph shows the distribution of monthly income in self-employment/entrepreneurship and in formal wage work in the export manufacturing sector. Nominal incomes are deflated to 2011 Mexican Pesos using the consumer price index (INPC) provided by the Mexican Central Bank.

ployment. Wage-employment is somewhat more stable: 83 percent of the wage-employed stay in the sector from quarter to quarter, and 93 percent of those in formal wage employment stay (not reported in Table 3).¹⁴

In this chapter I am mainly interested in transitions out of self-employment. Table 4 provides some correlates of the decision to leave self-employment to either non-employment or to a wage job. Entrepreneurs who take up a wage job are predominantly male – as opposed to entrepreneurs who transit to non-employment, three quarters of whom are female. Retail shop owners make up almost half of those leaving work altogether, whereas construction firm owners are largely overrepresented in the group who leave to wage

¹⁴ Workers in export manufacturing stay in wage employment at the same rate as formal workers in other industries.

TABLE 3
TRANSITION RATES IN THE MEXICAN LABOUR MARKET

Employment status	Employment status next quarter					Total
	OLF	WE	SE	OE	UE	
Out of LF %	69.8	8.0	5.4	3.2	13.6	100.0
Wage-employed %	5.7	82.8	5.7	1.2	4.5	100.0
Self-employed %	10.5	15.4	67.5	2.7	3.9	100.0
Other Employment %	24.4	14.2	11.0	43.4	6.9	100.0
Unemployed %	43.9	20.8	6.7	2.9	25.7	100.0
Total %	30.2	41.2	15.5	3.8	9.2	100.0

Source: ENOE, 2005-2014. Quarterly transitions for all respondents aged 15-64 with non-missing employment status in adjacent periods. OLF = Out of labour force; SE = self-employed; WE = wage employed; OE = unpaid employment; UE = unemployed.

employment. Finally, those who take up wage jobs are younger than other entrepreneurs. Migration status, hours of work, and even education seem to play close to no role in the propensity to taking up a wage job. One should bear in mind, though, that these aggregates hide substantial heterogeneity: those who take up manufacturing jobs, for example, are much more likely to have left school with secondary education.

2.3.4 How well does the model match the setting?

The section concludes the descriptive analysis by assessing some of the key model assumptions about the economic environment for entrepreneurs in Mexico against the evidence in the data.

The descriptive statistics above already shed some lights on these questions. A substantial fraction of entrepreneurs leave their occupation, and taking up a wage job is the most important reason. Wage jobs in the export manufacturing sector appear to be an almost disjoint subset of the labour mar-

ket from entrepreneurship: they are formal jobs in large firms, whereas entrepreneurs run small firms with a very high degree of labour informality. Export manufacturing jobs pay relatively well, and judging by the income distributions, there seems to be some wage jobs that are financially attractive for almost every entrepreneur. While there are some differences in demographics across occupations and switching behaviour in a way that we would expect, no single variable seems to be a strong predictor.

TABLE 4
CORRELATES OF THE DECISION TO LEAVE SELF-EMPLOYMENT

	Stay entrepreneur	To non-employment	To wage employment
<i>Individual characteristics</i>			
Male	0.60	0.28	0.75
Inter-state migrant	0.25	0.23	0.25
<i>Education</i>			
Primary	0.40	0.46	0.36
Secondary	0.45	0.45	0.49
Tertiary	0.15	0.10	0.15
<i>Age</i>			
15-29	0.11	0.20	0.22
30-39	0.26	0.24	0.31
40-49	0.31	0.24	0.27
50-64	0.30	0.31	0.20
Full time (min 35 hrs/week)	0.70	0.43	0.69
<i>Industry</i>			
Manufacturing	0.15	0.17	0.11
Construction & Transport	0.13	0.08	0.30
Commerce	0.38	0.44	0.24
Services	0.33	0.29	0.32
Observations	613,736	170,567	132,782

Source: ENOE, 2005-2014. Figures are obtained using sampling weights.

The correlates of the choice to leave self-employment are in line with what the model predicts. Entrepreneurs who take up a wage job experience an average quarter-on-quarter income gain of 3 percent, compared to stagnating incomes for those who stay. Relative growth in firm size in the previous quarter is lower for those who leave self-employment than those who stay.

My model relies on a set of four key assumptions about entrepreneurs, formally stated in Assumption 1.1. First, the choice between self-employment and wage-employment is a discrete one - it is not possible for workers to simultaneously hold a full-time wage job and to run a small business. Second, management of small firms cannot be transferred to another individual. One motivation for this is a missing market for managerial labour. Third, investments into the firm are specific, and cannot be liquidated for some time. Fourth, there is uncertainty about the arrival of wage offers; yet the distributions are known.

I use descriptive data in order to assess whether entrepreneurs in Mexico conform to these assumptions. Secondary activities are rare for Mexican workers: Only 6 percent of those who are self-employed in their main occupation, and 5 percent of those who are wage employed, report any secondary activity. To assess the second assumption, I look at evidence whether small firms continue to exist even though the initial firm owner has changed their main occupation to wage employment. The data allows us to examine two possible ways through which this could happen: entrepreneurs might put the firm into the trusted hands of a family member, or they might be able to continue managing their firm as a secondary, part-time activity. In order to quantify the magnitude of each of these two channels, I look at whether a firm within the same 2-digit NAICS sector continues to exist in the trimester after the initial entrepreneur has switched from self-employment to wage-employment, either with a different member of the same household, or as a secondary activity of the original entrepreneur. I find that 10.8 percent of households of previous firm owners continue to operate a business in the same sector, and that 2.5 percent of entrepreneurs previously self-employed as their primary activity continue to operate a business in the same sector as

their own secondary activity. Conversely, however, this implies that in 88 percent of all instances, there is no evidence that a firm will be continued after the entrepreneur leaves self-employment as their main activity.¹⁵

To assess the ability of exiting entrepreneurs to liquidate assets, I need to observe occupational histories and microdata on personal finance of the same individual. To my knowledge, the Mexican Family Life Survey (MxFLS) is the only dataset that fulfils these criteria. I compare the differences in incomes from major asset sales between individuals who recently left self-employment and entrepreneurs who stayed, controlling for industry and profits of entrepreneurs.¹⁶ The results will be included in a future revision of this paper. Finally, the large dispersion in export manufacturing wages documented in Figure 2 is consistent with uncertainty in the job search process. I will revisit the assumption that (changes in) the distributions governing wage offers are anticipated by entrepreneurs in section 2.5.3.

2.4 Empirical results

In this section I document the main empirical findings. I proceed in four steps. First, I show that my constructed variation in export manufacturing employment is a good predictor of the observed variation. Then I document the impact of export manufacturing shocks on the local labour market. Third, I show how occupational choices of entrepreneurs respond to the creation of jobs in export manufacturing. Finally, I show the effect of future

¹⁵ It is not possible to link the ownership of a particular business over time, so these figures should be considered an estimate. They could constitute an overestimate if household or secondary activity firms in the same sector are in fact different enterprises. Or they could be an underestimate if important ways of continuing the firm are not accounted for in the data, for example installing a manager who does not live in the same household. However, the figure I report can be corroborated for the subset of firms surveyed in the ENAMIN, where 87.5 percent of small business owners declare to run firms that they established themselves.

¹⁶ The MxFLS has no information on business assets or investments.

export manufacturing opportunities on entrepreneurial choices.

2.4.1 The predictive power of exports for employment

I construct the predicted change in export manufacturing jobs per thousand working-age population in each municipality according to equation (2.1). To remove the short-term cyclical and seasonal variation in exports that are evident from Figure 1, I apply a moving-average smoother with a bandwidth of half a year, and use the smoothed export volumes for prediction.¹⁷ For comparison, I also construct predicted export manufacturing employment *levels*.

TABLE 5
GOODNESS OF FIT OF PREDICTED EXPORT MANUFACTURING
EMPLOYMENT

	(1) Levels	(2) Levels	(3) Levels	(4) Levels	(5) FD	(6) FD	(7) FD
Pred. XME/WAPop	0.869*** (0.030) [29.247]	0.875*** (0.030) [29.183]	0.642*** (0.063) [10.131]	0.670*** (0.057) [11.728]			
Δ Pred. XME/WAPop					0.696*** (0.092) [7.547]	0.723*** (0.105) [6.913]	0.761*** (0.116) [6.577]
Constant	2.467*** (0.671) [3.675]	17.043*** (3.347) [5.092]	2.144*** (0.593) [3.615]	16.398*** (3.243) [5.057]	-0.036 (0.070) [-0.510]	0.290 (0.193) [1.505]	0.492*** (0.179) [2.752]
Initial XME	No	No	Yes	Yes	No	No	Yes
Trimester FE	No	Yes	No	Yes	No	Yes	Yes
State FE	No	Yes	No	Yes	No	Yes	Yes
Observations	97520	97520	97520	97520	95082	95082	95082
R^2	0.930	0.938	0.934	0.941	0.235	0.265	0.269

Note: This table reports regression of actual on predicted export manufacturing employment for all 97,520 unique local labour market constellations in the cross-section / 95,082 in the first difference (2,438 non-D.F. municipalities over 40/39 trimesters). Municipal employment variables are normalised per 1000 working-age population of the corresponding trimester and municipality, and winsorised at the top and bottom 0.1%. Two-way clustered standard errors by municipality and quarter in parentheses, t-statistics in square brackets.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

¹⁷ I tried different methods for smoothing, and moving averages turned out to be most effective in removing seasonal variation, and resulted in the instrument with the highest predictive power.

The regression reported in column (5) of Table 5 demonstrates the predictive power of my Bartik-style instruments. The change in export manufacturing employment (XME) predicted from initial exposure and changes in Mexican-US exports alone explains about a quarter of the variation in observed XME changes. The coefficient of 0.696 is highly significant. The following columns reveal that this does not change by much when the individual ingredients of the prediction and state trends are separately controlled for. This confirms that it is really the interaction of export volumes and initial exposure that provides for a strong instrument. The prediction in levels reported in columns (1)-(4) achieves an almost perfect fit, with an R^2 of basically unity. This is driven by the fact that predicted XME *levels* are akin to municipality fixed effects: some municipalities always have a high number of export jobs, while other municipalities always have zero export jobs. For this reason, I disregard the level instruments in my further analysis.

2.4.2 Export shocks and local labour markets

By construction, export-driven variation in manufacturing employment constitutes a shock at the level of the local labour market. I document market-level adjustments to these shocks by regressing average quantities (transition rates) and prices (incomes and wages) on predicted export manufacturing employment changes, using the specification from equation (2.2). Usually, we would expect labour demand to move up wages and thus create incentives for entrepreneurs to move to wage employment. But in a labour market with search frictions, the demand elasticity of wages may be small or even zero, at least in the short run. Higher labour demand might simply manifest itself in a higher offer probability. This is why I look at both margins of adjustment. I further document effects on wages and employment

in other sectors to gauge the potential for spillovers arising from market adjustments. I report results on broad outcomes here, and relegate further discussion of local equilibrium effects to section 2.5.

TABLE 6
EFFECT OF EXPORT SHOCK ON WAGES

Type of income:	Dependent Variable: $100 \times \Delta \log(\text{average income})$					
	(1) Profits	(2) Wages	(3) Profits	(4) Wages	(5) Profits	(6) Wages
Δ Pred. XME/WAPop	0.0446 (0.0678)	0.0429 (0.0273)	0.0593 (0.0721)	0.0440 (0.0288)	0.0675 (0.0749)	0.0207 (0.0269)
Standard controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	No	No	No	No	Yes	Yes
Firm controls	No	No	No	No	Yes	Yes
Locality controls	No	No	No	No	Yes	Yes
Municipality fixed effects	No	No	Yes	Yes	No	No
Dep. var. mean	-0.66	-0.18	-0.66	-0.18	-0.66	-0.18
R ²	0.01	0.01	0.01	0.01	0.02	0.03
# municipalities	1,482	1,491	1,482	1,491	1,481	1,483
Observations	37,240	38,321	37,240	38,321	37,002	37,543

Standard errors in parentheses

This table reports regression of the log average income difference on predicted quarterly export manufacturing employment difference. Unit of observation is an unbalanced panel of Mexican municipalities (apart from D.F.) over 40 quarters from 2005 to 2014. All employment variables are normalised per initial 1000 working-age-population of the corresponding trimester and municipality, and employment and outcome variables are winsorised at the top and bottom 0.05%. All specifications include a constant, and per standard control for state, trimester, 2005 export employment. Individual controls are municipal average levels of education, experience, and share of men. Firm controls are municipal baseline industry shares and firm size distributions. Locality controls are locality size distribution within municipality. Observations are weighted by working-age population. Standard errors clustered by municipality.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

I start by reporting the effect of export manufacturing shocks on municipality-level incomes in Table 6. I am interested in the effects on incomes in two sectors: profits for entrepreneurs, and total wages for employees. I report three specifications for these two outcome variables: a baseline specification with the standard set of instrument controls; a specification with municipality fixed effects; and a specification controlling for demographic and firm

characteristics of the municipality at baseline (i.e. at $t - 1$). The latter is my preferred specification. Municipality controls allow for differential trends depending on time-varying market structures. The estimates are small in economic magnitude and mostly insignificant: coefficients from columns (5) and (6) suggest that an additional export manufacturing job per 1000 workers in the local municipality increases entrepreneurs' profits by 0.07 percent and wages by 0.02 percent.

TABLE 7
EFFECT OF EXPORT SHOCK ON TRANSITION RATES

		(1) To NPE	(2) To WE	(3) To SE
From NPE	Δ Pred. XME/WAP	-0.0526** (0.0263)	0.0540*** (0.0199)	-0.000603 (0.0173)
	Dep. var. mean	83.38	10.92	5.70
	Observations	27,021	27,021	27,021
From WE	Δ Pred. XME/WAP	-0.0227 (0.0280)	0.00174 (0.0407)	0.0212 (0.0246)
	Dep. var. mean	11.48	83.67	4.85
	Observations	26,570	26,584	26,584
From SE	Δ Pred. XME/WAP	0.00569 (0.0681)	0.103** (0.0403)	-0.108 (0.0663)
	Dep. var. mean	19.14	14.61	66.33
	Observations	26,736	26,752	26,752

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses, clustered by municipality.

Note: This table reports regression of transitions (in percentage points, rowwise) on predicted quarterly export manufacturing employment difference. Unit of observation are an unbalanced panel of Mexican municipalities (apart from D.F.) over 40 quarters from 2005 to 2014. Export employment is normalised per initial 1000 working-age-population of the corresponding trimester and municipality, and employment and outcome variables are winsorised at the top and bottom 0.05%. All specifications include a constant, and per standard control for state, trimester, 2005 export employment. Individual controls are municipal average levels of education, experience, and share of men. Firm controls are municipal baseline industry shares and firm size distributions. Locality controls are locality size distribution within municipality. Observations are weighted by working-age population.

In Table 7, I report the effect of export manufacturing shocks on employment transition rates. I divide the labour force into three mutually exclusive groups: those not in paid employment (NPE), the wage-employed (WE), and the entrepreneurs/self-employed (SE). As expected, transitions to wage-

employment from other sectors increase when an export employment shock occurs. The effect on transitions out of self-employment is twice as high as on transitions out of non-employment. Conversely, transitions into NPE and SE are reduced by export shocks; but none of the individual effects are statistically significant. The net effect of all of these transition patterns is that employment shares rise for WE and fall for NPE, by a similar magnitude (see Appendix B.2). Overall, the findings in this section imply that local labour markets adjust employment rather than wages when faced with an export manufacturing shock.¹⁸

2.4.3 Export shocks and occupational choice

A driving force of the model is the notion that entrepreneurs respond to the incentives that wage employment opportunities provide. Descriptively, we have seen significant income gains for switchers from self-employment to wage-employment. Here I provide more causal evidence how export manufacturing employment opportunities affect occupational choices of entrepreneurs.

In Table 8, I report the effect of an export shock in the local municipality on the propensity of entrepreneurs to leave self-employment, using the regression model (2.3). In the model, the only occupational choice is between self-employment and wage employment. In the real world, the choice between non-employment and employment constitutes another important margin that I ought to examine empirically (with the variable denoted ‘Leave’).

The base specification in columns (1) and (3) includes the standard controls:

¹⁸ Using matched employer-employee data from individual-level IMSS records, Kaplan, Lederman, and Robertson (2011) arrive at a similar conclusion for local labour market adjustment during the 2008/09 U.S. crisis.

TABLE 8
EFFECT OF EXPORT SHOCK ON PROPENSITY TO LEAVE
SELF-EMPLOYMENT

Dependent Variable:	(1) Leave	(2) Leave	(3) Leave to WE	(4) Leave to WE
Δ Pred. XME/WAP	0.0846*** (0.0258)	0.105* (0.0580)	0.0931*** (0.0340)	0.108*** (0.0359)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	No	Yes	No	Yes
Firm controls	No	Yes	No	Yes
Locality controls	No	Yes	No	Yes
Dep. var. mean	33.46	33.57	17.68	17.41
R ²	0.01	0.07	0.01	0.08
# distinct LLM	25,875	25,522	25,460	24,985
# municipalities	1,491	1,490	1,479	1,477
# periods	39	39	39	39
Observations	918,160	727,539	747,411	590,176

Standard errors in parentheses


This table reports regression of probability of leaving self-employment on predicted quarterly export manufacturing employment difference in the local municipality. Unit of observation are entrepreneurs with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. Municipal export employment variables are normalised per 1000 working-age-population of the corresponding month and municipality, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by municipality and trimester. Municipality controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

trimester fixed effects, initial municipality XME exposure, and state fixed effects. In my preferred specification in columns (2) and (4), I add vectors of controls at the individual, firm and locality levels to capture potential differential trends. The coefficient estimate implies that an additional export manufacturing wage job per thousand working-age population in the local municipality – a 0.1 percent increase in jobs per adult – raises the probability that an individual leaves self-employment by 0.1 percent. There is virtually no difference when we look at leaving to any sector versus leaving only to wage employment, suggesting that the choice to leave to non-employment is orthogonal to export manufacturing employment shocks.

In a next step, I repeat the exercise from Table 8, but successively shift the export shock to the past and to the future, respectively. That is, I run separate regressions of model (2.3), each for a different value of s . How occupational choice reacts to future export shocks tells us about how individuals respond to anticipated changes; and how choices react to past export shocks is informative about the speed of labour market adjustments. In Figure 3a, I graphically report the results from these individual regressions. The coefficient at zero on the horizontal axis corresponds the estimate in column (4) of Table 8. To its left are coefficients for the effects of future shocks ($s < 0$), and to its right for past shocks ($s > 0$). The asymmetric figure shows that occupational choice of entrepreneurs does not respond to future, anticipated export manufacturing employment shocks. Instead, they leave self-employment in the three quarters following an XME shock.

FIGURE 3
TIMING OF ENTREPRENEUR'S RESPONSES TO EXPORT SHOCKS



Source: Results from regression of model (2.3), using the specification reported in Column (4) of Table 8 (left) and Column (2) of Table 9 (right). Horizontal axis documents the number of quarters s that the dependent variable has been shifted. Each coefficient corresponds to a separate regression. Reported 95 % confidence intervals use two-way clustered standard errors by municipality and quarter.

Taken together, the results from this section show that entrepreneurs are

more likely to shut down their businesses at the time when job opportunities in export manufacturing arrive – but not before. This adjustment continues for another two or three quarters, but not beyond.

2.4.4 Export shocks and entrepreneurial choice

Armed with evidence on what export shocks do to local labour markets, and how individual occupational choices of entrepreneurs react to these shocks, I now proceed to testing for whether entrepreneurial choices, too, respond to local export manufacturing employment (XME) shocks. In particular, I test for whether the growth of small firms is affected by export manufacturing shocks.

My main measure of firm growth is the change in firm size of small firms, expressed by the number of paid workers that small firms employ.¹⁹ I use two variants of this measure. *Absolute firm growth* refers to the raw change in number of workers, for all firms where employment is observed. This measure includes both the intensive margin of firm growth, and the extensive margin for firms who hire their first employee. It cannot be negative for the majority of firms that initially employ no workers. Therefore my preferred measure, *relative firm growth*, expresses growth as the percent change in the number of workers.²⁰ This variable is only defined for firms that initially employ workers.

In the model from Chapter 1, firm growth is affected by *future* wage employment opportunities. In the regression I therefore shift the export shock

¹⁹ The findings are robust to expanding the definition to include unpaid workers, who are mostly members of the same household.

²⁰ So relative firm growth is 100 for a change from 1 to 2 workers, 50 for a change from 2 to 3 workers, etc.

TABLE 9
EFFECT OF EXPORT SHOCK ON SMALL FIRM GROWTH

Dependent Variable:	(1) Relative Firm Growth	(2) Relative Firm Growth	(3) Absolute Firm Growth	(4) Absolute Firm Growth
F.Δ Pred. XME/WAP	-0.618** (0.304)	-0.990*** (0.221)	-0.170 (0.192)	-0.225 (0.199)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	No	Yes	No	Yes
Firm controls	No	Yes	No	Yes
Locality controls	No	Yes	No	Yes
Dep. var. mean	-17.32	-22.17	6.56	3.88
R ²	0.00	0.02	0.00	0.05
# distinct LLM	16,274	14,759	24,507	23,899
# municipalities	1,272	1,255	1,458	1,456
# periods	38	38	38	38
Observations	132,998	97,524	599,408	475,005

Standard errors in parentheses

This table reports regression of firm growth on predicted quarterly export manufacturing employment difference in the local municipality. Unit of observation are entrepreneurs with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. Municipal export employment variables are normalised per 1000 working-age-population of the corresponding month and municipality, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by municipality and trimester. Municipality controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

by one quarter into the future. Table 9 reports the results for the base and the full specification. The effect of future employment opportunities on relative firm growth is negative and statistically significant. The effect on absolute firm growth is negative but insignificant. My preferred estimate in column (2) implies that an additional future export manufacturing wage job per thousand adults in the local municipality reduces growth of small firms by about 4 percent of the sample mean.

The observation that *future* export employment opportunities matter for small firm growth holds more generally when we again examine the effects of different lags and leads of the export shock. I document this pattern in Figure

3b. XME shocks up to three quarters of a year into the *future* have a significant and negative effect on small firms job growth. Past XME shocks, however, do not appear to have any effect. This is in contrast to the finding on occupational choice, which is primarily affected by XME shocks of the recent *past*.

2.5 Mechanisms and Heterogeneity

These findings provide evidence that is consistent with the setup and the predictions of my theoretical model. In this section I present further evidence which suggests that the effects on entrepreneurial choices – reduced growth and investment – arise *through* the channel of anticipated occupational choice. First, I document that not only the main effects but also their heterogeneity exhibits a pattern that is consistent with the model. In a second step, I directly address possible alternative explanations. I show that the results hold for subsamples where the alternative explanations are not likely to apply. Third, I provide direct evidence on the role of expectations with an event study on announcements of large automotive plant expansions. In the appendix, I show that these conclusions are robust to (i) defining commuting zones as the relevant local labour market and (ii) changes in the definition of my set of export manufacturing industries.

2.5.1 Testing for heterogeneous effects

If the main effects I document in Tables 8 and 9 are due the mechanism from the model, then the coefficients on leaving *and* on firm growth should be stronger for those who expect to benefit more from job opportunities.

TABLE 10
HETEROGENEOUS EFFECTS - SPECIES ANALYSIS

Dependent Variable:	Below median Pr(Leave)		Above median Pr(Leave)	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth
Δ Pred. XME/WAP	0.0617 (0.0389)		0.130*** (0.0432)	
F. Δ Pred. XME/WAP		-0.730 (0.485)		-1.309*** (0.333)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	8.84	-22.00	25.46	-22.34
R ²	0.01	0.03	0.05	0.02
# distinct LLM	22,814	10,476	22,286	10,956
# municipalities	1,460	1,090	1,435	1,143
# periods	39	38	39	38
Observations	274,290	46,350	305,071	49,556

Standard errors in parentheses

This table reports regression of outcomes on predicted quarterly export manufacturing employment difference in the local municipality. Unit of observation are entrepreneurs with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. Municipal export employment variables are normalised per 1000 working-age-population of the corresponding month and municipality, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by municipality and trimester. Municipality controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

No single variable is a strong predictor of which workers are most likely to take up job opportunities.²¹ I employ a general statistical classification method that separates the type of entrepreneur who is likely to become a wage worker (in any industry, not just in export manufacturing) from the type who is likely to remain an entrepreneur. In particular, I conduct a species analysis similar to De Mel, McKenzie, and Woodruff (2010). To implement this, I construct the predicted probability from a logit model that

²¹ In Appendix B.3, I test for heterogeneous effects according to education level, gender, age, migration status, full-time versus part-time self-employment, the size of the locality and the sign of the XME change. No clear pattern arises from these distinctions.

regresses a dummy for leaving to wage-employment on a rich set of baseline variables in a flexible specification.²² In Table 10, I split the sample at the median predicted probability. I find that both coefficients are twice as large for the entrepreneurs with a high predicted probability of leaving. This suggests that type of entrepreneurs who are most susceptible to wage offers are also the ones that respond in a much stronger way to changes in export employment opportunities.

2.5.2 Alternative explanations

I consider various alternative explanations for my findings. These alternative explanations all operate through the mechanism that local export employment manufacturing shocks could affect margins of entrepreneurial choice other than through the expectation of leaving self-employment.

First, local export manufacturing shocks could affect the equilibrium of the local economy, for example through their effects on local wages.²³ This is a concern because higher labour costs – just like a higher probability of leaving self-employment – reduce the incentives for hiring workers. In section 2.4.2, I documented that the effect on average wages is small (see Table 6). However, this average effect may hide some important underlying heterogeneity, and the type of small firms that closely compete for workers with export manufacturers could indeed be affected. Table 11 reports the effect on wages separately for the manufacturing and non-manufacturing

²² The variables included in the regression are age categories, education levels, marital status, whether the entrepreneur works full time, the number of children, and a detailed classification of business types (e.g. corner shop with family helpers; construction firm with employees; sole professional working from home); all interacted with gender.

²³ This is the mechanism through which shocks to the exporting sector transmit to small domestic firms in the standard trade model with heterogeneous firms (Melitz (2003)).

sector, and for small and large firms.²⁴ The results suggest that an additional local export manufacturing job bids up wages in small manufacturing firms by 0.21 percent; this is significant at the 5 percent level. The effect is twice as large for starting wages of new workers (see Table A10 in Appendix B.4). These are precisely the workers that small firms need to compete the hardest for. The coefficients for wages paid by small non-manufacturing firms are much smaller and not statistically significant.

A second possibility is that XME shocks lead to local demand spillovers, through supply chain linkages, goods market competition, or Keynesian multiplier effects. Positive spillovers work in the opposite direction to the occupational choice channel: positive demand effects increase the incentives for small firm growth, whereas employment opportunities decrease them. Yet negative spillovers work in the same direction as employment opportunities. They could therefore constitute a confounding mechanism. In Table 11, we observe a negative (but not statistically significant) coefficient on profits of small manufacturing firms, which is suggestive of negative demand spillovers.

This evidence implies that these two alternative mechanisms should especially be of concern for small manufacturing firms. In Table 12, I report separate regressions for entrepreneurs operating manufacturing firms, and those with non-manufacturing firms. Manufacturing entrepreneurs reduce their firm size when faced with an XME shock. But they are, if anything, less likely to take up a wage job. This suggests that indeed mechanisms other than occupational choice are driving the decline in firm size for manu-

²⁴ In the appendix I provide an alternative split by formal and informal employment, instead of employment in large firms and in small firms.

TABLE 11
EFFECTS ON INCOMES IN MANUFACTURING AND
NON-MANUFACTURING FIRMS

Sector:	Dependent Variable: $100 \times \Delta \log(\text{average income})$					
	Manufacturing			Non-manufacturing		
Type of income:	(1) Profits	(2) Small firm wages	(3) Large firm wages	(4) Profits	(5) Small firm wages	(6) Large firm wages
Δ Pred. XME/WAPop	-0.113 (0.183)	0.211** (0.0901)	-0.0435 (0.0586)	0.0630 (0.0734)	0.0465 (0.0606)	0.0781** (0.0388)
Standard controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality fixed effects						
Dep. var. mean	-0.34	-0.12	-0.11	-0.68	-0.24	-0.14
R ²	0.01	0.01	0.01	0.02	0.02	0.01
# municipalities	1,322	1,172	1,026	1,464	1,475	1,369
Observations	21,753	19,437	20,183	35,714	36,069	30,996

Standard errors in parentheses

This table reports regression of the log average income difference on predicted quarterly export manufacturing employment difference. Unit of observation is an unbalanced panel of Mexican municipalities (apart from D.F.) over 40 quarters from 2005 to 2010. All employment variables are normalised per initial 1000 working-age-population of the corresponding trimester and municipality, and employment and outcome variables are winsorised at the top and bottom 0.05%. All specifications include a constant, and per standard control for state, trimester, 2005 export employment. Individual controls are municipal average levels of education, experience, and share of men. Firm controls are municipal baseline industry shares and firm size distributions. Locality controls are locality size distribution within municipality. Observations are weighted by working-age population. Standard errors clustered by municipality.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

facturers. The results for non-manufacturing entrepreneurs – for whom the data show no evidence that labour or product market competition are at play – are similar to the main results, and remain highly significant.

The timing of effects supports this interpretation. Figure 4 graphs the timing separately for manufacturing and non-manufacturing firms. Manufacturing firms reduce their size in the same quarter when export jobs arrive, as well as in the preceding quarter. This supports the interpretation that labour and product market competition – which kick in together with the appearance

TABLE 12
EFFECTS ON OCCUPATIONAL AND ENTREPRENEURIAL CHOICES:
MANUFACTURING VS NON-MANUFACTURING

Dependent Variable:	Manufacturing		Not Manufacturing	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth
Δ Pred. XME/WAP	-0.105 (0.103)		0.131*** (0.0435)	
F. Δ Pred. XME/WAP		-1.461*** (0.555)		-0.869*** (0.142)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	13.70	-19.16	18.05	-22.74
R ²	0.05	0.03	0.08	0.02
# distinct LLM	15,583	5,907	24,405	13,455
# municipalities	1,353	911	1,467	1,205
# periods	39	38	39	38
Observations	83,834	14,700	506,342	82,824

Standard errors in parentheses


This table reports regression of outcomes on predicted quarterly export manufacturing employment difference in the local municipality. Unit of observation are entrepreneurs with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. Municipal export employment variables are normalised per 1000 working-age-population of the corresponding month and municipality, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by municipality and trimester. Municipality controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$


of export jobs – drive the results for entrepreneurs in manufacturing. If the same channels were responsible for the decrease in firm size among non-manufacturing firms, then we should see the same timing. Instead, non-manufacturing firms reduce their firm size three to one quarters *before* the creation of export jobs, and before the entrepreneurs themselves eventually take up opportunities in wage employment.

FIGURE 4
TIMING FOR MANUFACTURING AND NON-MANUFACTURING FIRMS

Manufacturing firms




(a) Occupational choice: leave to WE




(b) Entrepreneurial choice: relative firm growth

Non-manufacturing firms



(c) Occupational choice: leave to WE



(d) Entrepreneurial choice: relative firm growth

Source: Results from regression of model 2.3, using the full set of controls. Horizontal axis documents the number of quarters s that the dependent variable has been shifted. Each coefficient corresponds to a separate regression. Reported 95 % confidence intervals use two-way clustered standard errors by municipality and quarter.

Finally, a third alternative mechanism originates in the fact that XME shocks increase the value of employment in occupations other than self-employment. This implicitly provides an insurance for entrepreneurs against the loss of income when their business venture fails and they are forced to shut it down and leave self-employment.²⁵ This implicit insurance can in turn encourage entrepreneurs to take greater risks, and we should expect to see an increase in the probability of leaving self-employment together with stronger growth in the surviving firms.²⁶ However, I find that leaving self-employment and firm growth move in opposite directions. This rules out the insurance mechanism.

2.5.3 The role of expectations

In order to have any effect on entrepreneurs' behaviour, future local export manufacturing employment opportunities must be anticipated by entrepreneurs. While comprehensive administrative data can tell us precisely when and where export manufacturing jobs were created, they are silent about when information about these jobs was revealed. Fortunately, publicly listed multinational companies centrally reveal information about major investment projects. The announcements are targeted mainly at the firm's shareholders. This revelation of information is well-documented in the financial and special interest industry press. I conduct an event study for the universe of Mexican automotive manufacturing plants, exploiting such central revelation of information on major plant expansions. Transport equipment manufacturing (NAICS sub-sector 336) is the largest among Mexican

²⁵ See Gottlieb, Townsend, and Xu (2016) for evidence on the implicit insurance of job-protected leave for women giving birth in Canada.

²⁶ However, increased entry might also create stronger competition and lead to job destruction in incumbent firms; see Hombert, Schoar, Sraer, and Thesmar (2014) for evidence from France in the context of a welfare reform aimed at insuring unemployed entrepreneurs if they enter new business ventures.

export manufacturing industries in terms of employment and export volume; and the one with the fastest growth rate (see section 2.3.1).

I construct a database on announcements of large plant expansions in Mexico from the leading special interest publication for the automotive industry, *Automotive News*. By the end of 2014, Mexico had 18 automotive manufacturing plants. Most of these were final assembly plants for cars and light trucks; a few establishments were dedicated to intermediate processes such as transmissions production, engine assembly, or stamping. In total, I identify 56 public announcements of major plant expansions: introductions of new product lines, capacity expansions, installation of major first tier suppliers, and two greenfield plant openings. Most of these expansions tend to take three quarters to a year, between the first announcement and the completed expansion. Within each local labour market indexed by z , I assign dummy variables γ_q for each quarter since the announcement in period t , which itself is assigned to event time $q = 0$.²⁷ I also assign dummies for quarters before the announcement. I corroborate these dates with municipality-level time-series of transport equipment manufacturing employment constructed from IMSS records.

My basic event study specification at the individual level is the regression model

$$\Delta y_{i,z,t} = \alpha + \sum_{q=Q_b}^{Q_e} \gamma_q + \mathbf{x}'_{z,t}\beta_1 + \mathbf{x}'_{i,z,t}\beta_2 + \Delta\epsilon_{i,z,t}, \quad (2.4)$$

where control variables are defined as before, and where quarterly event dummies capture a non-parametric trend of arbitrary shape. I additionally

²⁷ In the case of greenfield assembly plants – which tend to exhibit a lag of around 3 years between the first public announcement and the start of production – I work backward from the start of production. The first quarter generally seems to coincide with the start of major construction works. I allocate four quarters ($q = 0$ through $q = 3$) exclusively to a single expansion. That is, I allow for subsequent events to take place in $q > 3$.

estimate a parametric form of the evolution of outcomes following an announcement. That is, I place restrictions of a particular parametric shape on the time dummies, together with an assumption on the average effect before the announcement.²⁸ I do not place any restriction on the announcement period itself, since announcements can occur at any point during a quarter. For instance, in the case of restricting the shape to a third-order polynomial, I estimate equation (2.4) subject to:

$$\gamma_q = \begin{cases} 0 & \text{if } q < 0; \\ \phi_0 + \phi_1 \cdot q + \phi_2 \cdot q^2 + \phi_3 \cdot q^3 & \text{if } q > 0. \end{cases} \quad (2.5)$$


Most auto plants are located in suburbs of mid-sized metropolitan areas. Many such suburbs constitute their own municipality. In order to adequately capture the exposure area for major plant expansions, I define commuting zones (CZ) as the relevant local labour market for the purpose of this event study.²⁹

I describe the evolution of auto jobs in local labour markets using an adapted version of model (2.4), without controls. In Figure 5, I report the change in CZ-level auto jobs for the commuting zones with an auto manufacturing plant. Each coefficient can be interpreted as the change of auto jobs in quarter q following a major expansion announcement, relative to other periods in these commuting zones. I also plot a quadratic trend estimate of the post-announcement change in auto jobs. Expansion announcements do not occur

²⁸ This follows the estimation of trajectories of treatment effects of randomised control trials (Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn (2016)).

²⁹ I use two pieces of information to aggregate up municipalities to commuting zones. First, I use the 2005 definition of *Zonas Metropolitanas* maintained by a working group of Mexican public agencies, and published by INEGI. Second, I apply a further aggregation from Atkin (2016) based on cross-municipality commuting patterns in the 2000 Census. These aggregation steps seem sufficient to properly delineate the urban area around an auto plant.

FIGURE 5
AUTOMOTIVE CASE STUDY: JOB CREATION



Source: The graph shows the coefficients from a regression of labour-market level changes in auto jobs on quarterly dummies indicating the time from the announcement of a major plant expansion. Confidence intervals at 95 percent. The quadratic trend corresponds to a constrained CZ-level regression.

entirely randomly. As the figure shows, announcements are preceded by an above-average increase in auto jobs in the commuting zone. There are at least two possibilities for this pattern. First, the auto plant starts its expansion already before the public announcement. Second, local suppliers (some of which fall within the same NAICS industry code) are driving this job creation. In order to distinguish between these explanations, I additionally report the evolution of auto jobs only for the municipality where the auto plant is located. There is no pre-announcement trend within the municipality itself. This suggests that this pre-announcement growth occurs in suppliers located elsewhere in the CZ. It is consistent with anecdotal evidence that local supply chains are a major determinant of plant location choices.

2.5. MECHANISMS AND HETEROGENEITY

TABLE 13
AUTOMOTIVE CASE STUDY: SMALL FIRM GROWTH

	Dependent variable: Relative firm growth			
	(1)	(2)	(3)	(4)
1Q to announcement	0.287 (7.479)	-3.852 (6.996)	-2.334 (7.342)	-2.164 (7.864)
2Q to announcement	-2.598 (2.830)	-4.304** (2.072)	-2.784 (2.021)	-3.336* (1.999)
3Q to announcement	-1.429 (2.877)	-0.983 (3.559)	0.546 (3.342)	-0.165 (3.213)
4Q to announcement	0.125 (8.594)	0.525 (5.534)	2.138 (5.582)	-2.807 (2.720)
Expansion announcement	-4.401* (2.278)	-2.214 (2.913)	-0.788 (2.399)	-0.663 (1.999)
Q1 after announcement	-6.719*** (2.269)	-6.587** (2.629)	-5.060** (2.266)	-5.278** (2.112)
Q2 after announcement	-2.501 (3.702)	-11.09** (4.675)	-9.614* (5.130)	-9.426* (5.541)
Q3 after announcement	0.365 (6.225)	5.157 (9.887)	6.787 (9.375)	6.531 (10.64)
Q4 after announcement	11.08 (8.929)	-1.782 (3.847)	-0.379 (4.017)	0.635 (4.085)
Q5 after announcement	10.80*** (3.470)	12.27*** (4.066)	13.65*** (3.999)	12.76*** (4.017)
Q6 after announcement	1.775 (2.963)	8.065*** (2.395)	9.417*** (2.243)	9.336*** (2.101)
Q7 after announcement	12.50 (9.938)	1.402 (5.064)	2.697 (5.159)	3.120 (5.483)
Q8 after announcement	6.836* (3.972)	7.171 (5.978)	8.555 (6.050)	2.355 (3.209)
Auto CZ			-4.526** (2.007)	-4.933** (2.243)
F.Δ Pred. XME/WAP				-0.903*** (0.310)
Auto CZ × F.Δ Pred. XME/WAP				1.316 (1.015)
Expansion Q0-Q2 × F.Δ Pred. XME/WAP				0.368 (1.010)
CZ controls	No	Yes	Yes	Yes
Individual controls	No	Yes	Yes	Yes
Firm controls	No	Yes	Yes	Yes
Locality controls	No	Yes	Yes	Yes
Dep. var. mean	-17.02	-22.09	-22.09	-21.96
R ²	0.00	0.02	0.02	0.02
# distinct LLM	11,874	10,871	10,871	10,595
# CZ	1,038	1,024	1,024	1,008
# periods	39	39	39	38
Observations	139,394	101,674	101,674	99,622

Standard errors in parentheses

In this table, I explore the effects of an announcement of a major auto manufacturing plant event in the local commuting zone (CZ). Unit of observation are entrepreneurs with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. CZ export employment variables are normalised per 1000 working-age-population of the corresponding month and CZ, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by CZ and trimester. CZ controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$


Auto job creation rises strongly following an announcement. It peaks a year after the announcement; which is when most expansions would be completed. During the first post-announcement year, on average 2,200 additional auto jobs are created in the local commuting zone. This combines jobs created directly in the announcing plant and indirectly in local suppliers.³⁰ Excess local job creation continues for another two or three quarters before falling back to zero.

In order to test for whether this revelation of information on future auto job opportunities has an effect on the firm size choices of entrepreneurs, I estimate model (2.4) for my preferred outcome variable, relative firm growth. My comparison group are observations on entrepreneurs within auto plant CZ that are outside the event study window, and all observations on entrepreneurs not residing in auto plant CZ. I report the regression results in Table 13. In my preferred specification in column (2), I include the full set of control variables from earlier specifications. I plot the coefficients and their 95 percent confidence interval in Figure 6. I also plot the estimated cubic trend.

First, note that, unlike for auto jobs, there is no pre-announcement trend in the growth of small firms, conditional on controls. Even though expansion announcements are not random, this parallel trend provides some confidence that any subsequent changes in small firm growth will be a result of the announcement. Second, small firm growth diminishes strongly in the first and second quarter after the announcement. These short-term effects

³⁰ This figure is descriptive; it does not constitute an estimate of the employment impact. The overall job creation of an expansion is likely larger than the subsector-level employment impact. Casual evidence from *Automotive News* suggests employment multipliers in the order of 5-6 for auto plants. Many suppliers do not share the same NAICS subsector code with auto plants.

FIGURE 6
AUTOMOTIVE CASE STUDY: SMALL FIRM GROWTH



Source: The dots plot the coefficients from model (2) of Table 13 together with their 95 percent confidence intervals. The line shows a cubic trend, imposing the additional restriction (2.5) on the regression model, with 95 percent confidence intervals. The hollow circle shows the coefficient on the announcement period from the constrained model.

are economically large. The coefficient at Q2 implies a 50 percent reduction in firm growth relative to the mean. Third, small firm growth reverts back to normal a year after the announcement. It even increases in the post-expansion phase, once the expanded plant is established. This is consistent with positive local demand spillovers that now come in as the dominant channel driving small firm growth.

If plant expansion announcements contain all the relevant information on changes to job opportunities over the course of the next quarters, then the effect of any further variables expressing changes to export manufacturing employment should be zero once we condition on announcements. Furthermore, if the lack of an announcement in my data implies that other episodes of job growth are unanticipated, then other XME changes in these specific local labour markets should not have any effect either. I test these hypothe-

ses in column (4) of Table 13. I augment the event study specification with predicted future XME changes and their interactions with auto plant CZ and announcement indicators. I cannot reject the hypotheses that, conditional on an announcement, (i) predicted future XME changes have no effect in auto plant commuting zones ($p = 0.67$) and (ii) predicted XME changes have no effect in the immediate aftermath of an announcement ($p = 0.31$). However, predicted future XME changes in municipalities that are not included in the event study continue to have a negative effect on small firm growth. The estimated coefficient is very similar to the main effect reported in column (2) of Table 9.

Taken together, this provides strong suggestive evidence that it is really the information about employment opportunities in the immediate future which links export manufacturing job creation to reduced growth in entrepreneur-run small firms. This holds true in the case of the auto industry, the largest export manufacturing sector. More speculatively, we can compare the graphical pattern in Figure 6 from this case study to the more general pattern of timing from Figure 3b. Relative to the peak of auto job creation in Q3, we see firm growth reduced in the two or three quarters that precede this peak. This mirrors the general pattern that small firm growth reduces in the three quarters prior to an XME shock.

2.5.4 Robustness

I relegate other robustness checks to Appendix B.5. First, I report OLS results. The OLS bias seems to work in the opposite direction of my model predictions: a downward (or attenuation) bias for leaving, and an upward bias for firm growth. Next, I re-run the main regressions at the commuting

zone level. The results hold overall – some coefficients become stronger, some weaker. Finally, I check whether local production shocks in industries with a strong agglomeration of production affect the results. If one or few municipalities can affect aggregate industry exports, then exports are no longer driven purely by US demand, an external factor. This may bias results. To implement this check, I in turn leave out each constituent export manufacturing industry. Results suggest that, if anything, the magnitudes are larger when I leave out highly concentrated industries. This suggests that any such potential bias would work in opposite direction to the occupational choice mechanism.

2.6 Conclusion

Small-scale entrepreneurship creates self-employment opportunities for a large fraction of workers in many developing countries. As a group, these entrepreneurs employ a sizeable fraction of workers who hold a wage job. Yet individually, most of their firms are very small, seldom employ any workers, and can rarely be found on transformative paths towards developing into a larger enterprise. This is despite vast support programmes to alleviate an array of constraints for entrepreneurs known to policymakers and researchers. Recent changes in the structure of labour markets and in information technology have made these issues salient again even in many advanced economies.

In this chapter, I provide an empirical test of a novel mechanism influencing firm size of entrepreneur-run small firms: the anticipation of outside employment opportunities provides incentives for underinvesting into firms. I look at small firms in Mexico over the 2005-2014 decade. I exploit

shocks to employment in the Mexican export manufacturing sector brought by the bust-and-boom of exports to the US during the financial crisis of 2008/09 and the subsequent recovery. I document how these shocks transmit into local labour market adjustments. They increase the transitions of entrepreneurs out of self-employment into wage employment, and they induce entrepreneurs to reduce their firm size in anticipation. With a species analysis, a case study of large plant expansion announcements in the Mexican auto industry, and an analysis of differential timing of effects, I distinguish between competing mechanisms behind these findings.

My empirical and theoretical results point to frequent transitions out of self-employed entrepreneurship as a new and additional determinant of firm size, at least in the lower part of the size distribution. Jobs in the export sector are one source for these incentives. This points at a tension between job creation in large, ‘formal’ establishments and in small informal businesses. The flipside, however, is the implication that small firms can act as a sort of automatic employment stabiliser against negative shocks to large, exporting firms. The 2016 US presidential election, and the new potential shock to Mexican exports it constitutes, may provide a setting for future empirical research on this implication. It would also be useful to determine to which degree this model can explain heterogeneous findings in the recent experimental microenterprise literature.

The present analysis remains subject to a few caveats. By construction, changes in wage employment opportunities – available to those who are not yet wage-employed – constitute market-level shocks. My empirical analysis distinguishes the occupational choice mechanism from other market adjustments that work in same direction, and therefore could confound the


mechanism. This is sufficient for establishing the existence of the mechanism. However, determining the magnitude, and assessing the welfare consequences of non-separability, require knowledge of the counterfactual: what firm size would entrepreneurs choose under separability of occupational and entrepreneurial choices? Such a counterfactual would move local economies to a completely different equilibrium, if it applies to all entrepreneurs simultaneously. More structure would be needed to establish the full counterfactual, and to answer these questions.

Appendix B

B.1 Map of Mexican export manufacturing in 2005

FIGURE A1

EXPORT MANUFACTURING EMPLOYMENT ACROSS MEXICAN MUNICIPALITIES IN 2005



Source: The map shows the penetration (per 1,000 working-age population) of export manufacturing jobs at the municipality level in 2005. Employment data from IMSS, population data from 2005 Census, municipality and state boundaries from INEGI *Marco Geoestadístico Nacional*.

B.2 Additional results on sector shares

This appendix documents additional results on the effects of XME shocks on sector employment shares in Table A1. I split the former category ‘not in paid employment’ into its components: out of the labour force (not employed, not searching for a job), other employment (mostly unpaid) and unemployment. The results in the table indicate that the significant increase in wage-employment which results from an XME is mostly (to 80 %) offset by a reduction in unemployment. The bulk of the remainder is driven by a reduction in entrepreneurship / self-employment, although this effect is economically small and statistically not significant. In Table A2, I further split up the increase in wage-employment. Unlike in table A1, the columns do not need to add up. Column (1) reproduces the estimate on overall wage employment. In columns (2) and (3) I split up wage-employment into large firms (more than 10 workers) and small firms. All of the employment effects are driven by large firms. In columns (4) and (5) I offer an alternative split into formal and informal wage-employment. Again, most of the effects are driven by formal employment. In column (6), I separately estimate the effects on export manufacturing employment, a subset of formal wage-employment. The magnitude of the estimate indicates that while constituting the largest part, not all wage employment effects are created by the export manufacturing sector. In other words, some (around 0.02 percentage points within an overall effect of 0.065) of jobs are created in other sectors.

B.2. ADDITIONAL RESULTS ON SECTOR SHARES

TABLE A1
EFFECT OF EXPORT SHOCK ON EMPLOYMENT SHARES

Type of employment:	Dependent Variable: $100 \times \Delta$ sectoral employment share				
	(1) Out of labour force	(2) Wage Employment	(3) Entre- preneurship	(4) Other employment	(5) Unempl- oyment
Δ Pred. XME/WAP	-0.00431 (0.0122)	0.0651*** (0.0120)	-0.00912 (0.00895)	-0.000870 (0.00480)	-0.0515*** (0.0118)
Municipality controls	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes	Yes
R ²	0.02	0.04	0.06	0.03	0.01
# municipalities	1,496	1,496	1,496	1,496	1,496
Observations	38,331	38,331	38,331	38,331	38,331

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses, clustered by municipality.
Note: This table reports regression of employment shares on predicted quarterly export manufacturing employment difference. Sample is an unbalanced panel of Mexican municipalities (apart from D.F.) over 40 quarters from 2005 to 2014. Export employment is normalised per initial 1000 working-age-population of the corresponding trimester and municipality, and employment and outcome variables are winsorised at the top and bottom 0.05%. All specifications include a constant, and per standard control for state, trimester, 2005 export employment. Individual controls are municipal average levels of education, experience, and share of men. Firm controls are municipal baseline industry shares and firm size distributions. Locality controls are locality size distribution within municipality. Observations are weighted by working-age population.

TABLE A2
EFFECT OF EXPORT SHOCK ON TYPES OF WAGE EMPLOYMENT

Type of wage-empl.:	Dependent Variable: $100 \times \Delta$ sectoral employment share					
	(1) Any	(2) Large Firm	(3) Small firm	(4) Formal	(5) Informal	(6) XME
Δ Pred. XME/WAP	0.0651*** (0.0120)	0.0659*** (0.0105)	-0.0000653 (0.0162)	0.0632*** (0.0127)	0.00221 (0.00727)	0.0477*** (0.00564)
Municipality controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.04	0.12	0.12	0.03	0.04	0.04
# municipalities	1,496	1,495	1,495	1,496	1,496	1,496
Observations	38,331	37,186	37,186	38,331	38,331	38,331

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses, clustered by municipality.
Note: This table reports regression of employment shares on predicted quarterly export manufacturing employment difference. Sample is an unbalanced panel of Mexican municipalities (apart from D.F.) over 40 quarters from 2005 to 2014. Export employment is normalised per initial 1000 working-age-population of the corresponding trimester and municipality, and employment and outcome variables are winsorised at the top and bottom 0.05%. All specifications include a constant, and per standard control for state, trimester, 2005 export employment. Individual controls are municipal average levels of education, experience, and share of men. Firm controls are municipal baseline industry shares and firm size distributions. Locality controls are locality size distribution within municipality. Observations are weighted by working-age population.

B.3 Additional results on heterogeneity

This appendix documents additional heterogeneity analysis based on sample splits along individual variables. Overall, no clear pattern emerges. I speculate that the underlying reason is that each subgroup differs not only in their propensity to take a wage job but also in the type of business they run, etc. Therefore, theoretical predictions about differences across sub-groups are ambiguous. This makes interpretation of these results quite hard. Furthermore, subsamples can be of quite different size, which affects the power of the test. I nevertheless report results for completeness.

The most compelling heterogeneity result is when the sample is split by the entrepreneur's level of education (Table A4). While the coefficient is quantitatively similar (and not significantly different) across education groups, the coefficient on leaving to WE is only individually significant for mid-skilled workers with secondary education. This is the most prevalent qualification among export manufacturing workers. The coefficient on relative firm growth is large, highly significant, and significantly different from the coefficients on other education group, which are small and not significantly different from zero.

When splitting by gender (Table A4) or age (Table A5) no clear pattern or conclusion emerges at all. For splits by hours (A8), size of locality (Table A7), or migrant status of the entrepreneur (Table A6), coefficients are similar to the main results (in size and statistical significance) for full-time entrepreneurs, those in urban localities, and non-migrants. The mechanism should be expected to hold especially for these sub-samples. In particular, the fact that results uphold for non-migrants alleviate some concerns that selective migration confounds results. Results for part-time, rural and migrant entrepreneurs are more mixed and not very conclusive. In any case, the differences across subgroups are not significant.

Finally, I split the sample not according to personal or locality characteristics, but according to the sign of the shock. For this split, theory provides some guidance. In order to have any effect on firm growth, according to the model, future wage employment opportunities need to be anticipated, at least a few quarters in advance. This seems more likely in the case of expansions – new factories, or enlargements of existing plants, take time to build. The automobile case study in section 2.5.3 is case in point. Redundancies and hiring freezes, on the other hand, tend to take place without much advance warning. Table A9 documents

the baseline results when the sample is split according to the sign of the XME shock. For positive shocks, the effect on leaving is twice as large; but it is not individually significant. The coefficient on growth is basically the same as the main result, and significant at the 5 % level. However, for negative shocks, the coefficients are also of similar magnitude and/or not significantly different.

I also repeat the timing analysis for positive and negative shocks, reported in Figure A2. The results point in the direction that we would expect, but they are not very precise or clear-cut. For positive XME shocks, leaving to WE has a positive coefficient for a number of quarters before and after the XME shock; although only the first lead is individually significant at the 5 % level. This is consistent with a gradual hiring process and a time to build. Firm growth is significantly negative in the three quarters before an XME shock. However, the coefficients continue to be negative afterwards as well. For negative XME shocks, switches to WE are only occur one quarter after the predicted shock; and this coefficient is significant. This is consistent with an interpretation that negative employment adjustments are sudden, and likely to be not anticipated. For negative XME shocks, we do not see an effect on small firm growth that moves in the opposite direction to XME changes; neither before nor after the negative shock. Rather, small firm growth moves in the same direction than XME shocks, in particular in Q3 and Q4 after the negative shock. This is consistent with a local demand channel being at play.

B.3. ADDITIONAL RESULTS ON HETEROGENEITY

TABLE A3
HETEROGENEOUS EFFECTS BY EDUCATION LEVEL

Dependent Variable:	Low skilled		Mid-skilled		High skilled	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth	(5) Leave to WE	(6) Relative Firm Growth
Δ Pred. XME/WAP	0.0851 (0.0887)		0.109** (0.0470)		0.118 (0.129)	
F. Δ Pred. XME/WAP		-0.307 (0.407)		-1.608*** (0.489)		-0.329 (0.762)
Municipality controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes	Yes	Yes
Dep. var. mean	16.12	-32.66	18.58	-24.55	17.38	-3.59
R ²	0.10	0.02	0.08	0.02	0.04	0.02
# distinct LLM	22,046	8,979	21,880	10,416	11,591	5,988
# municipalities	1,452	1,119	1,405	1,067	976	705
# periods	39	38	39	38	39	38
Observations	227,059	28,042	276,389	43,317	86,514	26,125

Low-skilled: up to primary education; mid-skilled: lower and upper secondary education; high-skilled: tertiary education.

Two-way clustered standard errors by municipality and trimester in parentheses. For other notes, see Table 12.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A4
HETEROGENEOUS EFFECTS - BY GENDER

Dependent Variable:	Male entrepreneurs		Female entrepreneurs	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth
Δ Pred. XME/WAP	0.0449 (0.0500)		0.255** (0.101)	
F. Δ Pred. XME/WAP		-1.236*** (0.249)		0.0537 (0.414)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	21.35	-19.41	11.38	-33.12
R ²	0.08	0.02	0.04	0.05
# distinct LLM	23,039	13,508	21,881	6,709
# municipalities	1,427	1,213	1,451	934
# periods	39	38	39	38
Observations	361,395	77,023	228,781	20,501

Two-way clustered standard errors by municipality and trimester in parentheses. For other notes, see Table 12.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3. ADDITIONAL RESULTS ON HETEROGENEITY

TABLE A5
HETEROGENEOUS EFFECTS - BY AGE GROUP

Dependent Variable:	Age 15-29		Age 30-39		Age 40-49		Age 50-64	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth	(5) Leave to WE	(6) Relative Firm Growth	(7) Leave to WE	(8) Relative Firm Growth
Δ Pred. XME/WAP	0.380* (0.205)		0.129** (0.0576)		-0.00364 (.)		0.0917 (0.0826)	
F. Δ Pred. XME/WAP		-0.226 (1.161)		-0.424 (0.587)		-1.512** (0.641)		-1.132** (0.532)
Municipality controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Dep. var. mean	29.59	-32.34	19.49	-23.15	15.27	-20.12	11.84	-20.58
R ²	0.10	0.04	0.06	0.03	0.05	0.03	0.05	0.02
# distinct LLM	15,106	3,897	19,709	8,020	20,088	8,932	19,279	7,842
# municipalities	1,321	727	1,401	1,003	1,396	1,032	1,393	995
# periods	39	38	39	38	39	38	39	38
Observations	79,125	8,316	162,604	26,993	178,899	33,721	165,058	27,932

Two-way clustered standard errors by municipality and trimester in parentheses. For other notes, see Table 12.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A6
HETEROGENEOUS EFFECTS - BY MIGRATION STATUS

Dependent Variable:	Migrant		Non-migrant	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth
Δ Pred. XME/WAP	0.166 (0.114)		0.0979*** (0.0380)	
F. Δ Pred. XME/WAP		-0.537 (0.718)		-1.185*** (0.216)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	17.61	-19.19	17.35	-23.20
R ²	0.06	0.03	0.09	0.03
# distinct LLM	13,702	6,047	24,521	13,492
# municipalities	1,066	689	1,475	1,241
# periods	39	38	39	38
Observations	133,573	24,420	456,603	73,104

Two-way clustered standard errors by municipality and trimester in parentheses. For other notes, see Table 12.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3. ADDITIONAL RESULTS ON HETEROGENEITY

TABLE A7
HETEROGENEOUS EFFECTS - ENTREPRENEURS IN URBAN VS RURAL LOCALITIES

Dependent Variable:	Rural localities (<15k inhabitants)		Urban localities (>15k inhabitants)	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth
Δ Pred. XME/WAP	0.163* (0.0906)		0.116*** (0.0352)	
F.Δ Pred. XME/WAP		0.183 (0.651)		-1.276*** (0.314)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	15.75	-30.27	18.20	-19.43
R ²	0.11	0.03	0.07	0.02
# distinct LLM	18,676	8,256	12,868	9,085
# municipalities	1,383	1,133	441	426
# periods	39	38	39	38
Observations	149,030	18,664	441,146	78,860

Two-way clustered standard errors by municipality and trimester in parentheses. For other notes, see Table 12.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.3. ADDITIONAL RESULTS ON HETEROGENEITY

TABLE A8
HETEROGENEOUS EFFECTS - FULL-TIME VS PART-TIME ENTREPRENEURS

Dependent Variable:	Full-time		Part-time	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth
Δ Pred. XME/WAP	0.0990** (0.0411)		0.117 (0.0759)	
F. Δ Pred. XME/WAP		-1.062*** (0.302)		-0.583 (0.683)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	17.23	-20.14	17.70	-33.18
R ²	0.09	0.02	0.06	0.04
# distinct LLM	23,755	13,677	20,226	5,578
# municipalities	1,464	1,215	1,431	936
# periods	39	38	39	38
Observations	400,391	82,688	184,097	14,089

Two-way clustered standard errors by municipality and trimester in parentheses. For other notes, see Table 12.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A9
DIFFERENTIAL EFFECTS FOR XME CONTRACTIONS AND EXPANSIONS

Dependent Variable:	Positive XME shock		Negative XME shock	
	(1) Leave to WE	(2) Relative Firm Growth	(3) Leave to WE	(4) Relative Firm Growth
Δ Pred. XME/WAP	0.200 (0.138)		0.0971 (0.140)	
F. Δ Pred. XME/WAP		-0.983** (0.422)		-0.890 (0.866)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	18.09	-21.94	17.49	-19.89
R ²	0.07	0.03	0.08	0.02
# distinct LLM	14,890	8,871	10,095	5,888
# municipalities	888	815	810	733
# periods	36	35	34	32
Observations	354,760	57,982	200,084	35,920


Two-way clustered standard errors by municipality and trimester in parentheses. For other notes, see Table 12.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$


B.3. ADDITIONAL RESULTS ON HETEROGENEITY

FIGURE A2
TIMING BY SIGN OF XME SHOCK

Positive XME shocks




(a) Occupational choice: leave to WE




(b) Entrepreneurial choice: relative firm growth

Negative XME shocks



(c) Occupational choice: leave to WE



(d) Entrepreneurial choice: relative firm growth

Source: Results from regression of model 2.3, using the full set of controls. Horizontal axis documents the number of quarters s that the dependent variable has been shifted. Each coefficient corresponds to a separate regression. Reported 95 % confidence intervals use two-way clustered standard errors by municipality and quarter.

B.4 Additional results on wages

This appendix documents two further results on wages. Table A10 repeats the analysis from Table 6, substituting wages of all workers with wages paid to new workers (those not observed in the same sector in the previous quarter) upon entry. The effects are more pronounced: the coefficient on small firm manufacturing wages has doubled, the one on small-firm wages not in manufacturing has increased by 1.5. However, the latter effect continues to be non-significant, and much smaller than the corresponding coefficient in manufacturing. What is new here is that profits of new entrants into entrepreneurship increase with an export manufacturing shock. This is consistent with a selection effect: given more attractive alternatives in the labour market, potential entrepreneurs only start relatively more profitable projects. Table A11 splits wage employment into formal (IMSS-affiliated) and informal. Around 80 percent of small firm jobs are informal, and conversely around 80 percent of large-firm jobs are formal. Given this imperfect overlap, the results are not very conclusive. There is no statistically significant effect on average formal or informal wages either in manufacturing or elsewhere. The magnitudes are smaller than for the small firm / large firm split in Table 6.

B.4. ADDITIONAL RESULTS ON WAGES

TABLE A10
EFFECTS ON STARTING INCOMES IN MANUFACTURING AND NON-MANUFACTURING
FIRMS

Sector:	Dependent Variable: $100 \times \Delta \log(\text{average entry income})$					
	Manufacturing			Non-manufacturing		
	(1) Profits	(2) Small firm wages	(3) Large firm wages	(4) Profits	(5) Small firm wages	(6) Large firm wages
Δ Pred. XME/WAPop	0.104 (0.438)	0.381*** (0.118)	-0.0593 (0.0908)	0.270*** (0.0974)	0.0781 (0.0722)	0.123 (0.0944)
Standard controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes	Yes	Yes
Municipality fixed effects						
Dep. var. mean	-0.87	-0.28	-0.26	-0.79	-0.54	-0.32
R ²	0.01	0.01	0.01	0.01	0.02	0.01
# municipalities	1,188	1,025	862	1,426	1,449	1,202
Observations	14,634	12,526	13,052	29,465	31,445	19,396

Standard errors in parentheses

This table reports regression of the log average income difference on predicted quarterly export manufacturing employment difference. Unit of observation is an unbalanced panel of Mexican municipalities (apart from D.F.) over 40 quarters from 2005 to 2010. All employment variables are normalised per initial 1000 working-age-population of the corresponding trimester and municipality, and employment and outcome variables are winsorised at the top and bottom 0.05%. All specifications include a constant, and per standard control for state, trimester, 2005 export employment. Individual controls are municipal average levels of education, experience, and share of men. Firm controls are municipal baseline industry shares and firm size distributions. Locality controls are locality size distribution within municipality. Observations are weighted by working-age population. Standard errors clustered by municipality.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A11
EFFECTS OF OF EXPORT SHOCKS ON WAGES - BY FORMAL/INFORMAL SPLIT

Sector:	Dependent Variable: $100 \times \Delta \log(\text{average income})$							
	Manufacturing				Non-manufacturing			
Subsample:	All		Entrant		All		Entrant	
Type of income:	(1) Informal wages	(2) Formal wages	(3) Informal wages	(4) Formal wages	(5) Informal wages	(6) Formal wages	(7) Informal wages	(8) Formal wages
Δ Pred. XME/WAPop	0.0714 (0.107)	0.00460 (0.0592)	0.136 (0.160)	-0.136 (0.0833)	0.0480 (0.0487)	0.0429 (0.0295)	0.0477 (0.0489)	0.0241 (0.0607)
Standard controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality fixed effects								
Dep. var. mean	-0.05	-0.20	-0.21	-0.37	-0.12	-0.23	-0.19	-0.42
R ²	0.01	0.01	0.01	0.01	0.02	0.01	0.02	0.01
# municipalities	1,203	963	1,097	814	1,473	1,351	1,473	1,176
Observations	20,858	19,543	15,082	12,998	36,277	30,576	36,126	20,958

Standard errors clustered by municipality in parentheses. For other notes see Table A10.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5 Robustness checks

This appendix documents a number of robustness checks on the main results. In Table A12 I repeat the specifications from Table 8 and from columns (3) and (4) of Table 9, but I restrict the sample to entrepreneurs with at least one paid worker. This is the same restriction which is applied in order to create the relative firm growth variable. These entrepreneurs are just as likely to take a wage job than are the self-employed with no workers, but they are much less likely to leave employment altogether. An XME shock has the same effect on the transition to wage work than it does for the full sample, but the effect on leaving to any occupation becomes zero in column (2) with the full vector of controls. For this subsample of employer-entrepreneurs, the effect on absolute firm growth is large and significant. This suggests that the null effect for the whole sample is driven by a failure of solo entrepreneurs to reduce their firm size, which is likely due to the censoring of firm size at zero paid workers.

Table A13 reports OLS results using actual instead of predicted export manufacturing employment changes as the explanatory variable. The coefficients on firm growth are now positive, but small and not significant. The coefficients on leaving are still positive, but reduced in size. This suggests a downward bias of OLS for leaving, and an upward bias for growth. We would expect a bias in these directions if small and large (export manufacturing) firm growth are driven by a common, unobservable variable. Then, small firms and large firms grow together. Episodes of export manufacturing growth also provide less incentives for taking up wage jobs if the relative earnings are less responsive.

In Table A14, I report results from 2SLS instrumental variables regression. Predicted XME changes serve as an instrument for actual XME changes. I restrict the sample to non-manufacturing entrepreneurs. Remember that for this sub-sample, alternative mechanisms are less of a concern. Predicted XME changes are a strong instrument for actual XME changes in a sample of individual entrepreneurs. This mirrors the result I obtain in Table 5 at the local labour market level. The Kleibergen-Papp heteroscedasticity-robust test statistic is well over the the Stock and Yogo (2005) critical values for a 10% maximal IV size. The second stage regression coefficients are about 1.7 times as large as the associated reduced-form results in table 12. This is to be expected, given that the first-stage coefficient is around 0.585 – the 2SLS coefficient equals the reduced form divided by the first stage

coefficient. Caution needs to be taken when interpreting the IV results. In particular, they should not be interpreted as an ‘average treatment effect’ estimate for the magnitude of the occupational choice mechanism on firm size. First, even though caution has been taken to remove direct mechanisms of XME changes on firm profits which work in the same direction as the occupational choice channel, some direct mechanisms may still have a quantitative influence, for example if they go in the opposite direction. Second, as I allude to in the conclusion, a simple linear extrapolation will likely not be a good approximation to the magnitude of the effects.

I finally report robustness checks to alternative definitions of XME changes. In Table A15, I apply the commuting-zone aggregation from section 2.5.3 and redefine local labour markets as commuting zones, instead of municipalities. The coefficients on leaving are similar to the main results. The coefficient on relative firm growth is reduced in magnitude and now only significant at the 10 % level. On the other hand, the coefficient on absolute firm growth is now large and also significant. Next, I disaggregate my measure of XME changes, and change the composition of the industries that are included in my definition of export manufacturing. One potential concern with the identification strategy could be that production of exports is concentrated in a small number of local labour markets. Then, a local market level shock transmits to aggregate Mexican exports to the US. The ideal check to alleviate this concern would be to predict export employment in municipality m with exports from all other municipalities apart from m . However, municipality-level export data are not available. Instead, I check whether results are driven by exports and employment in any particular industry (defined by a NAICS sub-sector). In Tables A16-A25, I report results obtained with predicted XME changes, leaving out each of the 10 constituent export industries in turn. None of these changes affects the sign or significance of any result. Most of these leave-one-out results are very similar in magnitude. When the magnitude is different – such as when leaving out large industries such as transportation equipment or electrical and appliances – it appears to be even larger. The auto industry is one of the more concentrated export industries. If aggregate exports are not immune to local production shocks, this seems to attenuate rather than magnify the results (see also the OLS table).

TABLE A12
 ROBUSTNESS: EFFECTS ON LEAVING ENTREPRENEURSHIP FOR ENTREPRENEURS
 WITH PAID WORKERS

Dependent Variable:	(1) Leave	(2) Leave	(3) Leave to WE	(4) Leave to WE	(5) Absolute Firm Growth	(6) Absolute Firm Growth
Δ Pred. XME/WAP	0.0723* (0.0420)	0.000563 (0.0734)	0.159*** (0.0471)	0.0966*** (0.00733)		
F. Δ Pred. XME/WAP					-0.981** (0.450)	-1.627*** (0.486)
Municipality controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	No	Yes	No	Yes	No	Yes
Firm controls	No	Yes	No	Yes	No	Yes
Locality controls	No	Yes	No	Yes	No	Yes
Dep. var. mean	25.69	25.90	17.98	18.20	-34.38	-42.54
R ²	0.01	0.07	0.01	0.07	0.00	0.03
# distinct LLM	18,475	17,005	17,914	16,419	16,274	14,759
# municipalities	1,340	1,320	1,325	1,307	1,272	1,255
# periods	39	39	39	39	38	38
Observations	182,645	133,993	165,530	121,512	132,998	97,524

Standard errors in parentheses

This table reports regression of outcomes on predicted quarterly export manufacturing employment difference in the local municipality. Unit of observation are entrepreneurs with at least 1 and at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. Municipal export employment variables are normalised per 1000 working-age-population of the corresponding month and municipality, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by municipality and trimester. Municipality controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A13
ROBUSTNESS: OLS REGRESSION

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ XME/WAP	0.0467 (0.0290)	0.0496 (0.0306)		
F. Δ XME/WAP			0.101 (0.265)	0.0830 (0.137)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.61	17.39	-22.27	3.82
R ²	0.07	0.08	0.02	0.05
# distinct LLM	20,397	20,024	12,976	20,166
# municipalities	1,489	1,476	1,274	1,471
# periods	38	38	39	39
Observations	707,913	573,675	99,540	485,832

Standard errors in parentheses

This table reports OLS regression of outcomes on observed quarterly export manufacturing employment difference in the local municipality. Unit of observation are entrepreneurs with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. Municipal export employment variables are normalised per 1000 working-age-population of the corresponding month and municipality, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by municipality and trimester. Municipality controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5. ROBUSTNESS CHECKS

TABLE A14
ROBUSTNESS: IV-2SLS REGRESSION RESULTS

Dependent Variable:	IV Second Stage results			
	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ XME/WAP	0.323*** (0.107)	0.243*** (0.0714)		
F. Δ XME/WAP			-1.628*** (0.367)	-0.319 (0.350)
Standard controls	Y	Y	Y	Y
Individual controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Locality controls	Y	Y	Y	Y
Dep. var. mean	33.52	18.02	-22.74	4.191
K-P Wald F-statistic	26.42	26.03	24.49	24.17
IV First Stage results				
Δ Pred. XME/WAP	0.585*** (0.114)	0.590*** (0.116)		
F. Δ Pred. XME/WAP			0.585*** (0.118)	0.586*** (0.119)
R ²				
# distinct LLM	23,729	23,099	13,044	22,543
# municipalities	1,481	1,466	1,205	1,442
# periods	38	38	38	38
Observations	602,838	492,126	82,824	404,120

Source: This table reports IV-2SLS regression of outcomes on quarterly export manufacturing employment difference in the local municipality. Unit of observation are entrepreneurs not in manufacturing with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. Municipal export employment variables are normalised per 1000 working-age-population of the corresponding month and municipality, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by municipality and quarter. Municipality controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm controls include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used. *Stock-Yogo (2005) Critical Values:* 16.38 (at 10% maximal IV size), 8.96 (at 15 % maximal IV size), 5.53 (at 20% maximal IV size).

TABLE A15
ROBUSTNESS: CZ-LEVEL REGRESSION

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.138** (0.0650)	0.134*** (0.0458)		
F. Δ Pred. XME/WAP			-0.656* (0.364)	-0.324* (0.181)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.63	17.68	-21.96	4.27
R ²	0.07	0.08	0.02	0.04
# distinct LLM	18,187	17,815	10,596	17,062
# CZ	1,230	1,217	1,008	1,197
# periods	39	39	38	38
Observations	746,596	605,922	99,622	487,264

Standard errors in parentheses

This table reports regression of outcomes on predicted quarterly export manufacturing employment difference in the commuting zone (CZ). Unit of observation are entrepreneurs with at most 9 employees from 40 quarterly ENOE waves 2005-2014 who are observed in the following period. CZ export employment variables are normalised per 1000 working-age-population of the corresponding month and CZ, and employment and outcome variables are winsorised at the top and bottom 0.01%. Two-way clustered standard errors by CZ and trimester. CZ controls include wave dummies, state dummies, and 2005 export manufacturing employment. Individual controls include gender, education level, and experience quadratic. Firm control include firm size, industry, and profit quintiles. Locality controls include locality size category. All specifications include a constant, and ENOE sampling weights are used.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5. ROBUSTNESS CHECKS

TABLE A16

ROBUSTNESS: LEAVING OUT TEXTILE PRODUCTS MANUFACTURING (NAICS 314)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.109* (0.0621)	0.111*** (0.0374)		
F. Δ Pred. XME/WAP			-0.973*** (0.233)	-0.207 (0.206)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	25,378	24,843	14,700	23,763
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out textile product manufacturing (NAICS sub-sector 314). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A17

ROBUSTNESS: LEAVING OUT APPAREL MANUFACTURING (NAICS 315)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.113* (0.0584)	0.109*** (0.0373)		
F. Δ Pred. XME/WAP			-0.961*** (0.255)	-0.231 (0.213)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	23,962	23,491	14,245	22,480
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out apparel manufacturing (NAICS sub-sector 315). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5. ROBUSTNESS CHECKS

TABLE A18
ROBUSTNESS: LEAVING OUT LEATHER, HIDE AND ALLIED PRODUCTS
MANUFACTURING (NAICS 316)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.102* (0.0538)	0.104*** (0.0365)		
F. Δ Pred. XME/WAP			-0.981*** (0.241)	-0.214 (0.204)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	25,360	24,827	14,696	23,749
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out leather, allied and hide manufacturing (NAICS sub-sector 316). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A19
ROBUSTNESS: LEAVING OUT METAL PRODUCTS MANUFACTURING (NAICS 332)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.106* (0.0617)	0.106*** (0.0390)		
F. Δ Pred. XME/WAP			-0.991*** (0.236)	-0.217 (0.212)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	23,576	23,118	14,000	22,123
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out metal products manufacturing (NAICS sub-sector 332). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5. ROBUSTNESS CHECKS

TABLE A20
ROBUSTNESS: LEAVING OUT MACHINERY AND EQUIPMENT MANUFACTURING
(NAICS 333)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.0992* (0.0603)	0.108*** (0.0359)		
F. Δ Pred. XME/WAP			-1.010*** (0.228)	-0.225 (0.211)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	25,483	24,948	14,746	23,864
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out machinery and equipment manufacturing (NAICS sub-sector 333). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A21
ROBUSTNESS: LEAVING OUT COMPUTER EQUIPMENT MANUFACTURING (NAICS
334)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.109* (0.0618)	0.111*** (0.0384)		
F. Δ Pred. XME/WAP			-1.162*** (0.246)	-0.226 (0.227)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	25,520	24,983	14,758	23,898
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out computer equipment manufacturing (NAICS sub-sector 334). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5. ROBUSTNESS CHECKS

TABLE A22
ROBUSTNESS: LEAVING OUT ELECTRIC APPLIANCES, ACCESORIES, ETC.
MANUFACTURING (NAICS 335)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.177** (0.0758)	0.150*** (0.0434)		
F. Δ Pred. XME/WAP			-1.121*** (0.380)	-0.128 (0.296)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	25,391	24,856	14,694	23,777
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out electric appliances, accessories, etc. manufacturing (NAICS sub-sector 335). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A23
ROBUSTNESS: LEAVING OUT TRANSPORT EQUIPMENT MANUFACTURING (NAICS
336)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.0664 (0.0949)	0.127* (0.0661)		
F. Δ Pred. XME/WAP			-1.540*** (0.347)	-0.704*** (0.250)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	25,410	24,884	14,731	23,801
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out transport equipment manufacturing (NAICS sub-sector 336). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.5. ROBUSTNESS CHECKS

TABLE A24
ROBUSTNESS: LEAVING OUT FURNITURE, MATTRESS AND BLINDS MANUFACTURING
(NAICS 337)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.113** (0.0564)	0.118*** (0.0347)		
F. Δ Pred. XME/WAP			-0.972*** (0.232)	-0.216 (0.206)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	24,719	24,214	14,479	23,181
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out furniture, mattresses and blinds manufacturing (NAICS sub-sector 337). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE A25
ROBUSTNESS: LEAVING OUT OTHER MANUFACTURING INDUSTRIES (NAICS 339)

Dependent Variable:	(1) Leave	(2) Leave to WE	(3) Relative Firm Growth	(4) Absolute Firm Growth
Δ Pred. XME/WAP	0.108* (0.0582)	0.112*** (0.0375)		
F. Δ Pred. XME/WAP			-1.002*** (0.207)	-0.220 (0.205)
Municipality controls	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Locality controls	Yes	Yes	Yes	Yes
Dep. var. mean	33.57	17.41	-22.17	3.88
R ²	0.07	0.08	0.02	0.05
# distinct LLM	25,283	24,749	14,621	23,670
# municipalities	1,490	1,477	1,255	1,456
# periods	39	39	38	38
Observations	727,539	590,176	97,524	475,005

In this table, predicted local export employment is constructed leaving out other manufacturing industries (NAICS sub-sector 339). Two-way clustered standard errors by municipality and quarter in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 3

Occupational Choice under Search Frictions: A Roy Search Model

3.1 Introduction

In this chapter, I consider the empirical content of a joint model of occupational choice and earnings. The basic structure is built around the Roy (1951) model where occupational choice is driven by comparative advantage. The second building block are search frictions, which drive a distinction between desired occupational choice based on the relative *attractiveness* of a sector (the comparative advantage), and observed occupational *allocations*. The allocations are not only driven by comparative advantage, but also by the exogenous availability of a particular occupation. I call this model the Roy Search Model (RSM).

Roy search models are increasingly popular building blocks of empirical analysis of occupational choice and labour supply. While the foundational work on structural life-cycle models such as Keane and Wolpin (1997) treats occupational choice as a pure discrete dynamic choice problem, recent lit-

erature (e.g. Blundell, Costa Dias, Meghir, and Shaw (2016), Adda, Dustmann, and Stevens (2017), Low, Meghir, Pistaferri, and Voena (2017)) combines aspects of discrete occupational choices that follow a Roy-model structure with elements of job search models. The model in Chapter 4 of this thesis is another such example of occupational choice under search frictions. This chapter considers the identification of this building block.

I propose a set of assumptions under which the RSM is identified from panel data on earnings and occupations. My identification strategy allows for selection into wage-employment to be a function of three unobservable elements: persistent differences in the earnings potential across individuals; idiosyncratic draws from the productivity and wage offer distribution; and constraints to choice in the form of search frictions. None of the assumptions require a large support of covariates. This mirrors the data availability from most labour force surveys, including from developing countries. These surveys typically exhibit a modest cross-sectional sample size and geographical variation, and a short individual panel duration. I distinguish which parts of the model are identified only from variation in covariates, which parts rely on economic theory about the form of the choice mechanism, and which parts require distributional assumption for identification.

I illustrate this approach with an application using labour force panel data from urban Colombia. I estimate a fully parametric version of the RSM. I focus on the occupational choice between wage-employment and self-employment. My findings from the estimation of the joint model of selection and earnings are threefold. First, I find a considerable degree of search frictions, expressed by the estimated wage offer probability. Second, I find a correlation of unobserved productivity between wage-employment

and self-employment in the order of 0.5. Third, I find that wages are fairly concentrated and persistent over time, whereas entrepreneurial profits are much more volatile and less persistent. The RSM parameter estimates are robust to alternative assumptions about the selection mechanism.

Finally, the estimated selection term of the model, which depends on the value of job opportunities in wage employment, relative to being self-employed, serves as a natural measure for the attractiveness of wage-employment for Colombian entrepreneurs. This allows me to carry out an alternative test of the model in Chapter 1 in a setting where cross-market shifters of wage employment opportunities are not available. I find some support for the two key predictions of the model. Relatively more attractive wage opportunities are associated with a higher probability that entrepreneurs take up a wage job. They are negatively correlated with the probability that entrepreneurs grow their firm. The strength of these associations is sensitive to some assumptions about the structure of the RSM.

3.1.1 Related literature

The empirical methods that I develop in this chapter are particularly applicable to the analysis of labour markets in developing countries. This literature was reviewed in Chapter 1. The RSM embeds both the ‘competitive’ and the ‘segmented’ view of labour markets in developing countries. These views are embodied, respectively, by the attractiveness and the availability elements of occupational choice. Methodologically, this chapter contributes to the literature on identification of the Roy model, and of selection models more generally. Here I review the contributions of that literature that are most relevant for this chapter.

Identification of the Roy Model. Methodologically, this chapter is situated within the literature on identification of the Roy (1951) model and related selection models. Heckman and Honoré (1990) establish the conditions under which the Roy model is non-parametrically identified. Subsequent work deals with extensions of the ‘pure’ Roy model (e.g. Heckman and Vytlačil (2007), Eisenhauer, Heckman, and Vytlačil (2015), Adão (2015)). French and Taber (2011) provide a review.¹ The Roy model and its extensions are based on the assumptions that selection is driven by the particular value of the job offer. In the context of the search model, this seems to be the most appropriate assumption to make. An alternative line of literature takes as a starting point the Chamberlain (1984) decomposition of the fixed effects model, and builds a correlated effects / coefficient model around this (Card (1996), Lemieux (1998), Suri (2011)). However, an important limitation to the applicability of that methodology is that only permanent and not transitory components of payoffs are allowed to influence switching behaviour. This assumption is clearly too restrictive in the context of occupational transitions in the short run.

Semiparametric identification of selection models A second strand of related methodological literature is concerned with the semiparametric identification and estimation of sample selection models (Powell (1987), Ahn and Powell (1993), Powell (2001), Honoré and Powell (2005), Aradillas-Lopez, Honoré, and Powell (2007)). I apply a pairwise differencing approach – common to that literature – to the particular case of the Roy Search Model. This approach identifies the coefficients of the wage equation under strong exclusion restrictions. I also discuss an alternative argument that dispenses

¹ Keane, Todd, and Wolpin (2011) review the identification of discrete choice dynamic programming models. They also discuss the difference that has evolved in the literature between ‘search models’ and ‘occupational choice models’.

with the need for exclusion restrictions by imposing the choice structure implied by the simple canonical search model.

The chapter is structured as follows. In section 3.2 I present the empirical framework. I motivate the RSM by deriving the requirements for a model-based test of Chapter 1. I then present the assumptions, and discuss identification in four steps. I describe the data in section 3.3 and present and discuss the empirical results in section 3.4. Finally, I briefly conclude in section 3.6.

3.2 Empirical Framework

This empirical framework centers on identifying the elements of a joint distribution of earnings in two sectors, where some elements will be observable and some unobservable to the researcher. My identification strategy allows for selection into wage employment to be a function of three unobservable elements: persistent differences in the earnings potential across individuals, idiosyncratic draws from the productivity and wage offer distribution, and constraints to choice in the form of search frictions. I extend a basic Roy model to a two-period panel setting and allow for search frictions and exogenous job separations, like in the canonical job search model (Mortensen (1970)).

In the context of the previous chapter of this thesis, the empirical strategy in this chapter takes a very different approach. Instead of exploiting cross-market variation in wage employment opportunities like in Chapter 2, I directly estimate the incidence of attractive wage-employment opportunities for each individual, based on individual characteristics that shift the mean of earnings distributions. I propose a model-based measure which quan-

tifies the attractiveness of wage-employment over self-employment, based on earnings differentials across the two sector. Once obtained, I correlate these measures with the decision to leave entrepreneurship, and with the decision to grow the firm, using linear regression.

3.2.1 Expressing the attractiveness of wage-employment

I start by motivating the empirical approach with the theoretical model in Chapter 1. The model suggests a particular form for the attractiveness of wage-employment opportunities, given by: $\Pr(w > \pi) = (1 - F(\pi))$, the probability that a wage offer w exceeds the reservation wage. In the static model, the reservation wage is given by the current entrepreneurial profits, π . The availability of wage employment, on the other hand, is determined by the degree of search frictions parameterised by λ .

Recall that the theoretical model is expressed for an individual entrepreneur. Both the profit function and the wage offer distribution are individual-specific. I acknowledge this by conditioning the attractiveness of wage employment on an information set Ψ_i :

$$\Pr(w_i > \pi_i | \Psi_i) = 1 - F_w(\pi_i | \Psi_i) \quad (3.1)$$

The information set generally contains observable characteristics, but potentially also other information about the self-employed individual, for example unobservable components of their current operational profit. On average, better entrepreneurs will also make for more skilled wage workers. Profits and wages should be expected to be correlated, most likely in a positive way.

In absence of sufficiently rich cross-market data which allows to conduct a test like in Chapter 2, we could instead directly estimate the attractiveness of wage-employment opportunities across entrepreneurs. In order to accurately estimate the measure (3.1) we need to be able to identify the components of the distribution of wages and entrepreneurial profits. To this I now turn.

3.2.2 The selection problem

To set up the problem, suppose that potential wages and entrepreneurial profits are determined according to the following process:

Assumption 3.1 *Joint earnings distribution*

The log of potential wages w_{it} , and entrepreneurial profits, π_{it} , of individual i in time period t are given by

$$\ln w_{it} = \mathbf{x}'_{it}\boldsymbol{\beta} + \epsilon_{it}; \quad (3.2)$$

$$\ln \pi_{it} = \mathbf{q}'_{it}\boldsymbol{\gamma} + \nu_{it}, \quad (3.3)$$

where $(\mathbf{x}_{it}, \mathbf{q}_{it})$ is independent of $(\epsilon_{it}, \nu_{it})$.

Log earnings (wages and profits) contain two linear components: an observable $(\mathbf{x}'_{it}\boldsymbol{\beta}$ and $\mathbf{q}'_{it}\boldsymbol{\gamma}$, respectively) and an unobservable part $(\epsilon_{it}$ and ν_{it} , respectively), which are independent from each other. Independence of this form implies homoskedasticity of the errors across individuals and time.² I make no assumption on the mutual relation of the unobservable error terms.

They may follow a joint stochastic distribution of arbitrary form, which I

² That is, while observable variables may shift the mean of the wage and profit distributions when they are included among the x_i and q_i , the variance and higher moments of wages are the same for everyone.

denote F in general. I will include subscripts which denote the marginal or conditional distributions to resolve ambiguity.

In order to identify (3.1), we need to identify both the observable and the unobservable part. For the observable part, this means correctly identifying the predictive coefficient β from the earnings function of the wage-employed, which is possible under weak conditions.³

However, we observe incomes only for those individuals actually employed in the respective sector $S_{it} \in \{0, 1\}$. Wage offers are observed only when an individual is wage-employed ($S_{it} = 1$); and vice versa for profits.

The main empirical challenge arises from a correlation between ϵ_{it} and ν_{it} . This poses two concerns: first, the residual of entrepreneurial profits, ν_{it} , may contain some information about the counterfactual ϵ_{it} and therefore the mean of the individual-specific wage distribution. Ignoring this information causes us to underestimate wage employment attractiveness for high-ability individuals, and to overestimate for low-ability individuals. Secondly, a correlation on correlated unobservables could be a source of selection bias for estimates of β and γ . Selection bias implies that the self-employed and the wage-employed are fundamentally different in terms of their unobservables. This biases wage predictions for individuals that are self-employed with coefficients that we would have estimated on data for the wage-employed.

³ In other words, we do not need to know the causal effect of a wage predictor on wages. Take the example of schooling. We only need to know how schooling correlates with wages. We do not need to decompose the Mincerian schooling coefficient into the causal returns to education, and the correlation between intrinsic ability and schooling. In fact, picking up this correlation will give our model a higher predictive power.

3.2.3 A Roy Search Model

In this section, I propose a Roy Search Model (RSM) based on an extension of the Roy (1951) model that incorporates search frictions. Here I state and comment on the assumptions of the model that are required for identification. I introduce the assumptions on the data-generating process for a general number of time periods. In the next section, I will assume that we only observe data from two adjacent periods.

Selection mechanism

As a first step, we need to make assumptions about the nature of the selection mechanism. The canonical search model implies a particular form of the transition probability between the current labour market status (here: self-employment) and an occupation in wage-employment: such a transition occurs when an offer is made *and* it is better than the reservation wage. For simplicity and in analogy to standard search models, I assume that the reverse transition between wage-employment and self-employment is given by an exogenous separation probability. This implies that self-employment is a free-entry sector, similar to unemployment in standard applications of the search model.⁴

Assumption 3.2 *Selection mechanism*

1a. TRANSITION OUT OF SELF-EMPLOYMENT The probability of individual i switching from self-employment to wage-employment between periods t and $t + 1$ is given by the joint probability of receiv-

⁴ I make this assumption for the convenience of working with a two-sector model. The life-cycle model in Chapter 4 relaxes the free-entry assumption in self-employment. It includes non-employment as a third sector, which then becomes the free-entry sector.

ing an offer, and of the offer being attractive:

$$\Pr(S_{i,t+1} = 1 | S_{it} = 0) = \lambda \cdot \left(1 - F_{w_{it}}(\pi_{it} | \Psi_i)\right). \quad (3.4)$$

1b. TRANSITION OUT OF WAGE-EMPLOYMENT The probability of switching between wage-employment and self-employment is given by the exogenous separation probability:

$$\Pr(S_{i,t+1} = 0 | S_{it} = 1) = \delta. \quad (3.5)$$

Two further comments are in order. First, note that there is some ambiguity in the model about what the right reservation wage is. In particular, it could be either equal to profits in the base period, π_{it} ; or to expected profits in the following period ($E(\pi_{i,t+1} | \pi_{it})$); or any convex combination of the two. The issue here is that we observe earnings only at discrete points in time (with the data in this chapter, every year) whereas transitions may occur at any point in the interim period. I resolve this ambiguity as following. As a baseline case, I assume that the reservation wage equals current profits, and that agents compare it to current counterfactual wages w_{it} . Following the main analysis, I then explore robustness to the alternative formulation where expected profits $E(\pi_{i,t+1} | \pi_{it})$ are compared to distribution of future, conditional wages ($w_{i,t+1} | \pi_{it}$). By a continuity argument, any convex combination reflecting a mid-point should lie in between these two cases.

Second, while I generally let wage offers differ by observable characteristics x_{it} and q_{it} , I assume that λ and δ are the same for every individual. In principle, one could parameterise these functions. However, this requires richer data, in particular if one wishes to nonparametrically identify the respective

functions.

Distributional assumptions

From Assumption 3.2, we see that the selection mechanism is sufficiently described by the offer probability and a *conditional* wage offer distribution. As I will show, these components are identified without imposing any restrictions on the joint earnings distribution F . To identify F , however, we need to impose further distributional assumptions. In Assumption 3.3, I restrict the time-series properties of the earnings process, in a quite general and standard way. In Assumption 3.4, I make a much stronger assumption on the parametric form of the distribution of productivity shocks.

Assumption 3.3 *Markov stationary earnings distribution*

The cumulative distribution function F_{ϵ_t, ν_t} obeys the properties

$$\begin{aligned} F_{\epsilon_t, \nu_t} &= \Pr(\epsilon_t < a, \nu_t < b \mid \epsilon_{t-1} = c, \nu_{t-1} = d, \epsilon_{t-2} = e, \nu_{t-2} = f, \dots) \\ &= \Pr(\epsilon_t < a, \nu_t < b \mid \epsilon_{t-1} = c, \nu_{t-1} = d) \end{aligned}$$

and

$$F_{\epsilon_t, \nu_t} = F_{\epsilon_s, \nu_s} \quad \forall \quad t, s.$$

This is equivalent to an AR(1) formulation which takes the form:

$$\epsilon_t = \rho_\epsilon \epsilon_{t-1} + \sqrt{1 - \rho_\epsilon^2} \tilde{\epsilon}_t \quad -1 < \rho_\epsilon < 1 \quad (3.6)$$

$$\nu_t = \rho_\nu \nu_{t-1} + \sqrt{1 - \rho_\nu^2} \tilde{\nu}_t \quad -1 < \rho_\nu < 1 \quad (3.7)$$

where $(\epsilon_{t-1}, \nu_{t-1})$ is independent of $(\tilde{\epsilon}_t, \tilde{\nu}_t)$; where the distribution the random innovations $F_{\tilde{\epsilon}_t, \tilde{\nu}_t}$ is stationary; and where initial values $(\epsilon_0 = \tilde{\epsilon}_0, \nu_0 =$

$\tilde{\nu}_o$) are drawn from this stationary distribution.

Assumption 3.4 *Joint normality*

$$\begin{pmatrix} \tilde{\epsilon}_t \\ \tilde{\nu}_t \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\epsilon^2 & \rho \sigma_\epsilon \sigma_\nu \\ \rho \sigma_\epsilon \sigma_\nu & \sigma_\nu^2 \end{pmatrix} \right)$$

I denote with ρ the cross-correlation between wage and profit offers (i.e. the correlation between sector-specific ‘abilities’); and with ρ_ϵ and ρ_ν the respective autocorrelations.

3.2.4 Identification

Suppose that we observe data from two adjacent periods, which I denote period 1 and period 2, without loss of generality. This ensures the widest possible applicability of this methodology, including with data from developing countries, where long panel data sets are rare.⁵ In this section, I intuitively discuss identification of the conditional mean parameters β and γ and the marginal earnings distribution $F_{\epsilon_1|\nu_1}$ under Assumptions 3.1 and 3.2. I also discuss identification of the joint earnings distribution $F(\epsilon_1, \nu_1, \epsilon_2, \nu_2)$ under the additional Assumptions 3.3 and 3.4. My line of argument, at least initially, follows the structure in French and Taber (2011) and builds upon their exposition, as well as upon the seminal work by Heckman and Honoré (1990).

The identification argument proceeds in four steps. Step 1 shows which objects can be identified from the data on transitions. Step 2 discusses how to identify the conditional mean parameters from observed data on earnings,

⁵ For example, the data used in this section (see section 3.3) is in principle available for up to four periods. Changes to the sampling frame, the addition and removal of particular cities, as well as attrition, however, imply that in practice, I use a two-period balanced panel.

holding the degree of selection fixed. In step 3, I identify the conditional earnings distributions. Step 4 then deals with recovering the full distribution of unobservable earnings components.

Step 1: Identification of the conditional choice probabilities

The first step involves establishing what objects can be identified directly by the observed data on choices. From Assumption 3.1 about the earnings process and Assumption 3.2 about the selection mechanism, the conditional sector choice probabilities can be written as

$$\begin{aligned}
& \Pr(S_{i2} = 1 \mid S_{i1} = 0, \mathbf{q}_{i1}, \mathbf{x}_{i1}, \pi_{i1}) \\
&= \lambda \cdot \Pr(w_{i1} > \pi_{i1} \mid \mathbf{q}_{i1}, \mathbf{x}_{i1}, \pi_{i1}) \\
&= \lambda \cdot \left(1 - \Pr(\epsilon_{i1} < \mathbf{x}'_{i1}\boldsymbol{\beta} - \pi_{i1} \mid \nu_{i1})\right) \\
&= \lambda \cdot \left(1 - F_{\epsilon_1|\nu_1}(\mathbf{x}'_{i1}\boldsymbol{\beta} - \mathbf{q}'_{i1}\boldsymbol{\gamma} - \nu_{i1})\right) \\
&\equiv g(\mathbf{x}_{i1}, \mathbf{q}_{i1}, \nu_{i1})
\end{aligned} \tag{3.8}$$

and

$$\Pr(S_{i2} = 1 \mid S_{i1} = 0, \mathbf{q}_{i1}, \mathbf{x}_{i1}, \pi_{i1}) = \delta, \tag{3.9}$$

where $F_{\epsilon_1|\nu_1}$ denotes the distribution of ϵ_1 conditional on ν_1 . The separation rate δ is immediately identified by the conditional choice probability (3.9). This conditional choice probability $g(\mathbf{x}_{it}, \mathbf{q}_{it}, \nu_{it})$ is a composite function of the search process, the conditional earnings distribution, and the conditional means parameters. Even if $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ were known at this stage, equation (3.9) only identifies the conditional distribution up to a monotonic transformation,

given by a multiplication with the constant λ .⁶

Step 2: Identification of the conditional earnings means

In the second step, we identify the vectors β and γ that parameterise the conditional mean. This is a form of the classical selection problem. Since selection into wage-employment occurs on the basis of the specific wage offer ϵ_{it} and the entrepreneurial ability ν_{it} , the error term of the conditional expectation function $E(w_{it}|\mathbf{x}_{it}, S_{it} = 1)$ is not mean-zero. It is also (negatively) correlated with observables $\mathbf{x}_{it}$. Those who select into wage-employment with low observable skills $\mathbf{x}_{it}$ must have a high level of unobservable skills ϵ_{it} . Conversely, if the wage-employed are a selected sub-population, then so are the self-employed. Since this logic also applies to period 1, we start our observation window already with a non-randomly selected distribution of sector of employment.

The nature of the selection problem differs from the standard Roy model where selection is only based on comparative advantage. The search friction ensures that even some of the (relatively) worst entrepreneurs are still self-employed, because they were unlucky not to receive a job offer. This makes it hard to write down a closed-form solution of a selection correction term, even under distributional assumption.

The solution to the selection problem relies on the insight that the degree of selection bias under the selection mechanism of Assumption 3.2 is governed by the conditional choice probability.⁷ If we are able to pair two individuals

⁶ The caveat of identification up to a monotonic transformation, at this stage, also applies in the standard cross-sectional Roy model without search frictions (French and Taber (2011), p. 553).

⁷ Assumption 3.2 is somewhat stronger than what would be necessary for this part of the identification argument. At this stage, we do not need to invoke the particular functional form of the conditional choice probability. We only require that it is a general function $g(\mathbf{x}_{i1}, \mathbf{q}_{i1}, \nu_{i1})$.

in the data with the same conditional choice probability, then we can condition on the degree of selection bias. I discuss two types of comparisons that identify parameters of the wage equation under different conditions on covariates. First, we can compare two different individuals who are identical entrepreneurs and who have the same choice probability conditional on being identical entrepreneurs. Second, we can exploit the panel component of the data and compare the same individual at two different points in time, if selection has not changed in the interim.

Pairwise comparisons across individuals

Imagine that we observe two individuals i and j who are ‘identical’ entrepreneurs in our first period of observation, in the sense that:

$$\mathbf{q}_{i1} = \mathbf{q}_{j1} \quad \text{and} \quad \pi_{i1} = \pi_{j1}. \quad (3.10)$$

From this it follows that $\nu_{i1} = \nu_{j1}$. Individuals with the same observable covariates and the same outcome must have the same unobservable residual. Suppose further that i and j have exactly the same probability of becoming wage-employed in the next period. Their observable skills in wage-employment may differ ($\mathbf{x}_{i1} \neq \mathbf{x}_{j1}$) as long as the following holds:

$$g(\mathbf{x}_{i1}, \mathbf{q}_{i1}, \nu_{i1}) = g(\mathbf{x}_{j1}, \mathbf{q}_{j1}, \nu_{j1}) \quad (3.11)$$

Under (3.10) and (3.11), the selection bias terms for i and j are identical. Taking the difference between i and j in the second period (if and when both are observed in wage-employment) eliminates the bias term, without a need for the econometrician to estimate or even to specify its exact form.

Therefore, the regression

$$E(w_{i2} - w_{j2} \mid \mathbf{x}_{i2}, \mathbf{x}_{j2}) = (\mathbf{x}_{i2} - \mathbf{x}_{j2})' \boldsymbol{\beta} + E(\epsilon_{i2} - \epsilon_{j2} \mid \mathbf{x}_{i,2}, \mathbf{x}_{j,2}) \quad (3.12)$$

identifies the wage equation parameters $\boldsymbol{\beta}$ for each dimension k of $\mathbf{x}_{i2} = [x_{i2}^1, \dots, x_{i2}^k, \dots, x_{i2}^n]$ where $x_{i2}^k \neq x_{j2}^k$; subject to conditions (3.10) and (3.11).

This example with two observations neatly illustrates the power and the limitations of this approach to identification ‘holding selection constant’. Given our assumption of linear conditional means for the earnings functions, the support requirements are rather weak. At a minimum, to identify a single element $\beta^k \in \boldsymbol{\beta}$, it would be sufficient to observe just two individuals with $x_{i2}^k \neq x_{j2}^k$, without any further condition on the support of x_2^k . If the regressors include a constant, this immediately implies that we can only identify the slope of the conditional mean function $\boldsymbol{\beta}$, but not the intercept.⁸ At the estimation stage, this logic could be implemented by a semiparametric pairwise difference estimator (Powell (1987, 2001); Ahn and Powell (1993); Honoré and Powell (2005); Aradillas-Lopez, Honoré, and Powell (2007)).

However, since we always require condition (3.10) and in particular $\mathbf{q}_{i1} = \mathbf{q}_{j1}$ to hold, this immediately implies that the identification requirement $x_{i2}^k \neq x_{j2}^k$ in general is only met by covariates that are not in the profit equation. Such lack of identification for variables that are included to control for selection is a limitation common to pairwise-differencing methods (see Aradillas-Lopez, Honoré, and Powell (2007) for a range of examples).

⁸ The intercept can be identified by an ‘identification at infinity’ argument, see discussion below. Having identified all the slope parameters ‘holding selection constant’, the support requirements for identifying the intercept ‘at infinity’ are quite weak. In particular, to identify the constant in $\boldsymbol{\beta}$ we only need to find a single region in the vector space spanned by the regressors where $g(\mathbf{x}_{i1}, \mathbf{q}_{i1}, \nu_{i1}) \rightarrow 1$ holds. The constant is then identified by the normalisation $E(w_{i2} - \mathbf{x}_{i,2}' \boldsymbol{\beta}) = 0$.

Identification of the whole parameter vector along these lines requires a strong form of exclusion restriction. In particular, the set of covariates $\mathbf{q}_1$ in the profit equation would need to be disjoint from the set of covariates $\mathbf{x}_2$ in the wage equation.

For time-varying variables, the panel structure of our model allows for circumventing a strong exclusion restriction of this form. This possibility relies on the particular timing embedded in Assumption 3.2. Selection is based on the earnings-generating process in period 1, whereas β is identified from wages in period two. Therefore, any changes in covariates between selection and the moment when wages are observed can substitute for exclusion restrictions of the form above. In my notation, this allows for $\mathbf{q}_{i1} = \mathbf{x}_{i1}$ but requires $\mathbf{x}_{i2} \neq \mathbf{x}_{i1}$ to identify at least some of the elements of β .

Identification of the profit equation parameters γ follows exactly the same steps, again requiring exclusion restrictions or time variation.

Remark 1: Identification does not require all covariates to be defined on a continuous support. Some may be defined on a discrete set. However, the constant selection condition (3.11) puts some limits on the sparsity of covariates. In particular, there need to be sufficient variation in covariates such that the conditional choice probability can be driven to the required level. As the discussion has made clear, linearity in conditional means ensures that the slope is identified from covariate variation across any two observations for which selection is held constant. In special cases, identification could be achieved with only discrete covariates if a suitable arrangement of covariate can deliver the same selection probability for a finite number of observations with distinct covariates. In general, however, at least one of the covariates should be a continuous variable.

Within-individual comparisons over time

In many applications of this two-sector model of employment and earnings, neither exclusion restrictions nor variation over time will be realistic assumptions on covariates. The same sets of observable covariates – measures of human capital, of other signals to the labour market, or of market conditions – will affect earnings in both self-employment and wage-employment. Many of these covariates are invariant or highly persistent over time.

Such covariates can nevertheless be identified by a different sort of comparison. For this we need to invoke Assumptions 3.1, 3.2, and 3.3. To see how a within-individual comparison can identify the coefficients of time-invariant, common covariates, write wages in the second period (for the subset who were wage-employed in the first period) as

$$\begin{aligned} w_{i2}|(S_{i1} = 1, \epsilon_{i1}) &= \mathbf{x}_i' \boldsymbol{\beta} + E_{\epsilon_{i2}|\epsilon_{i1}}(\epsilon_{i1}) + u_{i2} \\ &= \mathbf{x}_i' \boldsymbol{\beta} + E_{\epsilon_{i2}|\epsilon_{i1}}(w_{i1} - \mathbf{x}_i' \boldsymbol{\beta}) + u_{i2} \end{aligned} \quad (3.13)$$

whereby the vector $\mathbf{x}_i$ is invariant over time and common for earnings in wage-employment and self-employment. (This will be the case in the empirical application later in this chapter.) The residual component u_{i2} is independent of $\mathbf{x}_i$ and w_{i1} by construction. Selection bias arises because entry into wage-employment depends on the error term ϵ_{i1} in a way that is correlated with $\mathbf{x}_i$. Under stationarity, and absent any other further selection of *stayers* between period 1 and – and this is what we assume in (3.9) – the autocorrelation of the error term ϵ_{i2} with ϵ_{i1} will be the only source of selection bias in the second period. Including the control function $E_{\epsilon_{i2}|\epsilon_{i1}}$ would remove the selection bias.

Assumption 3.3 further implies that the functional form of $E_{\epsilon_{i2}|\epsilon_{i1}} = \rho_\epsilon \cdot \epsilon_{i1}$ is linear.⁹ Then, (3.13) becomes

$$\begin{aligned} w_{i2}|(S_{i1} = 1, \epsilon_{i1}) &= \mathbf{x}_i' \boldsymbol{\beta} + \rho_\epsilon (w_{i1} - \mathbf{x}_i' \boldsymbol{\beta}) + u_{i2} \\ &= \mathbf{x}_i' \boldsymbol{\beta} (1 - \rho_\epsilon) + \rho_\epsilon w_{i1} + u_{i2} \\ &= \mathbf{x}_i' \tilde{\boldsymbol{\beta}} + \rho_\epsilon w_{i1} + u_{i2} \end{aligned} \quad (3.14)$$

where $\tilde{\boldsymbol{\beta}} = \boldsymbol{\beta}(1 - \rho_\epsilon)$. OLS of (3.15) identifies $\tilde{\boldsymbol{\beta}}$ and ρ_ϵ , which is sufficient to recover $\boldsymbol{\beta}$. Similarly, second-period profits for leavers (who randomly switch to self-employment) can be written as

$$\begin{aligned} \pi_{i2}|(S_{i1} = 1, \epsilon_{i1}) &= \mathbf{x}_i' \boldsymbol{\gamma} + \tilde{\rho} (w_{i1} - \mathbf{x}_i' \boldsymbol{\beta}) + v_{i2} \\ &= \mathbf{x}_i' (\boldsymbol{\gamma} - \boldsymbol{\beta} \tilde{\rho}) + \tilde{\rho} w_{i1} + v_{i2} \\ &= \mathbf{x}_i' \tilde{\boldsymbol{\gamma}} + \tilde{\rho} w_{i1} + v_{i2} \end{aligned} \quad (3.15)$$

where $\tilde{\boldsymbol{\gamma}} = \boldsymbol{\gamma} - \boldsymbol{\beta} \tilde{\rho}$, v_{i2} is independent of $\mathbf{x}_i$ and w_{i1} , and $\mathbf{x}_i$ is again the time-invariant and common vector of covariates (i.e., $\mathbf{q}_{i2} = \mathbf{x}_{i1} = \mathbf{x}_i$). From $\tilde{\rho}$, $\tilde{\boldsymbol{\gamma}}$ and our previously identified $\boldsymbol{\beta}$ we can recover $\boldsymbol{\gamma}$. If additionally the error terms are jointly normally distributed, then the auxiliary parameter $\tilde{\rho}$ has a structural interpretation as $\tilde{\rho} = \rho_\nu \rho_{\sigma_\epsilon}^{\sigma_\nu}$.

To sum up Step 2: the conditional mean parameter vectors are identified either under strong exclusion restrictions but general assumptions about the

⁹ If $E_{\epsilon_{i2}|\epsilon_{i1}}$ were non-linear, the shape restriction that this implies should also be sufficient for identification. This is analogous to the argument that non-linearity (of general form, not confined to a particular non-linear functional form) in the process governing the selection equation is sufficient for identification of the Heckman selection model made by Ahn and Powell (1993) and Lewbel (2017). Proving this conjecture might be an avenue for further research.

selection mechanism; or under stronger economic and statistical assumption on the selection mechanism and the structure of unobservables that dispense with the requirement for exclusion restrictions. Once the conditional mean parameters are known, further elements of the structure, in particular the form on earnings distribution and the degree of search frictions, can be identified.

Step 3: Identification of conditional earnings distributions

Given that transition from wage employment to self-employment is independent of any determinant of earnings by Assumption 3.2, the empirical observations on earnings from switchers identify the bivariate conditional earnings distribution:

$$F_{\nu_2|\epsilon_1}(\pi_{i2} - \mathbf{q}_{i2}'\gamma) = \hat{F}(\pi_{i2} - \mathbf{q}_{i2}'\gamma \mid \mathbf{q}_{i2}, \epsilon_{i1}) \quad (3.16)$$

where $\hat{F}$ is the empirical distribution function. By varying the point of evaluation and holding ϵ_{i1} fixed, we can trace out the entire array of conditional distributions. We cannot identify yet the joint bivariate distribution, given that the cross-sectional empirical distribution of ϵ_{i1} does not equal the marginal distribution $f(\epsilon_1)$. This is a consequence of the selection problem.

In a parallel manner, we can identify the conditional distribution

$$F_{\epsilon_2|\epsilon_1}(w_{i2} - \mathbf{x}_{i2}'\beta) \quad (3.17)$$

from the empirical distribution of wages of those who stay in wage-employment.

We know the selection rule for those who switch into wage-employment, we observe their earnings, and by now we are able to decompose profits and

earnings into observable and unobservable components. Therefore, we can trace out the function

$$\begin{aligned}
& \Pr(S_{i2} = 1, w_{i2} \leq W \mid v_{i1}) \\
&= \lambda \Pr(\mathbf{x}'_{i1}\boldsymbol{\beta} + \epsilon_{i1} \geq \mathbf{q}'_{i1}\boldsymbol{\gamma} + \nu_{i1}, \mathbf{x}'_{i2}\boldsymbol{\beta} + \epsilon_{i2} \leq W \mid v_{i1}) \\
&= \lambda \cdot F_{(-\epsilon_1, \epsilon_2 \mid \nu_1)}(\mathbf{x}'_{i1}\boldsymbol{\beta} - \mathbf{q}'_{i1}\boldsymbol{\gamma} - \nu_{i1}, W - \mathbf{x}'_{i2}\boldsymbol{\beta}) \\
&\equiv \Lambda(\epsilon_{i1}, \epsilon_{i2}, \nu_{i1})
\end{aligned} \tag{3.18}$$

by varying the points of evaluation. Note that in Λ , selection is driven by both the degree of search frictions and by the productivity differential $\epsilon_{i1} - \nu_{i1}$, conditional on the productivity in entrepreneurship ν_{i1} . In the language used by some of the literature on identification of Roy selection models (e.g. Suri (2011), Adão (2015)), we might call the productivity differential *comparative advantage*, and the absolute level of productivity in the reference sector *absolute advantage*. Only comparative advantage matters for selection. However, a crucial distinction of my model and these existing approaches is that the degree of selection alone does not identify the schedule of comparative advantage. Search frictions create a wedge between the degree of selection, and the distribution of absolute advantage.

In order to isolate comparative advantage from frictions, note that since we can identify ϵ_{i1} and ϵ_{i2} , we are able to evaluate the integral

$$\int_{-\infty}^{\infty} \Lambda(\epsilon_{i1}, \epsilon_{i2} \mid \nu_{i1}) d\epsilon_1 = \lambda \cdot F_{\epsilon_2 \mid \nu_1}(W - \mathbf{x}'_{i2}\boldsymbol{\beta}) \tag{3.19}$$

and thereby obtain a function containing the marginal (conditional) distribution of ϵ_2 . Given that (3.18) and (3.19) both contain the same transformation

by λ , this is enough to identify the conditional distribution

$$\begin{aligned}
 F_{(-\epsilon_1|\nu_1)}(x) &= \frac{\Lambda(\epsilon_{i1}, \epsilon_{i2}, \nu_{i1})}{\lambda \cdot F_{\epsilon_2|\epsilon_1, \nu_1}(W - \mathbf{x}'_{i2}\boldsymbol{\beta})} \\
 &= \frac{F_{(-\epsilon_1, \epsilon_2|\nu_1)}(x, W - \mathbf{x}'_{i2}\boldsymbol{\beta})}{F_{\epsilon_2|\nu_1}(W - \mathbf{x}'_{i2}\boldsymbol{\beta})} \\
 &= F_{(-\epsilon_1|\epsilon_2, \nu_1)}(x) \\
 &= F_{(-\epsilon_1|\nu_1)}(x)
 \end{aligned} \tag{3.20}$$

where the last step follows from the law of total probabilities. The conditional distribution $F_{\epsilon_1|\nu_1}$ corresponds to the schedule of comparative advantage, since by a change of variable, (3.20) also identifies $F_{(\epsilon_1 - \nu_1|\nu_1)}$.

Noting that we have now identified all the ingredients of the conditional choice probability (3.8) but λ , we can straightforwardly identify λ from the conditional choice probability.

Remark 2: At this point, we have also identified the relevant information set Ψ_i for the entrepreneur's evaluation of conditional choice probability (3.8), given by:

$$\Psi_i = \{\pi_{i1}, \mathbf{x}'_{i1}\boldsymbol{\beta}, \mathbf{q}'_{i1}\boldsymbol{\gamma}, F_{\epsilon_1|\nu_1}\}. \tag{3.21}$$

In words, the agent fully knows what their expected incomes given observables would be. Together with her observation of profits in period 1, this allows her to construct the productivity residual ν_{it} (which is unobservable to the econometrician). Furthermore, she also is able to construct the unobservable part of her expected wages, given her unobservable productivity as an entrepreneur. This requires knowledge of the conditional distribution of productivity in entrepreneurship and in wage work. This information set is sufficient for evaluating the attractiveness of wage-employment from the

perspective of the entrepreneur, assuming the entrepreneur is able to correctly evaluate her conditional choice probability.

Step 4: Identification of the joint earnings distribution

To recap the identification results so far: we have identified the conditional mean functions given by vectors β and γ , the search friction parameters δ and λ , and the conditional distributions $F_{\epsilon_1|\nu_1}$, $F_{\nu_2|\epsilon_1}$ and $F_{\epsilon_2|\epsilon_1}$. For this we relied on Assumptions 3.1 and 3.2, and partly on stationarity in Assumption 3.3. Even then, we did not put strong parametric restrictions on the distribution of unobservables.

Recovering the joint distribution of earnings requires us to know the marginal distributions. Even stationarity in Assumption 3.3 is not sufficient for this. The self-employed are a selected sub-population, and the selection correction is a complicated function of the underlying productivity distribution, the search frictions, and the productivity decay parameterised by ρ_ϵ and ρ_ν .¹⁰ To recover the joint distribution $F_{(\epsilon_1, \epsilon_2, \nu_1, \nu_2)}$, we need to make additional distributional assumptions that allow us to recover a marginal distribution (e.g. F_{ϵ_1}) from the conditional distribution (e.g. $F_{\epsilon_2|\epsilon_1}$). The assumption that F follows a normal distribution achieves this.

Under Assumptions 3.3 and 3.4, the joint earnings distribution across the

¹⁰ Perhaps some progress could be made with additional stationarity assumptions on selection, e.g. that the observed distribution of $(\nu_t|S_t = 0)$ is constant over time. It is outside the scope of this chapter to establish identification using this and most likely additional auxiliary assumptions.

two observed periods takes a particular parametric form:

$$\begin{pmatrix} \epsilon_1 \\ \nu_1 \\ \epsilon_2 \\ \nu_2 \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\epsilon^2 & \rho \sigma_\epsilon \sigma_\nu & \rho_\epsilon \sigma_\epsilon^2 & \rho_\nu \rho \sigma_\epsilon \sigma_\nu \\ \rho \sigma_\epsilon \sigma_\nu & \sigma_\nu^2 & \rho_\epsilon \rho \sigma_\epsilon \sigma_\nu & \rho_\nu \sigma_\nu^2 \\ \rho_\epsilon \sigma_\epsilon^2 & \rho_\epsilon \rho \sigma_\epsilon \sigma_\nu & \sigma_\epsilon^2 & \rho \sigma_\epsilon \sigma_\nu \\ \rho_\nu \rho \sigma_\epsilon \sigma_\nu & \rho_\nu \sigma_\nu^2 & \rho \sigma_\epsilon \sigma_\nu & \sigma_\nu^2 \end{pmatrix} \right)$$

The three known conditional distributions are then fully characterised by the conditional expectations and variances, given by:

$$E(\epsilon_2 | \epsilon_1) = \rho_\epsilon \cdot \epsilon_1 \quad (3.22)$$

$$Var(\epsilon_2 | \epsilon_1) = \sigma_\epsilon^2 (1 - \rho_\epsilon^2) \quad (3.23)$$

$$E(\nu_2 | \epsilon_1) = \rho_\epsilon \rho \frac{\sigma_\nu}{\sigma_\epsilon} \cdot \epsilon_1 \quad (3.24)$$

$$Var(\nu_2 | \epsilon_1) = \sigma_\nu^2 (1 - \rho^2 \rho_\nu^2) \quad (3.25)$$

$$E(\epsilon_1 | \nu_1) = \rho \frac{\sigma_\epsilon}{\sigma_\nu} \cdot \nu_1 \quad (3.26)$$

$$Var(\epsilon_1 | \nu_1) = \sigma_\epsilon^2 (1 - \rho^2) \quad (3.27)$$

These these six equations identify the five parameters that characterise the joint normal distribution. The entire structure of the RSM is now identified.

3.2.5 Relation to literature on identification of selection models

Identification of the Roy Search Model differs from other approaches because selection is not only driven by comparative advantage, but also by a search friction constraint which makes the arrival of a wage offer (and not just the value of the offer) a stochastic process. As the last section has shown, the identification argument still can follow some steps that are simi-

lar to the identification of the Roy model. In this section, I briefly relate the approach I take to other known approaches to the specification and identification of selection models.

My argument of “identification holding selection constant” is inspired by the “identification along a ray” argument in Heckman and Honoré (1990). A different approach for the same class of “pure” Roy models is “identification at infinity” along the lines of French and Taber (2011) and earlier contributions (going back to at least Chamberlain (1986)). There, the idea is to make the selection problem disappear, instead of holding it constant. This requires exclusion restrictions on a much larger support, which shift the conditional choice probabilities to zero or unity for *every* combination of the covariates. Then, the wage equation can be identified from those individuals for whom selection is not an issue. The advantage of this fully non-parametric approach is that the conditional expectation function can take any arbitrary, non-parametric form. The full support assumption is likely not met by the kind of data that I have in mind with this model (see section 3.3).

My discussion of holding selection constant bears resemblance to a strand of literature on semiparametric estimation that takes as its basis pairwise differences of individuals with the same selection probability (Powell (1987), Ahn and Powell (1993), Powell (2001), Honoré and Powell (2005), Aradillas-Lopez, Honoré, and Powell (2007)). This strand of literature tends to treat the mechanism determining selection as a non-parametric “nuisance” problem. As Aradillas-Lopez, Honoré, and Powell (2007) point out, selection models with a non-parametric control function (i.e. one that is not summarised by a single index model) require fairly strong exclusion restric-

tions. In the context of the classic Roy model which I consider here, such exclusion restrictions are hard to justify economically. I therefore discuss an alternative approach that dispenses with exclusion restriction by imposing the choice structure implied by the simple canonical search model, where self-employment takes the role of the free-entry sector.

By exploiting the discrete choice structure implied by economic theory, the RSM is able to uncover the full earnings distribution in a two-sector model under an endogenous switching mechanism with potential search frictions. Another alternative approach in the literature (Lemieux (1998); Suri (2011)) identifies correlated, and sector-specific, returns to individual unobserved heterogeneity using linear panel methods.¹¹ Identification using this alternative approach requires that selection is based only on permanent comparative advantage, and not on transitory components of the relative earnings across sectors, such as particular wage offers. This is arguably too restrictive an assumption for modeling entry into wage employment in many applications.

Finally, the “generalised Roy model” (Heckman and Vytlačil (2007)) constitutes a class of selection models that lie between the “pure” Roy model as in Heckman and Honoré (1990) – where selection is entirely driven by particular wage offers – and the correlated random coefficient model of Lemieux (1998) and Suri (2011) – where selection is quasi-random, conditional on permanent heterogeneity. Instead of being totally random, selection in the generalised Roy model is driven by an observable variable that is excluded from the earnings functions in either sector. Shifters of preferences, or of non-pecuniary job amenities in a particular sector, or costs to

¹¹ Applications of the correlated random coefficients model to occupational choice and earnings distributions in developing countries include Haywood and Falco (2016) and Imbert (2013).

participation are examples of such excludable selection shifters. Identification of the generalised Roy model is discussed in Heckman and Vytlačil (2007), French and Taber (2011), and Eisenhauer, Heckman, and Vytlačil (2015).

3.2.6 Maximum likelihood estimation of the RSM

I estimate the Roy Search Model via Maximum Likelihood, imposing Assumptions 3.1–3.4. As the discussion above has shown, these assumptions identify the structural parameters of the model: the conditional means parameters, the two parameters relating to search frictions (offer and separation rates) as well the five parameter governing the joint normal distribution of residuals.¹² I derive the likelihood function under the model in the Appendix.

My data are from a labour force survey in urban Colombia (see section 3.3 and the Appendix for more details and descriptives). I create a balanced panel of individuals who we observe over two adjacent annual panel waves in either self-employment or wage-employment. With this sample restriction criterion, I abstract from the participation decision. There are some empirical justifications for this choice, which I explain in section 3.3. In practical terms, a balanced panel is required for estimation using a likelihood approach. Unbalanced data would require a moment-based estimation approach.

I estimate different specifications of the model that differ in the variables (and interaction terms) included in the conditional means of wages and

¹² As stated elsewhere in this chapter, I do not attempt at formally proving identification. This would require additional technical regularity assumptions. Instead I focus on the empirical content, the key substantial assumptions that are important for empirical applications.

profits. In each specification, however, the covariates are common across the two sectors, and most covariates are constant over the two-year horizon over which I observe the data. Therefore, in my application, earnings equation coefficients are not identified through exclusion restrictions; but rather through assumptions on the economic structure (see section 3.2.4, Step 2). Absent natural experiments from policy or otherwise, exclusion restrictions are already hard to justify for the selection problem into work participation (Blundell, Duncan, and Meghir (1998)). They are even more challenging for the problem of choosing between occupations, especially in a developing country context with high informality, where taxes and regulations constitute only soft constraints for individual behaviour.¹³

I include measures of human capital (either years of schooling or dummies for education levels; and an experience quadratic), gender, as well as city-year dummies to capture differential shocks to earnings across local labour markets. The latter are important because differences in local labour market conditions can be a confounding factor in the study of labour supply (Heckman, Ichimura, and Todd (1997), Smith and Todd (2001)).

3.2.7 The Roy Search Model and the Lure of Wage Work

While interesting in its own right, in this chapter, the Roy Search Model (RSM) also serves as an intermediate step in the construction of a measure of the attractiveness of wage employment for entrepreneurs. This in turn allows me to conduct an alternative test of the model in Chapter 1. For

¹³ I tried to exploit strict programme assignment cutoff rules to subsidised health insurance based on the Colombian SISBEN system; see Espinosa (2012) for evidence on the incentives for labour informality this creates. The idea was that all self-employment is informal, while wage employment is at least partly formal work. In practice, though, this turned out to lead to little variation along the self-employment / wage-employment margin. Furthermore, such non-pecuniary benefits are difficult to reconcile with the classic Roy model from this chapter; they are more adequately represented by a generalised Roy model.

this test, I regress outcomes of interest on the attractiveness of wage employment. Similar to the empirical analysis in Chapter 2, the outcomes of interest are transitions out of self-employment into wage-employment, and firm growth. Here, the outcomes are measured as binary variables, and the latent variable is a function of alternative employment opportunities (3.1) and control variables:

$$y_i = \begin{cases} 1 & \text{if } \alpha_0 + \alpha_1 \cdot M_i + \mathbf{\Gamma}_i' \alpha_2 + \tau_i > 0 \\ 0 & \text{if } \alpha_0 + \alpha_1 \cdot M_i + \mathbf{\Gamma}_i' \alpha_2 + \tau_i \leq 0 \end{cases} \quad (3.28)$$

where y_i denotes a binary outcome for individual i , $M_i = M(\mathbf{x}_i, \mathbf{q}_i)$ is the measure of attractiveness of wage employment, and τ_i an error term. The main parameter of interest is α_1 . I vary the vector of variables $\mathbf{x}_i$ and $\mathbf{q}_i$ that enter the conditional means in the RSM. I further include a vector of individual-level controls $\mathbf{\Gamma}_i$. All variables are specified at baseline (that is, in the first observation period). In practice, I estimate the model (3.28) as a linear probability model using OLS.

Estimation Procedure

I combine all of this in a procedure that consists of three steps.

1. In the first step, I estimate the RSM using maximum likelihood, and I predict counterfactual wages for entrepreneurs.
2. In the second step, I construct the empirical counterparts to the measure of the attractiveness of wage employment (3.1):

$$\hat{\Pr}(w_{i1} > \pi_{i1}) = \Phi \left(\frac{\mathbf{x}_{i1}' \hat{\boldsymbol{\beta}} + \left(\hat{\rho} \frac{\hat{\sigma}_\epsilon}{\hat{\sigma}_\nu} \right) \cdot \nu_{i1} - \pi_{i1}}{\hat{\sigma}_\epsilon (1 - \hat{\rho}^2)} \right) \quad (3.29)$$

3. Finally, I estimate binary outcome equations using constructed measures as a regressor. In particular, I estimate the linear probability model (3.28) using as outcomes:

- *Leave to WE*: A dummy equal one if a business owner transits to wage-employment between periods 1 and 2, zero if stays self-employed.
- *Firm Growth*: A dummy equal one a business owner adds salaried employees from outside the household to the firm between periods 1 and 2, zero otherwise.

I bootstrap the standard errors in the first and third steps by drawing with replacement from my data, for 1,000 iterations. I stratify the bootstrap draws. I draw from a sample of size N_1 and N_2 , respectively, of wage-employed and self-employed individuals from my balanced panel on which I estimate the RSM. I further draw from a sample of size N_{Out} for the outcome equations.

3.3 Data

3.3.1 Background

The data used in this analysis consist of a panel survey of urban households in Colombia, the *Encuesta Social Longitudinal de Fedesarrollo* (Fedesarrollo Social Longitudinal Survey; ESLF).¹⁴ The survey was collected by the Bogotá-based foundation Fedesarrollo, supported by the Inter-American Development Bank (IADB) and the Bogotá, Bucamaranga and Calí Chambers of Commerce. The 2007-2010 waves include a unified module on the

¹⁴ For recent empirical work on occupational choice in Colombia, see Magnac (1991) Bosch and Maloney (2010) Mondragón-Vélez and Peña (2010), and Cuesta and Bohórquez (2011).

labour market. The 2007 wave covers only the Metropolitan Areas of Bogotá, Bucamaranga and Calí. In 2008, the survey was enlarged to include ten additional cities and Metropolitan Areas.¹⁵

The sampling frame consists of the 1993 Census population list of dwellings, households and individuals. Census lists are updated using information from the Colombian National Statistical Office DANE and municipal planning departments. Households are selected into the survey using a thrice-stratified random sampling design: the primary sampling units are Census enumeration areas, secondary sampling units are housing blocks (*'manzanas'*), and the final and tertiary sampling units are households. Households attriting from the panel are replaced by new ones. I retain individuals who are observed in adjacent waves of the survey, either in self-employment or in wage-employment.

3.3.2 Descriptive statistics

The universe of the survey from 2008 onwards comprises of the 13 cities and metropolitan areas constituting the 'total national urban population' as defined for statistical purposes by the Colombian National Statistical Office DANE. Table 1 illustrates some general properties of the urban labour market in Colombia. The labour force participation rate, that is the proportion of the working age population in the labour force, is fairly constant at about 55 percent. A little over 50% of the labour force are wage employed, about a third self-employed, and about 10% are classified as unemployed. Over the years, there is a slight but significant trend of decreasing rates of wage-employment, and increasing rates of unemployment and self-employment.

¹⁵ Barranquilla, Cartagena, Cúcuta, Ibagué, Manizales, Medellín, Montería, Pasto, Pereira, and Villavicencio.

TABLE 1
COMPOSITION OF THE URBAN LABOUR MARKET IN COLOMBIA
2007-2010

	2007	2008	2009	2010
LF Participation Rate	54.4	55.4	54.7	55.4
% of Labour Force				
Wage-Employed	58.3	54.7	53.0	50.9
Self-Employed	31.3	33.7	33.9	35.9
Unemployed	8.8	10.0	11.7	12.1

Source: Encuesta Social Longitudinal de Fedesarrollo (ESLF). Ratios are representative for the total urban national population of Colombia. *Wage-employed* includes private and government sector employees, domestic and casual workers. *Self-Employed* includes own-account workers and employers.

The period covered by the sample corresponds to the global financial crisis. GDP growth in 2007 was 6.9 percent, then dropped stepwise to 3.5 percent in 2008 and 1.7 percent in 2009 before recovering in 2010 with 4 percent and 2011 with 5.9 percent. This adverse macroeconomic environment may explain some of the trends. The global financial crisis constitutes a purely external shock to the Colombian economy. Like other South American economies, it had been on a high growth trajectory for a number of years. The global recession that ensued after the financial crisis in developed countries led to a sharp reduction in GDP growth, with a quick recovery.

I exploit the fact that the economic downturn brought with it a higher number of involuntary job separations. Labour market transitions across all years are shown in Table 2. The vertical dimension corresponds to the employment status of an individual in a given year t , and the horizontal dimension to the employment status in the following year $t + 1$. The table should be read row-wise. Every entry in the main matrix shows the percentage of those in the row category in year t that are in the column category in year $t + 1$. For example, 6.6 percent of those out of the labour force switched

to wage-employment in the following year. The main diagonal corresponds to inertia: individuals who stay in the same form of employment in two adjacent years.

TABLE 2
LABOUR MARKET TRANSITIONS IN URBAN COLOMBIA 2007-2010


Employment Status t	Employment Status $t+1$					Total
	Out of LF	WE	SE	Other	UE	
Out of LF	81.5	6.6	6.5	0.6	4.7	100
WE	8.3	72.9	11.9	0.5	6.4	100
SE	15.7	16.7	62.3	1.0	4.3	100
Other	44.6	20.1	23.4	5.4	6.5	100
UE	28.6	29.8	16.5	1.1	24.0	100
Total $t + 1$	45.1	29.0	18.9	0.7	6.3	100

Source: Source: Encuesta Social Longitudinal de Fedesarrollo (ESLF). *WE* includes private and government sector employees, domestic and casual workers. *SE* includes own-account workers and employers. *Other* includes unpaid labour and unspecified.

Several stylised facts emerge from the analysis of transitions. First, there is large flexibility in the labour market. Fewer than 75 percent of the wage-employed, and fewer than two thirds of the self-employed remain in their respective occupations. Second, bilateral movements between self-employment and wage employment are more prevalent than movements with other sectors, for example unemployment. This supports the idea of on-the-job-search of the self-employed as an important economic mechanism in urban labour markets of developing countries.

There is an overlap of profits in self-employment and wages in wage-employment across the entire distribution, as Figure 1 shows. However, while profits are quite dispersed, wages are much compressed, in particular in the middle of the distribution. There is a large asymmetric spike in wages at

FIGURE 1
DISTRIBUTION OF PROFITS AND WAGES



the level which corresponds to the minimum wage. Given the prevalence of the informal sector, there is a mass in the wage distribution below the minimum wage. Nevertheless, minimum wages seem to constitute an important distortion, and a focal point, in the wage distribution in Colombia.

Summary statistics of the variables used in the analysis, and precise variable definitions, can be found in Appendix C.2. On average, self-employed individuals are almost a decade older, have more than a year less of formal education, and are more likely to be married as well as head of their household. Interestingly, more than three quarters of both wage-employed and self-employed report being satisfied with their employment, although 41 percent report reasons for self-employment that imply involuntary self-employment back lack of opportunities. Just under 20 percent of those self-employed in a given year switch into wage employment in the following

FIGURE 2
FIRM SIZE DISTRIBUTION



year.¹⁶ A quarter of those who stay self-employed in consecutive periods add employees to their firm. Figure 2 reveals that more than half the firms in the sample (56 percent) are owner-operator firms without other employees. Approximately 21 percent have a single employee, 8 percent have two employees. The ninetieth percentile is at five workers. Less than three percent of the firms in the sample have more than ten workers.

3.4 Results

This section presents the results from an estimation of the Roy Search Model using Colombian panel data, and then continues into using the results for a test of the model of Chapter 1.

¹⁶ These transition rates are higher than in Table 2 because here I only consider individuals who are either wage-employed or self-employed in t and $t + 1$.

3.4.1 First step: An estimated Roy Search Model

I estimate the Roy Search model via Maximum Likelihood, imposing Assumptions 3.1–3.4. I derive the likelihood function under the model in Appendix C.1.

In Table 3, I report the results from an RSM estimation for urban Colombia. The dataset lacks a suitable variable which could reasonably be assumed to be excludable from either of the earnings functions. The small number of markets and time period further preclude using sectoral productivity shifters, such as the export manufacturing shock of Chapter 2.

I report six specifications that differ in the set of variables which are included in the mean wage and profit functions. In column (1), the conditional mean follows a standard Mincer specification, only augmented by gender. In column (2), I add city-year dummies. In absence of any observable demand shifters, city-year dummies are the best way to capture aggregate shocks across markets, for example external shocks due to spillovers of the US financial crisis. Column (4) features a saturated model of human capital and demographic interactions. I evaluate the specifications with a likelihood ratio test. The test reveals that the more saturated models do achieve a better fit. I also show an alternative specification where years of schooling are substituted by education level dummies, which is a coarser but more flexible specification. Although the models are not nested and a likelihood test can therefore not be performed, by informal comparison, years-of-schooling specifications achieve a higher likelihood with fewer parameters. They are therefore my preferred specifications.

Across the different conditional mean specifications, the estimated earn-

TABLE 3
ROY SEARCH MODEL

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Earnings distribution</i>						
σ_ϵ	0.524 (0.008)	0.521 (0.008)	0.517 (0.008)	0.514 (0.008)	0.525 (0.008)	0.521 (0.008)
σ_ν	3.096 (0.097)	3.078 (0.096)	3.078 (0.096)	3.077 (0.096)	3.096 (0.096)	3.079 (0.095)
ρ_ϵ	0.558 (0.022)	0.554 (0.022)	0.545 (0.022)	0.536 (0.022)	0.549 (0.022)	0.545 (0.022)
ρ_ν	0.089 (0.031)	0.080 (0.031)	0.082 (0.031)	0.082 (0.031)	0.090 (0.031)	0.082 (0.031)
ρ	0.531 (0.111)	0.499 (0.119)	0.527 (0.119)	0.545 (0.111)	0.600 (0.109)	0.565 (0.101)
<i>Search frictions</i>						
λ	0.276 (0.017)	0.278 (0.017)	0.276 (0.017)	0.275 (0.017)	0.273 (0.017)	0.275 (0.016)
δ	0.147 (0.006)	0.147 (0.006)	0.147 (0.006)	0.147 (0.006)	0.147 (0.006)	0.147 (0.006)
<i>Cond. mean function</i>						
Years of Schooling	✓	✓	✓	✓		
Educ. Dummies					✓	✓
Exp. Quadratic	✓	✓	✓	✓	✓	✓
Gender	✓	✓	✓	✓	✓	✓
Gender $\times$ yos			✓	✓		
Gender $\times$ educ.						✓
Gender $\times$ exp.			✓	✓		✓
yos $\times$ exp.				✓		
City-year dummies		✓	✓	✓	✓	✓
Log-likelihood	-17627	-17584	-17549	-17532	-17644	-17600
LR test statistic		86	70	34		74
LR test p-value		0.0000	0.0000	0.0000		0.0000
Comparison model		(1)	(2)	(3)		(5)
# of parameters	18	36	52	54	52	64

Each column corresponds to an estimation of the Roy Search Model (RSM) with a different vector of variables $\mathbf{x}_{it} = \mathbf{q}_{it}$ that describe the mean of wages and profits. Sample size are 4,590 Colombian workers observed for two adjacent periods in either wage-employment or self-employment. Bootstrapped standard errors with 1,000 replications in parentheses.

ings distribution and search friction parameters are remarkably stable. The relative magnitude of parameters is as expected. The dispersion of entrepreneurial productivity is much higher than that of wage productivity. Profits are only weakly correlated over time, conditional on observables, whereas wage residuals have an autocorrelation of about 0.55. The absolute advantages in wage-employment and entrepreneurship are positively cor-

related, with a fairly strong correlation of about 0.5. Around 15 percent of all wage workers lose their job and become self-employed every year. This parameter is directly identified by a single moment that is independent of earnings; it is therefore completely constant across specifications. The probability of receiving a wage offer while self-employed is about 27 percent. This suggests a stronger degree of search frictions than in developed countries, where estimates of λ are closer to 1.¹⁷


3.4.2 Second step: Measures of alternative employment opportunities

Using these results, I construct the sample measure of attractiveness of wage-employment (WE) defined in equation (3.1). Figure 3 summarises its distribution. I focus on three conditional mean specifications from Table 3: the bare-bones Mincer model (1), model (2) which includes city-year level earnings shifters, and the fully saturated specification (3). I plot a kernel estimate of the density function of the measure, as well as the cumulative distribution function.

The first immediate observation is that the conditional mean specification virtually makes no difference for the distribution of attractiveness. This is also true at the individual level: the pairwise correlation of the different measures is always above 0.98. Second, the distribution of the measure seems to be well approximated by a uniform distribution between roughly the 10th and 60 percentile. Third, there is a mass of entrepreneurs (around 20 per-

¹⁷ Jolivet, Postel-Vinay, and Robin (2006) report estimates of λ , for the unemployed, ranging from 0.33 in Belgium to 0.77 in Germany. Ridder and van den Berg (2003) report *monthly* estimates of λ of 0.13 for the UK and 0.61 for the US. On the other hand, Cahuc, Postel-Vinay, and Robin (2006) find annual job offer rates in the range of 0.05-0.32 for France, often viewed as an OECD country with little flexibility in labour market institutions.

FIGURE 3
DISTRIBUTION OF ATTRACTIVENESS OF WAGE-EMPLOYMENT



Source: Attractiveness of wage-employment for each period 1 entrepreneur in the sample ($N = 1,623$), constructed according to equation (3.1). Parameter estimates come from the different models that are presented in Table 3.

cent) with predicted attractiveness of 1. A similar proportion are assigned measures between 0.95 and 1. A measure of 1 implies that entrepreneurs would be certain to obtain a financially attractive wage offer. All of these observations correspond to entrepreneurs with profits below the lowest support point of the predicted wage offer distribution.

3.4.3 Third step: Outcome regressions

In the final step of the empirical analysis, I proceed to testing whether alternative wage employment opportunities matter for business decisions of owners. Following on the analysis of Chapter 2, I test again for the effects on two outcomes: leaving entrepreneurship, and small firm growth. Given that I have a much smaller sample with limited variation, I construct dummy variables for both outcomes. I defined them in section 3.2.7.

Table 4 reports the result results of this test. I estimate model (3.28) as

TABLE 4
RSM TEST OF THE LURE OF WAGE WORK

Roy Search Model specification: (1)						
Dependent Variable	Leave to WE			Firm Growth dummy		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Pr(w_{i1} > \pi_{i1})$	0.036 (0.027) [0.196]	0.047 * (0.028) [0.092]	0.040 (0.027) [0.138]	-0.071 ** (0.035) [0.045]	-0.066 * (0.035) [0.062]	-0.062 * (0.035) [0.079]
City-year dummies		✓	✓		✓	✓
Gender			✓			✓
Years of Schooling			✓			✓
Exp. Quadratic			✓			✓
Observations	1,623	1,623	1,623	1,316	1,316	1,316
R ²	0.001	0.012	0.063	0.003	0.024	0.004
Roy Search Model specification: (2)						
Dependent Variable	Leave to WE			Firm Growth dummy		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Pr(w_{i1} > \pi_{i1})$	0.037 (0.028) [0.188]	0.047 * (0.028) [0.091]	0.040 (0.027) [0.142]	-0.068 * (0.035) [0.055]	-0.064 * (0.035) [0.067]	-0.060 * (0.035) [0.085]
City-year dummies		✓	✓		✓	✓
Gender			✓			✓
Years of Schooling			✓			✓
Exp. Quadratic			✓			✓
Observations	1,623	1,623	1,623	1,316	1,316	1,316
R ²	0.001	0.012	0.063	0.003	0.024	0.039
Roy Search Model specification: (4)						
Dependent Variable	Leave to WE			Firm Growth dummy		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Pr(w_{i1} > \pi_{i1})$	0.038 (0.027) [0.167]	0.049 * (0.028) [0.076]	0.033 (0.027) [0.223]	-0.063 * (0.035) [0.075]	-0.060 * (0.035) [0.086]	-0.062 * (0.035) [0.076]
City-year dummies		✓	✓		✓	✓
Gender			✓			✓
Years of Schooling			✓			✓
Exp. Quadratic			✓			✓
Observations	1,623	1,623	1,623	1,316	1,316	1,316
R ²	0.001	0.013	0.062	0.003	0.024	0.040

Linear Probability Model regressions of binary outcomes. All specifications include a constant. The explanatory variable $\Pr(w_{i1} > \pi_{i1})$ constructed according to equation (3.1). Each panel corresponds to a different first step specification, as detailed in Table 3. Bootstrapped standard errors, stratified by step, with 1,000 replications, in parentheses. P-values in square brackets.

a linear probability model.¹⁸ I report the coefficient of interest, α_1 , and I indicate the set of control variables in the final step regression. Each of the three panels of the table corresponds to a different specification of the conditional earnings means in the RSM.

The results suggest a positive but weak correlation between the attractiveness of wage-employment and the probability of leaving to wage-employment. Theoretically, the multiplier on the attractiveness of wage employment should be the probability of receiving an offer, estimated as $\lambda = 0.278$. The estimated coefficient amounts to only a fraction of this, and is only significant at the 10 percent level in one of three specifications. The coefficient lies between 0.033 and 0.049 in the nine combinations of first and third step specifications tested. This clearly suggests that the attractiveness of wage-employment inferred from the RSM is only one among several mechanisms that explain occupational choice. Also, the measure might be too coarse to capture actual counterfactual job opportunities. Firm growth is somewhat more significantly related to the attractiveness of wage-employment. Again the coefficient is quite stable, between 0.060 and 0.071 across specifications. It is significant at the 10 percent level throughout.

Taken together, the findings are consistent with the predictions from the model, on both margins of choice: entrepreneurs with better alternatives in wage-employment are less likely to grow their firms, and they are (somewhat) more likely to take a wage job. However, these correlations are fairly weak, and not consistently significant.

¹⁸ Given that my regressor is a constructed variable, this precludes the use of a nominal model for discrete outcomes. Chesher (2010) discusses the problems that could occur in the context of instrumental variable models.

3.5 Robustness

In this section, I explore robustness of the empirical results to alternative assumptions about the nature of selection and the structure of the joint earnings distribution.

3.5.1 An alternative formulation of the RSM

In section 3.2.3, I discussed the theoretical ambiguity with respect to the appropriate point in time at which individuals can be assumed to evaluate their comparative advantage. The main analysis is under the assumption that this comparison is made at baseline, in period 1. Here, I present the results for an alternative model based on a comparison of expected period 2 earnings. Appendix C.3 states the likelihood function of this alternative model.

There are at least two takeaways from comparing the estimates in Table 5 to the benchmark model in Table 3. First, the parameter estimates are again remarkably similar across the different conditional mean specifications. Second, the estimated parameters are also very close to the ones from Table 3. This enhances the confidence in the benchmark results.

3.5.2 Sensitivity of outcome regressions to alternative RSM formulations

I also check the sensitivity of the outcome regressions in step 3 to alternative formulations of the RSM. In the top panel of Table 6, I use the alternative RSM estimates obtained in the previous table to construct the attractiveness measures. The results are quite different than the ones reported in Table

TABLE 5
ROY SEARCH MODEL: ALTERNATIVE SELECTION

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Earnings distribution</i>						
σ_ϵ	0.525	0.521	0.517	0.514	0.525	0.521
σ_ν	3.094	3.078	3.077	3.076	3.094	3.078
ρ_ϵ	0.565	0.560	0.551	0.543	0.558	0.551
ρ_ν	0.079	0.071	0.072	0.071	0.081	0.074
ρ	0.167	0.166	0.185	0.190	0.239	0.239
<i>Search frictions</i>						
λ	0.299	0.300	0.300	0.300	0.299	0.300
δ	0.147	0.147	0.147	0.147	0.147	0.147
<i>Cond. mean function</i>						
Years of Schooling	✓	✓	✓	✓		
Educ. Dummies					✓	✓
Exp. Quadratic	✓	✓	✓	✓	✓	✓
Gender	✓	✓	✓	✓	✓	✓
Gender $\times$ yos			✓	✓		
Gender $\times$ educ.						✓
Gender $\times$ exp.			✓	✓		✓
yos $\times$ exp.				✓		
City-year dummies		✓	✓	✓	✓	✓

Each column corresponds to an estimation of the alternative version of the Roy Search Model (RSM) with a different vector of variables $\mathbf{x}_{it} = \mathbf{q}_{it}$ that describe the mean of wages and profits. Sample size are 4,590 Colombian workers observed for two adjacent periods in either wage-employment or self-employment.

4. There is still a positive association between the measure and the probability of leaving to wage-employment, and a negative association with firm growth. The point estimates are much larger, but with much bigger standard errors, which render almost all of the coefficients highly insignificant.

One possible explanation for this difference lies in the fact that the range of the predicted $\Pr(w_{i2} > \pi_{i2})$ is much smaller than for the baseline RSM formulation. Whereas the baseline attractiveness measure covered the entire probability space between zero and one, the measure constructed with the alternative model only covers the range $[0.58, 0.77]$. The distribution is graphed in Appendix C.4 A main driver seems to be low estimated autocorrelation of entrepreneurial profits ($\rho_\nu = 0.071$). This implies that con-

TABLE 6
SENSITIVITY OF OUTCOME REGRESSIONS TO RSM FORMULATIONS

Dependent Variable	Earnings: Alternative RSM					
	Leave to WE			Firm Growth dummy		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Pr(w_{i2} > \pi_{i2})$	0.183 (0.265) [0.855]	0.489 (0.569) [0.391]	-0.080 (0.432) [0.852]	-0.495 * (0.288) [0.086]	-0.300 (0.500) [0.519]	-0.160 (0.500) [0.749]
City-year dummies		✓	✓		✓	✓
Gender			✓			✓
Years of Schooling			✓			✓
Exp. Quadratic			✓			✓
Observations	1,623	1,623	1,623	1,316	1,316	1,316
R ²	0.001	0.012	0.063	0.003	0.024	0.004
Dependent Variable	Earnings: OLS					
	Leave to WE			Firm Growth dummy		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Pr(w_{i1} > \pi_{i1})$	0.025 (0.032) [0.440]	0.035 (0.032) [0.272]	0.033 (0.031) [0.288]	-0.076 * (0.039) [0.053]	-0.069 * (0.039) [0.075]	-0.066 * (0.039) [0.090]
City-year dummies		✓	✓		✓	✓
Gender			✓			✓
Years of Schooling			✓			✓
Exp. Quadratic			✓			✓
Observations	1,623	1,623	1,623	1,316	1,316	1,316
R ²	0.001	0.012	0.063	0.004	0.025	0.04

Linear Probability Model regressions of binary outcomes, using different formulations of the RSM discussed in the main text. All RSM estimations use the most saturated specification corresponding to the explanatory variables in model (4) of (3) and (5). The explanatory variable $\Pr(w_{i1} > \pi_{i1})$ constructed according to equation (3.1), with RSM-specific parameters. Each panel corresponds to a different first step specification, as detailed in Table 3. Bootstrapped standard errors, stratified by step, with 1,000 replications, in parentheses. P-values in square brackets.

ditioning on the residual profits of period 1, ν_{i1} , does not carry much information content for future profits and wages. With the variables rescaled to a different range of units, the coefficients of this regression and Table 4 might not be immediately comparable.

As an alternative sensitivity check, I estimate the wage equation using OLS. In the context of the RSM, OLS identifies the wage coefficients and the distribution of wage residuals if the earnings distribution F is independent

across occupations and across time. Under normality, this corresponds to the assumption $\rho = \rho_\epsilon = \rho_\nu = 0$. The lower panel of Table 6 presents the results. The estimates are quite similar to the ones in Table 4.

3.6 Conclusion

Transitions between occupations are driven by two key forces: comparative advantage that creates a differential in earnings, and search frictions that constrain choice, or give rise to random transitions. In this chapter, I consider the identification of a basic Roy Search Model that includes these two forces. I show which parts of the structure of the earnings and transition process are identified under fairly weak assumptions, and for which parts stronger assumptions are required.

I illustrate the methodology by estimating a Roy Search Model for the choice between wage-employment and self-employment on panel data in Colombia. I find that both selection on comparative advantage and search frictions are important determinants of transitions. My estimates give rise to a straightforward measure of the attractiveness of wage-employment for entrepreneurs. Correlations of this measure and occupational and entrepreneurial choices of Colombian entrepreneurs are consistent with the lure of wage work theory.

The Roy Search Model is an important addition to the library of ‘simple’ selection models that abstain from a fully structural analysis of the worker’s choice problem. It overcomes some limitations to the application of existing approaches. It does come with limitations itself. First, exclusion restrictions in earnings equations can be very challenging to find in labour force data; absent policy changes or other natural experiments. Second, as I show, an

alternative approach relies in turn on fairly strong assumptions about the economic structure. We need to assume that separations from wage jobs are essentially random. In reality, they might partly reflect endogenous quits. It would be useful to extend the discussion of identification to accommodate this possibility. Third, I assume that the search friction parameters are constant and common across individuals. A direction for further work would be to relax these restrictions, and allow for search frictions to be a function of observable and unobservable heterogeneity.

Appendix C

C.1 Derivation of Likelihood Function

The likelihood function is derived as follows. First, it should be noted that the likelihood is not separable over time, but it is separable between individuals. The likelihood for every individual, which comprises an observation in each of the two periods, is dependent on the combinations of states, or labour market sectors, the individual is observed in. With 2 periods and 2 states, there are 4 possible cases.

Case 1: SE in period 1, WE in period 2

With $\pi_{it} = \mathbf{z}'_{it}\gamma + \nu_{it}$, the density of SE income is $f(\nu_{i1})$. All densities are conditional (on x_{it}), with the conditioning suppressed for easier reading. Similarly, with $w_{it} = \mathbf{x}'_{it}\beta + \epsilon_{it}$, and noting that the earnings observation in the second period is conditional on the error term of the first period, the density of wage in period 2 is $f(\epsilon_{i2}|\nu_{i1})$. Together, the product $f(\nu_{i1}) \cdot f(\epsilon_{i2}|\nu_{i1})$ is the joint density $f(\nu_{i1}, \epsilon_{i2})$ by Bayes' Rule. The probability mass of observing SE in period 1 is $(1 - \omega)$ and the probability distribution for wages in period 1 conditional on residual profits ν_{i1} is $\lambda \cdot \Pr(w_{i1} > \pi_{i1}|\nu_{i1})$. This implies that the joint likelihood for observing occupations and earnings of individual i is $(1 - \omega) \cdot f(\nu_{i1}) \cdot \lambda \cdot \Pr(w_{i1} > \pi_{i1}|\nu_{i1}) \cdot f(\epsilon_{i2}|\nu_{i1})$. It is manipulated in the following manner:

$$\begin{aligned}
 L_i^1 &= (1 - \omega) \cdot f(\nu_{i1}) \cdot \lambda \cdot \Pr(w_{i1} > \pi_{i1}|\nu_{i1}) \cdot f(\epsilon_{i2}|\nu_{i1}) \\
 &= f(\nu_{i1}) \cdot f(\epsilon_{i2}|\nu_{i1}) \cdot \Pr(\mathbf{x}'_{i1}\beta + \epsilon_{i1} > \mathbf{z}'_{i1}\gamma + \nu_{i1}|\nu_{i1}) \cdot \lambda(1 - \omega) \\
 &= f(\nu_{i1}, \epsilon_{i2}) \cdot \Pr(-\epsilon_{i1} < \mathbf{x}'_{i1}\beta - \mathbf{z}'_{i1}\gamma - \nu_{i1}|\nu_{i1}) \cdot \lambda(1 - \omega) \\
 &= f(\nu_{i1}, \epsilon_{i2}) \cdot \Pr\left(-\epsilon_{i1} + \rho \frac{\sigma_\epsilon}{\sigma_\nu} \nu_{i1} < \mathbf{x}'_{i1}\beta - \mathbf{z}'_{i1}\gamma - \nu_{i1} + \rho \frac{\sigma_\epsilon}{\sigma_\nu} \cdot \nu_{i1} \middle| \nu_{i1}\right) \lambda(1 - \omega) \\
 &= f(\nu_{i1}, \epsilon_{i2}) \Pr\left(\frac{-\epsilon_{i1} + \rho \frac{\sigma_\epsilon}{\sigma_\nu} \nu_{i1}}{\sigma_\epsilon \sqrt{1 - \rho^2}} < \frac{\mathbf{x}'_{i1}\beta - \pi_{i1} + \left(\rho \frac{\sigma_\epsilon}{\sigma_\nu}\right) \cdot \nu_{i1}}{\sigma_\epsilon \sqrt{1 - \rho^2}} \middle| \nu_{i1}\right) \lambda(1 - \omega)
 \end{aligned}$$

C.1. DERIVATION OF LIKELIHOOD FUNCTION

Under the normality Assumption 3.4, the conditional density of $\epsilon_{i1}|v_{i1}$ is normal with mean

$E(\epsilon_{i2}|v_{i1}) = \rho \frac{\sigma_\epsilon}{\sigma_\nu} \nu_{i1}$ and variance $Var(\epsilon_{i2}|v_{i1}) = \sigma_\epsilon^2 (1 - \rho^2)$. The variable

$$z = \frac{-\epsilon_{i1} + \rho \frac{\sigma_\epsilon}{\sigma_\nu} \nu_{i1}}{\sigma_\epsilon \sqrt{1 - \rho^2}} \sim N(0, 1)$$

then follows a standard normal distribution. Furthermore, the marginal bivariate distribution of ν_{i1} and ϵ_{i2} is a bivariate normal distribution. I denote the joint distribution of

standardized normal variables $z_1 = \frac{x_1 - \mu_1}{\sigma_1}$ and $z_2 = \frac{x_2 - \mu_2}{\sigma_2}$ with correlation $\tilde{\rho}$ as

$$\phi(z_1, z_2; \tilde{\rho}) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1 - \tilde{\rho}^2}} \exp\left\{-\frac{1}{2(1 - \tilde{\rho}^2)} [z_1^2 - 2\tilde{\rho} \cdot z_1 z_2 + z_2^2]\right\}.$$

The likelihood under normality in Case 1 is then

$$L_i^1 = \phi\left(\frac{\nu_{i1}}{\sigma_\nu}, \frac{\epsilon_{i2}}{\sigma_\epsilon}; \rho_\epsilon \cdot \rho\right) \cdot \Phi\left(\frac{\mathbf{x}_{i1}'\beta - \pi_{i1} + \rho \frac{\sigma_\epsilon}{\sigma_\nu} \cdot \nu_{i1}}{\sigma_\epsilon \sqrt{1 - \rho^2}}\right) \cdot \lambda(1 - \omega)$$

noting that $corr(\nu_{i1}, \epsilon_{i2}) = \rho_\epsilon \cdot \rho$.

Case 2: SE in period 1 and 2

The reasoning follows exactly Case 1. The density of SE income in period 1 is $f(\nu_{i1})$ and in period 2 $f(\nu_{i2}|\nu_{i1})$, and the joint density $f(\nu_{i1}, \nu_{i2})$. The probability mass of observing SE in period 1 is $(1 - \omega)$, the probability mass of observing SE in period 2 conditional on SE in period 1 and a residual ν_{i1} is $1 - \lambda \cdot \Pr(w_{i1} > \pi_{i1}|\nu_{i1}) = (1 - \lambda) + \lambda \cdot (1 - \Pr(w_{i1} > \pi_{i1}|\nu_{i1}))$. Then, the derivation for the likelihood of an individual in Case 2 is

$$\begin{aligned} L_i^2 &= (1 - \omega) \cdot f(\nu_{i1}) \cdot [(1 - \lambda) + \lambda \cdot (1 - \Pr(w_{i2} > \pi_{i1}|\nu_{i1}))] \cdot f(\nu_{i2}|\nu_{i1}) \\ &= f(\nu_{i1}, \nu_{i2}) \left[(1 - \lambda) + \lambda \Pr\left(\frac{-\epsilon_{i1} + \rho \frac{\sigma_\epsilon}{\sigma_\nu}}{\sigma_\epsilon \sqrt{1 - \rho^2}} < \frac{\mathbf{x}_{i1}'\beta - \pi_{i1} + \rho \frac{\sigma_\epsilon}{\sigma_\nu} \nu_{i1}}{\sigma_\epsilon \sqrt{1 - \rho^2}} \middle| \nu_{i1} \right) \right] (1 - \omega) \\ &= \phi\left(\frac{\nu_{i1}}{\sigma_\nu}, \frac{\nu_{i2}}{\sigma_\nu}; \rho_\nu\right) \left[(1 - \lambda) + \lambda \cdot \Phi\left(-\left(\frac{\mathbf{x}_{i1}'\beta + \rho \frac{\sigma_\epsilon}{\sigma_\nu} \cdot \nu_{i1} - \pi_{i1}}{\sigma_\epsilon \sqrt{1 - \rho^2}}\right)\right) \right] (1 - \omega). \end{aligned}$$

Case 3: WE in period 1 and 2

The density of the wage in period 1 is $f(\epsilon_{i1})$ and in period 2 $f(\epsilon_{i2}|\epsilon_{i1})$, jointly $f(\epsilon_{i1}, \epsilon_{i2})$.

The probability mass of observing WE in period 1 is ω , the probability mass of observing

WE in period 2 conditional on WE in period 1 and a residual ϵ_{i1} is $(1 - \delta)$. The likelihood of an individual in Case 3 is

$$\begin{aligned} L_i^3 &= \omega \cdot f(\epsilon_{i1}) \cdot [1 - \delta] \cdot f(\epsilon_{i2}|\epsilon_{i1}) \\ &= f(\epsilon_{i1}, \epsilon_{i2}) \cdot [1 - \delta] \cdot \omega \\ &= \phi\left(\frac{\epsilon_{i1}}{\sigma_\epsilon}, \frac{\epsilon_{i2}}{\sigma_\epsilon}; \rho_\epsilon\right) \cdot [1 - \delta] \cdot \omega \end{aligned}$$

Case 4: WE in period 1, and SE in period 2

The density of the wage in period 1 is $f(\epsilon_{i1})$ and the density of income in period 2 $f(\nu_{i2}|\epsilon_{i1})$, jointly $f(\nu_{i1}, \epsilon_{i2})$. The probability mass of observing WE in period 1 is ω , the probability mass of observing SE in period 2 conditional on WE in period 1 and a residual ϵ_{i1} is δ . The likelihood of an individual in Case 4 is

$$\begin{aligned} L_i^4 &= \omega \cdot f(\epsilon_{i1}) \cdot \delta \cdot f(\nu_{i2}|\epsilon_{i1}) \\ &= f(\epsilon_{i1}, \nu_{i2}) \cdot \delta \cdot \omega \\ &= \phi\left(\frac{\epsilon_{i1}}{\sigma_\epsilon}, \frac{\nu_{i2}}{\sigma_\nu}; \rho_\nu \cdot \rho\right) \cdot \delta \cdot \omega \end{aligned}$$

Joint Likelihood

Under assumption of independence across individuals, the joint likelihood is the product of the individual likelihoods:

$$\mathcal{L}(\beta, \gamma, \delta, \lambda, \omega, \sigma_\epsilon, \sigma_\nu, \rho, \rho_\epsilon, \rho_\nu) = \prod_{i=1}^N [L_i^1]^{D_i^1} [L_i^2]^{D_i^2} [L_i^3]^{D_i^3} [L_i^4]^{D_i^4}$$

D_i^j is a dummy variable indicating the case. The log-likelihood is then obtained by taking the natural logarithm of the respective expressions.

C.2 Data definitions and summary statistics

Table A1: DEFINITIONS OF ALL VARIABLES USED

Variable	Description
A: Main Variables	
Male	Dummy equal one if male, zero if female.
Age	Age of the individual in years.
Years of Schooling	Years of formal education. Constructed by combining information on highest level of education attended, and highest grade completed.
Experience	Labour market experience in years, constructed as $Age - Years\ of\ Schooling - 6$.
Monthly income	Real monthly income (in 2010 COP, inflated with CPI) for self-employed; real monthly wage for wage-employed, including overtime and extra pay, excluding fringe benefits
Firm Size	<i>(SE only)</i> Number of employees that are not household members.
Registered	<i>(SE only)</i> Dummy equal one if business registered with any authority or official organism.
C: Outcome Variables	
Leave WE to SE	<i>(WE only)</i> Dummy equal one if a wage-employed individual transits to self-employment in the next period, zero if stays wage-employed.
Leave SE to WE	<i>(SE only)</i> Dummy equal one if a self-employed individual transits to wage-employment in the next period, zero if stays wage-employed.
Grow Firm	<i>(Non-switching SE only)</i> Dummy equal one if self-employed individual adds employees to firm, zero otherwise.

C.3. LIKELIHOOD FUNCTION OF THE ALTERNATIVE MODEL

TABLE A2
SUMMARY STATISTICS OF VARIABLES USED

	Wage-employed			Self-employed		
	N	Mean	Stdev	N	Mean	Stdev
Panel A: RSM sample						
Male	5,806	0.57	.	3,374	0.60	.
Age	5,806	37.37	11.48	3,374	44.76	12.93
Years of Schooling	5,806	10.76	4.08	3,374	9.47	4.41
Experience	5,806	20.61	13.03	3,374	29.29	14.65
Monthly Income †	5,806	777.98	620.22	3,374	656.41	794.05
Monthly Income (log)	5,806	13.35	0.63	3,374	12.27	3.12
Firm Size	.	.	.	3,368	1.28	10.73
Registered	.	.	.	3,297	0.29	.
Panel B: Outcome regression sample						
Leave WE to SE	.	.	.	1,623	0.19	.
Leave SE to WE	2,967	0.15	.	.	.	.
Grow Firm	.	.	.	1,316	0.23	.

Source: Encuesta Social Longitudinal de Fedesarrollo (ESLF), 2007-2010. Standard deviations of binary variables omitted.

† In thousands of Colombian Pesos (COP) of 2010. In 2010 the exchange rate was COP 1,898 for one US dollar. Standard Deviation of binary variables not reported.

C.3 Likelihood Function of the alternative model

This section states the likelihood function of the alternative RSM estimated in section 3.5. Here, selection into wage-employment is based on a comparison of expected earnings in period 2.

Case 1: SE in period 1, WE in period 2

$$\begin{aligned}
 L_i^1 &= (1 - \omega) \cdot f(\nu_{i1}) \cdot \lambda \cdot \Pr(w_{i2} > \pi_{i2} | \nu_{i1}) \cdot f(\epsilon_{i1} | \nu_{i1}) \\
 &= (1 - \omega) \cdot \phi\left(\frac{\nu_{i1}}{\sigma_\nu}, \frac{\epsilon_{i2}}{\sigma_\epsilon}; \rho_\epsilon \cdot \rho\right) \cdot \lambda \Phi\left(\frac{\mathbf{x}_{i2}'\boldsymbol{\beta} - \mathbf{q}_{i2}'\boldsymbol{\gamma} + \left(\rho \rho_\epsilon \frac{\sigma_\epsilon}{\sigma_\nu} - \rho_\nu\right) \cdot \nu_{i1}}{\sqrt{\sigma_\epsilon^2(1 - \rho_\epsilon^2 \rho^2) - 2\rho\sigma_\epsilon\sigma_\nu(1 - \rho_\epsilon\rho_\nu) + \sigma_\nu^2(1 - \rho_\nu^2)}}\right)
 \end{aligned}$$

Case 2: SE in period 1 and 2

$$\begin{aligned}
 L_i^2 &= (1 - \omega) \cdot f(\nu_{i1}) \cdot (1 - \lambda \cdot \Pr(w_{i2} > \pi_{i2} | \nu_{i1})) \cdot f(\nu_{i1} | \nu_{i1}) \\
 &= (1 - \omega) \cdot \phi\left(\frac{\nu_{i1}}{\sigma_\nu}, \frac{\nu_{i2}}{\sigma_\nu}; \rho_\nu\right) \cdot \left(1 - \lambda \Phi\left(\frac{\mathbf{x}_{i2}'\boldsymbol{\beta} - \mathbf{q}_{i2}'\boldsymbol{\gamma} + \left(\rho \rho_\epsilon \frac{\sigma_\epsilon}{\sigma_\nu} - \rho_\nu\right) \cdot \nu_{i1}}{\sqrt{\sigma_\epsilon^2(1 - \rho_\epsilon^2 \rho^2) - 2\rho\sigma_\epsilon\sigma_\nu(1 - \rho_\epsilon\rho_\nu) + \sigma_\nu^2(1 - \rho_\nu^2)}}\right)\right)
 \end{aligned}$$

Case 3: WE in period 1 and 2

$$L_i^3 = \omega \cdot f(\epsilon_{i1}) \cdot [1 - \delta] \cdot f(\epsilon_{i2} | \epsilon_{i1}) = \phi\left(\frac{\epsilon_{i1}}{\sigma_\epsilon}, \frac{\epsilon_{i2}}{\sigma_\epsilon}; \rho_\epsilon\right) \cdot [1 - \delta] \cdot \omega$$


Case 4: WE in period 1, and SE in period 2

$$L_i^4 = \omega \cdot f(\epsilon_{i1}) \cdot \delta \cdot f(\nu_{i2} | \epsilon_{i1}) = \phi\left(\frac{\epsilon_{i1}}{\sigma_\epsilon}, \frac{\nu_{i2}}{\sigma_\nu}; \rho_\nu \cdot \rho\right) \cdot \delta \cdot \omega$$

C.4 Additional tables and figures


FIGURE A1

ATTRACTIVENESS OF WAGE EMPLOYMENT: RSM COMPARING EXPECTED EARNINGS



Source: Attractiveness of wage-employment for each period 1 entrepreneur in the sample ($N = 1,623$), constructed according to equation (3.1) for the alternative RSM formulation. Parameter estimates come from the different models that are presented in Table 5.

FIGURE A2
ATTRACTIVENESS OF WAGE EMPLOYMENT ASSUMING INDEPENDENCE OF EARNINGS
DISTRIBUTION



Source: Attractiveness of wage-employment for each period 1 entrepreneur in the sample ($N = 1,623$), constructed according to equation (3.1) under $\rho = \rho_\epsilon = \rho_\nu = 1$. Parameter estimates come from an OLS estimation of the different conditional mean models in Table 3.

Chapter 4

Self-Employment, Capital and Job Search: Evidence from a Structural Model

MICHAEL KOELLE AND SIMON QUINN

4.1 Introduction

In recent years, academic literature has made enormous advances in understanding earnings and employment dynamics in developing countries. A big part of this is due to our improved understanding of constraints faced by the self-employed, who constitute the majority of workers in developing countries (Banerjee, Karlan, and Zinman (2015), Blattman, Fiala, and Martinez (2014), De Mel, McKenzie, and Woodruff (2012)). We are also rapidly improving our understanding of what matters for firms to grow and to provide jobs (Bloom, Eifert, Mahajan, McKenzie, and Roberts (2013), Atkin, Chaudhry, Chaudry, Khandelwal, and Verhoogen (2017)) and what encourages potential entrepreneurs to start successful firms in the first place (McKenzie (2017), Fafchamps and Quinn (2016)). A long literature in development economics suggests that the process of matching workers with eventual jobs is ridden with important frictions, and new empirical evidence

seems to support this view (Lewis (1954), Harris and Todaro (1970), Groh, Krishnan, McKenzie, and Vishwanath (2016), Franklin (2016), Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn (2016)). However, we still know very little of which of these aspects are complements or substitutes when they come together as individuals choose between jobs and self-employment, when they acquire skills, and invest into search for opportunities – in short, when people in developing countries have a career. Our objective is to fill this gap.

In this paper, we provide a unified framework to think about career paths in urban labour markets of developing countries. We place a special emphasis on understanding the role of self-employment, and on labour market frictions. We do so by formulating a discrete dynamic choice model of career choices, job search as a costly process, and savings and capital accumulation over the life-cycle of individuals. We structurally estimate this model on a very rich long labour force panel survey from urban Ghana, the 2005-2013 Ghana Urban Household Panel Survey (GUHPS) using indirect inference.

In particular, we formulate a discrete dynamic model of occupational choice, investment, and job search. Our model has three sectors – self-employment, wage employment, and non-employment – and is stylised enough to be tractable. At the same time, the model is rich enough to capture a crucial distinction between self-employment and wage employment: self-employment depends on having a business idea and access to inputs such as capital, and wage employment depends on having found a job with an employer. We explicitly model how individuals accumulate assets, and how capital increases the outputs and profits of businesses. We do so by expressing business prof-

its as an outcome of a production process. We also model a costly search process, where job seekers can invest into search effort that might increase their chances of finding a job (and possibly of opening up a business, too).

We incorporate rich patterns of individual heterogeneity in our model, in three ways. First, some of the heterogeneity arises within the model itself. Experience and tenure – state variables of the model – affect wages through accumulation of human capital, and business profits through their effect on productivity. Experience and tenure also affect job offers and losses. Capital enables individuals to invest into business opportunities and into costly job search. Second, we let pre-determined characteristics affect relative payoffs, costs, and opportunities agents face. Specifically, we allow all parameters to vary by education and estimate the model separately for men with low and with high levels of education. Third, we allow for unobserved heterogeneity to shape absolute and comparative advantage in wage and self-employment.

By modelling the choices that shape the career paths of very heterogeneous agents, we aim to uncover potentially generalisable mechanisms that operate in labour markets with a large self-employment sector. We provide quantitative evidence for one such market, the urban labour market in Ghana. We believe that this is useful for our understanding of *how* labour markets work in developing countries; where large self-employment sectors usually are found. This fills an important gap in the existing body of research (Deaton, 2010a, 2010b). Moreover, we believe that this is complementary to recent empirical literature which, taken together, has produced a variety of stylised facts. Our model is useful for connecting these observations, and for unifying our understanding of the mechanisms that could lead to them.

We find that the model is able to reproduce well the overall patterns in the data, including the distributions of employment, earnings, and capital stock in microenterprises. We also obtain a reasonable approximation to moments that our estimation does not directly target, such as the different shapes of capital distributions across self-employment, wage employment, and non-employment. We replicate two actual field experiments to validate our model: a randomised control trial of capital grants to microenterprises in Ghana, and a job search subsidy in a different but comparable setting. Through the lens of our model, we examine whether a lack of opportunities for entrepreneurs can explain differences among microenterprises, and whether it constrains later career outcomes of individuals. Our results suggest that while some groups are significantly affected by a lack of opportunities – such as individuals who involuntarily lost a wage job – a stark categorisation of entrepreneurs into distinct types is probably not warranted. Finally, we simulate a job guarantee for the unemployed and find that this can provide a small stepping stone for future self-employment.

4.1.1 Related literature

Capital and job search constraints in developing countries. A literature on capital grants for microenterprises in developing countries (De Mel, McKenzie, and Woodruff (2008, 2012); Fafchamps, McKenzie, Quinn, and Woodruff (2014)) has found that capital shocks have effects that are large, lasting, and heterogeneous. Capital grants to aspiring entrepreneurs can encourage entry to self-employment (Blattman, Fiala, and Martinez (2014), Fafchamps and Quinn (2016), McKenzie (2017)). The literature describing occupational choice finds that people can choose self-employment both

through high potential and through lack of good alternatives (Banerjee and Duflo (2011), Günther and Launov (2012), Fafchamps and Woodruff (2016), Haywood and Falco (2016)). A growing literature (Franklin (2016), Groh, Krishnan, McKenzie, and Vishwanath (2016), Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn (2016)) finds new empirical evidence on the magnitude of frictions, and the costs of job search and matching in developing countries. This fills a longstanding gap in evidence than can speak to a classical theoretical literature on market frictions and labour market structure (Lewis (1954), Harris and Todaro (1970), Fields (1975)).¹ Our model provides a unified framework to all of these themes, and how these different aspects might be complements or substitutes for earnings and employment dynamics.

Dynamic models of job search. We contribute to the literature on dynamic models of job search with a model where success in self-employment not only depends on skills, but also on capital. We do so in a framework that incorporates asset accumulation under credit constraints. Our approach introduces ideas from the classical literature on selection into entrepreneurship (Evans and Jovanovic (1989), Evans and Leighton (1989), Blanchflower and Oswald (1998)) to a framework of occupational choice with job search (Eckstein and Wolpin (1990) Blundell, Costa Dias, Meghir, and Shaw (2016) Adda, Dustmann, Meghir, and Robin (2013)) that itself builds on discrete dynamic models of occupational choice (Eckstein and Wolpin (1989), Keane and Wolpin (1997)).² We also extend and develop a class of models that look at the combined implications of job search and savings (Rendon (2006), Lise (2013)). In our model, asset accumulation through

¹ A notable early exception is Magnac (1991).

² Eckstein and Van den Berg (2007) review the empirical literature on labour search.

saving serves three purposes: to smooth out consumption when income fluctuates, to finance costly search, and to build up capital stock for a business start-up. These extensions allow us to take this class of models to studying central aspects of urban labour market in developing countries. In contrast to other recent work on entrepreneurship in a framework of lifecycle occupational choices (Humphries (2017), Dillon and Stanton (2017)), we allow for market frictions such as capital constraints and search frictions to play a key role in constraining choices. Alleviating these frictions has been the focus of the large literature on self-employment in developing countries that motivates our work.

Structural development economics. Our work also contributes to a stream of literature in development economics which speaks to the symbiosis of field experiments and theory-based approaches. Like the seminal applications by Todd and Wolpin (2006) and Attanasio, Meghir, and Santiago (2011), we validate our structural model with a randomised control trial. Our work also contributes to recent attempts at characterising the behaviour of microenterprises with a structural model (e.g. Keniston (2011), Banerjee, Breza, Duflo, and Kinnan (2015), Cohen (2016)). We provide a unified framework which structurally links the decisions of microenterprises with occupational choices. At the same time, this introduces notions of capital accumulation and production technologies to structural models of self-employment in developing countries (e.g. Lopez-Garcia (2013), Narita (2014)). With a focus on analysing the implications of individual choice behaviour for informality, they do not include capital or production technologies as a factor driving self-employment. More generally, our paper relates to the structural analysis of firms in developing countries (Meghir, Narita, and Robin (2015), Ulyssea (2017)).

This chapter is structured as follows. In section 4.2 we present the data, and descriptive evidence on career paths in urban Ghana. We then describe the model in section 4.3. In section 4.4 we discuss identification of the model. We present estimation results in section 4.5. In section 4.6 we conduct policy simulations to interpret various questions of interest to labour markets in developing countries. Finally we conclude in section 4.7.

4.2 Data and descriptive evidence

We use data from the Ghana Urban Household Panel Survey (GUHPS), the longest annual panel survey of workers in Sub-Saharan Africa. The GUHPS sample was drawn using the 2000 Population and Housing Census, and is designed to be representative of the cities of Accra, Kumasi, Takoradi and Cape Coast. Overall the data span nine waves from 2004-2013. The survey interviews all labour force participants in the household, and puts effort into tracking individuals over time, including when they leave their original household.

We retain data for individuals of working age, which we define as 15-64 years of age, from seven waves between 2005 and 2013 that allow us to construct a consistent panel dataset. We focus on a fairly homogenous sub-population: men without an upper secondary (Senior High School) qualification.³ We obtain an unbalanced panel of 3,307 observations that are made up of 1,255 individuals. We have information on their sector of employment, occupation, incomes, labour market experience, assets, demographic characteristics and if they are unemployed, whether they undertook

³ It is beyond the scope of this work to do justice to modelling female labour supply in a developing country. We divide the sample by education; and eventually want to estimate the model also for men with higher levels of education.

any steps towards finding employment in either a wage job or by becoming self-employed.

The survey contains separate modules for eliciting the occupations and weekly incomes of self-employed and wage-employed workers. Data collection using hand-held devices allows for automated on-the-spot consistency checks of profits, incomes, and expenses. This improves the accuracy of measuring small business incomes in developing countries (Fafchamps, McKenzie, Quinn, and Woodruff (2012)). The survey is quite unique for a representative labour force survey in that it allows for accurate measurement of not only incomes but also capital stock of the self-employed. We can construct capital stock from a list of questions on ownership and value of common business and financial assets.

Respondents had the option to report more than one job, but few of them actually did.⁴ We focus on the main occupation. Ideally, we would observe the actual years of labour market experience for every individual. Unfortunately, for most respondents we do not observe this for the years prior to the survey. We reconstruct potential experience by counting the years that have passed since an individual obtained their highest educational qualification. We do observe the starting year and month of the current employment activity, although this date seems to be measured with a lot of noise, especially for longer recall periods. We deal with this in two ways. Firstly, we supplement the reported tenure by a constructed variable whenever we observe the start of a new employment spell in the panel data. Secondly, we truncate the tenure variable, and use the information in our model only in

⁴ About 5 percent of all wage workers, and 3 percent of all self-employed, report having any secondary activity.


the form of a binary variable that indicates more than two years of tenure. Given that tenure-wage profiles typically are steepest in the first two years in an occupation (Keane and Wolpin (1997)), we think this is a reasonable approximation. We treat the length of an unemployment spell in a parallel manner.⁵

A number of descriptive facts from the data motivate our thinking about wages, incomes, and employment in a developing country like Ghana in terms of a dynamic life-cycle framework. Figure 1 plots the distribution of employment over the life-cycle. We make a number of observations from this. Entry into employment takes a long time. Even 10 years after entering the labour force upon completion of their education, around half of all men in our sample are not employed; although, eventually, around 80% find themselves in employment. At first glance, delayed entry into self-employment seems to be the most prominent factor that might explain this. Entry into self-employment occurs very slowly over the life-cycle. Even though almost half of them eventually become self-employed, fewer than 20% of our sample are self-employed ten years after they enter the labour market. In fact, young workers are much more likely to be found in wage employment than running their own business. After around 12 years of potential experience, wage employment rates plateau at around 45 percent.

The employment transition matrices in Table 1 corroborate these observations. It is much harder for younger individuals to break into self-employment, than it is for older individuals. For instance, only 9% of the non-employed with up to 20 years of potential experience – that is up to around 35-40 years

⁵ Appendix D.1 provides the details of how each of our variables was constructed from the survey data.

FIGURE 1
EMPLOYMENT OVER THE LIFE-CYCLE



Source: GUHPS 2004-2013. Sample restricted to men with less than Senior High School (SHS) education. SE = self-employed; WE = wage-employed. Local mean regressions with optimal plug-in bandwidth.

of age – enter self-employment every year; whereas this rate is three times higher (at 26 %) for older non-employed workers. Once employed, there is a similar propensity of old and young workers to keep their jobs, and only a slightly higher propensity for older workers to stay in a business. Across the board the rate of churning is high: only around two thirds stay in a given type of employment in two consecutive years.

Turning to earnings in Figure 3, profits in self-employment are consistently higher than wages along the life-cycle, by about half of a log-point in the second half of the life-cycle. Figure 2 shows the cross-sectional distribution of profits and wages. Profits are more dispersed and there is a considerable share of the self-employed who earn profits that are either lower or higher

TABLE 1
EMPLOYMENT TRANSITION MATRIX

	Employment at t+1			
Employment at t	SE	WE	NE	Total
Fewer than 20 years of experience				
SE	63.1% (106)	22.6% (38)	14.3% (24)	100% (168)
WE	8.9% (25)	68.8% (194)	22.3% (63)	100% (282)
NE	8.5% (33)	27.2% (106)	64.3% (251)	100% (390)
More than 20 years of experience				
SE	67.6% (230)	22.9% (78)	9.4% (32)	100% (340)
WE	18.3% (59)	69.3% (224)	12.4% (40)	100% (323)
NE	26.1% (36)	29.7% (41)	44.2% (61)	100% (138)


Source: GHUPS. Employment transitions for male respondents aged 15-65 without SHS degree and with less (top panel) and more (bottom panel) than 20 years of potential experience and non-missing employment status in adjacent periods. SE = self-employed; WE = wage employed; NE = not employed. Number of transitions in sample in brackets.

than the distribution of wages in the market.

4.3 Model


In order to gain an understanding of the individual decisions and circumstances that can explain these observations of related to entry, transitions, and occupational choice we turn to formulating a dynamic model of employment choice in a developing country. We particularly focus on the importance of self-employment and costly search; and we allow for frictions that govern job offers and losses, and business opportunities and closures. Our model is a partial labour market equilibrium: we take the wage of-

FIGURE 2
INCOME OVER THE LIFE-CYCLE



Source: GUHPS. Sample restricted to men with less than Senior High School (SHS) education. Local mean regressions with bandwidth 1.5.

FIGURE 3
CROSS-SECTIONAL DISTRIBUTION OF INCOME



Source: GUHPS. Sample restricted to men with less than Senior High School (SHS) education. Kernel density estimation with plug-in bandwidth.

fer and firm productivity distributions as given and do not model aggregate labour demand and supply. This in turn lets us allow for rich patterns of heterogeneity between individuals.

The model is set in discrete time. Every period corresponds to a year. The model starts at the point of entry into the labour market, and ends after 50 periods. We take education at the point of labour market entry as given (and do not model education choices). The labour market consists of three sectors of employment: self-employment, wage employment, and unemployment. The timing of the model in any period t is as follows:

1. Agents earn income from work (if self-employed or wage-employed) or support from other members of their household (if not employed). Income adds to the current wealth that individuals bring into the period.
2. Agents decide how much of their wealth to consume, and how much to save for the next period. At this stage, unemployed agents may also invest into costly job search effort. Agents then receive a utility from consumption.
3. At the end of a period, agents may lose their jobs if they are employed. If they do not lose their job, they may experience a shock to their earnings (for example, a productivity shock). They may also encounter employment opportunities in other sectors: a wage job, or a business opportunity. Every opportunity comes with a wage or productivity level draw from a distribution.
4. Agents then choose the best employment option from the choice menu available to them in t . The option to become non-employed at a fixed

level of income support is always available. In addition, the choice menu may contain one or both of:

- A job offer at a specified wage.
- A business opportunity with a specified level of productivity, and an income that is determined by productivity and the capital stock. Agents are credit-constrained and can only invest up to their accumulated wealth.

5. Agents choose the optimal employment option amongst the menu of choices available to them. They choose the option with the highest discounted lifetime value.
6. Agents then enter the next period $t + 1$, take up their chosen employment, earn experience and – if they stay employed in the same sector as in t – tenure. Tenure otherwise resets to zero.

This dynamic discrete choice model has four state variables: experience t , tenure x_{it} , capital k_{it} , and the current wage or productivity level ϵ_{it} . The first three of these state variables are observed by the researcher; the productivity level is observed only by agents in the model. We specify below how these state variables affect payoffs and other functions. Agents make three choices in this model: investment, effort, and sector of employment.

These choices are forward-looking, and we model agents as maximising expected utility under a known distribution of shocks. At the beginning of the period – before observing information about future income, or employment shocks – agents must decide on job search and on future capital. Investing into search effort changes the probability with which employment shocks occur, whereas investing into capital changes the value of the employment

options agents choose from. At the end of the period, all uncertainty about the *next* period is revealed, and agents choose from a deterministic menu of employment and incomes. Agents again take an expectation over option values in periods that are *further* into the future.

We need to parameterise the functions that appear in the model in order to be able to solve and simulate it. We make parametric choices about three types of functions in particular: wage and productivity distributions, search frictions, and utility. We plan to eventually estimate the model separately for workers with low and high levels of education, so all the parameters should be thought of as specific to a level of education. In our exploratory work we built a strong sense that the labour market, in particular the market for wage labour, is segmented by the level of education; and that the completion of senior high school (SHS) in particular provides the watershed that segments the markets. In this version, we restrict our attention to the larger subgroup: men without a SHS degree.

4.3.1 Wages and incomes

We parameterise the log of profits π_{it} , wages w_{it} , and unemployment support b_{it} as:

$$\begin{aligned} \ln \pi_{it} = & \gamma_0 + \gamma_1 \cdot t + \gamma_2 \cdot t^2 + \gamma_3 \cdot \mathbb{1}(x_{it} > 2) + \alpha \cdot \ln(k_{it}^*) \quad (4.1) \\ & + \mu_i^s + \epsilon_{1,it}; \end{aligned}$$

$$\ln w_{it} = \beta_0 + \beta_1 \cdot t + \beta_2 \cdot t^2 + \beta_3 \cdot \mathbb{1}(x_{it} > 2) + \mu_i^w + \epsilon_{2,it}; \quad (4.2)$$

$$\ln b_{it} = \bar{\mu}^n + \mu_i^n. \quad (4.3)$$

The log of earnings in wage and self-employment depends on a constant, a quadratic function of potential experience t (which is synonymous with the time periods of the model), a dummy whether occupation-specific tenure $x_{it} \in \{0, 1, 2, \dots\}$ is larger than two years. It is also a function of permanent heterogeneity and a time-varying error term. Business profits also depend on capital stock k_{it}^* , through a Cobb-Douglas form. Income support deviates from its mean $\bar{\mu}^n$ as a function of unobserved heterogeneity.

Each period, entrepreneurs make an optimal portfolio allocation between business capital and cash. That is, they optimally allocate a share ϕ of their given financial capital k_{it} to business purposes.⁶ Business capital yields a productive return and is subject to depreciation. Cash earns the real interest rate r (which in our application is set to minus the inflation rate):

$$\max_{\phi \in [0,1]} A_{it} \cdot (\phi \cdot k_{it})^\alpha - d \cdot \phi \cdot k_{it} + r \cdot (1 - \phi) \cdot k_{it} \quad (4.4)$$

Total factor productivity in the production function is given by

$$A_{it} = \exp(\gamma_0 + \gamma_1 \cdot t + \gamma_2 \cdot t^2 + \gamma_3 \cdot \mathbb{1}(x_{it} > 2) + \mu_i^s + \epsilon_{1,it}). \quad (4.5)$$

We can think of tenure, experience, permanent heterogeneity, and the residual $\epsilon_{2,it}$ as observable and unobservable productivity shifters for incomes from self-employment. Note that we assume that entrepreneurs choose capital in full knowledge of the contemporaneous productivity shock. The firm's optimally chosen capital stock, given initial capital and credit constraints, is

⁶ This modelling device allows us to have a single capital state variable, which makes the model computationally feasible. The limitation is that this precludes us from modelling a mechanism such as in Chapter 1, where business capital choices are influenced by the probability of leaving self-employment. This would require an additional continuous state variable. In our model, intertemporal capital choices are determined by expected net returns and consumption-smoothing motives.

given by

$$k_{it}^* = \phi^* \cdot k_{it} = \min \left\{ k_{it}, \left(\frac{A_{it} \cdot \alpha}{d + r} \right)^{\frac{1}{1-\alpha}} \right\} \quad (4.6)$$

We assume that the residuals of wages and productivity follow an AR(1)-process in each sector j :

$$\epsilon_{j,it} = \zeta_j \epsilon_{j,i,t-1} + \nu_{j,it} \quad \nu_{j,it} \sim \mathcal{N}(0, \sigma_j^2) \quad 0 \leq \zeta_j < 1 \quad (4.7)$$

On the one hand, autocorrelation allows earnings to have a persistent component. For instance, this could be thought of as a match-specific component of wages, or a persistent level of productivity of a firm. On the other hand, the innovations $\nu_{j,it}$ can accommodate for earnings risks within a given job, which is particularly high in developing countries. The variance of earnings in either sector is flat over the life-cycle in our dataset. We therefore restrict the residuals to follow a stationary process over time.

4.3.2 Permanent heterogeneity

The previous section introduced variables μ_i^s , μ_i^w and μ_i^n , representing permanent heterogeneity in comparative advantage, unobservable to the researcher. Specifically, this structure allows for comparative advantage in wage versus self-employment, as well as individual heterogeneity in the advantage of either employment type over non-employment. In particular, we use a trivariate normal random effects structure:

$$\begin{pmatrix} \mu_i^s \\ \mu_i^w \\ \mu_i^n \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_s^2 & \rho_{sw} \cdot \sigma_s \cdot \sigma_w & \rho_{sn} \cdot \sigma_s \cdot \sigma_n \\ \rho_{sw} \cdot \sigma_s \cdot \sigma_w & \sigma_w^2 & \rho_{wn} \cdot \sigma_w \cdot \sigma_n \\ \rho_{sn} \cdot \sigma_n \cdot \sigma_s & \rho_{wn} \cdot \sigma_n \cdot \sigma_w & \sigma_n^2 \end{pmatrix} \right).$$

This parametric structure allows us to capture a rich structure of unobserved heterogeneity in a feasible manner with only few parameters. Having a rich structure of permanent heterogeneity is crucial for modelling well the precautionary savings behaviour, a key mechanism for asset accumulation in our model. Agents accumulate savings to insure against anticipated fluctuations in earnings, and differences in their personal earnings potential in either sector play an important role in the extent of these fluctuations. In practice, we discretise the trivariate support using $4 \times 4 \times 3$ grid points.

4.3.3 Consumption

We model consumption in a way that allows for integrating intra-household resource allocation in a parsimonious, reduced-form way. We recognise that not all the income that is left over after savings are made can be used for private consumption. Some will be shared with other household members. We assume that individuals consume a constant, equivalised share of household consumption that depends only on the potential experience (a function of age) of the individual. We estimate consumption equivalence scales using data on household composition and size from the GUHPS. Household consumption is shared equally by the adult-equivalent dependents denoted d_t . Agents only care about their own private consumption when making consumption-savings choices.

We assume a constant relative risk aversion (CRRA) utility function in private consumption. We also allow for a non-pecuniary utility differential across the alternative occupations, captured by a non-pecuniary utility benefit from being self-employed (compared to both other sectors). Utility is

defined by:

$$U^j(C_{it}) = \frac{1}{1-\xi} \cdot \left(\frac{C_{it}^j}{d_t} \right)^{1-\xi} + \bar{u}_s,$$

where consumption in sector j is given by

$$C_{it}^{SE} = \pi(k_{it}^*, t, x_{it}, \epsilon_{1,it}; \mu_i^s) + (1+r) \cdot k_{it} - (r+d) \cdot k_{it}^* - k_{i,t+1}; \quad (4.8)$$

$$C_{it}^{WE} = w(t, x_{it}, \epsilon_{2,it}; \mu_i^w) + (1+r) \cdot k_{it} - k_{i,t+1}; \quad (4.9)$$

$$C_{it}^{NE} = b(\mu_{it}^n) - \kappa \cdot e_{it} + (1+r) \cdot k_{it} - k_{i,t+1}. \quad (4.10)$$

The notation is a reminder that incomes are a function of observed and unobserved state variables: experience, tenure, capital, occupation-specific productivity, and permanent heterogeneity. We assume that wealth can earn a real interest rate r when it is not invested in a business. Capital invested into a business earns returns and depreciates at rate d . Consumption then equals income minus savings. In addition, individuals who are not employed can put effort into active job search e_{it} at a monetary cost of κ . Search effort is a binary choice here, capturing the extensive margin of whether individuals take *some* means to find wage employment or start a business.

4.3.4 Search frictions

Our occupational choice model allows for search frictions. These search frictions change the choice menu that a particular individual faces. The parameters of the process depend on variables that are endogenous to the model. We model two kinds of search frictions:

- Involuntary separation from a wage contract or shutdown of a business (which we refer to as *job loss*);

- The possibility of receiving an offer for a wage job, or encountering a business opportunity (which we refer to as *job offer*).

Job losses effectively remove an item (the current sector of employment) from the choice menu, job offers add an item (a sector different to the current one) to the choice menu. In general, neither offers in all sectors of employment nor continuation in the current job are guaranteed. The option of non-employment, however, is always available as a choice. Since the frictions are determined by binary stochastic processes, their distribution is necessarily Bernoulli, governed by a single parameter, the success probability. We parameterise this success probability as a logit transformation of parameters. For the job loss probabilities, we write

$$\delta^{SE} = \Lambda \left[\delta_C^{SE} + \delta_{T1}^{SE} \cdot t + \delta_{T2}^{SE} \cdot t^2 + \delta_X^{SE} \cdot \mathbb{1}(x_{it} > 2) \right]; \quad (4.11)$$

$$\delta^{WE} = \Lambda \left[\delta_C^{WE} + \delta_{T1}^{WE} \cdot t + \delta_{T2}^{WE} \cdot t^2 + \delta_X^{WE} \cdot \mathbb{1}(x_{it} > 2) \right]. \quad (4.12)$$

Experience enters this logit index in a quadratic form. This allows both for firmer job attachment after an initial early-career experimentation phase, as well as possible exogenous transitions to early retirement. The logit transformation restricts the probabilities to be within the unit interval. We include tenure as a two-piece step function to allow for layoff risk to change (possibly decrease) with job attachment; for example by building match-specific human capital.

We express the job offer probabilities (from non-employment) as a function

of experience, tenure, and additionally effort

$$\begin{aligned}\lambda_{SE}^{NE} = & \Lambda \left[\lambda_{C,SE}^{NE} + \lambda_{T1,SE}^{NE} \cdot t + \lambda_{T2,SE}^{NE} \cdot t^2 + \lambda_{X,SE}^{NE} \cdot \mathbb{1}(x_{it} > 2) \right. \\ & \left. + \lambda_{E,SE}^{NE} \cdot e_{it} \right];\end{aligned}\quad (4.13)$$

$$\begin{aligned}\lambda_{WE}^{NE} = & \Lambda \left[\lambda_{C,WE}^{NE} + \lambda_{T1,WE}^{NE} \cdot t + \lambda_{T2,WE}^{NE} \cdot t^2 + \lambda_{X,WE}^{NE} \cdot \mathbb{1}(x_{it} > 2) \right. \\ & \left. + \lambda_{E,WE}^{NE} \cdot e_{it} \right].\end{aligned}\quad (4.14)$$

We expect job offer probabilities to increase for individuals who put in more effort and take specific action to search for work. We expect them to decrease with unemployment duration – ‘tenure’ in unemployment – as a result of possible hysteresis effects. We allow for a background probability of receiving job offers which is independent of search effort, and that potentially varies over the life-cycle. The quadratic form allows for both the gradual build-up of opportunities as individuals advance in their careers, as well as a reduction in new employment and business opportunities in the later stages of the working life.

Finally, our model also allows for the possibility of on-the-job search that results in transitions from one sector of employment to another. We write the probability of being offered to switch sectors as

$$\lambda_{WE}^{SE} = \Lambda \left[\lambda_{C,WE}^{SE} + \lambda_{T1,WE}^{SE} \cdot t + \lambda_{T2,WE}^{SE} \cdot t^2 \right]; \quad (4.15)$$

$$\lambda_{SE}^{WE} = \Lambda \left[\lambda_{C,SE}^{WE} + \lambda_{T1,SE}^{WE} \cdot t + \lambda_{T2,SE}^{WE} \cdot t^2 \right]. \quad (4.16)$$

We express the switching offer probability as a function of experience only. Ideally, we would express it also as a function of on-the-job search effort.

Unfortunately, we do not measure this consistently in our dataset.⁷

4.4 Identification

We identify the model using Indirect Inference (Gouriéroux, Monfort, and Renault (1993)). This involves the following steps:

1. We run a series of descriptive regressions on our dataset, and calculate a number of statistical moments (means, variances, and quantiles). We save b resulting ‘auxiliary parameters’ in a vector $\mathbf{m}$.⁸
2. For a $k \times 1$ parameter vector $\boldsymbol{\theta}$, we solve the model recursively by backwards induction and simulate a dataset of 15,000 imaginary Ghanaians. We solve the model separately for each cell in the permanent heterogeneity grid. When we simulate the model, we use the same distribution of initial conditions (capital and the initial state of employment at zero years of experience; roughly age 15) as in the actual dataset. We assume that these initial conditions are independent of permanent heterogeneity.
3. For the simulated data, we run the same set of descriptive regressions, and calculate the same moments. We save these ‘simulated auxiliary parameters’ in a vector $\mathbf{m}_S(\boldsymbol{\theta})$.
4. We update $\boldsymbol{\theta}$ in order to solve:

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} [\mathbf{m}_S(\boldsymbol{\theta}) - \mathbf{m}]' \cdot W(\mathbf{m}) \cdot [\mathbf{m}_S(\boldsymbol{\theta}) - \mathbf{m}].$$

⁷ We see no reason to expect the switching offer probability to depend on tenure. The *relative value* of an outside offer, and the probability of accepting it will of course probably depend on tenure; but not the fact of an offer being made or not.

⁸ We adjust for potential macroeconomic trends and local labour market shocks by partialling out year and city dummies in our descriptive regressions and statistics on the actual data. We do not include the adjustment dummy coefficients in $\mathbf{m}$.

4.4.1 Identification of structural parameters

The structural model is identified if the binding function that links the auxiliary parameters to the structural parameters of the model is a one-to-one mapping (Gouriéroux, Monfort, and Renault (1993)). Here we detail how our auxiliary parameter vector relates to the structural parameters of interest, and intuitively discuss how it identifies the structural model. We offer a more formal analysis of local identification, and a numerical evaluation of the binding function, in our discussion of estimation results.

Our structural model contains 45 structural parameters, and the auxiliary regressions m contain 97 auxiliary parameters. Table 2 provides an overview of how our moments and structural parameters relate to each other. We have 8 structural conditional means parameters in the income functions (4.1)-(4.3). The conditional mean parameters in self-employment and wage employment are identified from descriptive Mincer regressions on experience quadratics and tenure dummies. The consumption value of unemployment, and the non-pecuniary utility of self-employment, are identified from observed occupational choices and incomes in the employment sectors. There are a total of 22 search friction parameters that appear in equations (4.11)-(4.16). We identify them by auxiliary regressions of sector dummies on experience; as well as dummies for sector switches on experience, tenure, and an individual's position in (depending on sector) the wage, profit, or capital distribution. We identify the benefit of search by including an effort dummy in the regressions for transitions out of non-employment.

Two aspects of our discussion of identification merit particular mention. Firstly, we back out the cost of search from the identified benefit of search

TABLE 2
INDIRECT INFERENCE: IDENTIFICATION

STRUCTURAL (45)	$\Leftrightarrow$	AUXILIARY (97)
<u>earnings functions</u> conditional means (8) variances (2)	$\Leftrightarrow$	<u>earnings functions</u> log income on experience & tenure (8) moments of residual distribution (8) log firm capital on experience (3) moments of firm capital distr. (4)
<u>preferences & benefits</u> mean NE income support (1) nonpecuniary SE utility (1)	$\Leftrightarrow$	<u>preferences & benefits</u> sector dummies on experience (6)
<u>sectoral frictions</u> sector switch opportunities (6) job loss probabilities (8) job search probabilities (10) cost of search (1)	$\Leftrightarrow$	<u>sectoral frictions</u> sector switches on exp., capital, effort & profit/ wage/ capital quantiles (42) effort on experience & tenure (4)
<u>persistence & heterogeneity</u> earnings autocorrelations (2) distribution of types (6)	$\Leftrightarrow$	<u>persistence & heterogeneity</u> moments of FD income distribution (8) log income at entry on experience (6) variance decomposition from income RE regressions (6)

(the increases of job offer probabilities together with the associated option value) and the observed proportion of non-employed individuals who actually search. Secondly, we combine different techniques to separately identify the different unobserved components of wage and productivity offers: permanent heterogeneity, persistence during an employment spell (the autocorrelation parameters), and innovations. We identify the joint variance of all unobserved components by moments of the distribution of residuals from Mincer regressions. We separate out productivity shocks from permanence and persistence using the distribution of first differences, and variance decompositions from random effects models; for each sector. Random effects regressions that combine sectors identify the cross-correlation of abilities. Finally, we include moments from wages and profits at entry into an em-

ployment spell, when a new AR(1) process of productivity shocks is initialised in our model.

While we solve the model for 50 periods, we trim the actual and simulated data for auxiliary regressions at 40 periods. This is for two reasons. First, we have not explicitly modeled retirement, so we do not expect the model to explain the data when individuals approach retirement. Second, we have very few observations beyond 40 years of potential experience, and moments of actual data become very imprecise and flutter a lot. Further, we sample an observation window from our perfectly balanced simulated data set that mimics the life-cycle profile and attrition rates of our actual, unbalanced data.

There are a number of parameters that we cannot identify. We hold them fixed. Table 3 lists the fixed parameters with their values. The most central of these parameters are the discount factor ψ and the coefficient of relative risk aversion ξ . We set the real annual interest rate r to the negative of the inflation rate under the presumption that low-educated individuals in a poor country such as Ghana find it difficult to access interest-bearing savings products. We set depreciation to 90 percent a year. While this may seem like an incredibly high value, the vast majority (86 percent) of the self-employed in our sample are trade and service businesses. Most of their capital is working capital (for the modal firm, a week's worth of business) which will depreciate rather quickly.⁹ We also fix the Cobb-Douglas coefficient on capital, α , and set it to 0.2.

⁹ In a structural estimation of a static model of the heterogeneous effects of microcredit, Banerjee, Breza, Duflo, and Kinnan (2015) similarly assume that microenterprises liquidate their entire capital stock every period.

TABLE 3
FIXED PARAMETERS

COEFFICIENT OF RELATIVE RISK AVERSION (ξ)	1.5
DISCOUNT FACTOR (ψ)	0.95
REAL ANNUAL INTEREST RATE (r)	-0.10
ANNUAL DEPRECIATION RATE (d)	0.90
MICROENTERPRISE PRODUCTIVITY (α)	0.20

We solve the model numerically, and approximate continuous by discrete values, and integrals by quadratures. Appendix D.2 explains some of the computational details of our estimation.

4.4.2 Inference

We follow recent implementations of indirect inference (e.g. Eckstein, Keane, and Lifshitz (2016), Blundell, Costa Dias, Meghir, and Shaw (2016)) and choose the inverse of the diagonal matrix of squared standard errors of the empirical moments $\mathbf{m}$ as the optimal weight matrix. The resulting weight matrix does not depend on the parameter vector θ . An alternative would have been to calculate the full variance-covariance matrix using a non-parametric bootstrap (see, e.g., Goettler and Gordon (2011)). However, such an approach has been shown to have potentially poor small-sample properties (Altonji and Segal (1996)).

We estimate the standard errors using the formula proposed by Gouriéroux, Monfort, and Renault (1993). Denote $\nabla \mathbf{m}_S(\theta)$ as a $b \times k$ matrix recording the slope of $\mathbf{m}_S(\theta)$ with respect to θ , which we calculate numerically. Then, for a simulated dataset with S simulations, our estimator converges

to a Normal limiting distribution:

$$\sqrt{N} \cdot (\hat{\theta} - \theta) \rightarrow \mathcal{N} \left(\mathbf{0}, \left(1 + \frac{1}{S} \right) \left[\nabla \mathbf{m}_S(\theta)' \cdot \text{cov}(\mathbf{m})^{-1} \cdot \nabla \mathbf{m}_S(\theta) \right]^{-1} \right). \quad (4.17)$$

We can then obtain standard errors in the usual manner from the diagonal of the covariance matrix of the limiting distribution.

4.5 Estimation results


In this section we present and discuss the results of our estimation of the structural model for the sample of Ghanaian men without an upper secondary (Senior High School) degree. We first show the fit of the model against the data. Second, we discuss the main parameter estimates. Third, we analyse the sensitivity of the model.

4.5.1 Goodness of fit

Figure 4 shows the fit of the model on the life-cycle evolution of employment, earnings, and search; and the cross-sectional distribution of microenterprise capital. All of these quantities correspond to moments which are targeted in the estimation, in a more parsimonious form, for instance with parametric auxiliary regressions. The full set of targeted empirical moments, and their simulated counterparts from the model, are listed in Appendix D.3.

The model is able to replicate well the overall patterns in the data. It matches very well the life-cycle evolution of employment in all three sectors. The

FIGURE 4
GOODNESS OF FIT




Life-cycle evolutions are generated with Kernel-smoothed local linear regressions of outcomes on potential experience. Densities are obtained with Kernel density regressions. Shaded areas are 95 % confidence intervals on actual data moments. Simulated moments generated with 15,000 simulations of the estimated model.

model decently fits the life-cycle evolution of mean earnings, although it somewhat overpredicts the experience gradient in the later career stages, both for wages and for profits. (The model also fits higher moments of the earnings distribution). One important input into profits is the capital stock of microenterprises. While we fix the factor share parameter α in the production function, our model lets individuals endogenously accumulate capital, and allows for an allocation choice between financial assets and firm capital. Our model is able to replicate the capital stock of Ghanaian

microenterprises remarkably well. In particular, the simulated distribution follows the empirical distribution very closely in the lower tail, as well as in the mean and median. The simulation does a decent job of fitting the capital stock distribution in the upper tail of the distribution.

There is one exception to our overall positive assessment of the goodness of fit of the model. This relates to the distribution of search effort in the data. While our simulations are able to pick up the hump-shaped pattern of search effort over the life-cycle, we underpredict the level of search activity of young Ghanaian job-seekers. We can think of at least two reasons for this. First, due to data constraints, our measure of effort is very coarse, and only picks up search on the extensive margin. Relatedly, our estimated search cost is non-divisible. It is possible, for a number of reasons, that search behaviour of middle-aged Ghanaian men is quite different on the intensive margin from search of younger men. A second, and related, reason is that our estimates of parameters governing search behaviour are not very precise. We will return to this in the next section.

FIGURE 5
GOODNESS OF FIT: NON-TARGETED CAPITAL DISTRIBUTIONS



Densities are estimated with Kernel density regressions. Shaded areas are 95 % confidence intervals on actual data moments. Simulated moments generated with 15,000 simulations of the estimated model. Measurement of capital is defined in Appendix D.1.

Figure 5 documents the fit on a family of moments that were not targeted in the estimation. We show the life-cycle evolution of capital (assets) in self-employment, wage employment and non-employment. The survey was designed to accurately represent microenterprises; the measurement tools for fixed and working capital were therefore carefully elaborated. Other assets are might be less well-measured, and we probably tend to systematically undermeasure capital in the data; a common problem in the literature.¹⁰ For example, we attribute an equal share of common household assets (durable goods such as refrigerators, mobile phones, or generators) to each household member, irrespective of their age, role in the household, or other variables. Identifying the share of household assets that each member can claim would require a fuller and richer model of intrahousehold bargaining.

Nevertheless, our simulated asset distribution in each of the occupations replicates the shape of the respective distribution in the data. We are able to replicate the bell-shaped distribution of capital for the self-employed (which includes the capital stock of microenterprises), the left-skewed distribution of wage workers' capital, and the bi-modal and right-skewed distribution of capital of the non-employed.

4.5.2 Parameter estimates

We report the main parameter estimates in Table 4. The productivity and wage offer functions have the shape that we expect. Potential experience enters wages through a concave form, and returns to experience in either sector are of plausible magnitude. Returns to experience in wage-employment are significantly higher than in self-employment. Productivity shocks have a

¹⁰ Our definition of capital includes personal financial assets, the firm's capital stock (if self-employed), and an implied share in household durables (see Appendix D.1).

larger variance in self-employment than for the wage-employed. The autocorrelation of productivity shocks is moderate, significantly positive, and very similar across sectors. Finally, the standard deviations of permanent heterogeneity are positive, but smaller than the productivity shocks. Fixed ability accounts for about half the variation of productivity shocks in self-employment; and three quarters in wage-employment. The correlation of ‘ability’ in all of the three sectors is significantly positive. That is, the best entrepreneurs, on average, are the best wage workers, and they also receive the most generous income support when they are non-employed. Together, this suggests that permanent heterogeneity is an important feature of our model. Indeed, a version of the model without permanent heterogeneity fails to replicate features of the data on many dimensions.

We also estimate parameters related to job search, as well as the implied value of income support for the non-employed, and the non-pecuniary utility value of being self-employed. On average, Ghanaian men value being self-employed positively, but the estimated value is very small, and the parameter is not statistically significant. Search effort is costly - our point estimate suggests an annual cost of effort of around 70 USD (102 cedis), compared to an estimated income support in non-employment of around 540 USD. This is in line with field evidence from urban Africa: Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn (2016) report that the median expenditure on job search for young job seekers in Addis Ababa is about 16 percent.

The estimated parameters for search efficacy are difficult to interpret. We therefore also graphically report the probability of encountering a business opportunity or finding a job in Figure 6. The corresponding coefficients

TABLE 4
MAIN PARAMETER ESTIMATES


		ESTIMATE	STD. ERROR	LOWER	UPPER
95 % CI					
Profit function					
Constant	γ_0	5.609	(0.094)	5.424	5.793
Experience	γ_1	0.032	(0.007)	0.019	0.045
Experience ²	γ_2	-0.042	(0.017)	-0.075	-0.010
Tenure > 2	γ_3	0.051	(0.080)	-0.106	0.207
Innovations (SD)	σ_1	0.818	(0.040)	0.740	0.896
Autocorrelation	ζ_1	0.616	(0.084)	0.451	0.781
Wage offer function					
Constant	β_0	6.688	(0.124)	6.445	6.930
Experience	β_1	0.068	(0.009)	0.051	0.084
Experience ²	β_2	-0.107	(0.021)	-0.149	-0.065
Tenure > 2	β_3	0.071	(0.060)	-0.046	0.188
Innovations (SD)	σ_2	0.498	(0.020)	0.459	0.537
Autocorrelation	ζ_2	0.610	(0.117)	0.381	0.838
Permanent Heterogeneity					
Std. dev. of SE ability	σ_s	0.412	(0.068)	0.278	0.546
Std. dev. of WE ability	σ_w	0.369	(0.042)	0.286	0.451
Std. dev. of NE support	σ_n	0.507	(0.142)	0.228	0.786
Correlation SE-WE	ρ_{sw}	0.871	(0.331)	0.222	1.521
Correlation WE-NE	ρ_{wn}	0.594	(0.129)	0.341	0.848
Correlation SE-NE	ρ_{sn}	0.712	(0.247)	0.228	1.196
Job search and preferences					
NE income support	$\bar{\mu}^n$	6.643	(0.133)	6.381	6.905
Non-pecuniary SE utility	$\bar{u}_s$	0.005	(0.004)	-0.003	0.013
Search cost	κ	101.900	(47.296)	9.200	194.600
Search efficacy SE	$\lambda_{E,SE}^{NE}$	-0.526	(0.512)	-1.531	0.478
Search efficacy WE	$\lambda_{E,WE}^{NE}$	-0.194	(0.120)	-0.430	0.041

Selected parameter estimates of structural model obtained via indirect inference estimation on sample of 3,307 panel observations of male Ghanaians aged 15-65 with education less than senior high school degree. Further parameter estimates are listed in appendix table A21.

are in Table A21 in the Appendix. The estimates suggest that search frictions are particularly important at the beginning of the life-cycle. Young workers face a very low probability of finding employment opportunities,

especially in self-employment. Our estimate of the probability of “encountering a business opportunity” likely captures multiple mechanisms: actual business ideas, the psychic or organisational barrier to starting a business (sourcing suppliers, finding clients, etc.). It also captures possible monetary entry costs in addition to the capital used in production, which we cannot separately identify. The probabilities rise towards the middle of the life-cycle and then fall again; but stay generally low. Effort increases the probabilities slightly, but not significantly so.

FIGURE 6
OFFER PROBABILITIES: FIRMS AND JOBS



Model implied probabilities of finding a business opportunity or a wage job when non-employed, for individuals with ≤ 2 years of tenure. Shaded areas are 95 % confidence intervals estimated by a parametric bootstrap, drawing 10,000 replications from the joint distribution of parameters of the offer distribution.

As the table and the figure show, our ability to estimate the parameters with precision varies. The conditional mean parameters and productivity shock variances of profits and wages are fairly precisely estimated; similarly the autocorrelation of profits, and the variances of permanent heterogeneity. We are only able to pin down fairly well the correlation between income support and wage employment ability. The other cross-correlations are positive, but very noisy: the confidence bounds cover the interval $[0.22, 1]$. The param-

eters governing the search process – the cost and efficacy of search – are similarly very imprecisely estimated. This may be the flipside of why effort provision implied by the model only matches the patterns in the data in very broad terms.

4.5.3 Sensitivity of estimates to moments


Having obtained our parameter estimates, we can use them to carry out a more formal analysis of the binding function which maps auxiliary moments into structural parameters. We calculate the *standardised sensitivity* measure proposed by Andrews, Gentzkow, and Shapiro (2017), where sensitivity is defined by

$$S = -\left(\nabla \mathbf{m}_S(\boldsymbol{\theta})' \cdot \mathbf{W} \cdot \nabla \mathbf{m}_S(\boldsymbol{\theta})\right)^{-1} \cdot \nabla \mathbf{m}_S(\boldsymbol{\theta})' \cdot \mathbf{W}; \quad (4.18)$$

and standardised sensitivity of structural parameter p to auxiliary parameter j is given by $S_{pj} \cdot \frac{\omega_j}{\sigma_p}$ where ω_j is the standard deviation of the auxiliary parameter j , and σ_p the standard deviation of the structural parameter estimate. We evaluate S at the estimates $\hat{\boldsymbol{\theta}}$ to obtain a local measure of sensitivity of parameters to perturbation in moments. We can interpret the standardised version as the elasticity of parameter estimates to moments.

In Figure 7, we graphically report standardised sensitivity of our 45 parameters to 97 moments. The darker a cell, the higher the sensitivity. While the reader may draw their own inferences from the vast amount of information we report in this figure, we highlight a couple of points of interpretation. First, each structural parameter is sensitive to at least one moment. This lends credence to the notion that our model is locally identified by the mo-

Figure 7: Standardised Sensitivity



This figure reports Andrews, Gentzkow, and Shapiro (2017) standardised sensitivity of 45 structural parameters to 97 auxiliary moments; as defined in the main text. The colour intensity of a cell is proportional to the square of standardised sensitivity, winsorised at the 99th percentile. Auxiliary parameters are listed in reverse order to the tables in Appendix D.3 – the bottom row represents the constant in the auxiliary regression of the proportion working in SE; the top row represents the between variation in a random effects regression of earnings in both WE and SE.

ments: each structural parameter estimate is constrained by the data. Second, sensitivity of parameters is not concentrated, in the sense that most structural parameters are sensitive to multiple moments. Conversely, each auxiliary parameter identifies multiple structural parameters. Our informal discussion of identification in section 4.4.1 could only reveal part of the full picture that we are seeing here. For example, the parameters of the wage mean ($\beta_0 - \beta_3$) are quite sensitive to the respective parameters in the auxiliary wage regression. This is something we had informally discussed. However, it turns out these parameters are also sensitive to other moments – higher moments of wages, or correlates of transition probabilities.

Third, all moments are informative for the model. Take, for instance, the moments of the distribution of residual wages. Even though these only identify a few parameters, they provide crucial information for identifying a particular parameter – in this case, the standard deviation and autocorrelation of wage shocks. This resonates with the importance we had attached to these particular moments in our informal discussion.

4.6 Interpretation

4.6.1 Structural replication of field experiments

Our paper is motivated by the desire to provide a unified framework to understanding career paths in urban labour markets of developing countries. Our existing body of knowledge about these markets mainly stems from the stylised facts that a rich and large literature using field experiments has put together in the last decade. If our framework is useful in this context, we should be able to use the model to replicate the findings from field experi-

ments.

We run two kinds of experiments. The first experiment is a capital grant to microenterprises. Fafchamps, McKenzie, Quinn, and Woodruff (2014) ran such an experiment in the same empirical context. In early 2009, they randomly allocated 150 Ghanaian cedis (ca. 120 USD) – either as cash or in-kind – to microenterprises in Accra. They conducted a total of six survey rounds, with the last round just under a year from the time the first microenterprise was treated. They report intent-to-treat effects on a range of outcome variables, including profits, capital, and household consumption expenditure. With only 2% non-compliance, the ITT effects are very similar to average treatment effects (ATE). We use the dataset and replication files provided by the authors. In Appendix Table A22, we show that we are able to replicate the ATE reported in the paper. We then adjust the data so that it corresponds with our setting (annualised flow variables; June 2010 as the base for monetary variables) and re-calculate the treatment effects on the adjusted data.

We use our estimated structural model to run the exact same experiment. For each of our simulated individuals, we randomly assign a life-cycle period for the treatment. For those individuals who are self-employed in that period, we then increase capital (assets) at the beginning of the model period by the capital grant. We use the policy functions of the model to determine how treated individuals adjust the capital stock of their firm, profits, and total household consumption. The flow variables in our model correspond to the 12-month period following the treatment. Since we have already simulated the model without the treatment, we obtain the treatment effect by taking the difference between the two counterfactual scenarios, for each simulated

TABLE 5
STRUCTURAL VS. EXPERIMENTAL TREATMENT EFFECTS OF CAPITAL
GRANTS

	Profits	Capital	Consumption
Structural model	90.21	44.10	62.38
Experiment: Cash treatment	69.11 (217.2)	2.48 (68.68)	96.20 (67.10)
Experiment: Pooled treatments	221.79 (176.85)	46.94 (59.41)	72.82 (60.81)

Source: Own calculation of experimental treatment effects using data from the field experiment in Fafchamps, McKenzie, Quinn, and Woodruff (2014). Appendix D.4 details the cross-walk between our estimates and the estimates reported in Fafchamps, McKenzie, Quinn, and Woodruff (2014). Robust standard errors of experimental parameter estimates from a fixed-effects regression in parentheses. Structural treatment effects are described in the main text.

individual. We average to obtain the ATE.

We report the structural and experimental treatment effects in Table 5. The structure of our model, which allows agents to freely allocate their assets between firm capital and cash after receiving the treatment, corresponds most directly to the cash treatment in Fafchamps, McKenzie, Quinn, and Woodruff (2014). However, since the authors are unable to reject the null hypothesis that the cash and in-kind treatments have the same effect for the specifications we use in our comparison, we also report the effects from the pooled treatment. Our structural model implies that the capital grant (168 cedis in 2010 money) increases microenterprise capital stock by 44 cedis, total household consumption expenditure by 62 cedis, and annual profits by 90 cedis, on average. The structural estimates for profits and capital are within the bounds formed by the experimental estimates for the two treatment variants. The structural estimate for consumption is within half a standard deviation of either experimental treatment effect. We conclude from this comparison that our structural model is able to replicate the findings

from an independent field experiment that took place in the same empirical context.

The second experiment is an active labour market policy to support young unemployed men in their job search. We simulate a policy to subsidize job search that provides vouchers to young unemployed men. These vouchers can only be redeemed for direct search costs. We implement this simulation by cutting κ in half. While we know of no similar experiment that was implemented in Ghana, we can compare our structural estimates to the findings from a recent field experiment that Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn (2016) conducted in Addis Ababa, Ethiopia. The authors provided a non-fungible transport subsidy to assist in job search, which corresponded to about two thirds of median expenditure on job search. From our conversations with some former participants in the GUHPS, transport expenses appeared to be an important part of monetary job search costs in Ghana.

Using our structural model, we find an average treatment effect of 8 percentage points on the proportion exerting search effort, and 11 percentage points for those under age 30. We find close to zero effects on employment in any or either sector. Abebe, Caria, Fafchamps, Falco, Franklin, and Quinn (2016) report that the transport subsidy increased the proportion of individuals doing any job search activity by a non-significant 5 percentage points, arguably from a higher baseline level. The authors do find much larger and significant effects on particular kinds of search activities that require physical transport to the city centre. They report that treated job seekers find jobs of higher quality, but no effects on the quantity of employment. Overall, and with the caveats that apply to comparison across different empirical

contexts, our structural model also is in line with experimental empirical evidence on job search.

4.6.2 Entrepreneurship: Necessity or opportunity?

These simulations highlight the potential of our structural model as a unifying framework for labour market policy in developing countries. Our model is able to simulate a large array of policy parameters of interest. Importantly, it tells us about behavioural responses that go beyond what is typically studied experimentally. We use our model to explore a long-standing question in the labour economics of developing countries: is entrepreneurship a manifestation of necessity, or of attractive business opportunities?

There are many reasons to start a business, yet the literature has often bifurcated entrepreneurs into two groups. These two groups, ‘necessity’ entrepreneurs and ‘opportunity’ entrepreneurs, appear in different variants: as ‘subsistence’ entrepreneurs versus ‘transformational’ entrepreneurs (Lerner and Schoar (2010)), or as ‘reluctant’ versus ‘gung-ho’ entrepreneurs (Banerjee, Breza, Duflo, and Kinnan (2015)). We ask our structural model two related questions: (i) How reasonable is such a bifurcation? (ii) Is the choice of ‘opportunity’ entrepreneurs about the person or about the circumstances?

To answer these questions through the lens of the model, we exploit the fact that the model simulates the full opportunity set of agents. That is, for every simulated choice, the model reports the choice set the agent was picking from. In Table 6, we look at different paths into self-employment. In particular, we categorise each simulated employment spell according to four mutually exclusive and collectively exhaustive types, according to the

TABLE 6
SIMULATION: PATHS INTO SELF-EMPLOYMENT

Type	‘Opportunity’ Entrepreneurs		‘Necessity’ Entrepreneurs	
	(1)	(2)	(3)	(4)
Share of new SE entrants	0.58		0.42	
Previous sector	NE	WE	NE	WE
Previous sector share	0.08	0.92	0.25	0.75
Overall share of type	0.04	0.54	0.10	0.32
Average spell length (years)	3.47	3.47	3.18	2.88
Median firm capital	411	711	545	365
Median annual net profits	2746	2830	2364	1765
Average number of job offers	1.77	1.72	1.56	1.39
Share leaving to NE	0.36	0.34	0.35	0.36
Share leaving to WE	0.64	0.66	0.65	0.64
Median starting salary	2452	2454	2428	2361


Source: Simulation of estimated structural model using 15,000 replications. After collecting all datapoints from a self-employment spell, we apply our observation window. We only retain spells where we observe the first spell period. Profits are net of the cost of capital (through depreciation).

path of entry: entry from NE while having a simultaneous wage offer (Type 1); entry from a continuing wage job (Type 2); entry from NE without a simultaneous job offer; and entry from WE immediately after being fired (Type 4). Types 1 and 2 correspond to what we would typically think of as opportunity entrepreneurs – individuals with a revealed preference for self-employment – whereas types 3 and 4 arguably entered entrepreneurship out of necessity, facing a lack of alternatives.

We make several observations. First, we classify around 60 percent of entrepreneurs as opportunity entrepreneurs. The vast majority of them entered self-employment out of an existing wage job. Second, even among necessity entrepreneurs, entry from wage-employment seems to be the more common path. Opportunity entrepreneurs stay in self-employment for longer than ne-

cessity entrepreneurs, tend to have somewhat larger firms, and earn higher profits. However, few of these differences are very pronounced. In fact, only type 4 entrepreneurs – those who become self-employed because they lost a wage job – appear to be a categorically different type. Their businesses make significantly lower profits and they stay in self employment for a shorter amount of time, on average. Even then, all types can expect to receive offers for wage work they choose not to accept. Finally, we also examine whether we continue to see differences across types after their self-employment spell has ended. This is not the case. All types face virtually the same job opportunities – both in employment outcomes and in salaries – after they leave their self-employment spell.

FIGURE 8
SIMULATION: CORRELATES OF JOB OFFER AND ACCEPTANCE RATES
OUT OF SELF-EMPLOYMENT



Model implied probabilities of receiving a job offer when self-employed, and of accepting such an offer.

In Figure 8, we take a different look at one manifestation of opportunity

in entrepreneurship: the opportunity to take up a wage job. We plot the gradient of job offer and acceptance probabilities with potential experience, and with capital. Expect for the very young or the very poor, our simulated entrepreneurs face quite similar opportunities to leave self-employment for a wage job. They differ more strongly in whether they choose to accept these job offers. The gradient with capital of the probability to accept an offer is very steep. The poorest entrepreneurs take all offers, the richest entrepreneurs virtually none. The slope of the acceptance probability with potential experience is also negative, but less steep.


Returning to our initial questions, we find that while there are differences in outcomes according to the mode of entry of entrepreneurs, these differences are probably not strong enough to warrant a bifurcation into qualitatively different ‘types’. Further, different experiences of entrepreneurship seem to leave no scarring effect for future careers. Admittedly, our model may not pick up all possible dynamic channels that unfold in the real world. We also find that episodes of self-employment are strongly linked with wage work. The one group that seems to be most strongly affected from lack of opportunities are people who involuntarily lost their wage job; about a third of all self-employed.

4.6.3 Wage work: Stepping stone to self-employment?

The findings from the previous experiment point to wage work as an important precursor to entry into self-employment. This raises the question of whether wage work is a stepping stone to self-employment in our context. A related question is: how persistent are occupational shocks?

We use our model to simulate a thought experiment. Suppose we were to offer every non-employed person a guaranteed wage job for a year. How would occupational trajectories change? This is of course not a real policy experiment that one could conduct on this scale; and even if it were, our model would be completely unsuited to address any general equilibrium adjustment. It is an instructive thought experiment we can use to study the career trajectory of the average Ghanaian man after a shock which puts him in wage work.

FIGURE 9
SIMULATION: JOB GUARANTEE



Employment trajectories implied by job guarantee experiment as described in the main text. Experiment in period 1 in red in red solid lines; experiment in period 10 in blue dotted lines; experiment in period 20 in green dash-dotted lines.

We run three different versions of this experiment: a job guarantee in the first period (defined by potential experience), a job guarantee in period 10, and a job guarantee in period 20. To implement this experiment, take all simulated individuals who are non-employed in the experimental period, and allocate them to a random wage job. The wage is drawn from the es-

timated offer distribution (which differs from the distribution of observed wages).¹¹ The job guarantee is implemented at the beginning of the period. After this, the model proceeds as usual, and shocks are drawn from the estimated distributions. Figure 9 shows the trajectories for wage-employment and self-employment from the experiment.

We find that the job guarantee fails to leave a lasting impact on wage employment. One period after a job guarantee to new labour market entrants, there is a 74 percent chance that the typical worker will have left his job. The probability of being wage employed five years after the experiment is indistinguishable from the baseline scenario. This holds for every timing of the experiment, but it is particularly pronounced for the experiments later in the life-cycle. Presumably by then, the pool of non-employed individuals contains only the least employable. In the right-hand panel of Figure 9, we examine the spillover effects to self-employment. For the job guarantee at entry and after ten years of experience, we do find some visible increases in the probability of being self-employed. For the ten-years job guarantee, this effect diminishes over time but lasts for an entire ten years.

Why do workers leave their new jobs so quickly? In Table 7, we list the reasons implied by the model, for the first three years of each of the experiments. Young workers in our sample face very high rates of job loss – so those whom the experiment places into a job experience the same fate. For more experienced workers, choosing to start a business instead becomes the leading reason for leaving behind a wage job. Finally, the fact that workers

¹¹ We do not implement this job guarantee for individuals who are already in a wage job or self-employed with a business. This form of job guarantee would uproot individuals from their current job; and any estimates we would obtain would mix the effects of the job guarantee with the effects from uprooting existing employment relationships.

TABLE 7
SIMULATION: REASONS FOR LEAVING THE WAGE JOBS

Job guarantee	Period since shock	Pr(Leave)	Lose Job	Choose SE	Choose NE
Period 1	$t = 1$	74.0%	86.8%	2.1%	11.0%
	$t = 2$	64.5%	90.4%	5.4%	4.2%
	$t = 3$	60.8%	89.8%	6.1%	4.1%
Period 10	$t = 1$	39.9%	67.8%	19.4%	12.8%
	$t = 2$	36.8%	69.0%	25.7%	5.3%
	$t = 3$	35.6%	64.5%	31.0%	4.5%
Period 20	$t = 1$	31.6%	42.6%	46.7%	10.8%
	$t = 2$	28.9%	47.0%	48.4%	4.6%
	$t = 3$	28.2%	47.0%	48.2%	4.7%

are placed in a random job – which they might not have chosen voluntarily – does not seem to play a large role in the short-lived wage employment effects. The importance of voluntary quits does not change with the timing of the experiment either.

In sum, our simulated thought experiment gives some credence to the argument that wage-employment provides a stepping stone to self-employment. This is particularly true for more experienced workers, who also have the ability to hold on to a business. However, even a guaranteed job does not leave a transformative impact on the probabilities that Ghanaian workers find themselves in either self-employment or wage work.

4.7 Conclusion

In this chapter we build and estimate a structural life-cycle model of sectoral choice in urban Ghana. We formulate a model that allows for endogenous sector choices, asset accumulation, and search effort. Our model

also accounts for rich patterns of unobserved heterogeneity. Our work extends existing dynamic models of occupational choice to a setting with a large self-employment sector. These extensions allow us to take this class of models to studying central aspects of urban labour market in developing countries. The applicability to such contexts is reinforced by the fact that, unlike very recent work modeling self-employment in developed countries, our model incorporates two central types of market frictions: credit constraints for entrepreneurs, and search frictions in the wage labour market. These types of frictions are studied by a rich field-experimental literature in development economics. Our model builds a unifying framework applicable to these settings.

We structurally estimate our model on the longest panel data set from Sub-Saharan Africa, the Ghana Urban Household Panel Survey (GUHPS). In this version, we estimate the model on data from men with low levels of education. Our model fits the data it is matched to well. It also gives a reasonable fit for moments that we do not match on, such as the distribution of accumulated savings. We validate our model by replicating two actual field experiments: a capital grant to microenterprises that took place in the same place and during the same time as our data, and a job search subsidy in a different but comparable setting. We apply our model to shed light on whether entrepreneurship is a manifestation of opportunities or necessities, and to which degree wage-employment is a stepping stone to self-employment in urban Ghana.

Our work demonstrates the usefulness of a structural model to complement experimental studies of firms, labour markets, and policies in developing countries. We believe that such an approach provides a fruitful area for fu-

ture research, including extensions to our framework in this chapter. We can think of at least two lines of extension. First, we model entry into self-employment in a parallel way to entry into wage employment, but this gives very high entry constraints in form of ‘lack of business ideas’. It would be useful to disentangle this lack of ideas from other forms of entry costs into self-employment, both monetary and psychological. Second, our treatment of capital in the stylised form of a standard production function may not capture all aspects of actual production perfectly. For example, most capital in microenterprises comes in the form of frequently replaced working capital.

Appendix D

D.1 Data and variable construction

From the 9 waves of the Ghana Urban Household Panel Survey (GUHPS) we retain a dataset of 1,255 men that contains 3,307 observations in an unbalanced panel and spans the waves 2005, 2006, 2008, 2009, 2010, 2012 and 2013. We do not include data from 2004 (where variable definitions were significantly different and several key variables were not elicited) and from 2007 (a recall wave).

We keep all individuals for whom information on the sector of employment (wage employment, self-employment or non-employment) and if they are employed, on earnings from the main job, are reported. We drop individuals if we observe them in an on-going education spell; and only retain them if they are in their final year of education. We use self-reported measures of weekly wages and incomes, create capital stocks by adding up assets owned by the respondent's business. We deflate all monetary variables by the monthly consumer price index from the Bank of Ghana. Inflation in Ghana is high, and in July 2010 the currency was redenominated and four digits were cut off (so that 1 Ghanaian Cedi replaces 10,000 New Cedis). People were still following the habit of quoting figures in the old currency years after the redenomination, and so we assume that any weekly income above 10,000 cedis (between 5,000-10,000 US dollar at market exchange rates) was quoted in old currency. All our monetary variables are expressed in June 2010 Ghanaian cedis (the new currency). The annual average exchange rate for 2010 (World Bank) is 1.431 cedis to the dollar, or approximately a factor of 0.7 for converting cedi to USD values. We remove remaining outliers by winsorizing at the first and 99th percentile.

We only create variables that measure employment transitions if we observe an individual

in two adjacent panel waves (so we record a missing observation for transition if an individual was in the survey, drops out for one wave, and rejoins the survey in the next wave in a different occupation). We fill in missing observations if there is no ambiguity about them. (For example, if an individual in the 2010 wave reports having been unemployed for 1 year and 3 months, we fill in ‘not employed’ for his employment status in 2009 if that observation is missing.)

Occupation variables

- **Self-employment:** All individuals who report an own-account / self-employment activity as their current main occupation (i.e. who answer the self-employment module of the questionnaire for their main activity).
- **Private wage employment:** All individuals who report a wage job (i.e. who answer the ‘jobs’ module of the questionnaire for their main activity).
- **Not employed:** All individuals who are not working - this includes individuals in training (e.g. schooling, vocational training, apprenticeships). This category is as broad as possible at the moment and includes many people that would otherwise be classified as ‘out of the labour force’. Given that this is about the life-cycle, we do not want to restrict attention only to those in the labour force (i.e. as a minimum we also want to capture entry into the labour force). There is no good indicator for people permanently out of the labour force, such as “not available for work”.

Education, tenure, experience, and effort variables

- **Years of Education:** Typical duration of schooling in order to obtain the highest educational attainment reported by the respondent.
- **High level of education:** Dummy for whether the highest educational attainment is Senior High School (SHS) or a tertiary qualification (university and college degrees, and certain professional qualifications such as nursing school, teachers training certificate, etc.); zero otherwise for lower levels of education.
- **Potential experience:** Calculated as $\min \{ \text{age at interview} - \text{years of education} - 6, \text{age} - 15 \}$. That is, potential experience starts counting either at age 15 or the age an individual leaves school, whichever is greater.

- **Tenure:** Time difference (rounded down to years) between interview date and the date (up to year and month) when the respondent started working for the current employer / in their current own business. For the unemployed who have worked at least once before, the time difference between interview date and the date they last had a job for pay, profit or gain. For those who never worked, set equal to potential experience. Truncated at 3 years.
- **Effort:** Dummy variable set equal to one if answered yes to “During the past 4 weeks, have you taken any ACTION to seek work/do business” (capitalised in original questionnaire), zero otherwise. In 2005, if agreed to being “Able to work and looking for a job”.

Income and Capital

- **Wage:** Answer to ‘how much do you CURRENTLY earn [per pay period] BEFORE TAXES (excluding bonuses/allowances). We calculate the real value as describe above.
- **Business Income:** Answer to ‘how much PROFIT does the business make in a normal week (after all expenses)?’ Respondents are explained the concepts of profit, revenues, and costs. They are asked for each of this components. A PDA consistency check prompts the enumerator to ask for revised figures if profit is not equal to revenues minus cost.
- **Capital / assets:** For the wage-employed and non-employed, the real value of their personal financial assets (cash, current account, formal and informal savings accounts), plus an equal fraction for each household member of household durables (elicited through a listing exercise). For the self-employed, additionally the real value of tools, equipment and other assets used in their business (elicited through a listing exercise), and the value of a week’s worth of working capital and unsold merchandise. The modal firm reports replacing their inventories on a weekly basis.

D.2 Computational details

We discretise the distributions of wage and productivity residuals, each with ten points of support. We follow the method of Tauchen (1986) and choose an equidistant grid, then

assign probabilities to each node. We use the same method to discretise the AR(1) process of residuals into a Markov transition matrix with a finite number of states; and to discretise the distribution of permanent heterogeneity. For the latter, we use a coarser grid with four nodes in each sector.

We discretise the capital into an equidistant grid in the log space. We use a grid of .1 log points up to 90 cedis ($\exp(4.5)$) and a grid of .05 log points thereafter.

We solve the model for each level of experience, tenure (up to including 3 years), wage and productivity residual, and effort. We map from the set of all combinations of frictions and earnings shocks to the set of all unique choice menus; and solve for occupational choice in each choice menu separately. We then numerically approximate the expectation by a sum over all choice menus, weighted by their probabilities. We repeat this solution by backwards induction for each type of the discretised distribution of permanent heterogeneity. The proportion of the overall 15,000 simulations that corresponds to each type is determined by the cross-correlation of abilities.

We solve the model in MATLAB using parallel processing on a workstation with 16 cores. A single evaluation of the objective function takes about 700 seconds.

D.3 Goodness of fit

In this section we detail the additional parameter estimates not reported in the main text. These parameters relate to the job offer probabilities, job loss probabilities, and probabilities of being offered to switch occupations. We also list all the 97 auxiliary parameters used in the estimation. We list the empirical moments and their standard errors (obtained analytically in the case of regressions, and with a bootstrap in the case of moments of a distribution). We then list the simulated moments that we obtain with our estimated parameters. For ease of comparing the goodness of fit, we also calculate the difference between the empirical and the simulated parameters in standard deviation units.

TABLE A1
GOODNESS OF FIT - PROPORTION WORKING IN SE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	-0.153	(0.027)	-0.118	1.264
Experience	0.032	(0.003)	0.030	-0.539
Experience ²	-0.048	(0.007)	-0.042	0.881

TABLE A2
GOODNESS OF FIT - PROPORTION WORKING IN WE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	0.014	(0.031)	0.021	0.205
Experience	0.043	(0.003)	0.039	-1.300
Experience ²	-0.086	(0.008)	-0.078	1.119

TABLE A3
GOODNESS OF FIT - TRANSITIONS SE TO WE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	0.394	(0.142)	0.048	-2.434
Experience	-0.008	(0.011)	0.008	1.365
Experience ²	0.016	(0.024)	-0.017	-1.334
Profits < P25	-0.107	(0.059)	0.060	2.842
Profits < P50	0.047	(0.059)	0.054	0.132
Profits < P75	-0.044	(0.058)	0.071	1.982

TABLE A4
GOODNESS OF FIT - TRANSITIONS SE TO NE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	0.366	(0.102)	0.349	-0.171
Experience	-0.022	(0.008)	-0.018	0.521
Experience ²	0.042	(0.017)	0.031	-0.664
Tenure > 2	-0.051	(0.032)	-0.064	-0.423
Profits < P25	0.001	(0.042)	0.028	0.658
Profits < P50	-0.008	(0.042)	0.011	0.469
Profits < P75	0.011	(0.041)	0.015	0.081

TABLE A5
GOODNESS OF FIT - TRANSITIONS WE TO SE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	0.113	(0.089)	0.014	-1.110
Experience	0.003	(0.007)	0.009	0.778
Experience ²	-0.000	(0.016)	-0.011	-0.649
Profits < P25	-0.037	(0.040)	0.038	1.900
Profits < P50	-0.012	(0.040)	0.029	1.014
Profits < P75	-0.048	(0.040)	0.031	1.983

TABLE A6
GOODNESS OF FIT - TRANSITIONS WE TO NE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	0.519	(0.099)	0.478	-0.418
Experience	-0.037	(0.009)	-0.032	0.542
Experience ²	0.071	(0.019)	0.062	-0.518
Tenure > 2	-0.025	(0.033)	-0.016	0.286
Profits < P25	0.065	(0.044)	0.013	-1.157
Profits < P50	-0.015	(0.045)	0.019	0.756
Profits < P75	0.066	(0.044)	-0.002	-1.527

TABLE A7
GOODNESS OF FIT - TRANSITIONS NE TO SE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	-0.172	(0.083)	-0.094	0.940
Experience	0.038	(0.008)	0.015	-3.020
Experience ²	-0.075	(0.019)	-0.032	2.266
Tenure > 2	-0.038	(0.040)	0.019	1.407
Effort	0.060	(0.044)	0.060	0.001
Profits < P25	-0.037	(0.083)	0.028	0.785
Profits < P50	-0.045	(0.071)	-0.003	0.586
Profits < P75	-0.014	(0.062)	0.005	0.290

TABLE A8
GOODNESS OF FIT - TRANSITIONS NE TO WE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	0.168	(0.111)	0.031	-1.238
Experience	0.021	(0.010)	0.029	0.737
Experience ²	-0.043	(0.026)	-0.049	-0.250
Tenure > 2	-0.068	(0.054)	-0.028	0.744
Effort	0.004	(0.059)	0.032	0.468
Profits < P25	0.045	(0.110)	-0.014	-0.535
Profits < P50	-0.176	(0.095)	0.042	2.286
Profits < P75	-0.024	(0.083)	-0.015	0.114

TABLE A9
GOODNESS OF FIT - PROFITS

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	7.335	(0.248)	7.032	-1.223
Experience	0.056	(0.023)	0.075	0.813
Experience ²	-0.105	(0.047)	-0.133	-0.602
Tenure > 2	0.210	(0.081)	0.207	-0.042

TABLE A10
GOODNESS OF FIT - WAGES

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	7.009	(0.102)	6.855	-1.521
Experience	0.062	(0.010)	0.066	0.399
Experience ²	-0.105	(0.023)	-0.104	0.040
Tenure > 2	0.140	(0.042)	0.148	0.186

TABLE A11
GOODNESS OF FIT - PROFIT RESIDUALS

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Std. deviation	1.020	(0.026)	1.054	1.327
P25	-0.638	(0.043)	-0.715	-1.781
P50	-0.015	(0.044)	-0.050	-0.784
P75	0.700	(0.052)	0.633	-1.296

TABLE A12
GOODNESS OF FIT - WAGE RESIDUALS

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Std. deviation	0.659	(0.014)	0.640	-1.361
P25	-0.439	(0.026)	-0.448	-0.348
P50	-0.015	(0.029)	-0.014	0.040
P75	0.455	(0.029)	0.468	0.472

TABLE A13
GOODNESS OF FIT - PROFITS FIRST DIFFERENCES

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Mean	0.190	(0.077)	0.118	-0.933
Std. deviation	1.239	(0.052)	1.074	-3.189
P25	-0.687	(0.099)	-0.687	0.002
P50	0.141	(0.079)	-0.007	-1.882
P75	1.070	(0.137)	0.677	-2.855

TABLE A14
GOODNESS OF FIT - WAGES FIRST DIFFERENCES

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Mean	0.099	(0.038)	0.061	-1.009
Std. deviation	0.723	(0.034)	0.749	0.758
P25	-0.244	(0.051)	-0.369	-2.474
P50	0.087	(0.026)	-0.011	-3.739
P75	0.489	(0.047)	0.461	-0.603

TABLE A15
GOODNESS OF FIT - PROFITS AT ENTRY

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	6.633	(0.452)	7.238	1.341
Experience	0.146	(0.045)	0.076	-1.562
Experience ²	-0.315	(0.101)	-0.136	1.768

TABLE A16
GOODNESS OF FIT - WAGES AT ENTRY

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	6.877	(0.171)	6.919	0.245
Experience	0.075	(0.019)	0.074	-0.047
Experience ²	-0.132	(0.045)	-0.121	0.258

TABLE A17
GOODNESS OF FIT - SEARCH EFFORT PROVISION

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	0.265	(0.059)	-0.143	-6.909
Experience	0.034	(0.006)	0.050	2.372
Experience ²	-0.084	(0.016)	-0.075	0.518
Tenure > 2	-0.120	(0.039)	-0.181	-1.549

TABLE A18
GOODNESS OF FIT - FIRM CAPITAL OVER LIFE-CYCLE

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
Constant	5.056	(0.522)	4.395	-1.265
Experience	0.121	(0.045)	0.145	0.526
Experience ²	-0.234	(0.094)	-0.234	0.005

TABLE A19
GOODNESS OF FIT - FIRM CAPITAL DISTRIBUTION

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
P25	4.640	(0.143)	4.800	1.125
P25	5.526	(0.096)	5.600	0.766
P50	6.525	(0.086)	6.450	-0.872
P75	7.511	(0.084)	7.226	-3.395

TABLE A20
GOODNESS OF FIT - RANDOM EFFECTS VARIANCE DECOMPOSITION

	ACTUAL	STD. ERROR	SIMULATED	DIFF. / SE
SE: Within variation σ_u	0.517	(0.054)	0.490	-0.494
SE: Between variation σ_e	0.885	(0.031)	0.936	1.666
WE: Within variation σ_u	0.338	(0.031)	0.384	1.511
WE: Between variation σ_e	0.571	(0.017)	0.515	-3.244
All: Within variation σ_u	0.388	(0.028)	0.396	0.269
All: Between variation σ_e	0.746	(0.015)	0.740	-0.407

TABLE A21
ADDITIONAL PARAMETER ESTIMATES

		ESTIMATE	STD. ERROR	LOWER	UPPER
				95 % CI	
Job loss probabilities in SE					
Constant	δ_C^{SE}	-0.756	(0.456)	-1.650	0.137
Experience	δ_{T1}^{SE}	0.196	(0.040)	0.118	0.275
Experience ²	δ_{T2}^{SE}	-0.377	(0.075)	-0.523	-0.230
Tenure > 2	δ_{X1}^{SE}	0.915	(1.040)	-1.124	2.954
Job loss probabilities in WE					
Constant	δ_C^{WE}	-0.782	(0.298)	-1.367	-0.197
Experience	δ_{T1}^{WE}	0.217	(0.028)	0.162	0.272
Experience ²	δ_{T2}^{WE}	-0.437	(0.071)	-0.576	-0.298
Tenure > 2	δ_{X1}^{WE}	0.096	(0.209)	-0.315	0.506
Job offer probabilities to SE					
Constant	$\lambda_{C,SE}^{NE}$	8.039	(1.520)	5.060	11.017
Experience	$\lambda_{T1,SE}^{NE}$	-0.559	(0.051)	-0.658	-0.460
Experience ²	$\lambda_{T2,SE}^{NE}$	1.204	(0.377)	0.465	1.943
Tenure > 2	$\lambda_{X1,SE}^{NE}$	-1.066	(0.829)	-2.692	0.559
Effort	$\lambda_{E,SE}^{NE}$	-0.526	(0.512)	-1.531	0.478
Job offer probabilities to WE					
Constant	$\lambda_{C,WE}^{NE}$	2.283	(0.286)	1.722	2.843
Experience	$\lambda_{T1,WE}^{NE}$	-0.169	(0.022)	-0.212	-0.127
Experience ²	$\lambda_{T1,WE}^{NE}$	0.310	(0.066)	0.181	0.439
Tenure > 2	$\lambda_{X1,WE}^{NE}$	-0.103	(0.329)	-0.749	0.542
Effort	$\lambda_{E,WE}^{NE}$	-0.194	(0.120)	-0.430	0.041
Sector switch offer SE to WE					
Constant	$\lambda_{C,WE}^{SE}$	1.106	(0.435)	0.254	1.958
Experience	$\lambda_{T2,WE}^{SE}$	-0.070	(0.024)	-0.117	-0.023
Experience ²		0.099	(0.074)	-0.046	0.243
Sector switch offer WE to SE					
Constant	$\lambda_{C,SE}^{WE}$	1.390	(0.439)	0.531	2.250
Experience	$\lambda_{T1,SE}^{WE}$	-0.115	(0.030)	-0.174	-0.055
Experience ²	$\lambda_{T2,SE}^{WE}$	0.221	(0.091)	0.042	0.400

Additional parameter estimates to complement estimates in table 4. Estimates obtained via indirect inference estimation on sample of 3,307 panel observations of male Ghanaians aged 15-65 with education less than senior high school degree.

D.4 Replication of the Flypaper Effect

We replicate the estimation results from Fafchamps, McKenzie, Quinn, and Woodruff (2014) (FMQW) using the replication code and dataset provided by the authors. Since our structural model not only models the main treatment effects but also the *mechanism* which give rise to the effect, our aim is to anchor our model in as many dimensions as possible. We therefore not only wish to reproduce the main treatment effects (which FMQW report in their Table 3) but also the treatment effects which relate to the underlying mechanisms in their Table 6. We follow the specification FMQW use to report the secondary treatment effects. This specification runs a fixed effects regression on a trimmed sample (where outliers in terms of profits are removed), and winsorises all monetary variables at the 99.5th percentile. This corresponds closely to the way we dealt with our data from the GUHPS, where we also remove extreme outliers and winsorise all monetary variables.

The fixed effects regression most closely mimicks our calculation of structural treatment effects. The latter is effectively a within-individual comparison. Since we are able to calculate both counterfactuals (under treatment and in absence of treatment) with our structural model, our structural ATE simply averages over individual treatment effects that result from a simple difference of counterfactuals.

In the first three columns of Table A22, we verify that we are able to replicate the treatment effects from FMQW. We apply all steps of data transformation that the authors make. Then, we restrict our sample to men only. In column (1), we replicate the coefficients for the male-treatment interaction terms from column (9) of Table 3 in their paper. In column (2), we similarly replicate the effect on truncated capital stock from column (2) of Table 6. Finally, in column (3) we report the coefficients for total annual spending. This corresponds to column (9) from Table 6. The coefficients are different because this reported regression in FMQW, unlike the previous two regressions, does not control for gender levels. (We are able to replicate the entire regression, not just the male coefficients, with the full sample which includes women).

Having satisfied ourselves that we can replicate the results in FMQW, we transform the data so that it matches with our simulated data from the model. We use the CPI from the Bank of Ghana to rebase monetary variables to June 2010, our nominal base. We further

D.4. REPLICATION OF THE FLYPAPER EFFECT

aggregate monthly to yearly profits. In columns (4)-(6) of Table 5, we repeat the same regressions on our transformed variables. In columns (7)-(9) we pool the treatments into a single indicator. We use the last two sets of results to benchmark our structural estimates in Table 5 in the main text.

Table A22: Treatment effects from the Flypaper Experiment

Dependent Variable:	May 2009 cedis				June 2010 cedis				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Monthly Profits	Capital Stock	Annual Consumption	Annual Profits	Capital Stock	Annual Consumption	Annual Profits	Capital Stock	Annual Consumption
Cash treatment	5.132 (16.13)	2.212 (61.20)	85.73 (59.79)	69.11 (217.2)	2.482 (68.68)	96.20 (67.10)			
In-kind treatment	27.83 (18.18)	83.74 (69.96)	43.12 (72.38)	374.7 (244.8)	93.97 (78.51)	48.39 (81.22)			
Treatment							221.8 (176.9)	46.94 (59.41)	72.82 (60.81)
R ²	0.01	0.02	0.02	0.01	0.02	0.02	0.01	0.02	0.02
Observations	1,599	1,602	1,705	1,599	1,602	1,705	1,599	1,602	1,705

Standard errors in parentheses

This table re-estimates treatment effects from Faichamps, McKenzie, Quinn, and Woodruff (2014) using fixed effects regressions on the trimmed sample from the paper, and following the truncation at the 99.5th percentile. We also use the same control variables for survey wave and randomisation group. Columns (1), (4) and (7) replicate the treatment effects on profits, capital and total household expenditure (consumption) reported in the paper. The other columns report the treatment effects in converted units (2010 cedis and annualised flow variables).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Bibliography

- ABEBE, G., S. CARIA, M. FAFCHAMPS, P. FALCO, S. FRANKLIN, AND S. QUINN (2016): “Curse of Anonymity or Tyranny of Distance? The Impacts of Job-Search Support in Urban Ethiopia,” Working Paper 22409, National Bureau of Economic Research.
- ACEMOGLU, D., D. AUTOR, D. DORN, G. H. HANSON, AND B. PRICE (2016): “Import Competition and the Great US Employment Sag of the 2000s,” *Journal of Labor Economics*, 34(S1 Part 2), S141–S198.
- ADAM, S., H. MILLER, AND T. POPE (2017): “Tax, legal form and the gig economy,” in *The IFS Green Budget: February 2017*, ed. by C. Emmer-son, P. Johnson, and R. Joyce, pp. 203–238. Institute for Fiscal Studies.
- ADÃO, R. (2015): “Worker Heterogeneity, Wage Inequality, and Interna-tional Trade: Theory and Evidence from Brazil,” Job Market Paper, MIT.
- ADDA, J., C. DUSTMANN, C. MEGHIR, AND J.-M. ROBIN (2013): “Ca-reer Progression, Economic Downturn, and Skills,” *NBER Working Paper No. 18832*.
- ADDA, J., C. DUSTMANN, AND K. STEVENS (2017): “The Career Costs of Children,” *Journal of Political Economy*, 125(2), 293–337.
- AGHION, P., AND P. BOLTON (1997): “A theory of trickle-down growth and development,” *The Review of Economic Studies*, 64(2), 151–172.
- AHN, H., AND J. L. POWELL (1993): “Semiparametric estimation of cen-sored selection models with a nonparametric selection mechanism,” *Jour-nal of Econometrics*, 58(1-2), 3–29.
- ALTONJI, J. G., AND L. M. SEGAL (1996): “Small Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business and Eco-nomics Statistics*, 14(3), 353–366.
- ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): “Measur-ing the Sensitivity of Parameter Estimates to Estimation Moments,” *The Quarterly Journal of Economics*, forthcoming.

- ANGELUCCI, M., D. KARLAN, AND J. ZINMAN (2015): "Microcredit Impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco," *American Economic Journal: Applied Economics*, 7(1), 151–82.
- ARADILLAS-LOPEZ, A., B. E. HONORÉ, AND J. L. POWELL (2007): "Pairwise difference estimation with nonparametric control variables," *International Economic Review*, 48(4), 1119–1158.
- ATKIN, D. (2016): "Endogenous Skill Acquisition and Export Manufacturing in Mexico," *American Economic Review*, 106(8), 2046–85.
- ATKIN, D., A. CHAUDHRY, S. CHAUDRY, A. KHANDELWAL, AND E. VERHOOGEN (2017): "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan," *The Quarterly Journal of Economics*, 132(3), 1101–1164.
- ATTANASIO, O. P., C. MEGHIR, AND A. SANTIAGO (2011): "Education choices in Mexico: using a structural model and a randomized experiment to evaluate Progreso," *The Review of Economic Studies*, 79(1), 37–66.
- AUTOR, D., D. DORN, AND G. H. HANSON (2013): "The China syndrome: Local labor market effects of import competition in the United States," *The American Economic Review*, 103(6), 2121–2168.
- BANERJEE, A., E. BREZA, E. DUFLO, AND C. KINNAN (2015): "Do credit constraints limit entrepreneurship? Heterogeneity in the returns to microfinance," Mimeo.
- BANERJEE, A., AND E. DUFLO (2011): *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*. PublicAffairs.
- BANERJEE, A., E. DUFLO, R. GLENNERSTER, AND C. KINNAN (2015): "The Miracle of Microfinance? Evidence from a Randomized Evaluation," *American Economic Journal: Applied Economics*, 7(1), 22–53.
- BANERJEE, A., D. KARLAN, AND J. ZINMAN (2015): "Six Randomized Evaluations of Microcredit: Introduction and Further Steps," *American Economic Journal: Applied Economics*, 7(1), 1–21.
- BANERJEE, A., AND S. MULLAINATHAN (2010): "The shape of temptation: Implications for the economic lives of the poor," Working Paper 15973, National Bureau of Economic Research.
- BANERJEE, A. V., AND A. F. NEWMAN (1993): "Occupational choice and the process of development," *Journal of Political Economy*, pp. 274–298.
- BARTELSMAN, E., J. HALTIWANGER, AND S. SCARPETTA (2013): "Cross-country differences in productivity: The role of allocation and selection," *The American Economic Review*, 103(1), 305–334.

- BENJAMIN, D. (1992): "Household composition, labor markets, and labor demand: testing for separation in agricultural household models," *Econometrica*, 60(2), 287–322.
- BIANCHI, M., AND M. BOBBA (2013): "Liquidity, risk, and occupational choices," *The Review of Economic Studies*, 80(2), 491–511.
- BLANCHFLOWER, D., AND A. OSWALD (1998): "What Makes an Entrepreneur?," *Journal of Labor Economics*, 16(1), 26–60.
- BLATTMAN, C., AND S. DERCON (2016): "Occupational Choice in Early Industrializing Societies: Experimental Evidence on the Income and Health Effects of Industrial and Entrepreneurial Work," Working Paper 22683, National Bureau of Economic Research.
- BLATTMAN, C., N. FIALA, AND S. MARTINEZ (2014): "Generating Skilled Self-Employment in Developing Countries: Experimental Evidence from Uganda," *Quarterly Journal of Economics*, 129(2).
- BLOOM, N., B. EIFERT, A. MAHAJAN, D. MCKENZIE, AND J. ROBERTS (2013): "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 128(1), 1–51.
- BLUNDELL, R., M. COSTA DIAS, C. MEGHIR, AND J. SHAW (2016): "Female labor supply, human capital, and welfare reform," *Econometrica*, 84(5), 1705–1753.
- BLUNDELL, R., A. DUNCAN, AND C. MEGHIR (1998): "Estimating labor supply responses using tax reforms," *Econometrica*, 66(4), 827–861.
- BOSCH, M., AND R. M. CAMPOS-VAZQUEZ (2014): "The trade-offs of welfare policies in labor markets with informal jobs: the case of the Seguro Popular program in Mexico," *American Economic Journal: Economic Policy*, 6(4), 71–99.
- BOSCH, M., AND W. MALONEY (2010): "Comparative analysis of labor market dynamics using markov processes: An application to informality," *Labour Economics*, 17(4), 621–631.
- BRUHN, M. (2013): "A tale of two species: Revisiting the effect of registration reform on informal business owners in Mexico," *Journal of Development Economics*, 103, 275–283.
- BRUHN, M., AND I. LOVE (2014): "The real impact of improved access to finance: Evidence from Mexico," *The Journal of Finance*, 69(3), 1347–1376.
- CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "Wage bargaining with on-the-job search: Theory and evidence," *Econometrica*, 74(2), 323–364.

- CARD, D. (1996): "The effects of unions on the structure of wages: A longitudinal analysis," *Econometrica*, 64(4), 957–979.
- CHAMBERLAIN, G. (1984): "Panel Data," in *Handbook of Econometrics*, ed. by Z. Griliches, and M. Intriligator, vol. 2, pp. 1247–1318. North Holland.
- (1986): "Asymptotic efficiency in semi-parametric models with censoring," *Journal of Econometrics*, 32(2), 189–218.
- CHESHER, A. (2010): "Instrumental Variable Models for Discrete Outcomes," *Econometrica*, 78(2), 575–601.
- COHEN, A. (2016): "Constraints on Labor and Land and the Return to Microcredit and Microinsurance," Mimeo.
- COSTA, F., J. GARRED, AND J. P. PESSOA (2016): "Winners and losers from a commodities-for-manufactures trade boom," *Journal of International Economics*, 102, 50–69.
- CUESTA, J., AND C. BOHÓRQUEZ (2011): "Labour market transitions and social security in Colombia," Discussion paper, World Bank Policy Research Working Paper No. 5650.
- DAUTH, W., S. FINDEISEN, AND J. SUEDEKUM (2014): "The rise of the East and the Far East: German labor markets and trade integration," *Journal of the European Economic Association*, 12(6), 1643–1675.
- DAVID, J. M., H. A. HOPENHAYN, AND V. VENKATESWARAN (2016): "Information, Misallocation and Aggregate Productivity," *The Quarterly Journal of Economics*, 131(2), 943–1005.
- DAVID, J. M., AND V. VENKATESWARAN (2017): "Capital Misallocation: Frictions or Distortions," Working Paper 23129, National Bureau of Economic Research.
- DE JANVRY, A., M. FAFCHAMPS, AND E. SADOULET (1991): "Peasant household behaviour with missing markets: some paradoxes explained," *The Economic Journal*, 101(409), 1400–1417.
- DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2008): "Returns to capital in microenterprises: Evidence from a field experiment," *The Quarterly Journal of Economics*, 123(4), 1329–1372.
- (2010): "Who are the microenterprise owners? Evidence from Sri Lanka on Tokman v. de Soto," in *International Differences in Entrepreneurship*, ed. by J. Lerner, and A. Schoar, pp. 63–87. University of Chicago Press.

- (2012): “One-time transfers of cash or capital have long-lasting effects on microenterprises in Sri Lanka,” *Science*, 335(6071), 962–966.
- DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2016): “Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises,” World Bank Policy Research Working Paper 7924.
- DEATON, A. (2010a): “Instruments, Randomization, and Learning about Development,” *Journal of Economic Literature*, 48(2), 424–455.
- (2010b): “Understanding the Mechanisms of Economic Development,” *Journal of Economic Perspectives*, 24(3), 3–16.
- DEL VALLE, A. (2014): “From Caring to Work: The Labor Market Effects of Noncontributory Health Insurance,” Mimeo.
- DILLON, E. W., AND C. T. STANTON (2017): “Self-employment dynamics and the returns to entrepreneurship,” Working Paper 23168, National Bureau of Economic Research.
- DIX-CARNEIRO, R. (2014): “Trade liberalization and labor market dynamics,” *Econometrica*, 82(3), 825–885.
- ECKSTEIN, Z., M. KEANE, AND O. LIFSHITZ (2016): “Sources of Change in the Life-Cycle Decisions of American Men and Women: 1962-2014,” Mimeo.
- ECKSTEIN, Z., AND G. VAN DEN BERG (2007): “Empirical Labor Search: A Survey,” *Journal of Econometrics*, 136(2), 531–564.
- ECKSTEIN, Z., AND K. WOLPIN (1989): “Dynamic Labour Force Participation of Married Women and Endogenous Work Experience,” *Review of Economic Studies*, 56(3), 375–390.
- (1990): “Estimating a Market Equilibrium Search Model from Panel Data on Individuals,” *Econometrica*, 58(4), 783–808.
- EISENHAUER, P., J. J. HECKMAN, AND E. VYTLACIL (2015): “The Generalized Roy Model and the Cost-Benefit Analysis of Social Programs,” *Journal of Political Economy*, 123(2), 413–443.
- ESPINOSA, S. (2012): “Social Protection and Incentives to Labour Informality: A Regression Discontinuity Design applied to Colombia’s Non-Contributory Health Regime,” Mimeo.
- EVANS, D., AND B. JOVANOVIĆ (1989): “An Estimated Model of Entrepreneurial Choice under Liquidity Constraints,” *Journal of Political Economy*, 97(4), 808–827.

- EVANS, D., AND L. LEIGHTON (1989): "Some Empirical Aspects of Entrepreneurship," *American Economic Review*, 79(3), 519–535.
- EVANS, D. S. (1987): "Tests of alternative theories of firm growth," *Journal of Political Economy*, 95(4), 657–674.
- FAFCHAMPS, M., D. MCKENZIE, S. QUINN, AND C. WOODRUFF (2012): "Using PDA Consistency Checks to Increase the Precision of Profits and Sales Measurement in Panels," *Journal of Development Economics*, 98(1), 51–57.
- (2014): "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana," *Journal of Development Economics*, 116, 211–226.
- FAFCHAMPS, M., AND S. QUINN (2016): "Aspire," *Journal of Development Studies*, forthcoming.
- FAFCHAMPS, M., AND C. WOODRUFF (2016): "Identifying Gazelles: Expert Panels vs. Surveys as a Means to Identify Firms with Rapid Growth Potential," *World Bank Economic Review*, forthcoming.
- FAMA, E. F., AND M. C. JENSEN (1983): "Separation of ownership and control," *The Journal of Law and Economics*, 26(2), 301–325.
- FIELDS, G. (1975): "Rural-urban Migration, Urban Unemployment and Underemployment, and Job-Search Activity in LDCs," *Journal of Development Economics*, 2(2), 165–187.
- FRANKLIN, S. (2016): "Location, Search Costs and Youth Unemployment: A Randomized Trial of Transport Subsidies in Ethiopia," *Economic Journal*, forthcoming.
- FRENCH, E., AND C. TABER (2011): "Identification of Models of the Labor Market," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4, Part A, pp. 537–617.
- FRÍAS, J. A., D. S. KAPLAN, AND E. VERHOOGEN (2012): "Exports and Within-Plant Wage Distributions: Evidence from Mexico," *American Economic Review Papers and Proceedings*, 102(3), 435–440.
- GARICANO, L., C. LELARGE, AND J. VAN REENEN (2016): "Firm Size Distortion and the Productivity Distribution: Evidence from France," *American Economic Review*, 106(11), 3439–3479.
- GOETTLER, R., AND B. GORDON (2011): "Does AMD Spur Intel to Innovate More?," *Journal of Political Economy*, 119(6), 1141–1200.

- GOLLIN, D. (2008): "Nobody's business but my own: Self-employment and small enterprise in economic development," *Journal of Monetary Economics*, 55(2008), 219–233.
- GOTTLIEB, J. D., R. R. TOWNSEND, AND T. XU (2016): "Experimenting with entrepreneurship: the effect of job-protected leave," Working Paper 22446, National Bureau of Economic Research.
- GOURIÉROUX, C., A. MONFORT, AND E. RENAULT (1993): "Indirect Inference," *Journal of Applied Econometrics*, 8, S85–S118.
- GROH, M., N. KRISHNAN, D. MCKENZIE, AND T. VISHWANATH (2016): "Do Wage Subsidies Provide a Stepping Stone to Employment for Recent College Graduates? Evidence from a Randomized Experiment in Jordan," *The Review of Economics and Statistics*, 98(3), 488–502.
- GÜNTHER, I., AND A. LAUNOV (2012): "Informal Employment in Developing Countries: Opportunity or Last Resort?," *Journal of Development Economics*, 97(1), 88–98.
- HAMILTON, B. H. (2000): "Does Entrepreneurship Pay? An Empirical Analysis of the Returns to Self-Employment," *Journal of Political Economy*, 108(3), 604–631.
- HARDY, M., AND J. MCCASLAND (2015): "Are small firms labor constrained? Experimental evidence from Ghana," Mimeo.
- HARRIS, J., AND M. TODARO (1970): "Migration, Unemployment and Development: A Two-Sector Analysis," *American Economic Review*, 60(1), 126–142.
- HAYWOOD, L., AND P. FALCO (2016): "Entrepreneurship vs Joblessness: Explaining the Rise in Self-Employment," *Journal of Development Economics*, 118, 245–265.
- HEATH, R., AND A. M. MOBARAK (2015): "Manufacturing growth and the lives of Bangladeshi women," *Journal of Development Economics*, 115, 1–15.
- HECKMAN, J., AND B. HONORÉ (1990): "The Empirical Content of the Roy Model," *Econometrica*, 58(5), 1121–1149.
- HECKMAN, J. J., H. ICHIMURA, AND P. E. TODD (1997): "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme," *The Review of Economic Studies*, 64(4), 605–654.
- HECKMAN, J. J., AND E. J. VYTLACIL (2007): "Econometric evaluation of social programs, part I: Causal models, structural models and econometric policy evaluation," in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. Leamer, vol. 6, pp. 4779–4874. Elsevier.

- HOMBERT, J., A. SCHOAR, D. SRAER, AND D. THESMAR (2014): “Can Unemployment Insurance Spur Entrepreneurial Activity?,” Working Paper 20717, National Bureau of Economic Research.
- HONORÉ, B. E., AND J. L. POWELL (2005): “Pairwise Difference Estimation of Nonlinear Models,” in *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, ed. by D. W. Andrews, and J. H. Stock, pp. 520–535. Cambridge University Press.
- HOPENHAYN, H. (1992): “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60(5), 1127–1150.
- HSIEH, C.-T., AND P. J. KLENOW (2009): “Misallocation and Manufacturing TFP in China and India,” *The Quarterly Journal of Economics*, 124(4), 1403–1448.
- HUMPHRIES, J. E. (2017): “The Causes and Consequences of Self-Employment over the Life Cycle,” Job Market Paper, University of Chicago.
- IACOVONE, L., F. RAUCH, AND L. A. WINTERS (2013): “Trade as an engine of creative destruction: Mexican experience with Chinese competition,” *Journal of International Economics*, 89(2), 379–392.
- IMBERT, C. (2013): “Decomposing the Labor Market Earnings Inequality: The Public and Private Sectors in Vietnam, 1993-2006,” *World Bank Economic Review*, 27(1), 55–79.
- JOLIVET, G., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): “The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the US,” *European Economic Review*, 50(4), 877–907.
- JOVANOVIC, B. (1982): “Selection and the Evolution of Industry,” *Econometrica*, 50(3), 649–670.
- KAPLAN, D. S., D. LEDERMAN, AND R. ROBERTSON (2011): “Employment and Wage Responses to Trade Shocks: Evidence from Mexico during the 2008-09 U.S. Recession,” Mimeo.
- KARLAN, D., R. KNIGHT, AND C. UDRY (2015): “Consulting and Capital Experiments with Microenterprise Tailors in Ghana,” *Journal of Economic Behaviour and Organization*, 118, 281–302.
- KARLAN, D., AND M. VALDIVIA (2011): “Teaching Entrepreneurship: Impact of Business Training on Microfinance Clients and Institutions,” *Review of Economics and Statistics*, 93(2), 510–527.

- KATZ, L. F., AND A. B. KRUEGER (2016): "The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015," Working Paper 22667, National Bureau of Economic Research.
- KEANE, M., P. TODD, AND K. WOLPIN (2011): "The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4, Part A, pp. 331–461.
- KEANE, M., AND K. WOLPIN (1997): "The Career Decisions of Young Men," *Journal of Political Economy*, 105(3), 473–522.
- KENISTON, D. E. (2011): "Experimental vs. Structural Estimates of the Return to Capital in Microenterprises," Mimeo.
- KIHLSTROM, R., AND J.-J. LAFFONT (1979): "A General Equilibrium Entrepreneurial Theory of Firm Formation Based on Risk Aversion," *Journal of Political Economy*, 87(4), 719–748.
- KILLINGSWORTH, M. R., AND J. J. HECKMAN (1986): "Female labor supply: A survey," in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and R. Layard, vol. 1, pp. 103–204. Elsevier.
- KRAAY, A., AND D. MCKENZIE (2014): "Do Poverty Traps Exist? Assessing the Evidence," *Journal of Economic Perspectives*, 28(3), 127–148.
- LAFAVE, D., AND D. THOMAS (2016): "Farms, Families, and Markets: New Evidence on Completeness of Markets in Agricultural Settings," *Econometrica*, 84(5), 1917–1960.
- LEMIEUX, T. (1998): "Estimating the Effects of Unions on Wage Inequality in a Panel Data Model with Comparative Advantage and Nonrandom Selection," *Journal of Labor Economics*, 16(2), 261–291.
- LERNER, J., AND A. SCHOAR (2010): "Introduction to International Differences in Entrepreneurship," in *International Differences in Entrepreneurship*, ed. by J. Lerner, and A. Schoar, pp. 1–13. University of Chicago Press.
- LEVINE, R., AND Y. RUBINSTEIN (2017): "Smart and Illicit: Who Becomes an Entrepreneur and Do They Earn More?," *The Quarterly Journal of Economics*, 132(2), 963.
- LEVY, S. (2010): *Good intentions, bad outcomes: Social policy, informality, and economic growth in Mexico*. Brookings Institution Press.
- LEVY, S., AND N. SCHADY (2013): "Latin America's social policy challenge: education, social insurance, redistribution," *The Journal of Economic Perspectives*, 27(2), 193–218.

- LEWBEL, A. (2017): "The Identification Zoo – Meanings of Identification in Econometrics," Mimeo.
- LEWIS, W. A. (1954): "Economic Development with Unlimited Supplies of Labour," *The Manchester School*, 22(2), 139–191.
- LISE, J. (2013): "On-the-Job Search and Precautionary Savings," *Review of Economic Studies*, 80(3), 1086–1113.
- LOPEZ-GARCIA, I. (2013): "Human Capital and Labor Informality in Developing Countries: A Life-Cycle Approach," Job Market Paper.
- LOW, H., C. MEGHIR, L. PISTAFERRI, AND A. VOENA (2017): "Marriage, Labor Supply and the Social Safety Net," Mimeo.
- LUCAS, R. (1978): "On the Size Distribution of Business Firms," *The Bell Journal of Economics*, 9(2), 508–523.
- MAGNAC, T. (1991): "Segmented or Competitive Labor Markets," *Econometrica*, 59(1), 165–187.
- MALONEY, W. F. (1999): "Does informality imply segmentation in urban labor markets? Evidence from sectoral transitions in Mexico," *The World Bank Economic Review*, 13(2), 275–302.
- MCCAIG, B., AND N. PAVCNIK (2014): "Export markets and labor allocation in a low-income country," Working Paper 20455, National Bureau of Economic Research.
- MCCALL, J. (1970): "Economics of Information and Job Search," *The Quarterly Journal of Economics*, 84, 113–126.
- MCKENZIE, D. (2017): "Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition," *American Economic Review*, 107(8), 2278–2307.
- MCKENZIE, D., AND C. WOODRUFF (2008): "Experimental evidence on returns to capital and access to finance in Mexico," *The World Bank Economic Review*, 22(3), 457–482.
- (2014): "What Are We Learning from Business Training and Entrepreneurship Evaluations around the Developing World?," *The World Bank Research Observer*, 29(1), 48–82.
- (2016): "Business Practices in Small Firms in Developing Countries," *Management Science*, forthcoming.
- MEGHIR, C., R. NARITA, AND J.-M. ROBIN (2015): "Wages and Informality in Developing Countries," *American Economic Review*, 105(4), 1509–46.

- MELITZ, M. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71, 1695–1725.
- MENDEZ, O. (2014): “Trade Shocks and Mexican Local Labor Markets in the Great Recession,” Job Market Paper, UC Davis.
- MONDRAGÓN-VÉLEZ, C., AND X. PEÑA (2010): “Business Ownership and Self-Employment in Developing Economies: The Colombian Case,” in *International Differences in Entrepreneurship*, ed. by J. Lerner, and A. Schoar, pp. 89–127. University of Chicago Press.
- MORTENSEN, D. (1970): “Job search, the Duration of Unemployment, and the Phillips Curve,” *American Economic Review*, 60(5), 847–862.
- NARITA, R. (2014): “Self-Employment in Developing Countries: A Search-Equilibrium Approach,” Mimeo.
- PIKETTY, T. (1997): “The dynamics of the wealth distribution and the interest rate with credit rationing,” *The Review of Economic Studies*, 64(2), 173–189.
- POLKOVNICHENKO, V. (2003): “Human Capital and the Private Equity Premium,” *Review of Economic Dynamics*, 6(4), 831–845.
- POWELL, J. L. (1987): *Semiparametric estimation of bivariate latent variable models*. University of Wisconsin–Madison, Social Systems Research Institute Working Paper No. 8704.
- (2001): “Semiparametric estimation of censored selection models,” in *Nonlinear Statistical Modeling: Proceedings of the Thirteenth International Symposium in Economic Theory and Econometrics: Essays in Honor of Takeshi Amemiya*. Cambridge University Press.
- RAMSEY, F. (1928): “A Mathematical Theory of Saving,” *The Economic Journal*, 38, 543–559.
- RENDON, S. (2006): “Job Search and Asset Accumulation Under Borrowing Constraints,” *International Economic Review*, 47(1), 233–263.
- RIDDER, G., AND G. VAN DEN BERG (2003): “Measuring Labor Market Frictions: A Cross-Country Comparison,” *Journal of the European Economic Association*, 1(1), 224–244.
- ROY, A. (1951): “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 3(2), 135–146.
- SCHULTZ, T. P. (2004): “School subsidies for the poor: evaluating the Mexican Progresa poverty program,” *Journal of Development Economics*, 74(1), 199–250.

- SMITH, J. A., AND P. E. TODD (2001): “Reconciling conflicting evidence on the performance of propensity-score matching methods,” *The American Economic Review*, 91(2), 112–118.
- SÖDERBOM, M., AND F. TEAL (2004): “Size and efficiency in African manufacturing firms: Evidence from firm-level panel data,” *Journal of Development Economics*, 73(1), 369–394.
- STIGLER, G. (1962): “Information in the Labor Market,” *Journal of Political Economy*, 70, 94–105.
- SURI, T. (2011): “Selection and Comparative Advantage in Technology Adoption,” *Econometrica*, 79(1), 159–209.
- TAUCHEN, G. (1986): “Finite state markov-chain approximations to univariate and vector autoregressions,” *Economics Letters*, 20(2), 177 – 181.
- TODD, P. E., AND K. I. WOLPIN (2006): “Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility,” *American Economic Review*, 96(5), 1384–1417.
- TOPALOVA, P. (2010): “Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India,” *American Economic Journal: Applied Economics*, 2(4), 1–41.
- ULYSSEA, G. (2017): “Firms, Informality and Development: Theory and Evidence from Brazil,” Mimeo.
- VERESHCHAGINA, G., AND H. A. HOPENHAYN (2009): “Risk taking by entrepreneurs,” *The American Economic Review*, 99(5), 1808–1830.
- VERHOOGEN, E. A. (2008): “Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector,” *The Quarterly Journal of Economics*, 123(2), 489–530.