Jeremy R. Magruder

Research Statement
Fall, 2012

The focus of my research is on employment and unemployment in developing countries, where labor markets are characterized by low formal sector employment, high levels of informal employment and entrepreneurship (which often exhibits low levels of productivity), and high levels of structural unemployment. My work has focused on three aspects of employment in this setting. The first is how the unemployed find work, and what aspects of job search may help or hurt different unemployed people. The second is what labor market policies influence labor demand, which affects the number of unemployed who will be able to find work. The third is understanding the factors which push firms and workers into long run success or back into unemployment.

From the perspective of the unemployed, I have studied the practical relevance of the use of social networks in finding work. My research on social networks has focused on understanding how people in networks make the choice of who to share job opportunities with, which people become excluded as a consequence of that decision, and what employers can gain from networks. Considering labor demand, I studied the relationship between labor regulations which are commonplace in developing countries and labor demand from firms, with care to explore differences between small and large, formal and informal firms. Finally, I have several projects which explore other determinants of firm and labor market success, both in developing countries and the US, including behavioral biases such as attention shortages; the role of Social Media; and the response of different workers to immigration.

Of course, social networks are used for other purposes as well as sharing job opportunities. Building on my research in social networks and job search, one additional project examines how networks can be used to disseminate information about new agricultural technologies, while a second asks what the job search literature can teach us in understanding the behavioral underpinnings of the HIV/AIDS epidemic in Africa.

A): Social Networks in Labor Markets in Developing Countries

About half of jobs globally are spread through social networks. A large literature in both economics and sociology has convinced us of this stylized fact through both observational data and through a variety of quasi-experimental research designs. Yet while we can have confidence that social networks are important, we know much less about how specifically they work in sharing job opportunities. To an unemployed person, the relevant question has to be not "do people find work through social networks," but rather "can I expect my network to work for me?" This question has in general been a challenge to answer through experimental or quasi-experimental means, as social networks are almost in principle complicated, hard-to-observe entities. My research focuses on empirical analyses of understanding how social networks choose to allocate job opportunities, and what the consequences of that allocation are for the unemployed and for employers.

In a first paper **"Intergenerational Networks, Unemployment, and Persistent Inequality in South Africa,"** which is solo-authored and published in the *AEJ: Applied Economics,* I examine how young adults find work in South Africa. South Africa is a country with tremendously high unemployment and scarce job opportunities. In such a context, we might imagine that having the right connections to a job would be quite important – and that individuals without the right connections may find themselves in long term unemployment.

I utilize a panel data set on young adults in Cape Town, and identify that fathers are particularly important network connections for their sons when they are capable (but not their daughters). Using a triple-differenced design which takes advantage of the facts that jobs are gender segregated (and therefore girls are unlikely to benefit from their fathers' network help) and that absent fathers can contribute money but not job opportunities, I find that sons with present fathers are substantially more likely to be working when their fathers' industries are growing, in contrast to sons with absent fathers and all daughters (who are less likely to be working, consistent with the hypothesis of wealth effects impacting labor supply). Moreover, sons of present fathers are more likely to be working in their fathers' industries, and these effects cannot be explained by a variety of hypotheses other than the use of social connections (for example, skills specific to an industry and correlated networks other than the father). I conclude that the use of networks is contributing to low mobility and inequality in South Africa, as these nepotistic uses generate a benefit for relatively well-off sons.

A second paper, "**Who gets the Job Referral? Evidence from a Social Networks Experiment**," which is co-authored with Lori Beaman and forthcoming in *The American Economic Review* uses experimental methods to ask directly who individuals choose for job opportunities. This paper is the first to test experimentally both whether people have useful information which allows them to screen on their employer's behalf when making a referral, and also whether they are willing to use that information. We propose that social incentives created by the network create tradeoffs between an employee's desire to refer the person he thinks is most qualified for a position and the person who gives him the highest social payments. To test this, we created a laboratory experiment in the field in Kolkata, India. The main idea is that once participants arrive at our laboratory, complete a task, and we pay them for it, we have essentially given them a casual day-job. We can then allow them to find referrals for our experiment and offer them different contracts for making the referral, so that some are (at random) offered a fixed fee which they receive for making any referral, while others are offered a performance incentive which they receive only if their referral performs sufficiently well. We document that individuals do face a tradeoff between socially-preferred referrals and incentives given them by their employers, as they refer relatives more frequently in the absence of performance-based incentives and more socially distant coworkers in the presence of performance incentives. However, while we find that all individuals on average change their referral choice, only high ability people do so in a way which results in more skilled referrals coming in. We therefore confirm that high ability people have the right information to screen on their employers' behalf, but they would prefer not to use that information. Low ability people, by contrast, show little evidence of having useful information.

Finally, a current working paper, "**Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi"** which is co-authored with Lori Beaman and Niall Keleher asks whether characteristics of the referral choice problem lead to the systematic exclusion of women. The potential of referral networks to exclude disadvantaged groups and lead to long-run unequal outcomes has been described in the theoretical literature, but there is to this point little direct empirical evidence. In this experiment, we worked with a survey firm in Malawi who were hiring enumerators to learn whether the use of referral systems excludes women. We advertised enumerator positions using conventional methods and then asked those applicants to make referrals. To learn both about the referrals of each gender that applicants wanted to make versus the referrals that they could make, we randomized along two dimensions: whether applicants were told they must refer a man, they must refer a woman, or they could refer someone of either gender, and also whether they would receive a fixed fee for making any referral or a premium for referring someone we would actually hire. We find stark results. First, the use of referral systems strongly disadvantages qualified women. This happens both because most men choose to refer other men when given a choice of which gender to refer, and while women show less evidence of gender bias, they are much less likely to refer qualified applicants of either gender than men are.   As a result, neither men nor women are likely to refer qualified women, even though we can also show that both men and women are perfectly willing to make female referrals when prompted. We then develop a model which allows networks to be heterogeneous in several characteristics and find that this disadvantage occurs because of social payments, rather than expected differences in ability. We also show that productive information is different between the genders, with network screening maximized when men refer men. As a result, we conclude not only that referral systems disadvantage women but that employers have incentives to encourage this disadvantage: allowing referrals to be systematically male both maximizes their employee's social payments and leads to the most useful screening. This result suggests the need for creative policy to overcome this labor market disadvantage.

B): How do Labor Regulations Influence Labor Demand?

Labor markets in developing countries are often characterized by relatively high labor standards in the formal sector. These standards are often presumed to play a strong role in the bimodal firm size distribution present in many developing countries (where a majority of firms are very small, and some firms are very large, but there are few medium-sized firms) and also the large informal sectors which employ many. However, labor regulations are often difficult to study in developing countries due to the limited potential for randomization and the fact that policies are enacted purposefully, often over large regions of the economy, creating a correlation between labor regulation and underlying labor market conditions. My work builds on a large literature from the developed world and smaller literature within developing countries in testing the effects of labor regulations on labor demand and firm size.

A first paper, "**High Unemployment Yet Few Small Firms: The Role of Centralized Bargaining in South Africa,**" which is solo-authored and was published in the *AEJ: Applied Economics*, focuses on understanding the implications of South Africa's Bargaining Council system on unemployment and entrepreneurship. Understanding the constraints which prevent employment and entrepreneurship is one of the most important policy questions to South Africa, as unemployment is very high and yet adults do not respond by engineering work for themselves through informal entrepreneurship. One of the frequent suspects in policy debate is the bargaining council system, whereby unions and firms can choose to extend labor standards that they agree upon to all workers in the same industry in a given political demarcation. These standards may be especially difficult for small firms to comply with, as small firms have low levels of capital and presumably lower marginal products of labor. Despite the importance of this policy in South Africa and similar policies in Latin America and Western Europe, there are no credible estimates of its effect on unemployment and firm size. Causal inference on this policy is challenged by the fact that unions and firms make a choice to implement the policy in particular places, which suggests that its existence may be correlated with a variety of unobserved underlying characteristics of local labor markets. In this paper, I suggest using a spatial discontinuity to identify these effects, and find that bargaining councils are increasing wages by 10-20% while reducing employment by about 10%, and that the entire burden of reduced employment is coming from small firms. Moreover, I note that the labor force is also becoming systematically more male and more tenured, suggesting that new workers and women are disadvantaged the most by this policy. Finally, I apply a variety of robustness techniques, illustrating that while firms are systematically resettling on the opposite edge of a bargaining council regime, this is not driving the results; that focusing on towns and industries which should have less ability to influence bargaining council policy increases the estimated effect sizes, and by introducing rings of weighted fixed effects to check whether estimates are robust to a broad variety of potential forms of spatial heterogeneity. In all cases they are robust. I conclude that employment effects are about 10%, larger than the effects of other labor regulations estimated in the literature, and potentially contributing 1 to 2 percentage points to South Africa's substantial unemployment.

A second paper, "**Can Minimum Wages Cause a Big Push? Evidence from Indonesia,"** which is solo-authored and forthcoming in *The Journal of Development Economics* also examines the effect of labor standards – in this case a rapid and varied increase in the minimum wage in 1990s Indonesia on local employment outcomes. However, broad employment trends associated with these minimum wages are not consistent with standard economic models: minimum wages are associated with increased formalization and formal sector employment, decreased informal sector employment, and increased consumption.

I demonstrate theoretically that minimum wages can provide a coordination mechanism to spur development. More specifically, when firms are forced to pay higher wages, it imposes an externality on other firms that supply the local market: higher wages can be associated with higher purchases, expanded local product markets, and increased labor demand. This mechanism can cause a "big push" where the strengthened local product market generates the potential revenues to encourage firms to pay high fixed costs and formalize. Of course, the potential of omitted variables or reverse causality to create spurious correlations where higher minimum wages are associated with more formalized labor markets without a causal effect is strong. I develop a difference in spatial differences estimator to address these concerns, which adjusts a spatial discontinuity estimator for time-constant differences. I find that the formalization response to minimum wages is robust to these controls. When a district sees minimum wages grow faster than its neighbors, it also sees formal sector employment grow faster than neighboring districts and informal sector employment shrink faster than its neighbors. Wages and consumption increase in the same way. Moreover, this response occurs entirely in the untradeable, industrializable industries suggested by the model, while tradable industry and non-industrializable industry also see employment trends conforming to model predictions. A set of robustness checks verifies that trends are not consistent with reverse causality, and that key potential omitted variables similarly do not drive estimates.

C): Other Constraints to Labor Market and Firm Success

While labor regulations are a much-studied and salient constraint on firm success, there are of course many other factors which influence profitability and firm growth. Of particular interest in developing countries is understanding which characteristics allow entrepreneurs to develop profitable businesses rather than merely a safety-net which allows a small amount of income in times of otherwise minimal earnings. Thus far, I have written two papers on determinants of firm success, one focusing on the success of entrepreneurs in Kenya, A second, which is less central to my overall research agenda but similarly focuses on firm success estimates the role of social media on restaurant success in the US. I have also written an additional paper on the role of migrants in labor market opportunities of natives.

A current working paper, **"Minding Small Change: Limited Attention Among Small Firms in Kenya"**, which is co-authored with Lori Beaman and Jonathan Robinson, tests whether limited attention, which has elsewhere been shown to be a behavioral bias in decision-making, bleeds over into business decisions to impact profits in a meaningful way. To test attention directly, we focus on a small, but important business decision: how much small change to keep on hand in order to make sales when customers have large bills. Our survey shows that Kenyan entrepreneurs are losing 5-8% of profits due to lost sales, which occur both because change cannot be made and because these entrepreneurs miss out on sales while spending 2 hours per week away from the store hunting for change. We use two interventions designed at redirecting entrepreneurial effort towards small change to test whether limited attention is leading to lost sales and missed profits in meaningful ways. More specifically, we randomize enrollment into a study which frequently surveys entrepreneurs about recent lost change experiences, and which may make change problems more salient without otherwise impacting the entrepreneurial decision problem. We also randomize a more intensive intervention, where we work with the entrepreneur to calculate the profits lost due to changeouts. We use a mix of self-reported and objective measures to confirm that both interventions are successful at reducing changeouts. This reduction seems to occur in part because entrepreneurs bring in more change in the morning but more strikingly because they reduce their willingness to share change with other nearby businesses. We also document that more treated entrepreneurs have more cash on hand when we come to visit; that they exhibit different patterns in willingness to pay for change, and that they have higher sales and profits. We conclude that small firms are constrained by limited attention, which may have implications for understanding the mixed results in the business training literature.

A second paper, "**Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database**," which is co-authored with Michael Anderson and was published in the *Economic Journal*, represents the sole paper in my research portfolio which does not take place in a developing country. This paper asks whether restaurants in San Francisco see increases in patronage when they receive better crowd-sourced reviews on Yelp, a popular on-line review website. Social media, like Yelp, represent a new mechanism for consumers to learn about the quality of various businesses, though the practical relevance of yelp reviews – as unsolicited reviews of peers, rather than experts – remains unclear. The challenge in estimating the relationship between review quality and restaurant business is that restaurants with better reviews will likely experience higher consumer demand for reasons unrelated to customer learning from the review itself. We solve this problem by taking advantage of the fact that Yelp rounds its reviews to the half-star level, so that restaurants with very similar average ratings (say, those with an average rating of 3.24 stars versus 3.26 stars) experience different displayed ratings (3 versus 3.5 stars). We match the database on yelp reviews to a database which registers whether we could make a reservation at peak hours to test whether there is a similarly discontinuous change in the availability to make a reservation. We find that Yelp reviews are very important to restaurants, with an extra half-star causing a restaurant to sell out 17 to 19 percentage points more frequently, with potentially huge increases in profits. This suggests that restaurants have a strong incentive to game the system by posting fake reviews, and a rich set of robustness checks conclude that gaming behavior does not vary discontinuously at star thresholds.

Finally, a paper which focuses on alternate constraints to labor market success is **"Labor Market Changes in Response to Immigration: Evidence from Internal Migration Shocks**" which is joint work with Marieke Kleemans. In this paper, we develop rainfall shocks as an instrument for the supply of migrants to a destination city in Indonesia, similar to Munshi's (2003) seminal work on Mexican migrants. Using the Indonesia Family Life Survey, we document migration patterns to determine rainfall shocks amongst the sending communities to a particular city. As more migrants move to the city in periods of bad rainfall, we can test directly the role of migrants in labor market outcomes while controlling for the rainfall shock experienced by the destinations themselves. We find that a one percentage point increase in migrants decreases average income by 1.9 percentage points and reduces the employment rate by 0.4 percent, with both effects concentrated in the informal sector and amongst low-skilled natives. Previous studies with similar findings have been concentrated in the US and attribute these effects to substitution between natives and migrants with similar skill levels; however, in the Indonesian context, the migrants are systematically relatively high ability. We demonstrate that the disemployment of the low skilled in response to an influx of highly skilled immigrants is inconsistent with a labor market model based on perfect competition and full employment similar to Card and Lemieux (2001) and Borjas (2003), but that it can be understood if the disadvantaged group faces chances of disemployment or employment to an informal sector. We furthermore find that young people are disproportionally affected by an influx of immigrants and that men face the negative consequences in terms of lower income, while women are less likely to be employed.

D): Job Search and Networks in Other Contexts

Networks and search models are relevant to many contexts outside of the labor market. While my primary research interests have focused on labor market outcomes in developing countries, I remain interested in how interactions between different types of social relationships effect other outcomes as well. A key candidate here is learning about new technologies. Building on my work on understanding how information flows through job networks – and the consequences of that decision, a first long-run field project which began in 2010 examines whether network theory can be used to improve the efficiency of agricultural extension. A second paper asks what the theory of job search, applied instead to searching for a spouse, can contribute to understanding some consistent patterns in the spread of HIV which accord poorly with many existing epidemiological models.

"**Making Networks Work for Policy"** is an on-going data collection effort funded by the Bill and Melinda Gates Foundation's Agricultural Technology Adoption Initiative (ATAI) and 3ie (the International Initiative for Impact Evaluation), and which is joint work with Lori Beaman, Ariel BenYishay, and A. Mushfiq Mobarak, examines the role of social connections in agricultural technology adoption. Two common criticisms of the literature on social connections in a variety of contexts are that empirical analysis remains fairly divorced from the advances in network theory and that, for all of the studies which examine social connections in a variety of contexts, policy prescriptions have been few and far between. In this project, we expect to contribute on both of these dimensions. Agriculture is a natural context for this study, as it determines the livelihood of a tremendous fraction of people in developing countries.

Learning through social connections has been shown to be important for agricultural technology adoption. In this project, we evaluate whether we can explicitly design agricultural extension systems around this fact in a way which achieves higher take-up of a pair of new agricultural technologies. In particular, in winter 2010-2011, we conducted a full network census within 200 villages in Malawi to map out existing patterns of communication. We then assisted the Malawi Ministry of Agriculture in selecting extension partners who are situated in social networks within villages in places which make them "optimal" partners; that is, those whose adoption would induce the greatest social learning multiplier and thus result in the greatest village level adoption. Which partners are optimal depends on assumptions on how people learn, and we randomized villages into treatments where different partners were selected depending on whether learning happens easily (so that partners should be spread out in the social network) or with greater difficulty (so that best partners would know many of the same people, treating them extensively)[1], as well as villages where more observable proxies for network location were utilized. In the end, we expect to both observe which partners were most effective in increasing village-level adoption and adoption amongst their friends and neighbors, and develop policy implications which inform the government's future extension policy.

Secondly a paper **"Marital Shopping and Epidemic AIDS"**, which is solo-authored and was published in *Demography*, applies a classic job search model to the problem of finding a suitable spouse, and asks about the consequences of spousal search for the HIV epidemic. This paper observes that there is a broad disconnect between two well-known facts related to Africa's HIV/AIDS epidemic and the behavioral epidemiological modeling which has dominated the policy debate. Specifically, it is by now well known that the transmission rate in unprotected intercourse is very low (about 1/1000 per sex act), and also that systematically the young most frequently

[1] In the network theory literature, these two cases are referred to as a "simple" versus a "complex" contagion.

contract the virus. While this latter fact may be unsurprising, it is inconsistent with many existing models for the epidemic. In this paper, I adapt a classic model of job search to create a model of serially-monogamous dating and marriage. The matching behavior in this model creates behavior where individuals have clusters of short relationships in between a few much longer-term boyfriends and girlfriends before eventually marrying. I show that this type of behavior can create epidemic HIV prevalence levels, because a short window of high infectiousness at the beginning of infection creates high infection rates in these clusters of short relationships. Moreover, I show that the predictions of this model on the age-pattern of infection line up well with a number of characteristics of South Africa's HIV epidemic, most strikingly predicting the actual deaths-by-age observed in South Africa. Finally, I estimate the effectiveness of two potential policies: reducing transmission rates, potentially through male circumcision or STI treatment; or compelling individuals to use condoms at the beginning of new relationships. Both are found to be effective at combating the epidemic, and used together eliminate the potential of this dating behavior to create an epidemic.