ESSAYS IN EMPIRICAL MICROECONOMICS

A Dissertation
submitted to the Faculty of the
Graduate School of Arts and Sciences
of Georgetown University
in partial fulfillment of the requirements for the
degree of
Doctor of Philosophy
in Economics

By

David C. Phillips, M.A.

Washington, DC
April 3, 2012

ESSAYS IN EMPIRICAL MICROECONOMICS

David C. Phillips, M.A.

Thesis Advisor: William G. Jack, Ph.D.

ABSTRACT

This dissertation contains three essays in empirical microeconomics aimed at measuring the causal effects of policies related to poverty both in the U.S. and internationally. First, I examine the whether raising wages at home can encourage skilled health workers to remain in their country of training. In particular, I investigate this question in the context of Ghana and find that wage increases resulting from a new salary structure induced health workers receiving the largest wage increases to remain in Ghana. Second, I study the role of spatial mismatch in generating persistently poor labor market outcomes for urban, minority workers. In particular, I describe the results of a novel randomized field experiment that provided public transit subsidies to active job seekers in Washington, DC. I find that receiving these subsidies does improve labor market outcomes, indicating that spatial mismatch plays an important role in this low-wage labor market. Using the same experimental data, I also estimate and validate a structural model of job search which can then be used to measure the welfare impact of receiving treatment. I find transit subsidies are cost-effective, generating significant surplus relative to the cost of the intervention.

For the people, too many to name,
who helped make this possible;

For the clients of Jubilee Jobs,
whose persistence and grace inspire me;

For Janelle, who always stays with me in love and support
and insisted that I do something worthwhile;

I give thanks to God for all of you and for
God's continual call and guidance.

Many thanks,
David C. Phillips

TABLE OF CONTENTS

INTRODUCTION

0.1. Overview

Much recent research in empirical microeconomics has taken on a "what works" mentality, attempting to measure the effects of well-defined policy interventions. Particular focus has been placed on rigorous experimental and quasi-experimental research design that allows for estimating the causal effects of a policy without contamination from selection bias and omitted variables. Some have gone as gone as far to say that this change has led to a "credibility revolution" in empirical microeconomic research as improved research designs have allowed results to have a plausible causal interpretation, perhaps for the first time (Angrist and Pischke, 2010). However, this approach has had its drawbacks. Some have criticized this trend toward experimentation as atheoretical and unable to capture the most important economic phenomena (Keane, 2010). In focusing on narrow populations and partial equilibrium settings, some argue that the sheen has worn off the new attraction of experimental designs.

In this dissertation, I take on a policy-focused mentality that attempts to benefit from both of these perspectives. I aim mainly to measure the causal effects of well-defined and important policy interventions. In particular, I focus on issues of poverty in the context of labor and development economics, addressing issues related to international migration of skilled health workers from developing countries and persistent unemployment in urban, U.S. labor markets. While this dissertation largely flows from the existing literature on economic experiments, it attempts to bring together relevant tools from both experimentation and theoretically-driven structural estimation to answer these questions. To this end, in the first chapter I consider the effects of a non-experimental policy change through the lens of a natural quasi-experiment; in the second chapter I describe results from a novel randomized field experiment; and in the final chapter, I use estimation of structural job search models to interpret the experimental data from chapter 2. While these chapters represent an eclectic mix of topics and tools, through them I intend to address the broad question of "what works" in addressing poverty both domestically and internationally while demonstrating that diverse tools are necessary to fully answer that question.

0.2. Summary of Dissertation Chapters

In the first chapter, I ask whether governments in developing countries retain skilled health workers by raising public sector wages? With my co-author, James Antwi, I take a natural experiment approach to this question using data from Ghana. We investigate this question using sudden, policy-induced wage variation, in which the Government of Ghana restructured the pay scale for health workers employed by the government. We find that a ten percent increase in wages decreases annual attrition from the public payroll by 1.5 percentage points (from a mean of 8 percentage points) among 20-35 year-old workers from professions that tend to migrate. As a result, the ten-year survival probability for these health workers increases from 0.43 to 0.52. The effects are concentrated among these young workers, and we do not detect effects for older workers or among categories of workers that do not tend to migrate. Given Ghana's context as a major source of skilled health professional migrants and high correlation of our attrition measure with aggregate migration, we interpret these results as evidence that wage increases in Ghana improved retention mainly through reducing international migration.

In the second chapter, I consider the issue of how access to public transit and spatial mismatch of workers from jobs affect low-wage, urban labor markets in the U.S. In urban areas job vacancies often exist but poor, minority residents tend to be concentrated in neighborhoods with limited geographic access to these jobs. Using a randomized field experiment with public transit subsidies, I test whether this spatial mismatch of workers from jobs causes poor labor market outcomes. Randomly selected clients of a non-profit employment agency received a public transit subsidy to assist in applying to jobs and attending interviews with potential employers. I find evidence that the transit subsidies have a large, short-run effect in reducing unemployment durations with treatment causing the probability of finding employment within 40 days to increase by 9 percentage points, from 0.26 to 0.35. After 90 days, this difference narrows to a large but statistically insignificant 5 percentage points. I find weaker evidence that this decrease in unemployment duration results from more intense search behavior, with the transit subsidy group applying to more jobs and jobs further from home. To my knowledge, these results provide the first experimental confirmation that spatial mismatch of workers from jobs can cause adverse labor market outcomes for poor, urban individuals.

In the third chapter, I use data from the same randomized field experiment to test the validity of a structural model of job search in the low wage labor market. I find that a simplified version of the Flinn (2006) incorporating search frictions, binding minimum wages, and Nash bargaining over wages fits the experimental results well. These results indicate that a structural model tailored to a given environment can successfully predict out of sample experimental results. Having validated the model, I use it to estimate the impact of treatment on an inherently unobservable outcome: welfare of the recipient. I find that receiving subsidized public transportation during job search leads to more than $500 of net benefits.

CHAPTER 1

WAGES AND HEALTH WORKER RETENTION: EVIDENCE FROM PUBLIC SECTOR WAGE REFORMS IN GHANA[1]

1.1. Introduction

High attrition of skilled employees can generate under-staffing in the public health care systems of developing countries. Wage differentials between domestic public employment and other options are one factor that could be driving doctors, nurses, and other skilled health workers to leave the public health sector, often for jobs in high-income countries. Though the impact of this trend on health outcomes and the best policy response are oft-debated (e.g. Clemens, 2007; Bhargava and Docquier, 2008), many developing countries have decided to take policy positions discouraging such migration. Increased salaries represent one possible but expensive option for improving retention. However, the cost-effectiveness of this policy depends heavily on how elastically attrition responds to higher salaries. This is especially important as policy makers consider the costs and benefits of raising salaries as opposed to mandatory public service, improved facilities, shifting from higher to lower skilled health workers, and other options.

Despite the importance of this issue to policy makers and the centrality of wages in the basic economic model of migration, little strong evidence exists on the causal impact of home country wages on attrition of health workers. This chapter aims to isolate the causal effect of wages on attrition by use of a natural experiment. In the ideal econometric situation, wages

[1] This essay is co-authored with James Antwi, Greenwich School of Management, University of Wales.

would be set randomly so that any correlation between wages and attrition would be causal. Lacking this situation, we exploit a sudden change in the compensation of different groups of public sector health workers in Ghana. In 2006, the Government of Ghana enacted a new wage schedule for publicly paid health workers. This policy creates sudden and wide variation in wages across time, grade[2], and step (seniority within a given grade). Paired with administrative data, we use this natural experiment to identify plausibly exogenous wage variation and measure the causal impact of wages of attrition.

We employ a fixed effects strategy, controlling for effects common to workers in a given grade-seniority group as well as common time effects, to test whether the groups of health workers who received the largest raises had their attrition rates fall the most. Using this fixed effects strategy, we find evidence that wage increases do cause lower attrition rates. A ten percent increase in wages decreases annual attrition by 1.7 percentage points (from a mean of 8 percentage points) among 20-35 year-old potential migrants. During this time period, Ghana was a major source of high-skilled health worker migrants, and measures of health worker migration correlate strongly with our measure of attrition at aggregate levels. Thus, we interpret the results as most plausibly capturing the effect of wages in reducing international migration. This empirical strategy controls for all time-invariant differences across occupations, grades or seniority groups as well as time shocks common to all health workers. However, different groups of workers may in fact face differing time-varying shocks to attrition over time. We address this

[2] Throughout, 'grade' refers to the wage grade of a worker. These generally indicate both occupation as well as large differences in seniority. For example, doctors and nurses are in different grades, but there are also 7 different grades of nurses: Staff Nurse, Senior Staff Nurse, Nursing Officer, Senior Nursing Officer, Principal Nursing Officer, Dep. Dir. of Nursing Services, and Chief Nursing Officer.

in two main ways. First, we allow for time fixed effects that differ across three groups: doctors, nurses, and other health workers. This allays concerns that our results are being driven by shocks to the two largest occupations in our sample, using instead variation within these occupational groups. Second, we control for a variety of observed individual demographics as well as concurrent policies affecting migration of health workers both out of Ghana and into the UK. Our preferred specification incorporates the set controls as well as the more general time fixed effects. With this setup, we continue to find strong negative effects of wages on attrition with a 10 percent wage increase leading to a 1.45 percentage point decrease in annual attrition. This implies an increase of the 10-year survival rate from 43 percent to 51 percent.

Across workers, we find evidence that the effect of wages on attrition is concentrated among early-career workers with no effects on older health workers. We also find evidence that the impact of wages on attrition is strongest for health workers living in urban areas but no evidence that impacts differ by gender. The effect is concentrated among workers in occupations that tend to migrate. We take this as further evidence that wages affect attrition mainly through reducing migration.

We also present results controlling for linear time trends that are specific to each grade-seniority group, addressing the concern that large wage increases may have been targeted at grade-seniority groups with attrition rates that were already falling. This specification provides qualitatively similar results, though there is not enough variation in our data to measure the effects at standard levels of statistical significance. Attempts to allow for general time effects at a finer level than broad occupational classifications face similar issues. While this remains a drawback, our approach provides new, credible causal estimates of the impact of wages on health worker attrition using micro data and plausibly exogenous variation in wages.

Our main results confirm the predictions of simple models of migration and employee retention. In the simplest migration models, an increase in home country wages unambiguously reduces the probability that an individual will migrate by reducing the ‘push’ effect of a large wage gap (Borjas, 1987). However, if a binding credit constraint prevents migration, higher wages can increase migration (Lopez and Schiff, 1998). The cross-country empirical literature generally has been rather mixed about which of these effects dominates (Clark, et. al., 2007; Pederson, et. al., 2008; Mayda 2010). A growing micro-empirical literature has paid serious attention to identifying the causal impact of higher income on out-migration. Some evidence for classic push-pull effects of wages has been found by Yang (2006) who uses exchange rate variation to measure the impact of real wages abroad on return migration of Filipino emigrants. Meanwhile, evidence for the role of income in relaxing credit constraints to migration has been documented. Yang and Choi (2007) find evidence that rainfall shocks in the Philippines induce migration by generating income that relaxes credit constraints, and Angelucci (2005) finds evidence that income receipt from Progresa cash grants results in higher migration. Finally, Gibson and McKenzie (2011) find very little role for wage differentials and credit constraints in the migration choices of very highly-skilled individuals from New Zealand, New Guinea, and Tonga. To the extent that our results for wages and attrition can be interpreted as the effect of wages on migration, we find strong support for a classic push effect of low home wages for highly-skilled health workers in Ghana. While we know of no other studies attempting to measure the causal effect of wages on health worker migration using micro data, our results are broadly similar to those of Okeke (2009), who finds that aggregate physician flows from Africa respond to rainfall shocks.

If we interpret our results more directly as the impact of wages on employee retention, this study relates to the literature of retention and job search. On-the-job search models of employee retention predict that an employer with higher wages will retain workers for longer by competing more effectively with other employers (Burdett and Mortensen, 1998). In a context with variable labor supply, though, higher wages could have wealth effects that encourage retirement or other decreases in labor supply. In this setting, our results indicate that for young workers the substitution effect of making remaining employed with the government of Ghana outweighs any wealth effect on labor supply.

1.2. Background

1.2.1. Migration of Health Workers from Ghana

Ghana has long been a major source of migrants in the health sector. Likely due their high quality training, low wages, and English proficiency, many Ghanaian health workers have left for jobs abroad. Bhargava and Docquier (2008) provide cross country data on physician migration into OECD countries. As shown in Figure 1.1, in an average year from 1991-2004 three to four percent of Ghana's physicians migrated annually, easily outpacing the African average. Prior to this time period, migration rates were even higher, with Dovlo and Nyonator (2003) reporting annual migration rates of 10 to 20 percent for graduates in the 1985-1994 classes of the University of Ghana Medical School. As shown in Table 1.1, these migrants mainly leave for English-speaking, high-income countries. Data from the Ghana Nurses and Midwives Council indicate that a full 71 percent of nurses leaving during 2002-2005 went to the

UK, with most of the remainder leaving for the US. Data from Dovlo and Nyonator (2003) indicate a similar pattern for physicians.

After multiple decades of extensive migration by Ghanaian health workers, flows of health workers out of Ghana have slowed in recent years. Figure 1.2 demonstrates this fact for nurses using Nurses and Midwives Council data and administrative payroll data. Ghana's Nurses and Midwives Council keeps statistics on the number of requests by domestically trained nurses to have their credentials verified for international employment. As the figure indicates, migration of nurses from Ghana plateaued in the early 2000's, dropped precipitously in 2006, and then leveled off at a reduced rate. Attrition of nurses under age 40 from the public payroll shows a similar pattern. In the same figure, attrition rates show a large drop in 2006 and subsequently stabilize, closely following the NMC migration data. Recent decreases in migration are also apparent for physicians, as depicted in Figure 1.3. Attrition from the public sector and data on new Ghanaian registrants to the UK's General Medical Council show a strong correlation with each other as well as a drop in 2006. The close correspondence between migration data and attrition from the public payroll will be important later for the interpretation of our results. Because individual-level migration data is unavailable for our sample, we will use attrition from the public payroll as our dependent variable. The time-series correlation of our dependent variable with migration measures provides an indication that attrition in our data is best interpreted as migration.

The recent, sudden decline in migration of health workers from Ghana begs the question as to its causes. As Figure 1.4 shows, health workers migrating to the UK can roughly double their earnings, even after adjusting for purchasing power differences. For example, doctors in

Ghana earned about 1,000 Ghana Cedis per month but could earn about 2,500 Ghana Cedis per month (PPP) in the UK. Many point to such wage gaps as the main cause of migration of skilled health workers to high-income countries. In 2006, at precisely the same time as the fall in migration, the government of Ghana introduced a new wage structure for health workers that increased earnings significantly for many health workers. While many other factors and policies in Ghana and abroad likely influenced the decline in migration, we will focus on isolating the role that wages played.

1.2.2. Public Health Sector Wage Changes in Ghana

In 1998, the Ministry of Health introduced the Additional Duty Hours Allowance (ADHA) for health workers. As its name implies, the ADHA's explicit purpose was to compensate doctors, nurses, and other core clinical workers for unusually long hours. However, shortly after its creation in 1998, the ADHA became a simple salary supplement and was extended to other cadres of health workers. The Ministry of Health (MOH) assigned a fixed number of notional hours to each cadre (doctor, professional nurse, etc.) of employee, and all employees in the same cadre received the same number of hours. Since these hours were paid at the worker's usual hourly rate, the ADHA amounted to a percentage bonus of a health worker's base salary. Within a cadre all employees received the same percentage bonus from ADHA, while different cadres received different bonuses due to differences in notional hours assigned. This system, supplemented by common percentage pay raises among all employees, ensured that the relative pay of all healthworkers was stable from 2000-2006.

In 2006 due to budgetary pressure, the Government of Ghana desired to fold the ADHA into regular pay. Health workers also pushed for ADHA to be converted into basic salary because the ADHA payments were often delayed and were not taken into account when determining pensions. Since ADHA had grown to be a very large proportion of many health workers' pay, the government decided to adopt an entirely new salary structure, the Health Salary Structure (HSS). Pay rates under the new schedule were defined based on a job evaluation that arranged different grades in the new salary structure according to the skills and tasks of the job performed by that grade. As a result, the new salary structure gave nominal wage increases of varying degrees to all health workers, but due to inflation and the loss of ADHA, some workers saw their real total earnings rise slowly or even drop. Importantly, due to differences in the raises assigned to various groups of workers and due to the fact that some cadres benefitted more from ADHA than did others, the new salary structure completely rearranged the relative pay of many workers. Finally, as a part of the arrangement made in adopting the HSS, workers' nominal wages were frozen from 2006-2009.

Figure 1.5 displays these wage changes. In the wage schedule, a health workers' pay is determined by 'grade' and 'step.' Grades differentiate large promotions (principal nursing officer, medical officer, senior medical officer, specialist, etc.) while steps embody smaller promotions within a grade. In the figure, each line represents real log wages (inclusive of ADHA) for each possible gradeXstep combination, normalized to zero in 2003. Thus, following an individual line over time traces the wages of a worker that is never promoted from 2003-2009. From 2003-2005 the lines generally move together, demonstrating that all groups of workers received common percentage wage increases. But from 2005 to 2006, the wages of different groups of workers diverge. Some workers received real wage increases of up to 10 percent while others even saw their real wages decrease. We will exploit the variation in wages across professions, seniority, and time generated by this policy change to measure the impact of wages on attrition.

1.3. A Simple Theoretical Framework

1.3.1. A Simple Model of Migration

Consider an individual choosing between continuing to work in the public health system and leaving for another job. We will interpret this other option as migrating for a job outside the country, but in principle the outside job could be in the private health sector or outside the health profession. Assuming a linear indirect utility function, an individual i will attrite at time t iff:

$$\alpha_0 + \beta_1 w_{it}^* + \delta_0 x_{it} \geq \alpha_1 + \beta_2 w_{it} + \delta_1 x_{it} + c_{it}$$

$$\Leftrightarrow c_{it} \leq \beta_0 + \beta_1 w_{it}^* - \beta_2 w_{it} + \delta x_{it} \quad (1.1)$$

where c_{it} is the cost of migration, w_{it}^* is the log wage abroad, w_{it} is the log wage at home, and x_{it} is a vector of individual characteristics that are valued differently at home and abroad (δ is the marginal value of an attribute abroad relative to home). If F is the distribution of c_{it} then the probability of attrition A_{it} is:

$$A_{it} = \Pr[Attrition] = F(\beta_0 + \beta_1 w_{it}^* - \beta_2 w_{it} + \delta x_{it})$$

In this simple model, the impact of home wages on the probability of attrition is unambiguously non-positive. Assuming F is differentiable with density f:

$$\frac{\partial A_{it}}{\partial w_{it}} = -f(\beta_0 + \beta_1 w_{it}^* - \beta_2 w_{it} + \delta x_{it}) * \beta_2 \leq 0$$

However, even in this model the magnitude of the impact of wages depends greatly on the functional form and support of F. In particular, if the wage gap is very large between the home and foreign countries, then the impact of wages will be likely be small. Intuitively, large wage gaps move us into the 'tail' of the distribution of migration costs.[3] In the case of health workers migrating from Ghana to their main destination in the UK, this could very well be the case. In fact, some policy-focused research discourages salary increases as a method for decreasing health worker migration from sub-Saharan Africa due to the perception that salary increases will likely be ineffective due to the large wage gaps (Vujicic, et. al., 2004).

1.3.2. Credit Constraints

The unambiguous negative impact of home wages on attrition disappears if a simple credit constraint is added to the model. In an extreme case, suppose that a worker receives the public sector wage at time t. Then, the individual can choose whether or not to leave, expecting that future wages will be the same as today. Finally, suppose that the cost of migration must be financed out of current wages. Then, for an individual to migrate, they must be able to finance migration:

$$c_{it} \leq w_{it} \qquad (1.2)$$

Thus, an individual attrites iff (1.1) and (1.2) both hold, i.e. the probability of attrition is

$$A_{it} = \Pr[Attrition] = F(\min\{w_{it}, \beta_0 + \beta_1 w_{it}^* - \beta_2 w_{it} + \delta x_{it}\})$$

For individuals with low wealth, wage increases may actually lead to higher migration rates:

$$\frac{\partial A_{it}}{\partial w_{it}} = f(w_{it}) \geq 0 \; if \; w_{it} < \beta_0 + \beta_1 w_{it}^* - \beta_2 w_{it} + \delta x_{it}$$

Thus if a change in home wages reflects both an increase in current wages and a similar increase in expected future wages, the sign of the marginal effect of home wages on migration is an empirical question as well.

---

[3] Formally, as long as $\lim_{w^* \to \infty} f(\cdot)$ exists, then it must be zero. As a result, the limit of $\frac{\partial A_{it}}{\partial w_{it}}$ must be zero as well.

1.4. Identification Strategy

1.4.1. Main Identification

Finding exogenous variation in wages is important for a study of migration and home wages because the correlation between wages and migration can rarely be interpreted as the causal impact of wages. Individuals with high ability generally receive higher wages and migrate more frequently (Hanson, 2008). As a result, the correlation across individuals between home wages and migration will not reflect the causal impact of wages on migration. Meanwhile, the correlation between migration and wages across different locations will also not generally reflect the causal impact of wages because causality also runs the other direction: migration is a supply shock potentially affecting wages.

The wage reforms described above help alleviate these difficulties. The scene depicted in Figure 1.5 closely mimics an experiment with variation in the intensity of treatment. This new wage schedule was certainly set intentionally with the aim of bringing relative pay in line with the tasks required for a particular job, but fixed effects for each salary group should control for these time-invariant unobserved characteristics of different jobs. Since salaries are uniform for workers in the same grade of the same step (i.e. seniority), we condition on fixed effects for each gradeXstep group and time fixed effects to exploit variation in wages resulting from policy-induced wage changes from 2005 to 2006. In this basic setup, we take the variation in these wage changes across different gradeXstep combinations as plausibly exogenous. As such, we estimate the impact of wages on attrition from the public payroll using a difference-in-differences[4] estimator:

$$A_{it} = \alpha + \beta w_{it} + gradeXstep_i + \eta_t + \epsilon_{it} \quad (1.3)$$

where A_{it} is an indicator of attrition from the payroll; w_{it} represents wages paid in the public sector in Ghana; η_t is a common time fixed-effect; and $gradeXstep_i$ is a fixed effect for which grade-step group an individual is in when they first enter the data. We will estimate equation (1.3) by instrumental variables, using wages that an individual would have received if never promoted as an instrument for actual wages. Formally, if log wages are defined as:

$$w_{it} = f_t(grade_{it}, step_{it})$$

where changes in $f_t(\cdot)$ embody policy changes to the public health sector wage schedule. Then we define an instrument $\hat{w}_{it}$ as:

$$\hat{w}_{it} = f_t(grade_{i0}, step_{i0})$$

where $grade_{i0}$ and $step_{i0}$ denote the grade and step of individual *i* when we first observe them on the public payroll. Importantly, this instrument eliminates variation in wages that comes from promotions because this variation is likely to reflect ability and thus be endogenous. When combined with time fixed effects and group effects, the variation remaining is that caused by enactment of the new wage schedule.

1.4.2. Potential Confounding Factors

The sudden reforms of 2006 and otherwise stable wage environment create a reasonable natural experiment in which to measure the impact of wages. However, this identification

[4] We prefer a fixed effects linear probability model. The results are similar if we explicitly model the fact that our dependent variable is a hazard probability by using a Cox-PH model.

strategy relies on the assumption that wage changes from 2005 to 2006 for different groups of workers can be taken as exogenous. We control for time and group effects, which remove the influence of time-variant factors that affect all groups of workers similarly and time-invariant differences across different groups. However, our identification could fail if the government of Ghana directed the largest wage increases to groups of workers based on observed characteristics correlated with attrition. Additionally, unobserved time-varying shocks could affect attrition differently across the various occupations in our data. To address these issues, we estimate the following equation using a similar IV strategy:

$$A_{it} = \alpha + \beta w_{it} + \delta x_{it} + gradeXstep_i + \eta_{ot} + \epsilon_{it} \quad (1.4)$$

where x_{it} is a vector of controls for individual demographics, labor market conditions abroad, and other domestic policies and the time fixed effects are now allowed to vary by occupational groupings.

To estimate the more general structure of time fixed effects, we group grades of workers into three broad occupational classifications: nurses, doctors, and other health workers. We then allow for time effects that vary across these three groups. Workers in these groups may differ widely in education, ability to migrate, and political bargaining power, which could conceivably cause them to receive widely different time-varying shocks while also influencing the size of their pay increases in 2006. This empirical strategy avoids these problems by eliminating variation across these groups and focusing attention on variation in policy-induced wage changes within these occupational groups. Education, labor union, and migration options are generally homogeneous within these groups, making variation in the 2006 wage increases more likely to be exogenous.

While more general than the common time-trend assumption, these occupation-specific time effects do not control for all possible violations of exogeneity. In particular, workers within the same broad occupational category differ on several observable dimensions. For example, more senior workers may receive differing wage increases from younger workers while simultaneously being unlikely to migrate regardless of wage levels. As a result, controlling for worker age is important. We control for the demographic and job characteristics available in our data including a polynomial in age, gender, region of job placement, and department of job placement.

We also control for several concurrent policies. In Ghana, several measures were taken with the goal of reducing migration of health workers. The Ghana College of Physicians and Surgeons opened in 2004, becoming the first medical specialist training school in Ghana. It is thought to have decreased migration of doctors who would have otherwise migrated for training purposes. Also during this time period, the Ministry of Health and the Nurses and Midwives Council collaborated with the service delivery agencies to enforce a bonding scheme for nurses. Under this program, publicly-trained nurses were required to complete a term of public service or pay a bond before they could be given verification by NMC to practice abroad. Other policies to reduce migration included a subsidized car loan scheme as well as increased availability of fellowships for continuing professional education.

Outside of Ghana, major policy changes also occurred, particularly in the UK. In 1999, the UK National Health Service adopted a Code of Practice limiting recruitment from developing countries. This policy strengthened considerably in subsequent years as the UK moved to limit not just the NHS, but also recruitment agencies working on behalf of the NHS. Meanwhile, wages and domestic supply of health workers in the UK were changing and could also coincide with the wage reforms that we are studying in Ghana.

For both foreign and domestic policies, if they affect all health workers within an occupational group equally, then our identification will still be valid. However, if any policy disproportionally affects particular groups of health workers, then this could bias our results. Figure 1.6 demonstrates that this does not appear to be a problem with regard to changes in UK demand for foreign health workers. Overall inflows of health workers from the world into the UK have decreased dramatically, but the change has affected doctors and nurses similarly. In any case, we will demonstrate that our identification strategy is not affected by controlling for the impacts of both foreign and domestic policy changes.

While the combination of occupation-specific time effects and extensive controls eliminates many threats to the exogeneity of these policy induced wage changes, we cannot control for all possibilities. While in theory our data contain enough variation in wages to include group-specific time effects at a finer level (e.g. splitting doctors into residents, medical officers, and specialists), in practice insufficient wage variation in our data prevents us from allowing for time fixed effects with more groups than the broad ones we include in our analysis. So, while we are able to identify plausibly exogenous variation for health workers not currently available in the literature, this remains a drawback of our approach. Given that we allow for differing time shocks to the two largest occupations in our sample, doctors and nurses, we are able to address the most likely cause of bias from this source.

Finally, we are able to explore one further generalization of the common time trends assumption. Our final specification also controls for linear time trends that are unique to each gradeXstep group. This specification ensures that our results are not driven by wage increases being targeting at groups of workers whose attrition rates are already decreasing. While the linear functional form poses some restrictions, it allows for a very general structure of time trends across the different groups of workers and contributes to the evidence from our main specification that the negative impact of wages on attrition is robust to a wide variety of specifications.

A final concern relates our instrumental variable approach to unobserved ability of the workers. While controlling for group effects at a fine level and removing promotions from our instrument should reduce bias generated by ability and positive selection, unobserved ability could still generate endogeneity. For example, our specification focuses on the contemporaneous effect of wages on attrition, but if individuals make migration decisions based on lifetime expected wages, then our instrument will systematically underestimate home wages for high ability individuals who expect to be promoted. Since migration likely selects on ability, this would induce correlation between our instrument and the error term. We expect that the magnitude of bias created by this type of endogeneity will be small because of how finely our groups are defined; however, we cannot rule out such a possibility.

1.5. Data

1.5.1. Administrative Wage Data

The main data source used in this study is individual-level payroll data obtained from the Controller and Accountant General’s Directorate of the Government of Ghana. This data contains payroll records from 2003 to 2009 for each health worker classified under MOH paid by the central government including employees of the Ministry of Health, Ghana Health Service, Christian Health Association of Ghana (CHAG), the teaching hospitals, and MOH training institutions. This panel data provides individual-level observations over several years, and yields a rich picture of the health labor market in Ghana. Additionally, for each individual this data provides detailed information on employee grade (i.e. Senior Medical Officer, Chief Lab Technologist, etc.) and salary step for each individual. Information on age, gender (equals 1 if male), department (CHAG, GHS, Headquarters, etc.), and region of posting are also available.

We use health sector public wage schedules to map grade and step into a basic salary for each worker. As described above, a major part of the 2005-2006 salary changes hinged on the Additional Duty Hours Allowance. Thus, it is important to consider not just the basic salary but also the Additional Duty Hours Allowance for each worker. Since ADHA was allocated according to a fixed formula that depends on a worker's category and their base pay, the payroll data along with data on ADHA hours allotments from GHS allow us to estimate their ADHA earnings. From 2003-2005 the formula for total wages is:

$$w_{it} = \log(E_{it} + \frac{E_{it}}{160} * 1.75 * h_{g_{it}})$$

where E_{it} is the basic salary of individual *i* at time *t* and $h_{g_{it}}$ is the number of ADHA hours allotted to a worker in grade g_{it}, which is the grade of worker *i* at time *t*. From 2006-2009, we simply use the log of basic salary. Table 1.2 provides summary statistics. The first column describes potential migrants (defined below), which is our main sample. Real (measured at 2004) monthly wages for this group average 410 Ghana Cedis.[5]

[5] About 450 USD.

As described above, our instrument for log wages will be log wages that would be earned by worker *i* if he or she was never promoted. In particular, the instrument for a worker *i* at time *t* takes the value from the time *t* wage schedule for a worker with the grade $grade_{i0}$ and step $step_{i0}$ where $grade_{i0}$ and $step_{i0}$ are the grade and step of worker *i* when he or she first enters the data. One complication with this definition is that from 2005 to 2006, the entire wage schedule changed from the Ghana Universal Salary Structure (GUSS) to the Health Salary Structure (HSS). Recall that 'step' refers to small differences in seniority within a grade. As a result of the change in salary structures, the number of steps within a given grade changed in some instances. In these cases, a step from before 2006 cannot be trivially mapped to a step from after 2006. For concreteness, consider an observation in 2007 for an individual who first enters the data in 2003. From the data, we can measure the grade and step of the individual in 2003, and these data are from the GUSS system. Call the grade and step $grade_{i0}$ and $step_{i0}^{GUSS}$. Our instrument should indicate the wages a person in $grade_{i0}$ and $step_{i0}^{GUSS}$ would receive in 2007. However, because of the change from GUSS to HSS, the actual wage schedule in 2007 for $grade_{i0}$ may have fewer steps within it in 2007 than it did in 2003. So, we approximate the initial GUSS step by an HSS step reflecting the same percentage progress up the grade. Formally, consider an individual at a time after 2006 who first entered the data before 2006. For observations after 2006, we define the intitial step in the HSS system, $step_{i0}^{HSS}$, as:

$$step_{i0}^{HSS} = step_{lowest}^{HSS} + \left(step_{i0}^{GUSS} - step_{lowest}^{GUSS}\right) * \frac{totalsteps_i^{HSS}}{totalsteps_i^{GUSS}}$$

where $step_{i0}^{GUSS}$ is an individual's first-observed step under the GUSS system; $step_{lowest}^{GUSS}$ is the entry-level step for individual i's grade and $totalsteps$ refers to the total number of steps in i's grade under a particular salary schedule.

The payroll data is monthly but due to technical challenges does not cover all months prior to 2006. As a result, it is not possible to study attrition over monthly intervals. Thus, we analyze the data at intervals approximating one year. To this end, we only use data from selected months: November 2003, July 2004, December 2005, December 2006, December 2007, December 2008, and July 2009. Since we use time fixed effects, the varying lengths of time between observations should not be an important issue.

1.5.2. Measuring Attrition

Individual records can be matched from year to year based on identifiers in the data. In particular we match records on first name, last name, gender, and date of birth.[6] We use these matched records over time to measure attrition of health workers from the payroll. In particular, if an individual is in the sample at time t but never after time t, we say the individual attrited at time t. In this study, we usually interpret the impact of wages on attrition from the payroll as the impact of wages on migration. The two are not, of course, generally equivalent. Workers could potentially leave the public payroll for employers outside the dataset (i.e. the private sector or military hospitals), or they could retire from working in the health field. We emphasize migration, but we cannot explicitly separate migration from other forms of attrition in the data.

[6] Employee numbers are available in the data and in most circumstances would be ideal. However, the treatment of some cadres' employee numbers changed from 2005 to 2006. Given this re-definition of employee numbers comes at the same time as the natural experiment, we choose to use a consistent method of matching over all time periods. Since these demographic identifiers are nearly always unique, this method is preferred.

1.5.3. Choosing the Sample

Interpretation of the results depends heavily on what part of the sample we use because the data covers workers of all ages and occupations with widely varying skill sets, education requirements, and responsibilities. For older workers, retirement is a major consideration. In our data, the probability of remaining after 20-40 years of experience is still positive and decreasing. This indicates that a significant minority of health workers still remains and attrite from the public sector after long careers. Retirement is an obvious explanation. To focus on migration rather than retirement, we truncate our main sample to those 35 years of age and younger, though we will expand this scope when we investigate the impact of wages on workers of different ages.

We also split our sample based on whether workers' occupations allow them to be a 'potential migrant.' Classification as such was determined after conversations with administrators in the Ministry of Health and Ghana Health Service. The first column of Table 1.3 provides a list of all worker categories classified as potential migrants along with their prevalence in the 2003-2006 sample. As expected, nurses and physicians compose much of the 'potential migrant' sample, though there are also many medical assistants (nurse practitioners) and skilled allied health workers, such as laboratory and x-ray officers. The second column indicates those excluded from this group, mostly low-skilled workers (orderlies, watchmen, drivers, etc.), as well as clinical workers whose skills are in low demand in developed countries (midwives, community health nurses, etc.). Nurses in training are available in the database but excluded from both groups because they are in school and receive only small stipends. The second column of Table 1.2 shows that, as expected, the non-potential migrants have an attrition rate about half that of the potential migrants, and they are paid less than half as much. However, mean age (conditional on being younger than 35) is similar to our main sample at around 30. While just over half of individuals are female in the main sample, the non-potential migrants are more likely to be female.

Splitting the sample in this way provides two services. First, it again focuses analysis on categories of workers that were known to have high rates of migration during the sample period. This helps focus the analysis on a group where there are external reasons for believing that migration is a main cause of attrition. Second, it also provides the opportunity to test whether wages affect these two groups in similar or different ways. Since retirement and the private sector are open to both the potential migrants in our main sample and to those excluded from our main sample, if we observe similar effects of the wage changes on both groups, then this would indicate that we are observing the effect of wages on some other form of attrition rather than migration. However, if the wage reforms show no effect on attrition of non-migrants, then we are likely measuring the impact of wages on migration.

Aside from limiting the sample in reasonable ways, circumstantial evidence also indicates that the large scale attrition seen in the data can most plausibly be explained by migration. As Table 1.2 shows, in our data roughly 8 percent of all potential migrant health workers leave the public payroll each period. As we have already seen, this coincides with well-known large-scale migration of health workers from Ghana. So, migration is at least consistent with the attrition rates in our data. Additionally, other forms of exit seem implausible. Large-scale retirement for young workers seems unlikely. Due to central payment of not only GHS but also CHAG and others, our data cover the vast majority of all health workers in Ghana. Ghana MOH estimates that 81.9 percent of all healthworkers work in the institutions covered by our sample, with most of the remainder in the private for-profit sector or prison and military hospitals. Also, for many categories of potential migrant health workers, coverage of our data frequently surpasses 90 percent (Ministry of Health (2007)). Given the small proportion of health workers in for-profit and military medicine, these sectors' human resource usage would have to increase by about 40 percent per year to absorb all of the health workers leaving the public payroll. What data exists on the private sector is not consistent with this story. For example, the World Health Organization's National Health Accounts indicate that the private sector's share in health sector spending actually fell from 58.6 percent to 48.4 percent from 2000 to 2007 (World Health Organization (2010)). Thus, external evidence seems to indicate that attrition from the public payroll in our sample is more plausibly attributed to migration than to other causes.

1.5.4. Additional Data Sources

The payroll data is supplemented by data from other sources. For inflation we use the GDP deflator from the World Bank's World Development Indicators. All values are reported in real 2004 Ghana Cedis. In 2004, the average exchange rate was USD 1.12 per Ghana Cedi. Foreign wages are drawn from the UK Annual Survey of Hours and Earnings. Individuals are matched to UK wages based on their cadre (doctor, professional nurse, etc.) at the SIC 4-digit level, when possible. Others are matched to SIC 3-digit and 2-digit codes when necessary. Necessarily, these data are not nearly as precisely measured as home wages because they are not individual specific. UK wages are changed into Ghana Cedis using PPP exchange rates from the Penn World Tables. As Table 1.1 shows, wages in the UK average 1158 Ghana Cedis per month at PPP rates. This is almost 3 times the purchasing power of the average domestic salary.

In addition to wages, other labor market factors in the UK likely affected migration during this time period, particularly adoption and strengthening of the NHS Code of Practice. We control for this using a measure of the openness of the UK to migrants from different professions: the log of the total number of new migrants to the UK for a particular profession in a given year from all source countries. We have this variable available from registration data in the UK with physicians from the UK General Medical Council, nurses from the UK Nurses and Midwives Council, and others (Art Therapists, Biomedical Scientists, Clinical Scientists, Dieticians, Occupational Therapists, Orthopists, Physiotherapists, Radiographers, and Audiologists) from the UK Health Professions Council. In our data, we also match Ghanaian dentists as physicians in the UK; medical assistants and anesthetist assistants as nurses; lab assistants and lab technical officers as clinical scientists; nutrition officers as dieticians; and x-ray officers as radiographers. For other categories, we code the variable as zero.

Two concurrent domestic policies affecting health worker migration are included in some specifications. The Ghana College of Physicians and Surgeons opened in 2004, becoming the first specialist training school in Ghana. It is thought to have decreased migration of doctors who would have otherwise migrated for training purposes. We measure this using enrollment in the Ghana College of Physicians and Surgeons, with data drawn from the Ghana College's newsletter as well as media reports. We include enrollment in the Ghana College interacted with a dummy for non-specialist doctors as a control. Also, during the sample period the Nurses and Midwives Council, in cooperation with GHS, began enforcing a public-sector service requirement for nurses. They began withholding certification from nurses who wished to migrate until they served their bond period. We model this as a dummy that is one for nurses starting in 2006 and zero otherwise.

1.6. Results

1.6.1. Main Effects

We estimate the instrumental variables regression of equation (1.3) with $\hat{w}_{it}$ as an instrument for w_{it}. Table 1.4 shows the results for this regression. Coefficients on wages are normalized so that they can be interpreted as the percentage point change in migration resulting from a 10 percent wage increase. Each column represents a different specification. First stage F-statistics are reported at the bottom of each column and indicate that a weak first stage should not be a problem. Column (1) of Table 1.4 shows the results of the simple difference-in-difference approach for the 'potential migrants' sample. There are no covariates other than the time effects and gradeXstep effects. The coefficient is negative and statistically significant at the 1 percent level, indicating that the 'push' effect of wages outweighs any credit constraint effect. The coefficient of -1.72 indicates that a 10 percent wage increase would lead to a 1.72 percentage point decrease in attrition. This is fairly substantial relative to the average attrition rate of about 8 percent. This first specification relies on a common time shock assumption for different groups of health workers. Column (2) relaxes this assumptions by allowing for doctors, nurses, and other health workers to have different time effects. The coefficient declines slightly to -1.45 and remains significant at the 1 percent level. Column (3) introduces a full set of controls for UK wages, total migration to the UK by profession, the two concurrent domestic policies, gender, a quartic polynomial in age, a set of dummies for department of posting, and a set of dummies for region of posting. These controls have almost no effect on the estimated coefficient and standard error. This combination of occupation-specific time effects and a full set of control variables represents our preferred specification. Thus our preferred estimate is that a ten percent wage increase reduces annual attrition by 1.45 percentage points, improving the ten-year survival probability from 0.43 to 0.52.

Our preferred specification controls for a large class of violations of the common time trends assumptions. Ideally, though, we would like to relax this assumption further. While variation in our wage data limit this somewhat, we do estimate a further specification in column (4) which allows for linear time trends that vary across the finest level of occupational classification in our data: gradeXstep groups. This allows for differing time trends in attrition for all but the most similar workers, workers who are not only in the same grade (e.g. Senior Medical Officers) but who also have the same seniority within their grade. This estimation leads to a much larger point estimate of -3.07 but also much larger standard errors. While clearly including a large amount of uncertainty, these results suggest that the negative impact of wages on attrition is robust to many violations of the common trends assumption.

1.6.2. Heterogeneous Effects

Recall that our main sample includes only workers under age 35 from occupations we classify as potential migrants. Further investigation indicates that the effects we detect are limited to this group of health workers. Table 1.5 replicates Table 1.4 but with the sample of workers from non-potential migrant occupations.[7] The effects for this sample are if anything positive, much smaller in magnitude, precisely estimated, yet for the most part statistically

[7] The one difference is that the first two columns of Table 1.5 are identical because this sampleincludes neither doctors nor nurses.

insignificant. For example, our preferred specification yields a coefficient on log wages of 0.42. Thus, we have evidence that higher wages decrease the probability of attrition for professions that do migrate but no evidence that such an effect exists for professions that do not tend to migrate. Since other interpretations of attrition such as retirement and moving to the private sector are available to both groups, it seems unlikely that these drive the results. Migration appears to be more consistent with the evidence.

We also explore heterogeneous effects along the dimensions of age, gender, and location of posting. Of these, age is perhaps the most important because we have focused our sample on workers under the age of 35. The first columns of Table 1.6 explore the importance of this restriction. The first column duplicates our preferred specification for the main sample of workers under age 35. The second column then applies this same specification to the entire sample of workers under the age of 65. We do not detect an impact of wages on attrition in the broader sample, with a small coefficient of 0.14 that is not distinguishable from zero. From the results in these two samples, we infer that wages appear to have no average effect across the entire population of health workers but a large negative effect among younger workers. Column (3) formalizes this result, where a positive coefficient on the interaction of wages and age, significant at the 10 percent level, suggests that among older workers the impact of wages on attrition is less negative.

The effect among early-career workers is consistent with the simplest model of migration that focuses on how increasing home wages reduces a push effect. It is perhaps not surprising that these effects diminish for older cohorts. In a life cycle model of migration, a fixed cost of migration finances the opportunity to receive higher annual wages. With fewer years remaining, older workers face a similar cost but lower benefits to migrating. Additionally, if individuals differ idiosyncratically in their preference for migration, the health workers that have chosen to remain in Ghana for several years may be a selected group that has such a low probability of migrating that marginal wage changes have no impact. Finally, ageing itself may be correlated with other factors (marriage, having children, building a home, etc.) that may make individuals very unlikely migrants, again causing wages to have no measurable impact on this very small probability. These explanations are all observationally equivalent in our data, but they provide ample explanation for the fact that the effect diminishes in older cohorts. In any case, it appears that these factors outweigh the impact of credit constraints, for in a model of credit constraints we would expect the youngest individuals to have low wealth and thus be credit constrained. Then, it would be these individuals for whom the 'push' effect would be most offset by a more relaxed credit constraint. Clearly, we do not observe this effect as dominant in this data.

We also explore other forms of heterogeneous effects in the final column of Table 1.6. Again, younger workers appear to have a more negative impact of wages on attrition. Perhaps surprisingly, gender does not appear to be a source of heterogeneity. Finally, we also explore how a rural posting may moderate the impact of wage increases on attrition. To this end we match the district of posting to data from the Core Welfare Indicators Questionnaire (CWIQ) on the percent of population who must travel over an hour to arrive at a health facility. When we interact this measure with wages, we obtain a positive coefficient, statistically significant at the ten percent level. This provides some suggestive evidence that the size of the effect of wages is larger (more negative) in urban areas (where travel time is lower). While only suggestive, it provides some interesting evidence that reductions in attrition are most effective in urban areas where staffing is already better than in rural areas.

1.7. Conclusion

This chapter measures the impact of home wages on attrition of skilled health professionals from the public sector in Ghana by exploiting variation in wages caused by a policy-induced natural experiment. We find that a 10 percent wage increase reduces the annual attrition rate by about 1.45 percentage points. This corresponds to an eight percentage point increase in the 10 year survival probability. This effect is concentrated solely among young workers age 20-35 who come from professions that tend to migrate. We take this as evidence that the effect of wages on attrition from the public sector mainly runs through migration. Unlike previous attempts at measuring the impact of home wages on migration, we find this negative effect of home wages on migration to be economically significant and robust to specification. While we do not have truly random variation in wages available, our use of sudden, policy-generated variation in wages allows us to plausibly estimate a causal effect of wages on health worker attrition using microdata. To our knowledge, this has not been possible previously. These results support the most basic economic models of migration in which individuals choose to migrate based on wage differentials between home and foreign countries. These results run counter to expectations that the impact of marginal wage changes may be small when wage gaps are very large or that higher home wages might relax credit constraints causing more migration.

As with all experiments, natural or designed, high internal validity does not guarantee external validity of the results. Context is important. The individuals in this study are highly-skilled health professionals who tend to migrate permanently with their families. The results will not necessarily apply to low-skilled or low income individuals who may face tighter credit constraints. However, the sample studied here is both under-studied and of policy significance. While our use of attrition from public payroll rather than an explicit measure of migration creates uncertainty about whether we are measuring other forms of attrition, it does allow us to approach the phenomenon of permanent migration on the individual level with a broad sample. In doing so, we find evidence for a vital element of migration theory in a sample characterized by long-term, permanent migration.

The migration behavior of skilled health workers is of particular policy importance. Health policy and strengthening health systems in particular have gained great notoriety recently. In this context, migration of skilled health professionals away from developing countries has been widely debated. To many, it is a notorious ‘brain drain’ of needed health workers from places where skilled health professionals are already scarce (e.g. Chen and Boufford, 2005). As a result, many paths have been taken toward reducing such migration. For example, the UK National Health Service’s has voluntarily imposed restrictions on foreign recruitment via its Code of Conduct, and this idea has been taken up recently by the World Health Assembly. To others, restrictions on migration are seen as violating human rights of migrants, ignoring more important problems of health worker performance and urban-rural distribution, or driven by recouping misguided education subsidies (e.g. Clemens, 2009). While important, deciding this debate is certainly out of reach for this article.

However, we do demonstrate that health workers can be retained, not just by restrictions on leaving but also by rewards for staying. If policymakers in developing countries desire to retain more health workers, as most do, then our results indicate that increasing salaries is one effective option. If policymakers in developed countries desire to redistribute health workers, subsidizing health worker salaries in sending countries is one means to this end that does not involve restrictions on the movement of individuals. Our results also indicate that migration of health workers should be an important consideration as policymakers in developing countries contemplate public sector wage reforms. As Ghana considers transitioning all public workers to a Single Spine Salary Structure, future policy research needs to more closely examine how such a change in public sector compensation would affect migration of health workers.

Of course, salaries for health workers are expensive, and other factors enter the cost-benefit calculus. For example, when dealing with fairly rigid public pay schedules, raising wages for health workers can sometimes lead to calls for wage increases in other sectors of the public payroll. A complete cost-benefit analysis is beyond the scope of this paper, but we do make progress toward measuring the cost-effectiveness of using higher salaries to reduce attrition of skilled health workers. This opens the way to compare the effectiveness of salary increases with other policy options.

Table 1.1. Destinations of Migrant Health Workers from Ghana

Destination	Nurses	Physicians
UK	71%	56%
US	22%	35%
South Africa	--	6%
Canada	3%	1%
Australia	2%	--
Other	2%	2%
Source	Ghana Nurses and Midwives Council	Dovlo and Nyonator (1999)

Table 1.2. Sample Statistics

	Potential Migrants	Others
Attrition	0.08	0.04
	(0.28)	(0.19)
Real Ghana Wage	410	174
	(206)	(82)
Real UK Wage	1158	577
	(788)	(299)
Age	30	29
	(2.8)	(3.4)
Male	0.48	0.38
	(0.50)	(0.48)
Nursebonding	0.31	0.00
	(0.46)	(0.00)
COPS	23.5	0.00
	(47.3)	(0.00)
Log UK Migrants	8.0	0.04
	(2.3)	(0.41)
N	14751	35872

Source: Administrative payroll data. Standard deviations are in parentheses. The sample includes only individuals under age 35.

Table 1.3. Defining Potential Migrants

Potential Migrants Worker Category	Obs (2003-2006)	Others Worker Category	Obs (2003-2006)
PROFESSIONAL NURSE	25,761	ORDERLY	18,138
MEDICAL OFFICER	3,648	COMMUNITY HEALTH NURSE	16,149
MEDICAL ASSISTANTS	1,773	HEALTH/WARD ASSISTANT	11,941
TECHINCAL OFFICER	940	MIDWIFE	8,724
TECHNICAL OFFICER (LAB)	858	ENROLLED NURSE	7,582
SPECIALIST	805	WATCHMAN	3,413
MEDICAL OFFER - HOUSEMAN	790	ACCOUNT OFFICER	3,057
BIOMEDICAL SCIENTIST	669	DRIVER	2,836
HEALTH SERVICE ADMINISTRATOR	538	ARTISAN	2,733
ANESTHETIST ASSISTANT	511	PHARMACY TECHNICIAN	2,420
TECHNICAL OFFICER (XRAY)	488	BIOSTATISTICS ASSISTANT	2,174
TECHNICAL OFFICER (NUTRITION)	482	FIELD TECHNICIAN	2,154
FACILITY/DISTRICT DIRECTOR	134	LABOURER	1,940
HEALTH RESEARCH OFFICE	117	TYPIST	1,887
HEALTH EDUCATOR	109	ACCOUNTANT	1,697
PHYSIOTHERAPIST	106	DISPENSING ASSISTANT	1,595
TECHNICAL OFFICER (ORTHOPEDIC)	102	KITCHEN ASSISTANT	1,588
PHYSIOTHERAPY ASSISTANT	98	SCAVENGER	1,588
DENTAL SURGEON	96	TECHNICAL OFFICER (CDC)	1,511
MINISTRY-LEVEL DIRECTOR	94	STOREKEEPER	1,417
OCCUPATIONAL THERAPY ASSISTANT	89	CONSERVANCY LABOURER	1,407
RADIOGRAPHER	75	MEDICAL RECORD ASSISTANT	1,396
HEALTH PLANNER	64	TECHNICAL OFFICER (BIOSTATS)	1,315
OPTOMETRIST	58	PHARMACIST	1,099
DENTAL TECHNOLOGIST	49	SECURITY GUARD	1,014
NURSING TUTOR	30	LABORATORY ASSISTANT	899
HEALTH EDUCATION OFFICER	28	WASHMAN/IRONER	868
OCCUPATIONAL THERAPIST	4	STENOGRAPHER	856

Table 1.4. Main Effects

Sample Dependent Variable	(1) Potential Migrants Attrition	(2) Potential Migrants Attrition	(3) Potential Migrants Attrition	(4) Potential Migrants Attrition
Log Ghana Wage	-1.72***	-1.45***	-1.45***	-3.07
	(0.34)	(0.50)	(0.49)	(2.34)
Log UK Wage	--	--	0.33	-0.26
			(0.41)	(0.89)
Log UK Migrants	--	--	0.03	-0.02
			(0.03)	(0.07)
Nursebonding	--	--	-0.13*	0.02
			(0.07)	(0.15)
COPS	--	--	0.0040**	0.012***
			(0.0017)	(0.004)
Gender	--	--	-0.03***	-0.03***
			(0.01)	(0.01)
Age Quartic	NO	NO	YES	YES
Year Fixed Effects	YES	NO	NO	NO
OccupationXYear Fixed Effects	NO	YES	YES	YES
GradeXStep Fixed Effects	YES	YES	YES	YES
Department Dummies	NO	NO	YES	YES
Region Dummies	NO	NO	YES	YES
GradeXStep Specific Time Trends	NO	NO	NO	YES
Obs	14,696	14,696	14,675	14,675
Number of Clusters	243	243	241	241
First Stage F	496	268	288	357

Statistical significance at the 1, 5, and 10 percent levels is denoted by ***, **, and * respectively. Standard errors are clustered at the gradeXstep level.

Table 1.5. Non-Potential Migrants

	(1)	(2)	(3)	(4)
Sample	Non-Potential Migrants	Non-Potential Migrants	Non-Potential Migrants	Non-Potential Migrants
Dependent Variable	Attrition	Attrition	Attrition	Attrition
Log Ghana Wage	0.37	0.37	0.42*	0.98
	(0.25)	(0.25)	(0.25)	(0.65)
Log UK Wage	--	--	0.03	-0.10
			(0.04)	(0.07)
Log UK Migrants	--	--	0.06	0.04
			(0.04)	(0.05)
Nursebonding	--	--	--	--
COPS	--	--	--	--
Gender	--	--	0.011***	0.012***
			(0.003)	(0.003)
Age Quartic	NO	NO	YES	YES
Year Fixed Effects	YES	NO	NO	NO
OccupationXYear Fixed Effects	NO	YES	YES	YES
GradeXStep Fixed Effects	YES	YES	YES	YES
Department Dummies	NO	NO	YES	YES
Region Dummies	NO	NO	YES	YES
GradeXStep Specific Time Trends	NO	NO	NO	YES
Obs	35,446	35,446	35,319	35,319
Number of Clusters	822	822	815	815
First Stage F	2260	2260	2303	525

Statistical significance at the 1, 5, and 10 percent levels is denoted by ***, **, and * respectively. Standard errors are clustered at the gradeXstep level.

Table 1.6. Heterogeneous Effects

	(1)	(2)	(3)	(4)
Sample	Under 35	Under 65	Under 65	Under 65
Dependent Variable	Attrition	Attrition	Attrition	Attrition
Log Ghana Wage	-1.45***	0.14	0.22	0.21
	(0.49)	(1.4)	(1.1)	(1.6)
WageXAge	--	--	0.015*	0.015*
			(0.009)	(0.009)
WageXGender	--	--	--	0.05
				(0.26)
WageXRural	--	--	--	0.00028*
				(0.00016)
Controls	YES	YES	YES	YES
OccupationXYear Fixed Effects	YES	YES	YES	YES
GradeXStep Fixed Effects	YES	YES	YES	YES
GradeXStep Specific Time Trends	NO	NO	NO	YES
Obs	14,675	58,112	58,112	58,112
Number of Clusters	241	774	774	774

Statistical significance at the 1, 5, and 10 percent levels is denoted by ***, **, and * respectively. Standard errors are clustered at the gradeXstep level.

Figure 1.1. Historical Migration of Physicians from Ghana

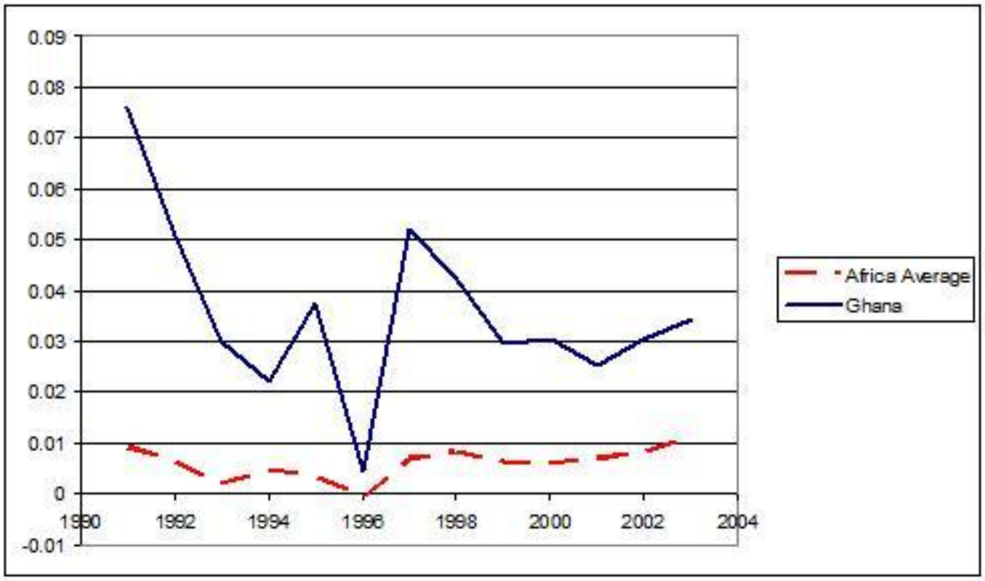

Source: Bhargava and Docquier (2007); the data have been converted into flows

Figure 1.2. Migration and Attrition of Nurses from Ghana

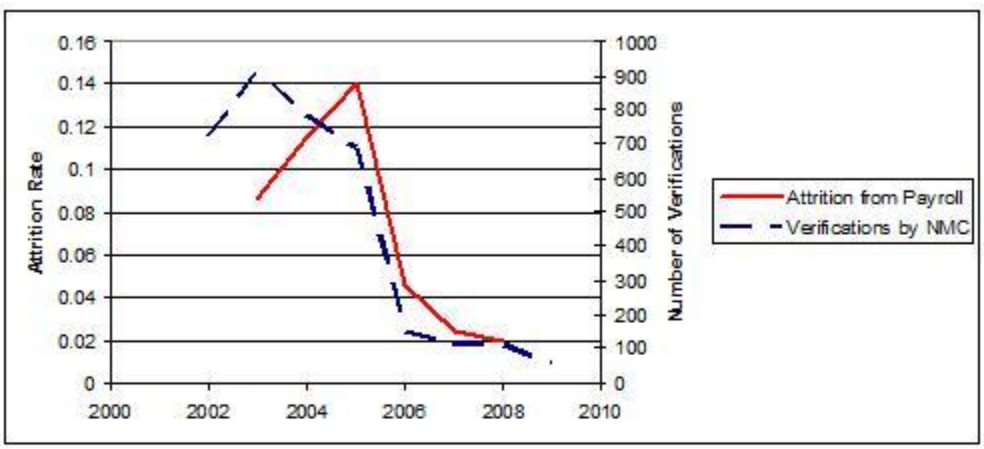

Source: Nurses and Midwives Council; IPPD Database

Figure 1.3. Migration and Attrition of Physicians from Ghana

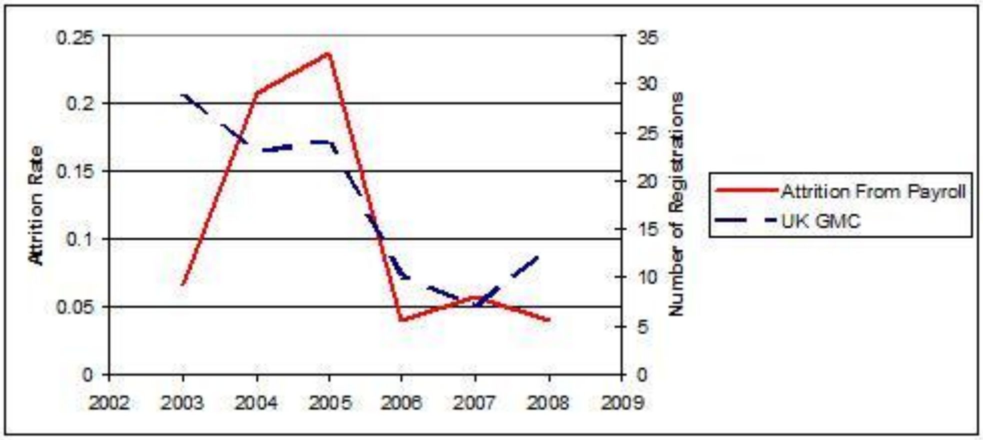

Source: UK General Medical Council and IPPD Database

Figure 1.4. Salaries of Health Workers in Ghana and the UK

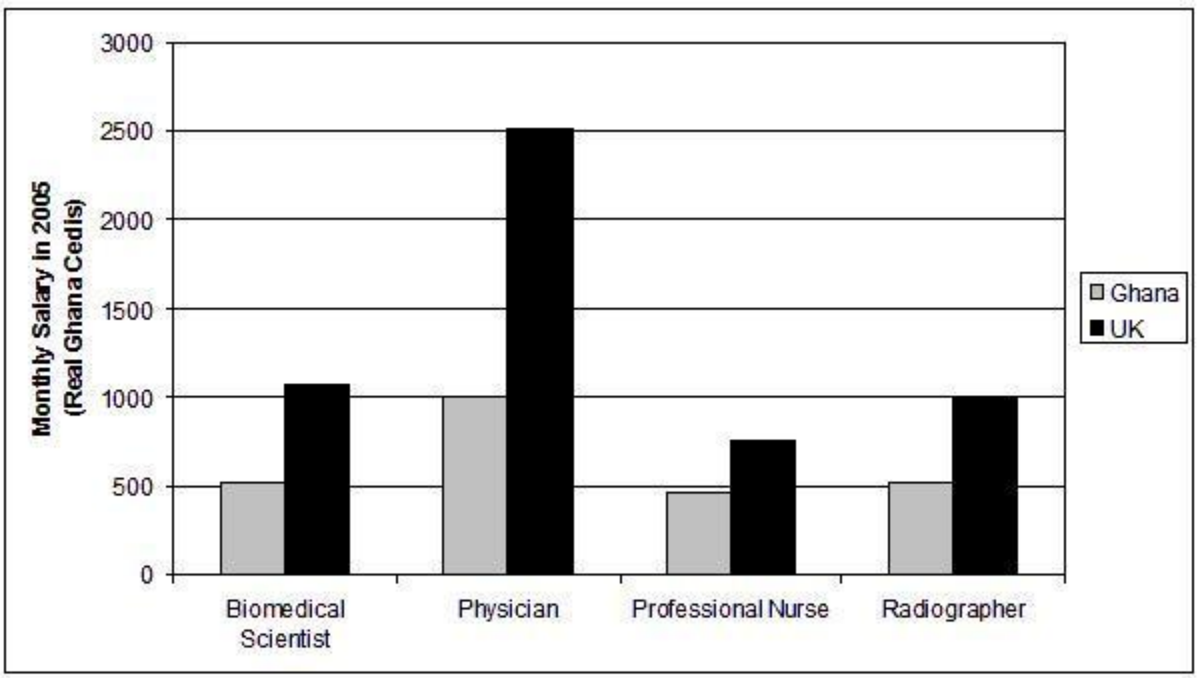

Source: IPPD Database and UK Annual Survey of Hours and Earnings; UK figures are PPP

Figure 1.5. Wages for Health Workers in Ghana, 2003-2008

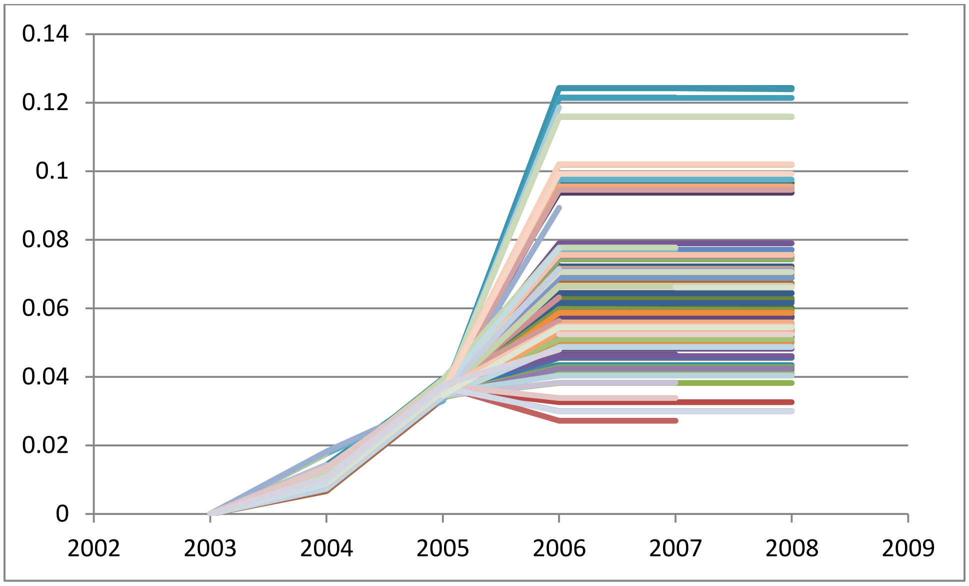

Each line indicates the real log wages of a particular gradeXstep group (e.g. senior medical officers on step 5); each group's wages are normalized to zero in 2003; Source: IPPD Database

Figure 1.6. Migration of Health Workers to UK from All Source Countries

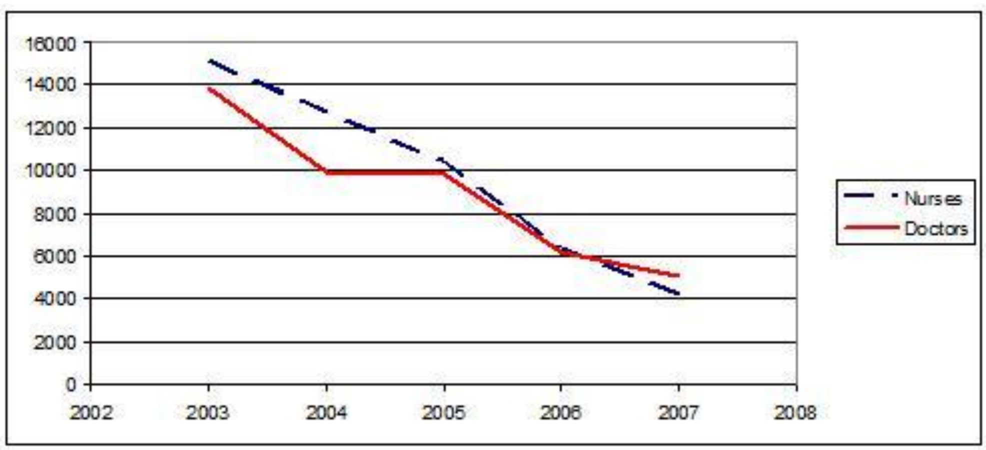

Source: UK General Medical Council; UK Nurses and Midwives Council

CHAPTER 2

GETTING TO WORK: EXPERIMENTAL EVIDENCE ON JOB SEARCH AND TRANSPORTATION COSTS[8]

2.1. Introduction

In the urban United States, de-facto residential segregation results in minority, poor individuals living in areas with few available jobs. Kain (1968) and Wilson (1997) have argued that this 'spatial mismatch' of workers from jobs contributes to adverse labor market outcomes for these individuals. In this chapter, I analyze spatial mismatch using an experiment of randomly provided public transit subsidies. In cooperation with a local, private, non-profit employment agency, I operate an experiment that provides public transit subsidies consisting of a fee-reducing public transit card as well as $50 of in-kind transportation subsidies to unemployed, low-wage job-seekers in Washington, DC. Individuals in the sample are disproportionately African-Americans from economically disadvantaged neighborhoods, providing a relevant sample in which to test the predictions of spatial mismatch. Additionally, this transportation intervention allows me to focus narrowly on spatial mismatch and access to local labor markets, separating these effects from other neighborhood characteristics and other impacts of residential mobility.

In analyzing the sample of 468 individuals, I find strong evidence that being assigned to treatment for subsidized transit reduces unemployment durations. This effect mainly operates

[8] This project was supported by the National Poverty Center using funds received from the U.S. Department of Health and Human Services, Office of the Assistant Secretary for Planning and Evaluation, grant number 1 U01 AE000002-03. The opinions and conclusions expressed herein are solely those of the author and should not be construed as representing the opinions or policy of any agency of the Federal government.

through reducing relatively short unemployment durations with the treatment group 9 percentage points more likely to have found a job within 40 days, narrowing to a statistically insignificant 5 percentage points after 90 days. Similarly, quantile regression results indicate that being assigned to receive the transit subsidy package decreases the 30th percentile of unemployment durations by two weeks, from 49 to 35 days, and the 35th percentile by three weeks, from 58 to 37 days. These results point to a large short-run effect of treatment on employment prospects that fades through time. On the other hand, I am not able to detect effects on earnings or wage rates, though this 'zero' is subject to considerable uncertainty.

The data provide some evidence that these lower unemployment durations result from increased intensity and scope of job search for individuals receiving the transit subsidy. Individuals in the treatment group search more intensely, completing 15 percent more applications and interviews. They also increase their scope of search, searching at locations that are on average 7 percent further from home than those of the control group, resulting in employment 20 percent further from home. However, these results for search actions are either marginally statistically significant or statistically insignificant due to the relatively small sample. Nonetheless they indicate potentially large changes in search behavior. Together, the results suggest that the causal mechanism for reducing unemployment duration likely runs through individuals in the subsidy group applying to more jobs and jobs further from home.

The results confirm the predictions of the 'spatial mismatch hypothesis,' which seeks to explain poor labor market outcomes using the distance of urban workers from available jobs (Kain, 1968). If minority individuals tend to live in areas with few job vacancies, then they are at a disadvantage in looking for work. A large empirical literature has used observational data and natural experiments to test the theory of spatial mismatch, with most of the literature confirming the existence of spatial mismatch (Raphael and Stoll, 2000; Zax and Kain, 1996; Holzer, Quigley, and Raphael 2003; Holzer, Ihlanfeldt, and Sjoquist, 1994; Raphael and Rice, 2002) A sizable theoretical literature interprets spatial mismatch as the outcome of a labor market search model where searching over distant is costly and minority workers reside in areas with fewer employment opportunities (Colson, Laing, and Wang, 2001; Gautier and Zenou, 2010; Gobillon, Selod, and Zenou, 2007). Building on this approach, I formulate a standard random search model of the labor market augmented to include variable search intensity and search over geographic space. My empirical results are consistent with this mechanism, whereby cutting transportation costs leads to greater search intensity which can translate into lower unemployment durations.

Despite its basis in theory and empirical studies using observational data, detecting the spatial mismatch hypothesis in experimental studies has been much more elusive. Most prominently, the Moving to Opportunity (MTO) experiment (Kling, et. al. 2007) tested the effect of neighborhood characteristics on labor market outcomes using a large-scale housing experiment. This program provided housing vouchers to a treatment group of families living in public housing projects with the condition that the vouchers could only be redeemed in lower-poverty areas. The experiment found no significant impacts of moving to richer neighborhoods on economic self-sufficiency for treatment families.[9] This has been interpreted by some as evidence against the importance of spatial mismatch, since change in residential location had no measurable impact on labor market outcomes. On the other hand, the housing mobility induced by MTO potentially changes an entire basket of neighborhood attributes including social

[9] The experiment did, however, document large improvements in mental health.

networks, vulnerability to violence, and access to labor markets. As a result, isolating the role of access to labor markets from other effects proves difficult, and a great deal of debate has surrounded how to interpret the non-effects for labor market outcomes (e.g. Clampet-Lundquist, S. and D. Massey, 2008; Ludwig, et. al., 2008).

The goals and scope of my study are less ambitious than MTO but also more focused on spatial mismatch. By experimentally subsidizing public transit for active job-seekers, I am able to isolate the impact of reducing the effective distance of individuals from available jobs, while leaving social networks and other factors related to residential location constant. In this light, my results lend support to the spatial mismatch hypothesis, indicating that an intervention focused solely on increasing geographic access to employment opportunities can improve labor market outcomes.

In the literature, the Bridges to Work (BTW) demonstration and evaluation is most similar to the present study. In BTW, the U.S. Department of Housing and Urban Development created a new shared van ride system in four cities and then evaluated it experimentally by randomly selecting individuals to receive free access to the system. However, in contrast to my results, BTW found negligible impacts on employment outcomes. The reason for this difference may stem from the difference in context and the design of treatment. Both Bridges to Work and the present study focus on using a randomized control trial to isolate the effect of better transportation on labor market outcomes, but while I focus on subsidizing access to an existing, extensive public transit system, Bridges to Work created a new transit system where such a system was lacking. As the evaluation of Bridges to Work describes, compliance of treatment individuals with the program was low, largely due to the enormous difficulty of operating a new, extensive transportation project in a timely and efficient manner. The program ultimately produced negligible intent-to-treat impacts on employment outcomes, but low compliance with treatment prevents measuring the impact of actually being treated with any precision (Roder and Scrivner, 2005). In the present study, I focus on increasing access to an existing transit system through simple individual-level subsidies, which is much less ambitious but also less prone to operational difficulties. In light of this previous literature, my results can be interpreted as an indication that alleviating spatial mismatch is easier in an environment with an extant mass transit system.

Aside from MTO and BTW, the recent literature on field experiments in labor economics does not appear to address the issue of transportation or spatial mismatch.[10] An extensive literature covers the use of randomized field experiments to study job search more generally. For example, randomized trials have been used to study the impact of job-finding bonuses (Meyer, 1995), active job search assistance (Dolton and O'Neill, 1996; Dolton and O'Neill 2002), and making re-employment services mandatory (Black, et. al., 2003). More recently these efforts have been expanded to examine spillovers of job placement services on non-treated individuals (Crepon, et. al., 2011). Yet the literature that applies an experimental approach to the role of transportation in job search, either in general or with the urban poor in particular, remains small. Thus, my study contributes to the literature by providing the first experimental evidence in support of the spatial mismatch hypothesis.

[10] See for example a recent review of field experiments by List and Rasul, 2010

2.2. Theoretical Predictions

Though this chapter focuses on empirical results from a randomized experiment, theory guides the study questions and experimental design. So, consider a standard partial equilibrium random search model, augmented to include variable search intensity and a spatial dimension. Unemployed job seekers receive job offers, (w, δ), which are a pair including a wage and a commute distance. Offers are received in continuous time from an exogenous distribution at a Poisson rate $\lambda(\delta)$, which can vary by distance, and in this model the arrival rate will be a choice variable. So, individuals will be choosing search intensity over all geographic locations with agents choosing the function $\lambda(\delta)$. Search intensity is costly for these potential workers, and the cost of increasing the arrival rate at location δ is $c(\lambda(\delta), \delta)$, which I will assume to be Cobb-Douglas, $\gamma\lambda(\delta)^{\alpha}g(\delta)$, where costs and marginal costs are increasing in δ (i.e. $g(\delta)$ is positive and increasing). As a result, the flow value of unemployment is $b - \int_{\delta} \gamma\lambda(\delta)^{\alpha}g(\delta)d\delta$. I will assume $\alpha > 1$ to ensure that the cost of search intensity is positive, increasing, and convex.

For simplicity, I will assume that the wage distributions at all locations are identical, i.e there is a common wage distribution $F(w|\delta) = F(w)$ at all locations with density $f(w)$. When job offers arrive, individuals choose whether to accept the offered wage-distance pair or continue searching. If they accept, they receive w at location δ forever after, discounting the future at rate r. If they reject, they continue searching, receiving the flow value of unemployment. In sum, I simply consider a standard random search model with variable search intensity, where a notion of distance influences search.

For the moment, take $\lambda(\delta)$ to be given. Given linear preferences over wages, the value function for a worker employed at wage w and commuting distance δ is:

$$rV(w,\delta) = w$$

For simplicity, commuting costs are not taken into account for employed workers.[11] The value of being unemployed is:

$$rU\big(\lambda(\delta)\big) = \; b - \int_{\delta} \gamma\lambda(\delta)^{\alpha} g(\delta) d\delta + \int_{\delta} \; \lambda(\delta) \int_{w} \max\big[V(w,\delta), U\big(\lambda(\delta)\big)\big] f(w) dw d\delta$$

For any given value of λ, the well-known result is that the optimal strategy is to set a reservation wage $w^{*}(\lambda(\delta),\delta)$, which here can depend on distance from home. Offers above the reservation wage are accepted. The reservation wage is defined as the wage at which the value of unemployment equals the value of employment, $V(w^{*}(\lambda(\delta),\delta),\delta) = U\big(\lambda(\delta)\big)$, or:

$$w^{*}(\lambda(\delta),\delta) = rU(\lambda(\delta))$$

Because there are no commuting costs, this is true for all δ. As a result, in this model the optimal reservation wage w^{*} is the same for all distances. Then, this reservation wage can be expressed as[12]:

$$w^{*} = \; b - \int_{\delta} \gamma\lambda(\delta)^{\alpha} g(\delta) d\delta + \int_{\delta} \; \lambda(\delta) \int_{w^{*}}^{\infty} \frac{w - w^{*}}{r} f(w) dw d\delta \qquad (2.1)$$

The optimal choice of search distance for an unemployed worker consists of maximizing the value of unemployment, or equivalently, the reservation wage. Given the Cobb-Douglas cost function, search intensity is positive everywhere. In this interior solution, the optimal choice of $\lambda(\delta)$ can be characterized by a first-order condition:

[11] This is only for expositional simplicity. The comparative static predictions of the model are unchanged if I include commuting costs.

[12] This equality is the result of noting that the reservation wage must satisfy $rU = w^{*}$, substituting in this equality twice, and substituting in for $V(w)$.

$$\gamma\alpha\lambda(\delta)^{\alpha-1}\, g(\delta) = \int_{w^*}^{\infty} \frac{w - w^*}{r} f(w)dw \qquad (2.2)$$

This model provides predictions for how transportation subsidies could affect job search. If transportation costs are considered to be part of job search costs, then a public transit subsidy will reduce γ, leading to lower marginal search costs for the treatment group.[13] The comparative static effects of a decrease in γ on the worker's strategy then follow from (2.1) and (2.2). The model predicts that in response to decreasing the marginal cost of search intensity, workers search more intensely at all locations while also becoming more selective, raising their reservation wage.

$$\frac{\partial\lambda(\delta)}{\partial\gamma} = \frac{-(1-F(w^*))\int_{\delta}\,(\alpha-1)\lambda(s)^{\alpha}g(s)ds - r\alpha\lambda(\delta)^{\alpha-1}\,g(\delta)}{[r+\int_{s}\,\big(1-F(w^*)\big)\lambda(s)ds\,]\,[\alpha(\alpha-1)\gamma g(\delta)\lambda(\delta)^{\alpha-2}]} < 0 \qquad (2.3)$$

$$\frac{\partial w^*}{\partial\gamma} = \frac{-r\int_{s}\,\lambda(s)^{\alpha}g(s)ds}{r+\int_{s}\,\big(1-F(w^*)\big)\lambda(s)ds} < 0 \qquad (2.4)$$

The first result is quite intuitive with lower marginal search costs leading to greater search intensity. This action results from equation (2.2) where a lower γ causes the worker to choose a higher $\lambda(\delta)$ for all values of δ.[14] Thus search intensity increases at all locations. The impact of lower search costs on the geographical concentration of search is less clear. Since workers are searching more everywhere, the average distance searched by workers could either

[13] In practice, the experiment both reduces marginal transit fees and provides an in-kind subsidy. I will simplify the exposition by considering only the shift in marginal costs. The comparative static effect of an in-kind transfer on wages, search intensity, and unemployment durations is qualitatively similar to a shift in marginal costs, provided that the transfer is not infra-marginal. If it is infra-marginal, then the transfer becomes an income supplement during unemployment, leading to higher wages and lower search intensity.
[14] Technically, the right side of (2.2) is also increasing due to the rising reservation wage. Given the Cobb-Douglas functional form, the impact on marginal costs dominates the effect through wages. See Appendix for a proof.

increase or decrease. So, when observing an experimental shock to search costs, we should expect to see greater search intensity with unclear expectations for where that search occurs.

The model also predicts that workers will raise their reservation wage in response to lower search costs. This occurs because decreasing γ increases the flow value of being unemployed, leading workers to become more selective about which wages they accept. One might expect that the simultaneous increase in search intensity might also have an impact on the reservation wage. However, this does not occur in this model because $\lambda(\delta)$ is being set to maximize w^*. As a result, this potential effect disappears with equation (2.4) following directly from applying an envelope condition to (2.1).[15] As a result, this model predicts that an experimental decrease in search costs should lead to higher reservation wages and thus higher observed wages.

The effect on unemployment duration is less clear. The hazard probability of leaving unemployment is:

$$h = \int_{\delta} \lambda(\delta) d\delta [1 - F(w^*)] d\delta \qquad (2.5)$$

Due to increased search intensity, job offers arrive at a faster rate, increasing the first term, but job-seekers also raise their reservation wage. This increased patience decreases the probability of accepting an offer, $[1 - F(w^*)]$. Together, the impact of lower marginal search costs on the hazard probability of leaving unemployment, and thus unemployment duration, depends on the parameters of the model. In models where search intensity is exogenous, certain general conditions can be assumed to get the result that a higher search intensity will decrease

[15] See Appendix for proof.

unemployment duration (e.g. van den Berg (1994)). In this situation, though, with endogenous search intensity, I cannot make a prediction about unemployment duration. However, equation (2.5) provides some guidance. If decreasing search costs result in a small change in the reservation wage but a large increase in search intensity, unemployment durations should decrease. In a case with a binding minimum wage, the wage effect will be zero and unemployment durations must decrease. On the other hand, a large wage effect should be associated with increased durations. In the experimental results, we should expect this sort of inverse relationship between wage effects and duration effects.

This model provides a tractable theory of how distance from employment opportunities can influence labor market outcomes. It predicts that a public transit subsidy which reduces the effective distance of workers from jobs should result in greater observed search intensity and higher observed wages. Finally, depending on which of these effects dominates, unemployment durations may increase or decrease. Given the experimental variation in transportation costs, these predictions can be empirically tested.

2.3. Experimental Design

2.3.1. Treatment

In this experiment, treated subjects received a fee-reducing fare card and an in-kind subsidy to ride Washington Metropolitan Area Transit Authority (WMATA) buses and trains, which cover the entire DC metropolitan area. In particular, they received a reusable SmarTrip card with a free $25 balance. When this balance had been depleted, clients who were still looking for work could return the empty card for a second card with a fresh $25 balance. Given that the card itself costs $5, the total cost of the package is $60. The balance functions similarly to a gift card in that the balance is debited with each trip, except that the card ceases to work only after the balance is negative. This $50 subsidy represents about 33 bus trips ($1.50 each) or 10 to 31 train trips ($1.60 to $5.00 each), depending on distance.

In addition to the $50 card balance, the SmarTrip cards also provide price subsidies for subsequent trips relative to using cash because WMATA charges different base prices to users of SmarTrip cards and those who pay with cash. Bus trips receive a 20 cent discount and train trips receive a 25 cent discount relative to riders using cash. Finally, the SmarTrip card also allows passengers to transfer between buses for free or between bus and rail at a 50 cent discount so long as the rides are within two hours of each other. This service is not available to passengers using cash and can represent a considerable subsidy to job-seekers who can make multiple successive bus trips in one day on a single fare.

Altogether, treatment provides a significant package of transportation subsidies that both decreases the marginal cost of public transit through a fee-reducing transit card and also provides an in-kind subsidy. In the model above, I simplify the theory by only considering a shift in the marginal cost of search. However, the in-kind subsidy will have effects similar to a drop in the marginal cost of search. Using labor market histories of individuals in the sample, I estimate that in a typical year an average member of the sample would have a monthly income of about $552.[16] Even this average is likely a significant overestimate of the sample's recent income level since the median individual has been unemployed for nearly one year. Together these facts indicate that the participants in the experiment likely face binding credit constraints that prevent spending on public transit even if the return to search is high. In this context, the $50 in-kind

[16] See summary statistics below. An average individual is employed for 30 hours per week at $10 per hour and is only employed 46 percent of the time.

subsidy is unlikely to be infra-marginal, and thus can also be interpreted as reducing the cost of search. Finally, control subjects did not receive transportation assistance but continued to receive job placement assistance from the partnering organization. Thus, I will measure the impact of transportation contingent on also receiving job placement services.

2.3.2. Study Design

I cooperated with a local, non-profit, private job placement assistance organization, Jubilee Jobs, to implement this experiment. Jubilee Jobs has two sites located in Washington, DC with one in Anacostia, a predominantly low-income and African-American area in southeast Washington, and one in Adams Morgan, a racially and economically diverse area in north-central DC. They provide job skills training and job placement assistance free of charge to all interested individuals, but Jubilee focuses on only low-wage employment. As a result they typically assist low-income individuals, often receiving public assistance, re-entering the workforce after incarceration, or recovering from substance abuse.

Applicants come through Jubilee Jobs in orientation cycles. A particular cohort begins with an orientation session, followed by a week of job skills workshops. Then, applicants search for employment with assistance from a job counselor, who actively markets them to prospective employers and meets with them regularly to set-up interviews. Every two weeks, a new cohort begins this process. Due to the timing of services provided by Jubilee Jobs, the experiment also runs in a sequential manner. After the initial week of workshops, applicants were recruited into the study through an informed consent process. This process was repeated roughly every two weeks from November 2010 through June 2011.

2.3.3. Recruitment

Table 2.1 summarizes participation in the experiment. Of the potential participants, 60 percent consented to participate in the study. Lack of informed consent prevents me from directly collecting data on reasons for non-participation; however, observation and volunteered information point to lack of time for informed consent as the main reason. Nonetheless, for some individual baseline characteristics, Jubilee Jobs was able to provide summary statistics on all individuals who were enrolled in their program at the time of recruitment for the experiment. Combining these summary statistics with baseline data on individuals who participate, I am able to calculate summary statistics for the non-participants in a manner that does not require obtaining their individual-level data.

Table 2.2 serves both to report baseline characteristics of individuals participating in the experiment and to compare them to average characteristics of those who did not participate. The first column reports mean baseline characteristics for those who were recruited. Nearly the entire sample, 98 percent, identifies as black or African-American and the average age is 40 years. The baseline data confirm that these low-wage job seekers face substantial disadvantages in the labor market. Only 24 percent of the sample attended any schooling after high school, with only 5 percent completing college, while 20 percent did not complete high school. Half of the sample is listed by their job counselors as having a criminal history and two thirds are receiving some type of public assistance. Only 11 percent of the sample is employed at baseline.[17] Strikingly, the average individual in the sample has been unemployed for 1.8 years (with a median of just under

[17] Most of these individuals are employed part-time, and in the analysis that follows, I analyze outcomes related to finding a new job.

1 year) and been employed for only 46 percent of the past 5 years. When employed, the most recent wage was about 10 dollars per hour and turnover was high with a median job duration of 1 year. While nearly all of the sample identify as black, 10 percent also identify as immigrants, over half of whom are from Ethiopia. The baseline data also identify this sample as a group of people reliant on public transit. More than half of the sample lacks a valid driver's license, and only 9 percent have access to a car. Finally, 82 percent of the sample resides in Washington, DC with most of the remainder living in the Maryland suburbs east of the city. As covered in more detail below, the sample inside DC disproportionately comes from Wards 7 and 8 in the Southeast part of the city, with very few applicants coming from the more affluent western part of the city.

For most of these characteristics, summary statistics can also be computed for those who did not participate. These are reported in the second column of Table 2.2.[18] Differences between mean characteristics for recruited and non-recruited individuals, along with p-values testing if this difference is different from zero, are reported in the final two columns. For the most part, individuals recruited into the experiment are similar to those who did not participate with no average differences in age, gender, educational attainment, ex-offender status, immigrant status, ability to drive, or residential location. However, individuals do appear to have selected into the experiment on a couple dimensions. The experimental sample is 17 percentage points more likely to be receiving public assistance and 5 percentage points less likely to have access to a car. Those recruited are also more likely to be black, though this difference is relatively small. Altogether, this baseline comparison indicates that individuals who selected to participate in the

[18] Unlike other baseline characteristics, the labor market history variables are stored in individual paper files and could not be obtained in a way that protected the privacy of those who refused to participate.

experiment had less access to transportation and a lower level of assets at baseline relative to individuals who decided not to participate. While selection into participating in the experiment has no impact on the internal validity of the treatment effects that I measure, it indicates that this study will be most externally valid for populations with heavy participation in public assistance programs and limited access to private vehicles. In practice, a scaled-up policy of transit subsidies for urban job-seekers would likely attract people with similar characteristics. While there is selection into the experiment, the sample is the one of most interest for measuring the impact of the transit subsidies.

2.3.4. Compliance with Treatment

Applicants who decided to participate are randomly assigned to treatment and control groups using a random number generator with treatment stratified by cohort, Jubilee Jobs site (two separate sites), and ex-offender status. In an initial pilot from November to January, the treatment probability was one quarter. Starting in February, as more funding became available, half of all subjects were offered treatment. As Table 2.1 documents, this results in slightly less than half of the sample being treated with 208 of 468 individuals assigned to treatment.[19]

Treatment group individuals were provided with a transit card (WMATA SmarTrip). Each selected applicant was offered the card from Jubilee Jobs staff prior to being sent out on their first interview and could obtain a second card if the first was exhausted prior to finding employment and they continued to work with Jubilee Jobs.[20] Table 2.1 also describes compliance of individuals with treatment, i.e. take-up of the SmarTrip cards. Take-up is high with 89 percent

[19] Note that since the probability of treatment varies over time, re-weighting the data or including strata dummies as control variables will be necessary in the analysis.
[20] This exchange helped protect against re-sale of the cards and connected the cards more closely with job search.

of the treatment group picking up at least one SmarTrip card and half receiving two cards.[21] The control group, meanwhile, did not receive any transportation assistance from Jubilee Jobs, as indicated in the table. Altogether, this points to high but imperfect compliance with treatment. Less than full take-up of the second card mainly occurred if individuals found employment prior to obtaining a second card or dropped out of Jubilee's program. To deal with imperfect compliance, all of the analysis will focus on Intent-to-Treat effects based on original treatment assignment.

2.4. Data and Follow-Up

Data come from three sources: Jubilee Jobs administrative records, a phone survey, and the WMATA SmarTrip on-line card manager.

2.4.1. Baseline and Randomization Test

Jubilee Jobs collects background and demographic information on all clients in intake interviews prior to randomization. This data provides baseline information on demographics, labor market history, ex-offender status, education attainment, and more. I make use of these data to characterize the baseline characteristics of the sample as detailed in Table 2.2. I also check the validity of randomization by testing the balance of baseline characteristics including all of the characteristics described above as well as an indicator variable for individuals with no work history and dummies for which of the nine job counselors an individual is working with. While the first two columns of Table 2.3 report mean baseline characteristics for treatment and control, the third column reports the coefficient on a treatment dummy in an un-weighted regression of baseline characteristics on treatment and indicators for all of the cohort-site-ex-

[21] Two individuals were able to obtain 3 cards due to administrative errors.

offender strata of the randomization. Of the 29 baseline characteristics, only two (Virginia residence and one of the job counselor dummies) have significant differences at the ten percent level and none are significant at the five percent level. This is consistent with the number of false positives expected from randomly assigned treatment when testing 29 characteristics.

2.4.2. Follow-Up and Attrition

I focus on two types of outcomes, those regarding employment and those regarding aspects of the search process. Administrative records from Jubilee Jobs track job search and employment outcomes for most of the sample. These records allow me to measure wages, unemployment duration, and an application/interview history that can be used to measure search intensity. These data also include home addresses and job locations that can be used to measure distance to jobs, interviews, and applications. I collect these records 3 months after enrollment of an individual in the study.

For employment outcomes, I combine administrative data with a phone survey. After 3 months, individuals fall into three groups: those who have reported a job placement to Jubilee Jobs, those who continue to search for work with Jubilee, and those who have lost contact with Jubilee Jobs. The first group has complete data on employment outcomes in the administrative data. For the second group, durations are censored at 90 days and employment characteristics cannot be observed. I take account of this censoring in analyzing durations. When considering earnings, one might be concerned that selection into employment biases estimates of the treatment effect on wages and earnings. To account for this, I examine both wage rates conditional on being employed at 90 days and weekly earnings, where earnings are assumed to be zero for those who are unemployed. For outcomes regarding the location of employment, I analyze both distance to work for the employed sub-sample and for the whole sample with unemployed individuals coded as a zero value.

The final data follow-up group, those who drop out of Jubilee's program and thus the administrative data, provide a more important problem. About 40 percent of the sample drops out of the administrative data between baseline and follow-up; however, there is no statistical difference between attrition in the treatment and control groups. Nonetheless, I attempt to complete the missing employment data using a phone survey of all individuals who drop out of the administrative data. The survey asks questions on employment and job characteristics. As listed in Table 2.4, attrition is reduced to 9 percent by the phone survey, and differential attrition between treatment and control remains negligible. Thus, the combination of administrative and phone survey records provide employment outcomes data with relatively low attrition.

As an additional check to ensure that selection into attrition does not introduce bias into the estimates for employment characteristics, I test baseline balance for only those individuals with follow-up data, either in the administrative records or the phone survey. Table 2.5 reports these results, which show no difference in baseline characteristics between the treatment and control groups. It appears that the validity of randomization is not violated by attrition from the employment data.

For all individuals, I also measure search outcomes, including the number, timing, and location of job applications. For these outcomes I use only Jubilee Jobs administrative data on the application and interview history of the individual as recorded by their job counselor. Selection out of this dataset occurs either when the individual finds employment or loses contact with Jubilee Jobs. To prevent bias from selecting out of job search, I focus on search outcomes from the first two weeks after randomization. Up to that point employment rates are under 10 percent and the difference between treatment and control is less than 1 percent. Meanwhile, dropout rates are similarly low. While this approach reduces the information available in the data, I focus on this narrow window to eliminate sample selection bias in job search outcomes.

2.4.3. Transit Card Data

Finally, I also collect data on usage of the transit cards. Electronic tracking of the cards provides time-stamped data on card transactions and balances. While I only have this information for the treatment group, these data assist in confirming that applicants do in fact use the cards and also provide some detail as to whether applicants appear to be using the cards for job search or for other purposes.

2.5. Context

2.5.1. Spatial Mismatch in Washington, DC

Washington, DC provides a prime example of how the de-facto residential segregation which has taken the place of legal segregation can lead to spatial mismatch of workers from available jobs. Figure 2.1 uses 2008 data on establishment-level employment per square mile across zip codes in the DC metro area from the U.S. Census Bureau's Zip Code Business Patterns, with darker areas indicating more employment in that area. As is apparent, employment in the city is concentrated downtown, and in the suburbs, jobs are more abundant to the southwest in Northern Virginia and to the northwest centered on Bethesda. Firms employ far fewer workers in the eastern parts of the city or the Maryland suburbs to the east.

For the same zip codes, Figures 2.2 and 2.3 map the the population density of African-Americans and White-Americans using data from the 2000 Population Census. These figures demonstrate the de facto residential segregation that characterizes the DC metro area, with the black population living mostly in the eastern part of the city with significant presence in the eastern suburbs of Prince George's County Maryland and the white population concentrated in the western part of the city as well as the western suburbs of Northern Virginia and Montgomery County Maryland. Together with Figure 2.1, these images demonstrate the distance between many minority job-seekers and available jobs in the DC labor market. While jobs are abundant downtown and to the west, individuals living in largely African-American communities to the east must travel extensively to access these employment opportunities.

Baseline data for my sample matches these qualities of the census data. At the time of recruitment, 83 percent of the sample lives inside the boundaries of the district, almost entirely in the eastern half of the city. The remainder largely lives in the Maryland suburbs east of the city. Figure 2.4 maps the residential addresses in the sample across DC area zip codes. The residential patterns of this sample closely follow those of the broader African-American population shown in Figure 2.2. Through the course of its regular operations, Jubilee Jobs collects information on available job vacancies that are appropriate for their clients. Figure 2.5 maps administrative data on these vacancies for the period April 2010 to April 2011. Not surprisingly, Jubilee targets vacancies somewhat closer to the homes of their applicants than the average job; however, even these vacancies, which are targeted as best suited for these particular individuals, tend to be in the western half of the city. As a result, the individuals participating in this experiment face a situation in which job vacancies exist but are often far from their residences. If the spatial mismatch hypothesis is true, taking steps to reduce the effective distance between potential workers and available jobs should be relevant for labor market outcomes in precisely this sample.

2.5.2. Comparability of the Sample

While this sample provides an interesting setting in which to study transportation subsidies, the external validity of the results will depend on how this context compares to other settings. To provide some sense of context, I compare the characteristics of the sample to data on respondents to the Current Population Survey in Table 2.6. The first column of Table 2.6 lists the educational attainment, age distribution, and gender distribution of my sample. The next three columns summarize 2010 CPS data for progressively selective groups. The second column describes all 18-65 year-old respondents; the third column 18-65 year-old unemployed individuals; and the fourth column 18-65 year-old, unemployed, black individuals living in metropolitan areas. Not surprisingly, the experimental sample has lower education attainment than the average working age person, with for example only 5 percent achieving at least a Bachelor's degree as opposed to 28 percent of the general population. More interestingly, the sample is less educated even when compared only to black, unemployed individuals in metro areas, of whom 11 percent completed a Bachelor's degree.

There are also major differences in age distribution between the experimental sample and a random sample of unemployed, black individuals living in metropolitan areas. While young people age 18-25 make up 31 percent of the CPS sample, only 11 percent of my sample is this young. Instead, the experimental sample has a disproportionate number of middle aged individuals. As such, the results in this study are much more relevant for mid-life to older individuals rather than urban youth. Finally, there are no major differences in gender distribution between the sample and the CPS data (when restricted to the unemployed).

2.6. Results

2.6.1. Transit Card Usage

As noted above, the vast majority of individuals assigned to treatment take up at least one free transit card. I use electronic data on usage of the cards to more precisely measure uptake of treatment. One first question when providing transit subsidies is the extent to which subsidies are used for job search versus other activities, including lending or selling the card to someone else. While tying the receipt of a second card to continued job search provides one check against these possibilities, I can also check the card usage data for consistency with job search. I find that trips are disproportionately concentrated on weekdays and during regular business hours. Figures 2.6 and 2.7 demonstrate these facts. While far from definitive, these findings indicate that observed card usage is consistent with transit subsidies being used for job search activities.

Given the transaction-level data, I am able to calculate both total spending on transit using the cards and deposits by individuals onto the cards. Figure 2.8 depicts three variables over time: total spending, total deposits, and spending net of deposits. Usage of the cards is extensive, with average spending surpassing $50 within two months of first using the card. Also of interest, individuals receiving the cards do not simply use the cards for the $50 subsidy, but continue depositing their own money onto them. Deposits are slow to pick up at first as the subsidy is used up, but within four months, the average applicant has added an additional $50 to the balance of their card(s). Spending net of these deposits rises quickly and flattens out within the first 50 days.[22]

These results indicate a pair of important facts about treatment. First, applicants make use of the lump sum subsidy relatively quickly, with individuals exhausting the subsidized balance with the first 50 days. Second, applicants demonstrate a revealed preference for using the transit card even after the subsidized balance has been exhausted. This seems to indicate that the value of treatment is not exclusively due to the balance but also because of the fare reductions that the card provides. In this sense, it is most accurate to think of treatment as a package intervention that both lowers the price of transit and provides a lump sum, in-kind subsidy.

2.6.2. Unemployment Durations

The spatial mismatch hypothesis posits that geographic isolation diminishes labor market outcomes. I will measure these outcomes through unemployment durations, weekly earnings, and wage rates. A standard search model predicts that wages should rise in response to decreased search costs, while the impact on unemployment duration is ambiguous. Using the available administrative and phone survey data paired with the random variation in transit subsidies, I can examine this question experimentally.

Figure 2.9 graphs the Kaplan-Meier survivor functions for the data. These functions non-parametrically describe the probability of remaining unemployed after a given number of days for both treatment and control groups. The survivor function for the treatment group clearly lies below that for the control group, indicating that for nearly all duration lengths the probability of remaining unemployed until that time is lower for the treatment group. As is apparent from the

[22] The average subsidized balance is less than $50 because of imperfect compliance with treatment.

graph, this effect is strongest from about 30 to 60 days after randomization after which the gap between treatment and control narrows. The difference between the survivor functions can be tested non-parametrically using a Kolmogorov-Smirnov test of the censored distribution of unemployment durations. A one-sided test of the null that the distribution of durations for the treatment group first order stochastically dominates that of the control group rejects the null at the 10 percent level (p-value 0.086). This indicates that there is some unemployment duration length for which the treatment group has a lower probability of being unemployed[23] and provides some confirmation of the graphical intuition that the treatment group appears to have a lower unemployment duration over a particular range.

To further investigate potential effects on employment, I test the difference between treatment and control in the (unconditional) probability of employment at each possible duration. In particular I estimate:

$$I[d_i \leq X] = \alpha + \beta_X T_i + s_i + \epsilon_i \qquad (2.6)$$

where d_i is the unemployment duration of person i, X is an integer, T_i is an assignment to treatment dummy, and s_i is a strata fixed effect. So, β_X gives the ITT effect on the probability that unemployment duration is less than or equal to X, i.e. testing the impact of treatment on employment probability at each point on the survivor function. Figure 2.10 graphs the results of repeating this regression for all X between 1 and 90. This formalizes the intuition above that treatment appears to have no effect on the probability of employment during the first month after treatment but then has a large effect, increasing the probability of employment by up to 10 percentage points in the 35-55 day range, a difference that is statistically significant. After two

[23] A two-sided test of the equality of the unemployment durations for treatment and control yields a p-value of 0.157

months, this difference shrinks and becomes statistically insignificant. While care needs to be taken in testing a large number of hypotheses in this manner, the results have a clear pattern indicating that treatment has a strong effect in reducing durations in the second month after treatment. Figure 2.11 displays the impact of treatment starkly by plotting the density of unemployment durations for the treatment and control groups. Treatment causes a leftward shift in probability density, with members of the treatment group much more likely to experience durations of 20 to 50 days and much less likely to experience durations of 50 to 80 days.

Given that the impact of treatment appears to differ considerably across the distribution of unemployment durations, examining the results of quantile regressions is useful. If the distribution of unemployment durations is $G(d)$ then the τ quantile of unemployment durations can be defined as:

$$Q(\tau) = G^{-1}(\tau)$$

In other words, $Q(\tau)$ is the unemployment duration length corresponding to the τ percentile, so for τ equal to 50, $Q(50)$ is the median unemployment duration. To test the impact of treatment on any quantile, I consider the conditional quantile function:

$$Q(\tau|T) = \alpha_\tau + \beta_\tau T$$

where α_τ is the τ quantile for the control group and β_τ measures the impact of treatment on the τ quantile. I estimate this conditional quantile function for every 5th quantile up to the 45th quantile using a bootstrap with 1,000 repetitions to estimate standard errors.[24] Table 2.7 provides the results of these estimations with each row representing a different quantile. For example, the first row shows the estimates for the 5th quantile, indicating that 5 percent of the control group

[24] I cannot estimate impacts over the 45th quantile because of censored durations.

finds employment within 8 days and treatment increases this by a statistically insignificant 1 day. As expected, there are no major differences between the treatment and control groups at low quantiles. However, in the middle range treatment has a large impact in reducing unemployment durations. Treatment causes the 30th percentile of the unemployment duration to drop by two weeks, from 49 to 35 days, and the 35th percentile drops by three weeks, from 58 to 37 days. Effects for higher quantiles are likewise large but larger standard errors make the estimates statistically insignificant. Altogether, these results indicate that receiving the transit subsidies has a strong effect on unemployment durations, reducing mid-length durations by about two weeks. Thus, this modest public transit subsidy appears to have a strong short-run effect on unemployment durations.

For employment as well as other outcomes, I will test differences in the means as well. To analyze the impact of treatment on these mean outcomes, I consider the following regression:

$$Y_i = \alpha + \beta T_i + \delta Z_i + s_i + \epsilon_i \qquad (2.7)$$

where all variables are as before except that Y_i is an outcome variable and Z_i is a vector of individual-level demographic and labor market history control variables. For all outcomes, I include results with and without Z_i. The first three rows of Table 2.8 replicate previous analysis of unemployment durations. The first two rows estimate equation (2.7) using a dependent variable dummy for being employed at 40 days and at 90 days as the outcome. The first two columns provide (weighted) means for the treatment and control groups while the third and fourth columns report on the coefficient on treatment from (2.7). The results confirm the previous discussion, that the impact of treatment on employment probability is strong, 9 percentage points, in the second month but diminishes to a statistically insignificant 4 to 5 percentage points thereafter. Similarly, comparison of mean unemployment duration, conditional on duration less than 90 days, yields a negative point estimate of about 3 days, though it is statistically insignificant.

2.6.3. Wages and Earnings

Both the search model discussed above and the spatial mismatch hypothesis predict effects not only on unemployment duration but also on wages. The next two rows of Table 2.8 examine the impact of treatment on weekly earnings and wage rates. Earnings are computed for the whole sample with hours assumed to be zero if an individual is unemployed.[25] Meanwhile, wage rates are measured only conditional on being employed within 90 days. In testing the impact on earnings, I find a relatively large but statistically insignificant impact of treatment on earnings with a point estimate of 16 dollars per week. However, if this effect exists, wage rates for the sub-sample of employed individuals do not appear to drive the difference in earnings. The treatment group receives wage rates that only marginally differ from the control group, and even this difference shrinks when controlling for demographics that may be related to selection into employment. Together, the results for earnings and wage rates indicate that treatment does not appear to have large effects on reservation wages. This is perhaps not so surprising. First, the job search model predicts that impacts in reducing unemployment duration will be greatest when wage effects are smallest. Given the strong impact of treatment in decreasing durations, theory would predict that the wage effects should be small. Second, by far the modal wage in the sample is the DC minimum wage of $8.25. If the minimum wage binds so that individuals are

[25] Actual hours are used when available. For much of the sample, though, they are estimated by assuming 20 hours for part-time employment and 40 hours for full-time employment.

willing to accept all employment offers, an increased reservation wage may not affect the distribution of accepted wages.

2.6.4. Search Intensity

From the theory, we would expect two possible mechanisms for a treatment that reduces the marginal cost of search to decrease unemployment durations. First, it should generate more intense search over a given geographic area. Second, it could cause individuals to search over a broader area. I measure these search outcomes using administrative records from Jubilee Jobs on jobs applied to and interviewed for by each individual using data on the number, timing, and location of these potential employers. Given non-random selection of individuals out of job search (into employment) and likely poor recall by individuals of job application histories, I use only administrative data for these outcomes and focus on job search during the two weeks after treatment assignment. This also has the merit of examining a time when the in-kind treatment subsidy was not yet exhausted.

I use three different complementary measures of search intensity: an indicator of whether any search action was taken, the number of search actions in the two weeks after treatment assignment, and the number of days until the first search action is taken. For all of these, applying to a job and interviewing for a job are combined and generically referred to as search actions. Given the experimental random assignment, treatment effects can be ascertained from a simple comparison of means.

Table 2.9 provides these results, which offer some evidence that treatment leads to greater search intensity. Individuals assigned to treatment take 0.34 more actions, an increase in search intensity of 15 percent, though this difference is only significant at the 10 percent level.

While treatment has no effect in increasing the fraction of people who take at least one action, the treatment group starts searching sooner, taking their first search action about one half of a day sooner on average, conditional on taking at least one action, though this is statistically insignificant. Together, these results provide some evidence that treatment does in fact encourage greater search intensity by decreasing marginal search costs.

2.6.5. Search Scope

Another path toward increasing search intensity lies in increasing the geographic scope of search. To test this channel, I measure the impact of treatment on distance travelled to search for work and distance travelled to work. I measure search distance as the distance between the home address of the individual and the address of the first job for which the individual applies or interviews, using travel time by public transit as the notion of distance.[26] Distance to work is measured similarly using recorded home and work locations. For both variables, I handle those who have no observed search/work location in two ways. I first consider the impact of treatment with the entire sample, assigning a distance of zero to those who do not search/work. Second, I examine the impact of treatment in the sub-sample that searches/works.

Table 2.10 reports the impact of being assigned to treatment on search distance. The point estimate for the impact of being assigned treatment is large but statistically insignificant with the treatment group travelling 2 to 3 minutes further, an increase of 4 to 7 percent. The point estimates for the impact on commute distance to work are larger, with the treatment group travelling 6 to 20 percent further to work, depending on the means of measuring. However, the

[26] Travel time is computed as the shortest trip on public transit from the individual's address to the employer's address on Monday at 8:00 a.m. according to WMATA's on-line Trip Planner. For some individuals, such routes do not exist at 8:00 a.m. on Monday. For these individuals I measure the value as close to 8:00 a.m. as possible.

point estimates are statistically insignificant. So, these results provide some suggestive but not definitive evidence that transit subsidies remove a binding constraint on applying to and engaging in employment further from home.

2.7. Extensions

2.7.1. Spillovers and Treatment Externalities

I rely on individual-level randomization of public transit subsidies to identify the impact of these subsidies on job search and labor market outcomes. However, this approach can be biased if treatment of particular individuals spills over and affects other members of the sample. In particular, if the spillovers differ systematically between treatment and control individuals, a simple comparison of means gives a biased estimate of treatment effects.

For example, each individual in the sample coordinates their job search with a staff member of Jubilee Jobs, referred to as their job counselor, who is aware of their treatment status. If transit subsidies and job counselor effort are complementary inputs in making job placements, job counselors with a large caseload and scarce time may optimally shift effort toward those receiving treatment and away from those not receiving treatment. In this situation, we would expect individuals to be worse off when they share a job counselor with other people who are receiving the subsidies. If job counselors are in fact exerting less effort on the control group than they would in the absence of the experiment, the estimated effect of treatment could be overstated. Alternatively, if transit subsidies do reduce unemployment durations for the treatment group, the existence of the experiment could reduce the total workload of a job counselor mechanically by removing individuals from their caseload. This would benefit those individuals who are still looking for work, who would disproportionately be members of the control group.

In this case, individuals would benefit when others working with the same job counselor were assigned to treatment, and positive spillovers from freeing up staff time could cause treatment effects to be understated.

To estimate the direction and magnitude of treatment spillovers, I exploit the cohort timing of the study and the fact that each member of the sample was assigned to a single Jubilee Jobs job counselor. As described above, the most plausible treatment externalities occur when two individuals share the scarce time of the same job counselor. So, I measure the fraction of treated individuals that share the same job counselor in the same cohort and examine whether having a large fraction of these 'neighbors' treated has any impact on unemployment duration and other outcomes. Formally, for person i with job counselor j at time t, I consider the linear regression:

$$Y_{ijt} = \alpha + \beta T_i + \gamma \bar{T}_{(i)jt} + \delta Z_i + s_i + \epsilon_{ijt}$$

where $\bar{T}_{(i)jt}$ is the fraction of individuals from job counselor j in cohort t who were assigned to treatment, excluding individual i.[27] All other variables are as before. In this regression, the coefficient γ measures the direction and magnitude of treatment spillovers on individuals sharing a job counselor in the same cohort, and β estimates the individual-level treatment effect of receiving subsidized transit without contamination from spillovers. This specification is similar in spirit to Miguel and Kremer (2004) who examine treatment spillovers of a public health intervention using a geographic area as the neighborhood in which treatment externalities exist. If individuals are affected by whether their job counselor-cohort 'neighbors' are receiving transit subsidies, it would disproportionately affect members of the control group whose 'neighbors,'

[27] I drop the fifteen observations with no other individuals in their cohort-job counselor group.

due to the stratified randomization, are more likely to be from the treatment group. So, controlling for $\bar{T}_{(i)jt}$ should allow me to measure a treatment effect unbiased by spillovers.

The results from these regressions are displayed in Table 2.11. For various outcome variables, I investigate the impact of treatment with and without controlling for spillovers. Each pane of Table 2.11 corresponds to a different outcome variable and each pane contains the results of two regressions. In the first column, I replicate the previous results from equation (2.7), which assume no treatment spillovers. In the second column, I analyze equation (2.8) which controls for spillovers within a job counselor-cohort. For example, the top-left pane of Table 2.11 provides results where the outcome is an indicator of employment within 40 days. The first column of the first pane replicates previous results, indicating that under the assumption of no spillovers treatment increases the probability of employment by 9 percentage points. As the second column indicates, controlling for potential spillovers increases, if anything, the estimated individual-level effect of receiving transit subsidies. After controlling for these spillovers, I estimate the treatment effect to be a slightly larger, 11 percentage point increase in the probability of being employed. This occurs despite the fact that large treatment externalities appear to exist because the spillovers are *positive spillovers*. The coefficient of 0.26 indicates that for a given individual, if all other individuals working with the same job counselor receive subsidized transit, this individual's probability of employment will *increase* by twenty-six percentage points relative to a situation where no one working with that job counselor receives the subsidy. In other words, positive treatment externalities exist whereby having close ‘neighbors’ treated improves one's own employment outcomes.

These results indicate that the stronger assumption of no treatment spillovers has little impact on measured treatment effects. The previously measured individual-level effect of treatment on reducing unemployment durations is robust to controlling for treatment spillovers and if anything increases when doing so. This occurs because control group members are mechanically more likely to have treated `neighbors' and thus benefit more from the positive treatment externalities, leading the original estimation to have a slight bias toward zero. This evidence is consistent with the idea that when some individuals find employment more quickly due to the subsidies, this reduces the workload of job counselors, allowing them to focus more effort on everyone else, and those remaining are disproportionately from the control group.

The remaining panes of Table 2.11 demonstrate similar results for other important outcome variables. As before, the treatment effect on employment within 90 days is larger in magnitude. The treatment externalities for it show a pattern similar to what was found before, with the previous estimates providing a lower bound for the individual-level treatment effect. For wage rates, search intensity, and search scope similar results are obtained when controlling for treatment spillovers. If anything, measuring treatment effects while disregarding spillovers appears to produce a lower bound for the actual impact of treatment.

2.7.2. In-Kind Subsidies and Infra-marginality

As discussed above, the treatment group in this study was offered a packaged bundle of transportation subsidies including both a fee-reducing public transit card and $50 of in-kind subsidies to pay for bus and train fares. As commonly discussed with in-kind subsidies like housing vouchers and food stamps, the in-kind transit subsidy could be infra-marginal resulting in effects identical to an equal size cash grant. Standard microeconomic theory indicates that in-kind subsidies will 'stick' only if the recipient would optimally spend less than the subsidy amount in the absence of subsidies, but they will be infra-marginal, and thus equivalent to cash, if the individual would already choose to consume more public transit than the subsidy value even in the absence of the transfer. The model introduced in section 2 can be extended to model in-kind subsidies, and if this transfer is not fungible at the margin, the in-kind transit subsidies have an effect similar to reducing the marginal cost of search, resulting in greater search intensity and a higher reservation wage. But if the transfer is infra-marginal, the subsidy acts as income, resulting in lower search intensity and a higher reservation wage. Thus, if the potential infra-marginality of in-kind transfers is important, we should see different effects on search intensity and similar effects on wages for these two groups.[28]

Ideally, one would test for heterogeneous effects across individuals with different ex-ante desired levels of transit spending. Unfortunately, I have only transit spending data for the treatment group and no baseline information on transit spending. Lacking these direct measurements, I investigate how treatment effects differ for individuals who deposit larger balances of their own money onto their transit cards and those who deposit little or no balance.[29] In the treatment group, depositing additional balances onto the transit cards indicates that an individual has a demand for public transit exceeding the $50 balance provided to them. According to the theory, the in-kind portion of the treatment should be infra-marginal for this group. On the other hand, individuals who do not deposit their own money onto the card can be considered those for whom the in-kind constraint binds, so that receiving the transfer results in greater transit spending than would occur its absence. Of course, there could be other reasons for

[28] Theoretical results available upon request.
[29] I divide the sample at the median of $10.35.

deposit behavior, but deposit behavior should provide a decent proxy for testing whether infra-marginality plays an important role in the results.

One technical problem prevents simply comparing treatment effects between those who deposit more versus those who deposit less. In my data, the members of the control group do not receive cards and thus cannot be classified according to deposit behavior.[30] This prevents a clean comparison of treatment individuals in the depositing and non-depositing groups to the relevant control group. To generate a control group, I turn to non-experimental methods, using kernel-weighted propensity score matching to generate matched control groups for both the depositing and non-depositing groups. To use this method, I assume that selection occurs only on observable demographic and labor market history variables. While requiring stronger identifying assumptions, this method provides a reasonable way to see if treatment effects vary in a way consistent with infra-marginal in-kind subsidies.

Table 2.12 reports the results of measuring treatment effects first for the sub-sample of individuals who deposit less than the median amount onto their cards and then for those over the median. Behavior consistent with infra-marginal in-kind subsidies would predict that the sign on job search intensity should switch from positive in the left pane to negative in the right pane and that wage effects should be positive in both panes. The results do not provide support for these predictions. Treatment effects are not statistically different for the number of search actions or employment probabilities. The one statistically significant difference (10 percent level) between the two groups comes in the effect on wages. In the potentially infra-marginal group, treatment decreases wages by $1.05 and in the binding group they increase by $0.59. While these tests

[30] In any case, deposit behavior would be affected by treatment.

have low power, the results are inconsistent with what we would expect from the theory if the fungibility of in-kind transfers is a major issue. Search behavior of depositers and non-depositers respond to subsidies similarly. This result suggests that treatment either operates through the fee-reducing card itself or that in-kind subsidies are not fungible in this sample.

2.7.3. Cost-Benefit Calculations

As has been demonstrated, individuals receiving subsidized public transit through this experiment benefitted from reduced unemployed durations. In this section, I consider whether these benefits outweigh the direct costs of the subsidies as well as the value of staff time required to oversee the program. In particular, I will examine whether the additional earnings gained from reduced unemployment durations offsets program costs. Using the reduced form results, I am unable to give a complete cost-benefit analysis because they do not allow me to address, for instance, the welfare value of leisure time lost while working. In chapter 3, I will return to some of these issues more formally. However, this analysis provides a first estimate of the cost-effectiveness of public transit subsidies. Given a follow period of t days, the earnings gained from receiving treatment can be expressed as:

$$\Delta E = \left(t - \bar{d}_T\right)\bar{w}_T - \left(t - \bar{d}_C\right)\bar{w}_C = t(\bar{w}_T - \bar{w}_C) - \left(\bar{d}_T * \bar{w}_T - \bar{d}_C * \bar{w}_C\right) = -\left(\bar{d}_T - \bar{d}_C\right)\bar{w}$$

where $\bar{d}_i$ is average unemployment duration, $\bar{w}_i$ is average daily earnings, and i denotes the treatment and control groups. Given the results above, I assume that treatment has no effect on wages, leading to the second equality.

Censored unemployment durations provide the main challenge in measuring the total change in earnings due to receiving transit subsidies. While the results above indicate a strong impact of treatment on unemployment, censoring prevents me from expressing this in terms of a difference in uncensored means, $(\bar{d}_T - \bar{d}_C)$. To estimate this difference, I make the simplifying assumption that durations follow an exponential distribution where the hazard rate of leaving employment depends on treatment.[31] In other words:

$$F(d_i) = 1 - h_i e^{-h_i d_i}$$

where h_i is the hazard rate of leaving unemployment that differs between treatment and control groups $(h_i = \alpha + \beta T_i)$. Given these assumptions the effect of treatment on total earnings can be written:

$$\Delta E = -\left[\bar{d}_C \left(\frac{h_C}{h_T} - 1\right)\right] \bar{w}$$

Table 2.13 shows the parameters that I use to estimate ΔE. To estimate average daily earnings, I simply assume the value in the data for the sub-sample that has gained employment within 90 days. At an hourly rate \$9.82 and for an average 30.6 hours per week, this works out to \$42.91 per day. To estimate the uncensored mean of unemployment durations in the control group, $\overline{d_C}$, I assume an exponential distribution and estimate the extended mean of the censored data. As shown in Table 2.13 I find it to be 147.5 days.

The remaining parameter, $\frac{h_C}{h_T}$, is the inverse of the hazard ratio of finding employment in the treatment group versus the control group. I estimate the hazard ratio by maximum likelihood, assuming an exponential distribution and taking into account the strata of randomization:[32]

$$F_{d_i}(d) = 1 - (\alpha + \beta T_i + s_i) e^{-(\alpha + \beta T_i + s_i)}$$

[31] Note that this is consistent with the search model in Section 2.2, which implies a constant hazard rate and thus an exponential distribution of durations.
[32] Estimating the hazard ratio using a Cox proportional hazards assumption provides a nearly identical estimate of 1.24.

This provides an estimate of the hazard ratio of 1.25, indicating that individuals in the treatment group leave unemployment at a rate 25 percent faster than the control group.

Altogether, these estimates imply that being assigned to the transit subsidy groups causes the mean unemployment duration to decrease by an average of 30 days. At the given wage rate and number of hours, this implies a $1,280 increase in earnings as a result of being assigned to receive transit subsidies. This valuation of the benefits of treatment exceeds the cost by an order of magnitude. The direct cost of treatment was, at most, $60 due to purchasing two transit cards for $5 each and adding a $25 balance to each card. Even a very generous accounting of administrative costs still leaves costs far below benefits. For instance, assuming 20 hours of time per week devoted to administering treatment over the 8-month course of the program at a cost of $25 per hour leads to a total administrative cost of $16,000, or about $80 per treated individual. Even a very large hourly rate of $50 per hour leaves administrative costs at $160 per treated individual. While many effects are left out of this estimate, including the within-sample treatment externalities, general equilibrium effects of increased search intensity, and the value of leisure to the job-seekers, these back-of-the-envelope calculations provide a strong indication that public transit subsidies pass a cost-benefit test in this case.

2.8. Conclusion

This chapter reports results from a randomized experiment that provided a transportation subsidy package consisting of a fee-reducing public transit card and a $50 in-kind transit subsidy to low-wage job seekers in Washington, DC. I find that being assigned to receive transit subsidies decreases unemployment durations for a sample of individuals with limited access to jobs. In particular, there is a large effect in the short-run with the 30th percentile of unemployment durations decreasing by two weeks, from 49 to 35 days. Meanwhile, I find no impact of treatment on wage rates. The unemployment effect is quantitatively large, and I estimate that by ending unemployment durations sooner, subsidies increase earnings for the treatment group by approximately $1,280 per person at a program cost of $140, including administrative costs. While these simple calculations cannot provide a complete cost-benefit accounting, they indicate that transit subsidies can be a cost-effective way to improve labor market outcomes for low-wage, urban job-seekers.

I also find some suggestive evidence that increased search intensity is the mechanism by which these transit subsidies reduce unemployment durations. This evidence suggests that members of the treatment group apply to more jobs and jobs further from home. While I find that unemployment durations decrease in response to this uptick in search intensity, I find no evidence that job seekers respond to the broader pool of feasible jobs by becoming more selective regarding which wages to accept. These results are consistent with job search theory, which predicts that lower search costs will lead to greater search intensity and that individuals can choose to translate more search intensity into either lower unemployment durations or higher wages.

This chapter provides the first experimental evidence supporting the theory of spatial mismatch, the idea that living in a neighborhood far from available jobs can reduce employment opportunities due to high search costs. This evidence indicates that subsidizing public transit costs for job seekers from such areas can stimulate employment by enabling them to search for work more intensely. In particular, I find these results for a group of individuals who are largely disconnected from the labor market, with a mean baseline unemployment duration of almost 2 years, and who have many barriers to employment. These results are particularly relevant in the current economic environment, where long unemployment durations have become dramatically more common. At a cost of $60 per treated subject, this intervention represents a cost-effective means for getting such individuals back to work more quickly.

As with any small-scale experiment, context matters a great deal, but it is simultaneously beyond the reach of the experiment to delineate what parts of the context were necessary for the measured results. I find that transit subsidies have a large labor market impact in this sample, but these results will not translate to all situations. It is impossible to know for certain without further research, but Washington's high degree of residential segregation and its extensive (relative to other U.S. cities) public transit system are likely important components in the outcome of this experiment. Additionally, individuals in this study were already actively seeking low-wage employment and receiving job search support services. As such, the sample in this study form a selected sample of motivated job-seekers who were receiving other professional support. As a result, it is not clear how a representative sample of the urban unemployed would respond to transit subsidies. As a result, my findings are most pertinent to individuals who are already seeking employment with the assistance of public or private employment agencies. Nonetheless, I find that a modest transit subsidy can encourage greater job search and lead to significantly better labor market outcomes for minority individuals seeking low wage employment in an urban, U.S. labor market.

Figure 2.1. Employment Density (by Firm Location) in Zip Codes

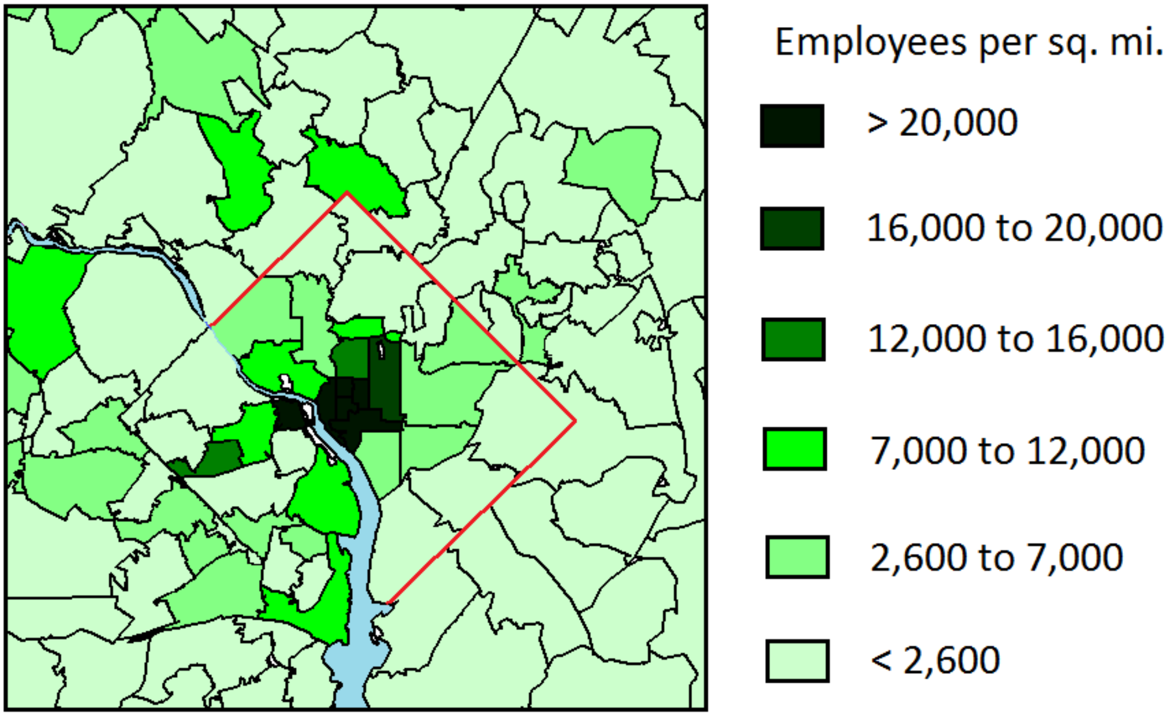

Source: US Census Zip Code Business Patterns, 2008

Figure 2.2. Population Density of African-Americans by Zip Code

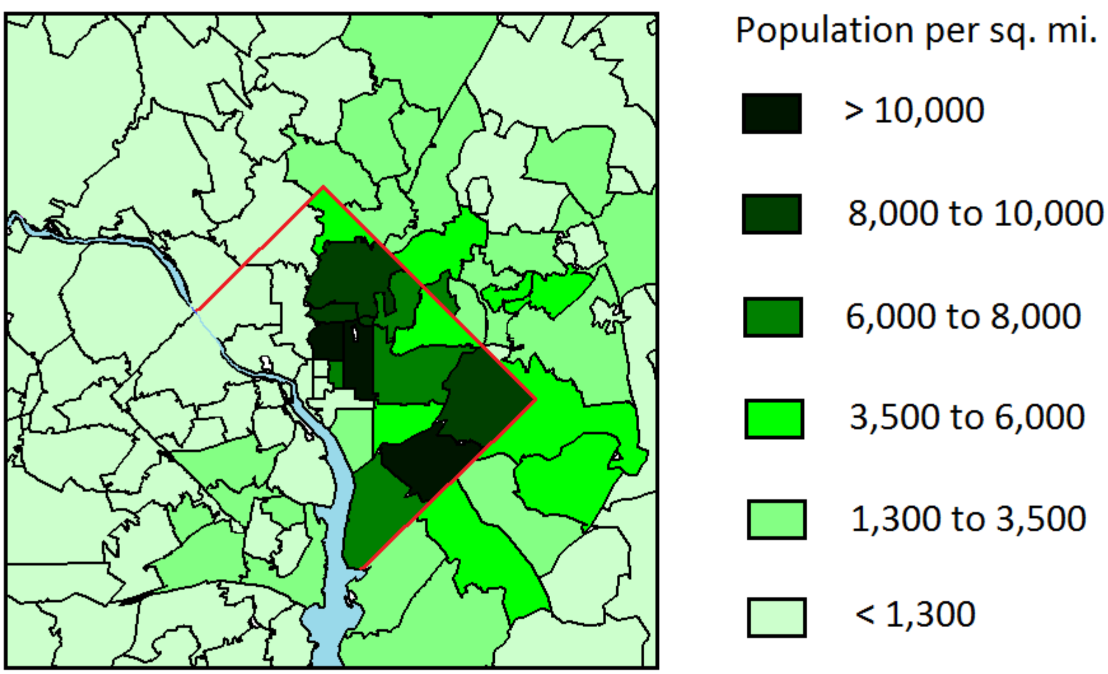

Source: US Census, 2000

Figure 2.3. Population Density of White-Americans by Zip Code

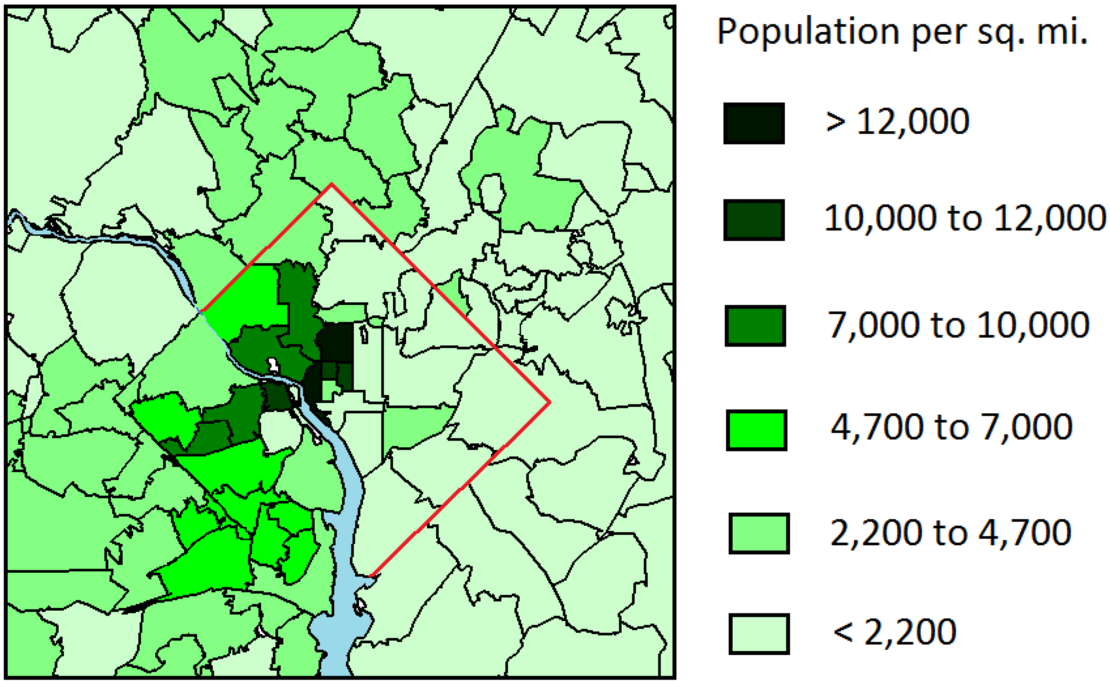

Source: US Census, 2000

Figure 2.4. Residence of Experimental Sample by Zip Code

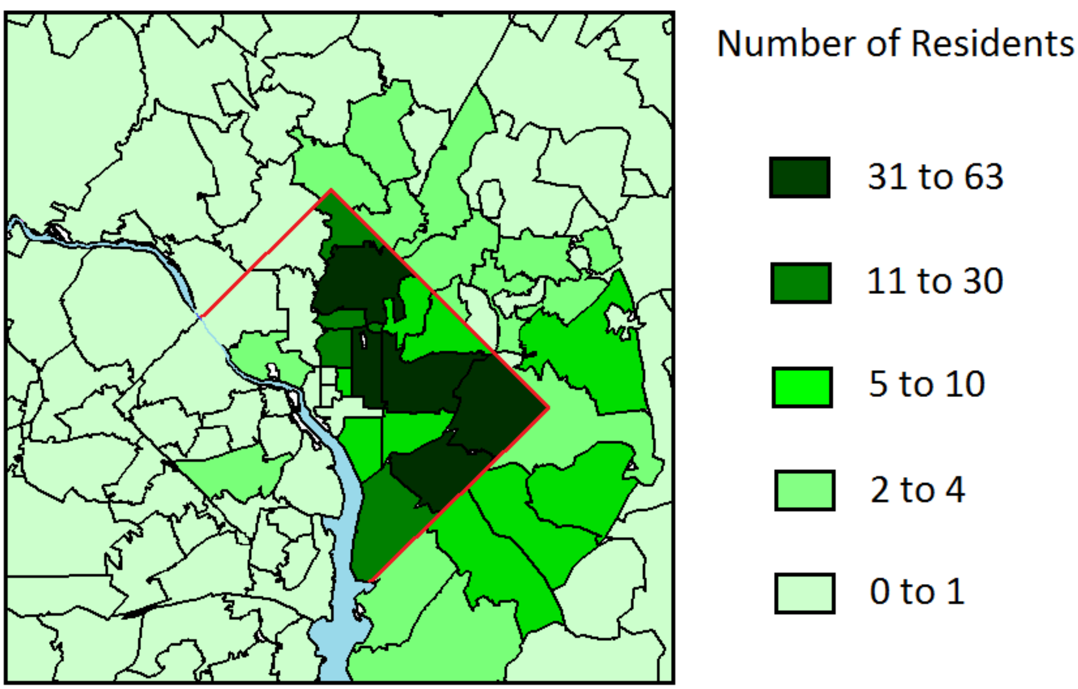

Source: Jubilee Jobs Administrative Data

Figure 2.5. Available Vacancies by Zip Code

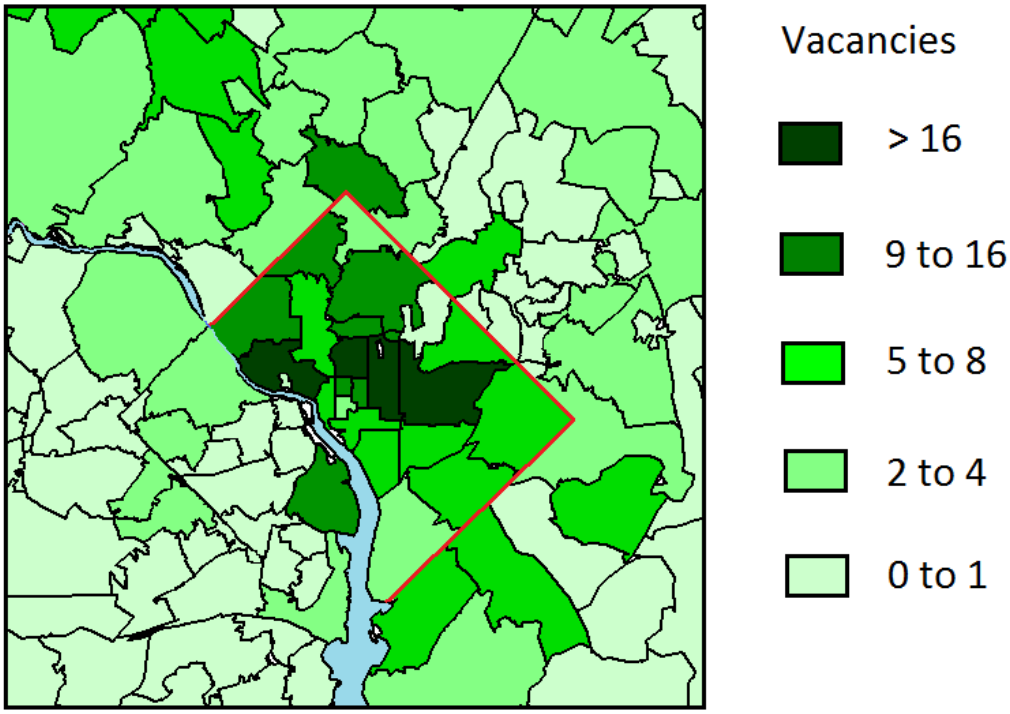

Source: Jubilee Jobs Administrative Data, April 2010 to April 2011

Figure 2.6. Card Usage by Day of the Week

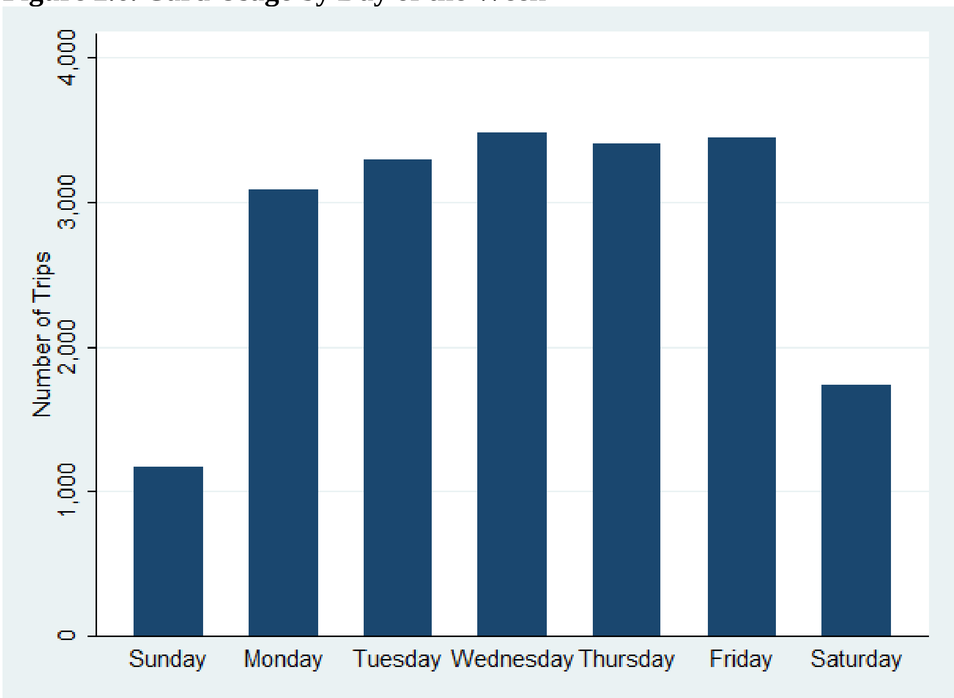

Figure 2.7. Card Usage by Hour of the Day

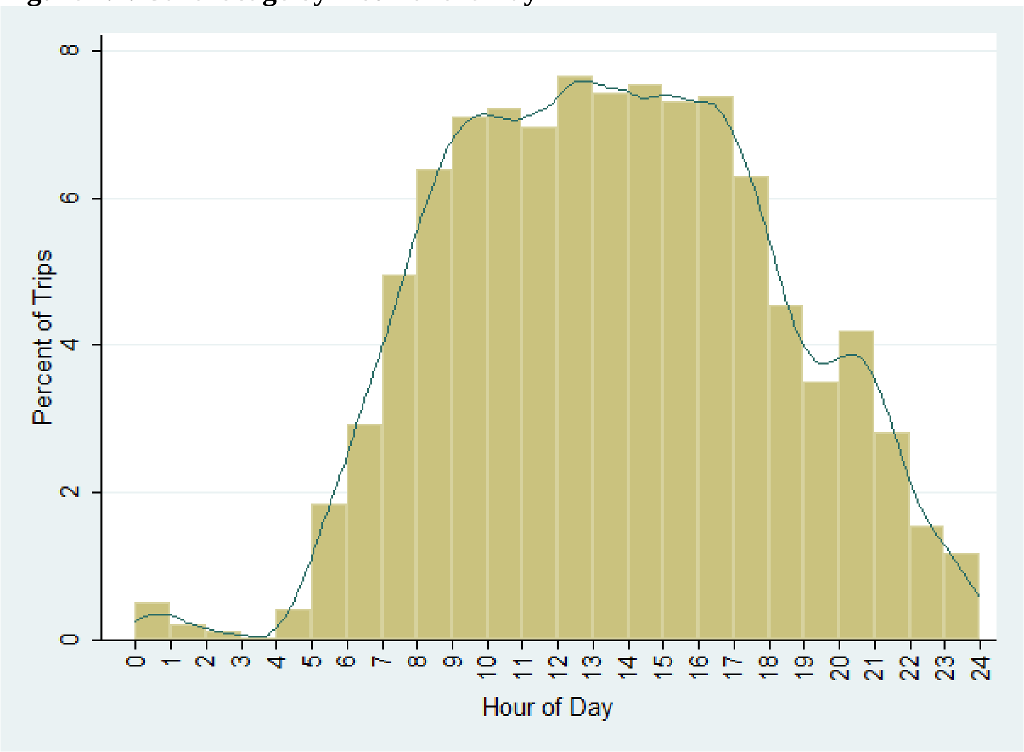

Figure 2.8. Gross and Net Usage of SmarTrip Cards

Figure 2.9. Effect of Treatment on Unemployment Duration

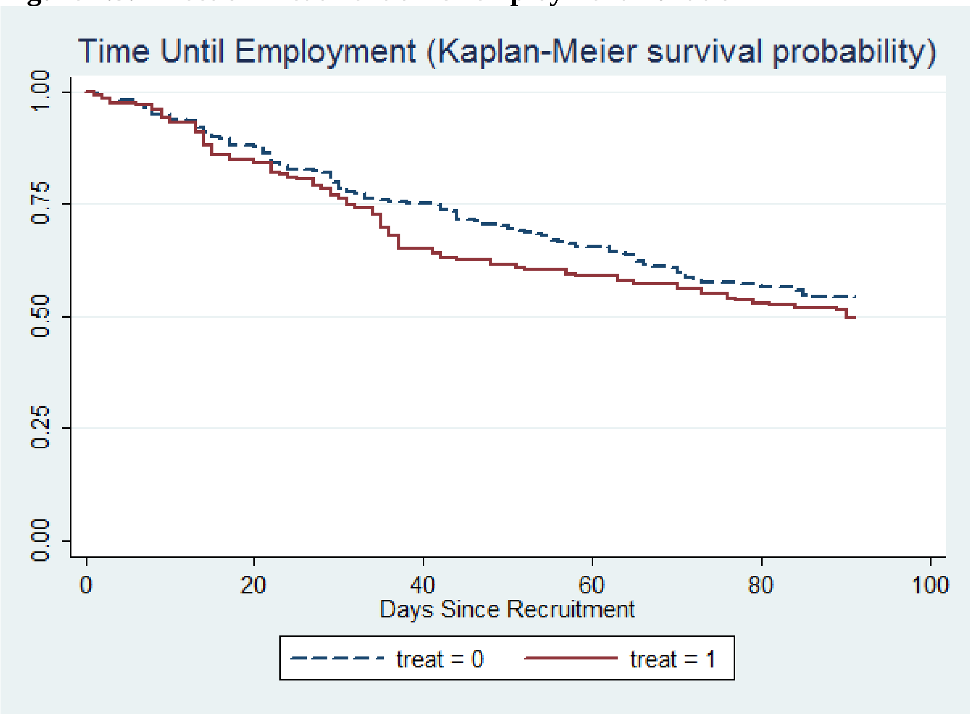

Figure 2.10. Effect of Treatment on Unemployment Duration: By Duration Length

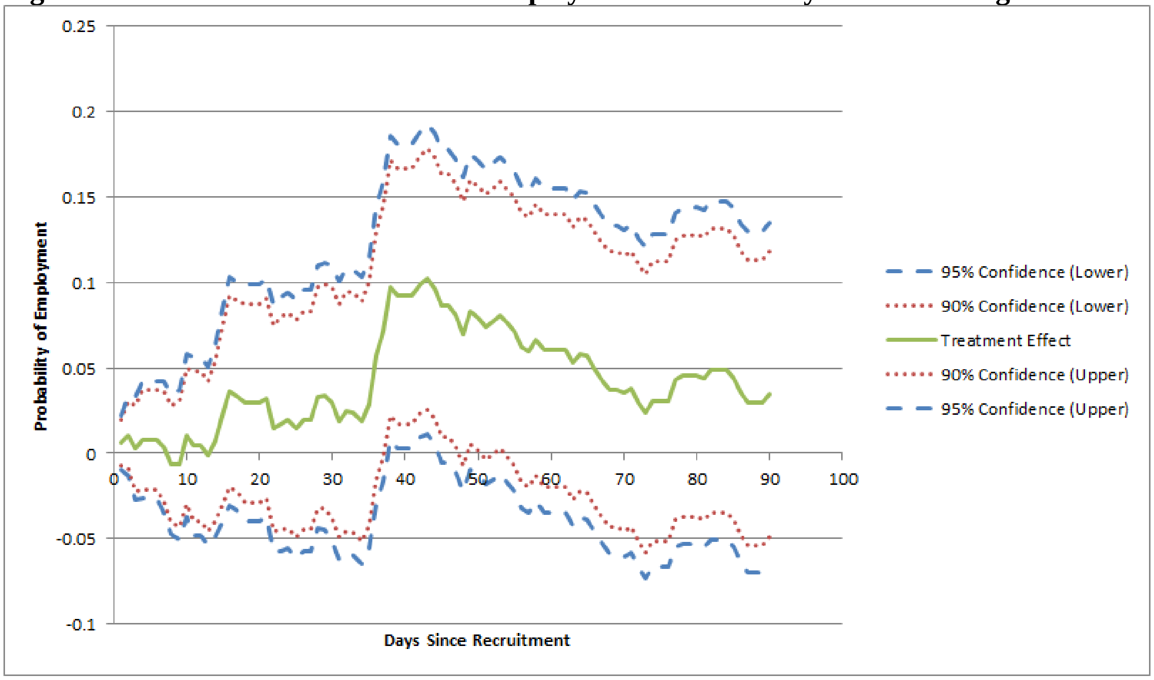

Figure 2.11. Effect of Treatment on Unemployment Duration: Kernel Densities

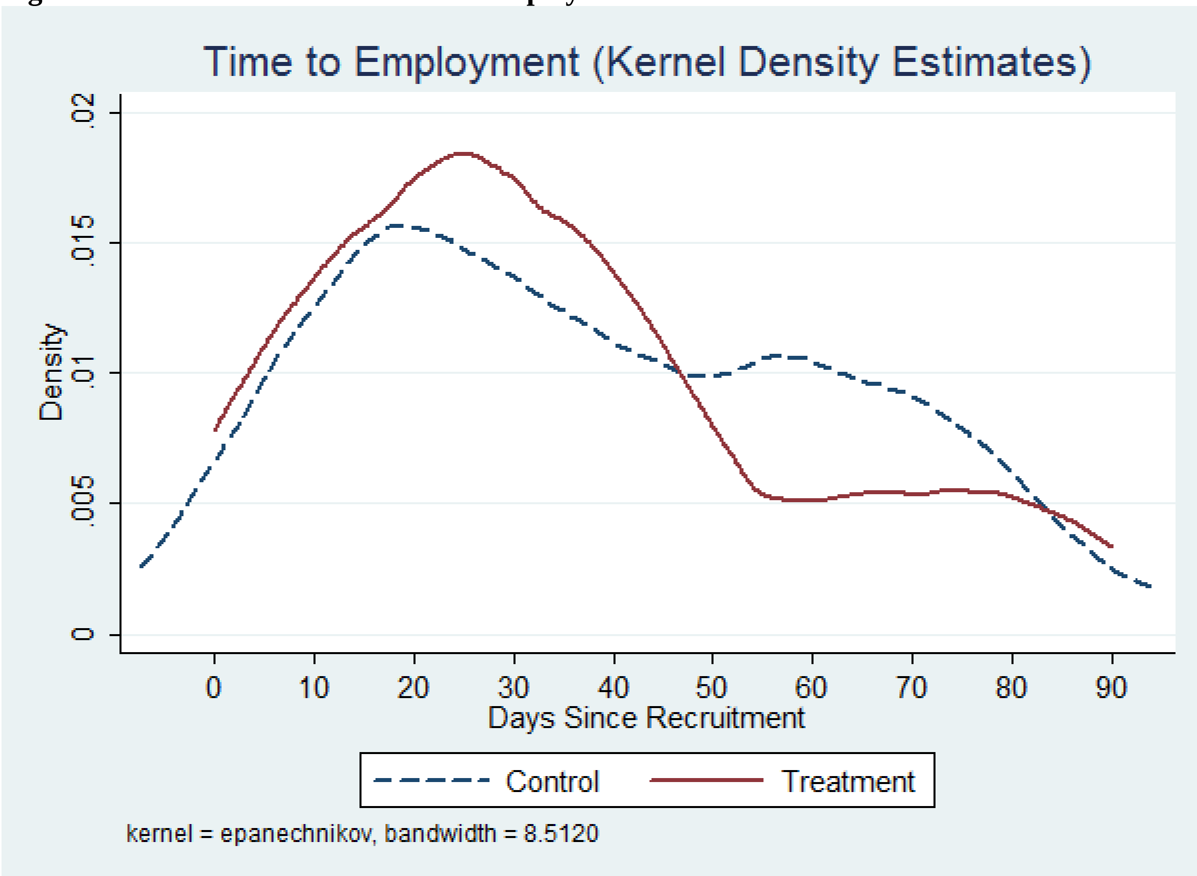

Table 2.1. Recruitment and Compliance

			Number	Proportion
Total			775	1.00
Recruited			468	0.60
	Treatment		208	0.44
		Cards Received		
		0	21	0.10
		1	80	0.38
		2	105	0.50
		3	2	0.01
	Control		260	0.56
		Cards Received		
		0	260	1.00
Not Recruited			307	0.40
	Not Eligible		17	0.06
	Refused		290	0.94

Table 2.2. Baseline Characteristics, Recruited and Non-Recruited

	Recruited		Not Recruited		Difference	p-value
	Mean	SD	Mean	SD		
Black	0.98	0.13	0.96	0.16	0.02	0.02**
Age	40.47	11.5	40.28	11.02	0.19	0.78
Male	0.58	0.49	0.61	0.45	-0.03	0.31
No HS Diploma	0.20	0.40	0.18	0.36	0.02	0.38
HS Diploma	0.56	0.5	0.54	0.45	0.02	0.57
Some College	0.19	0.39	0.16	0.34	0.03	0.24
College Graduate	0.05	0.21	0.05	0.20	0.00	0.76
Ex-Offender	0.50	0.50	0.46	0.45	0.04	0.16
Public Assistance	0.66	0.47	0.49	0.46	0.17	0.00***
Employed	0.11	0.31	--	--	--	--
Duration of Most Recent Job	2.24	3.24	--	--	--	--
(Median)	(1.03)					
Current Unemployment Duration	1.76	3.16	--	--	--	--
(Median)	(0.90)					
Pct of Last Five Years Employed	0.46	0.30	--	--	--	--
Most Recent Wage	10.04	5.44	--	--	--	--
Most Recent Wage Missing	0.10	0.30	--	--	--	--
Immigrant	0.10	0.30	0.10	0.27	0.00	0.96
Driver's License	0.45	0.50	0.45	0.45	0.00	0.99
Access to a Car	0.09	0.29	0.14	0.30	-0.05	0.01***
Maryland Residence	0.16	0.37	0.14	0.32	0.02	0.30
Virginia Residence	0.02	0.14	0.02	0.14	0.00	0.71

Source: Baseline administrative data. Difference is an unweighted comparison of means. 1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.

Table 2.3. Randomization Test for All Individuals

	Treatment		Control		Treatment-Control	
	Mean	SD	Mean	SD	Difference	p-value
Age	40.05	11.2	41.01	11.5	-0.94	0.4
Male	0.58	0.50	0.56	0.50	0.01	0.79
No HS Diploma	0.19	0.39	0.20	0.40	0.01	0.89
HS Diploma	0.55	0.50	0.57	0.50	-0.03	0.54
Some College	0.20	0.40	0.19	0.39	0.00	0.96
College Graduate	0.06	0.24	0.04	0.18	0.03	0.20
Ex-Offender	0.50	0.50	0.49	0.50	0.01	0.89
Public Assistance	0.66	0.47	0.67	0.47	-0.01	0.87
Employed	0.10	0.30	0.11	0.31	0.00	0.89
Duration of Most Recent Job	2.37	3.38	2.10	3.17	0.28	0.40
(Median)	(1.17)		(1.00)			
Current Unemployment Duration	1.99	3.50	1.81	3.12	0.15	0.64
(Median)	(0.90)		(0.97)			
Pct of Last Five Years Employed	0.49	0.29	0.47	0.29	0.01	0.76
Most Recent Wage	10.86	6.05	10.11	4.56	0.61	0.29
Most Recent Wage Missing	0.08	0.27	0.05	0.22	0.04	0.21
Immigrant	0.13	0.33	0.08	0.27	0.04	0.15
Driver's License	0.42	0.49	0.49	0.50	-0.07	0.12
Access to a Car	0.08	0.27	0.12	0.33	-0.04	0.11
Maryland Residence	0.15	0.36	0.17	0.38	-0.02	0.58
Virginia Residence	0.01	0.09	0.03	0.17	-0.02	0.08*

Source: Baseline administrative data. Difference is from a regression of the baseline characteristic on a treatment dummy and dummies for the strata of randomization. Means are computed by weighting individuals by the inverse probability of treatment for their cohort. 1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.

Table 2.4. Attrition

	Whole Sample		Treatment		Control	
Cards	Number	Proportion	Number	Proportion	Number	Proportion
Non-Attriters	427	0.91	190	0.91	237	0.91
-Administrative Data	289	0.62	133	0.64	156	0.60
-Phone Survey	138	0.29	57	0.27	81	0.31
Attriters	41	0.09	18	0.09	23	0.09
Total	468	1.00	208	1.00	260	1.00

Table 2.5. Randomization Test for Analysis Sample Only

	Mean	SD	Mean	SD	Difference	p-value
Age	40.00	11.12	41.12	11.39	-1.09	0.34
Male	0.57	0.50	0.57	0.50	0.00	0.95
No HS Diploma	0.19	0.39	0.19	0.40	0.01	0.74
HS Diploma	0.56	0.50	0.58	0.50	-0.02	0.68
Some College	0.19	0.39	0.19	0.39	-0.01	0.75
College Graduate	0.06	0.24	0.04	0.19	0.02	0.36
Ex-Offender	0.50	0.50	0.48	0.50	0.02	0.72
Public Assistance	0.66	0.47	0.68	0.47	-0.01	0.83
Employed	0.10	0.30	0.11	0.31	-0.01	0.65
Duration of Most Recent Job	2.23	3.24	2.18	3.25	0.05	0.88
(Median)	(1.08)		(1.00)			
Current Unemployment Duration	2.03	3.65	1.76	3.10	0.24	0.49
(Median)	(0.81)		(0.93)			
Pct of Last Five Years Employed	0.49	0.30	0.48	0.29	0.00	0.99
Most Recent Wage	11.17	6.03	10.34	4.43	0.68	0.25
Most Recent Wage Missing	0.06	0.24	0.03	0.18	0.04	0.20
Immigrant	0.13	0.34	0.08	0.27	0.05	0.11
Driver's License	0.43	0.50	0.50	0.50	-0.07	0.17
Access to a Car	0.08	0.27	0.12	0.33	-0.04	0.13
Maryland Residence	0.14	0.35	0.16	0.37	-0.03	0.42
Virginia Residence	0.01	0.10	0.03	0.16	-0.02	0.20

Source: Baseline administrative data. Difference is from a regression of the baseline characteristic on a treatment dummy and dummies for the strata of randomization. Means are computed by weighting individuals by the inverse probability of treatment for their cohort. 1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.

Table 2.6. CPS Comparison

Sample	Experimental Data All	2010 CPS All Working Age	2010 CPS Unemployed	2010 CPS Unemployed, Black Metro Area
Less Than High School	20	13	19	17
HS Diploma	56	30	39	43
Some College	19	29	27	28
Bachelor's or More	5	28	15	11
18-25	11	18	25	31
25-30	11	11	13	16
30-35	13	10	11	11
35-40	9	11	11	10
40-45	14	11	11	10
45-50	18	12	11	9
50-55	13	11	9	7
55-60	7	10	7	4
60+	1	8	4	3
Percent Female	42	51	39	41
N	468	128845	9473	1511

All data are listed in percentages. Means in the experimental sample are unweighted.

Table 2.7. Employment Duration Quantile Regressions

Quantile	Coefficient on Constant (Control Group Quantile)	Coefficient on Treatment
5	8	1
		(3)
10	15	-1
		(3)
15	22	-5
		(5)
20	29	-2
		(6)
25	38	-7
		(7)
30	49	-14**
		(7)
35	58	-21**
		(10)
40	70	-13
		(15)
45	85	-12
		(15)

1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.
Standard errors are in parentheses and bootstrapped with 1000 repetitions.

Table2.8. ITT Estimates for Employment Outcomes

	Obs	Control Mean	Treatment Mean	Difference	Difference (with baseline control variables)
Found Job within 40 Days	427	0.26	0.35	0.09**	0.09**
				(0.05)	(0.05)
Found Job within 90 Days	427	0.45	0.50	0.05	0.04
				(0.05)	(0.05)
Unemployment Duration	205	37	34.7	-2.5	-2.5
(Conditional on Working)				(4.0)	(3.9)
Weekly Earnings	424	134	154	21	16
				(18)	(18)
Wage Rate	202	9.59	10.03	0.26	0.07
				(0.44)	(0.41)

Standard errors are in parentheses in fourth and fifth columns. Treatment-control differences come from a regression of the outcome on treatment. Both columns include a vector of dummies for each randomization strata. The final column adds demographic controls: age, age squared, gender, education dummies, and labor market history variables. 1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.

Table 2.9. ITT Estimates for Search Outcomes

	Obs	Control Mean	Treatment Mean	Difference	Difference (with baseline control variables)
Any Search Action Taken	465	0.90	0.88	-0.01	0.00
				(0.03)	(0.03)
Number of Actions	465	2.33	2.72	0.34*	0.35*
				(0.20)	(0.20)
Time to First Action	418	13.31	12.38	-0.58	-0.65
				(0.75)	(0.80)

Standard errors are in parentheses in fourth and fifth columns. Treatment-control differences come from a regression of the outcome on treatment. Both columns include a vector of dummies for each randomization strata. The final column adds demographic controls: age, age squared, gender, education dummies, and labor market history variables. 1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.

Table 2.10. ITT Estimates for Location Outcomes

	Obs	Control Mean	Treatment Mean	Difference	Difference (with baseline control variables)
Distance to Search Location (Zero for Non-Searchers)	465	34.2	36.3	2.4 (2.6)	2.5 (2.7)
Distance to Search Location (Conditional on Searching)	396	40.9	42.3	1.9 (2.7)	1.7 (2.7)
Distance to Work (Zero for Non-Working)	422	19.2	22.5	3.9 (3.0)	3.3 (3.0)
Distance to Work (Conditional on Working)	200	42.5	44.8	2.7 (4.5)	3.8 (4.8)

Standard errors are in parentheses in fourth and fifth columns. Treatment-control differences come from a regression of the outcome on treatment. Both columns include a vector of dummies for each randomization strata. The final column adds demographic controls: age, age squared, gender, education dummies, and labor market history variables. 1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.

Table 2.11. Treatment Spillovers

	Employed Within 40 Days		Employed Within 90 Days	
	(1)	(2)	(1)	(2)
T_i	0.09**	0.11**	0.04	0.06
	(0.05)	(0.05)	(0.05)	(0.05)
$\bar{T}_{(i)jt}$		0.26***		0.17
		(0.09)		(0.11)
N	427	414	427	414
R^2	0.18	0.20	0.12	0.11
	Wage Rate		Number of Search Actions	
	(1)	(2)	(1)	(2)
T_i	0.07	0.28	0.35*	0.41*
	(0.41)	(0.45)	(0.20)	(0.21)
$\bar{T}_{(i)jt}$		1.37		-0.12
		(0.94)		(0.50)
N	202	193	465	450
R^2	0.50	0.51	0.15	0.16
	Distance to Search Location		Distance to Work	
	(1)	(2)	(1)	(2)
T_i	2.5	3	3.3	3.4
	(2.7)	(2.8)	(3.0)	(3.0)
$\bar{T}_{(i)jt}$		0.7		5.4
		(5.8)		(5.9)
N	464	449	422	410
R^2	0.14	0.16	0.15	0.15

Standard errors clustered at the job counselor-cohort level are in parentheses. All regressions control for a vector of dummies for each randomization strata and demographic controls: age, age squared, gender, education dummies, and labor market history variables. 1, 5, and 10 percent statistical significance are denoted by ***, **, and * respectively.

Table 2.12. Inframarginality - Propensity Score Matching Estimates

	Non-Depositors			Depositors		
	Treatment	Matched Control Group	Difference	Treatment	Matched Control Group	Difference
Found Job within 40 Days	0.31	0.24	0.08 (0.06)	0.40	0.26	0.14** (0.06)
Found Job within 90 Days	0.49	0.43	0.06 (0.07)	0.53	0.47	0.06 (0.07)
Job Search Actions	2.60	2.20	0.40 (0.30)	2.84	2.55	0.29 (0.27)
Wage Rate	10.41	9.83	0.59 (0.69)	9.35	10.4	-1.05** (0.51)

Matched control group mean generated by kernel-weighted propensity score matching. Kernel is Epanechnikov and the first-stage uses randomization strata, education dummies, gender, and labor market history variables to predict treatment. Standard errors are in parentheses. 1, 5, and 10 percent statistical significance are denoted by ***, **, and *.

Table 2.13. Cost-Benefit Calculations

Parameter	Estimate
Mean Wages	9.82
Average Hours	30.6
Average Daily Earnings	42.91
Control Group Mean Unemployment Duration	147.5
Hazard Ratio	1.25

Cost-Benefit Analysis	
Benefits	
--Earnings Increase	$1,280
Costs	
--Subsidy	$60
--Administrative	$80
Net Benefits	$1,140

CHAPTER 3

VALIDATING A STRUCTURAL JOB SEARCH MODEL AND ESTIMATING PROGRAM WELFARE EFFECTS WITH EXPERIMENTAL DATA

3.1. Introduction

Randomized field experiments and quasi-experiments have become much more common in economics in recent years, largely due to an appeal to clean identification and internal validity (Angrist and Pischke, 2010). However, increased use has also led to greater understanding of their limitations, including considerations of external validity (Deaton, 2009) and sometimes an atheoretical interest in the treatment effect of interventions rather than how those treatment effects relate to economic theory (Card, et. al., 2011). In response, designers of field experiments have increasingly taken theory into account when constructing experiments, using designs that allow researchers to distinguish between competing theoretical mechanisms (e.g. Karlan and Zinman, 2009). Additionally, in a small but prominent and growing number of cases, researchers have directly applied theory to experimental data ex-post. Most famously, Todd and Wolpin (2006) estimate a structural model of human capital formation using the PROGRESA experiment data. In particular, they use the PROGRESA control group to estimate their model, validate the model with experimental variation in the conditional cash transfer program, and then use the validated model to make predictions about policy interventions that were not included as treatments in the original PROGRESA program. In the context of job search models, a similar approach has been taken by Lise, Seitz, and Smith (2005) using the Canadian Self-Sufficiency Project experiment. More recently, Attanasio, et. al. (2011) also estimate a structural model on the PROGRESA data but instead use instead experimental variation directly in the estimation and identification of structural parameters.

In this chapter, I also combine structural estimation with experimental data, using data from the transportation subsidy experiment described in chapter 2 in which low-wage active job seekers were randomly selected to receive public transit subsidies during their job search period. In particular, I take a structural model of job search to the experimental data. Because the sample contains job-seekers looking for low wage jobs with a large modal wage equal to the minimum wage, I consider a model with a binding minimum wage. In particular, I estimate a simplified version of the Flinn (2006) job search model where wages are determined by Nash bargaining in the presence of a minimum wage.

I consider two exercises. First, I simply estimate the structural model on the whole sample and then on the treatment and control groups separately. This process provides a first pass at seeing how the experiment affects the parameters of the model without restricting the impact of treatment. Second, I conduct an exercise similar to that of Todd and Wolpin (2006). In the experiment, the treatment group was observed to increase its search effort, as measured by number of applications and interviews, by 15 percent relative to the control group. Taking the structural parameters estimated for the control group, I then simulate treatment as a 15 percent increase in the arrival rate of job offers, which is itself a structural parameter, and observe the impact that this treatment would have on unemployment durations and wage rates, according to the estimated model. This simulation generates predictions from the model regarding the size of treatment effects to be expected from the experiment using only control group data on wages and unemployment durations in addition to search intensity data from both the treatment and control groups. These independent predictions of treatment effects can then be compared to those actually observed in the experiment.

The estimated model precisely predicts not only the control group's distribution of wages but also the impact of the experimental treatment on wages. In particular, it fairs well in this environment where low wages make the minimum wage bind. For unemployment durations, I find that the model underestimates the impact of treatment in shortening unemployment durations somewhat, though the predictions are relatively close. The Nash bargaining model predicts treatment effects at two-thirds of the observed effect. While not perfect, I show in appendix B that this proves a major improvement over a simple partial equilibrium model with an exogenous, continuous distribution of posted wages. Together, these results largely validate the use of a search model with Nash bargaining and binding minimum wages to explain job search behavior in an urban, low-wage labor market.

Having validated the bargaining model, I use it to measure the welfare effects of treatment. Welfare, of course, cannot be directly observed and thus overall welfare effects cannot be measured in a typical experiment. This is particularly true in a job search setting with labor market frictions and minimum wages because the value of job seekers' non-market time proves difficult to impute. However, a validated structural model of job search provides simple estimates for the welfare of participants in the experiment explicitly taking into account the value of non-market time. Using the Nash bargaining model with binding minimum wages, I estimate that welfare gains from receiving the transportation subsidies far outweigh their costs, generating an excess value of over $500 per person.

3.2. Background on the Experiment

The data used in this paper come from an experiment described in detail in chapter 2. To summarize, clients of a non-profit employment services organization were invited to participate in a field experiment in which just under half of participants were randomly selected to receive two fee-reducing public transit cards that both reduced the price of bus and train trips and provided $25 worth of fares per card. Selected participants were provided with these subsidies while they actively looked for employment in the low-wage labor market of Washington, DC. For the study, data were collected on unemployment durations and wage rates for those who found employment within a 90-day follow-up period. Importantly for the estimation below, durations greater than 90 days and the corresponding wages are censored. Data were also collected on search behavior of applicants, particularly the number of interviews and applications undertaken by each individual in their first 2 weeks of search. These three pieces of data (unemployment duration, wage rate, and search intensity) will allow me to both estimate the structural models of job search and simulate the impact of treatment.

For purposes of comparison with the results of the theoretical models, I summarize the results of the experiment in Table 3.1. As can be seen, receiving the subsidized transportation had a large impact on the likelihood of being employed. In particular, the treatment group is 9 percentage points more likely to be employed within 40 days and 5 percentage points more likely to be employed within 90 days. Though censoring of unemployment durations prevents measuring the total drop in average unemployment duration, these figures indicate that receiving the subsidies decreased average durations considerably. With regard to wages, the effects are much smaller (and statistically insignificant), indicating that members of the treatment group did have higher wages, but only by $0.44 per hour. These estimates for the impact of treatment on unemployment durations and wage rates will provide the benchmark against which to validate the theoretical models.

A final variable will be useful in the simulation exercises below. As the final row of Table 3.1 shows, increased search intensity appears to be the mechanism by which the treatment group obtained shorter unemployment durations. The average number of job search actions (i.e. interviews or applications) in the first two weeks of searching increased from 2.33 in the control group to 2.72 in the treatment group, an increase of about 15 percent. This fact will become important in the present study when simulating the impact of treatment in the structural models. Throughout, I will consider this change in search intensity to be the change in structural parameters generated by the experiment.

3.3. Search with Nash Bargaining and a Binding Minimum Wage

3.3.1. A Canonical Search Model

In a standard search model (e.g. McCall, 1970), an unemployed job seeker looks for employment in continuous time. Job offers are completely characterized by a wage rate, w, and workers receive offers from firms at a Poisson rate λ. If the worker and firm agree on a match, the worker is employed at wage w forever. Given that she discounts the future at a constant rate ρ, the value of being employed at wage w is:

$$\rho V(w) = w \qquad (3.1)$$

The sole decision to be made by the job seeker is whether to accept the job and work at wage w forever or to reject the job and continue looking for work. To make this decision, she must compare the value of being employed to the value of unemployment. Unemployment has a flow value b which can be positive or negative and comes with the option value of potential future employment. In the model I will consider, wages are determined by an asymmetric Nash bargain (described below), but from the worker's perspective, this is equivalent to making an independent draw from the distribution of bargained wages $F(w)$ any time she comes in contact with a firm. Thus, the standard setup is that the value of being unemployed is:

$$\rho U = b + \lambda \int \max\big(U, V(w)\big)\, dF(w) \qquad (3.2)$$

In the absence of a binding minimum wage, comparison of (3.1) and (3.2) leads directly to the standard reservation wage rule: the job seeker sets a reservation wage w^*, accepting any wage above the reservation wage and rejecting all wages below w^*. In particular, the reservation wage will be the wage at which she is indifferent between working and looking for work, i.e. where $V(w^*) = U$. Substituting this and (3.1) into (3.2) leads to the standard reservation wage equation:

$$w^* = b + \lambda \int_{w^*}^{\infty} \frac{w - w^*}{\rho}\, dF(w) \qquad (3.3)$$

3.3.2. A Search Model with Binding Minimum Wage and Nash Bargaining

I follow Flinn (2006) and complicate this standard job search model in two ways. First, as in Pissarides (2000), wages result from an asymmetric Nash bargain between workers and firms rather than simply a draw from an exogenous distribution of wages. Second, a binding minimum wage exists. There are many approaches to incorporating a minimum wage into equilibrium search and matching models. For example, Eckstein and Wolpin (1990) and van den Berg and Ridder (1998) estimate equilibrium search models based on wage posting that incorporate a legal minimum wage. In a more recent example, Eckstein, Ge, and Petrongolo (2011) estimate a wage posting model with on-the-job search and lifecycle considerations. I choose to follow the Nash bargaining approach because it matches important features in the data. In particular, interacting the mandated minimum wage with a wages set by bargaining generates a realistic distribution of wages with a mass point at the minimum wage.

Thus, in the model wage offers arise through a process of Nash bargaining. Whenever a firm and worker meet (at Poisson rate λ), nature draws the productivity of the match, θ, from an exogenous distribution $G(\theta)$. The firm and the worker then they engage in an asymmetric Nash bargain over the wage rate, w. If they stay in the match, the worker receives w and the firm receives $\theta - w$. The firm's outside option is normalized to zero and the outside option of the worker is unemployment, valued at U. As with any Nash bargain, both sides will wish to participate as long as there is positive surplus relative to their outside options. In the absence of a minimum wage, the match endures if $\theta \geq \theta^* = \rho U$ and otherwise the worker returns to searching. If they decide to keep the match, they split the surplus by choosing a wage to solve:

$$\max \; [V(w) - U]^{\alpha} \left[\frac{\theta - w}{\rho}\right]^{1-\alpha} \qquad (3.4)$$

It can be shown that this results in a wage of:

$$w = \alpha\theta + (1 - \alpha)\theta^* \qquad (3.5)$$

With no minimum wages, productivities map linearly into wages so that the distribution of wages, $F(w)$, is a simple transformation of the distribution of productivity, $G(\theta)$. Note from (3.5) that the reservation productivity and the reservation wage are identical, $\theta^* = w^*$. Then, (3.3) can be used to find the optimal reservation wage.

However, the situation changes if there is a legally enforced minimum wage, m, such that the Nash bargain now consists of solving (3.4) subject to the wage lying above the minimum wage. I will consider the case where the minimum wage binds, i.e. $m > \theta^*$. If it does not bind, the model reduces to the simpler one considered in section 3.3.1. First, note that if the minimum wage binds the worker's optimal strategy will be to accept all offers because the binding minimum wage forces all offers to exceed the reservation wage. On the other side, the firm will refuse a match if the productivity of the worker is below the minimum wage, i.e. $\theta < m$. Only matches with productivity at or above the minimum wage will stick, and this will eliminate some previously profitable matches. Some profitable matches will also be affected by the minimum wage. As Flinn (2006) argues, consider the productivity value that would lead to a wage of m in the absence of regulation:

$$\hat{\theta} = \frac{m - (1-\alpha)\theta^*}{\alpha}$$

For productivity values at $\hat{\theta}$ and above, the minimum wage constraint does not bind and wages are determined by (3.5). However, in the binding minimum wage case, there is a range of productivities $m \leq \theta \leq \hat{\theta}$ for which the constraint binds but there is still positive surplus form the match. In this case, the firm will find it profitable to keep the match but raise the wage above the unconstrained level to m. In sum, workers with productivity above $\hat{\theta}$ will receive wages according to (3.5); workers with productivity $m \leq \theta \leq \hat{\theta}$ will receive m; and workers with productivity lower than m will not make a match and continue searching. So, the minimum wage will eliminate some matches, generating unemployment, but will also give workers bargaining power, raising wages for some workers.

3.3.3. Estimation

I will estimate the model by maximum likelihood using duration and wage data, relying on the identification results of Flinn and Heckman (1982). Throughout I will assume that the minimum wage binds, i.e. $\theta^* \leq m$, both because the interesting features of the model require this and because it will be born out in the experimental data. With a binding minimum wage, the hazard probability of leaving unemployment is constant but depends on the minimum wage rather than a reservation wage. Matches are made when a worker and firm meet and then draw a productivity greater than the minimum wage, leading to a match rate of $\lambda[1 - G(m)]$. As a result, the distribution of unemployment durations is exponential:

$$\Pr[t \leq t_i] = 1 - e^{-\lambda[1-G(m)]t_i}$$

which leads to a density of:

$$f_t(t_i) = \lambda[1 - G(m)]e^{-\lambda[1-G(m)]t_i}$$

This bears great similarity to a simple, off-the-shelf search model. The main change resulting from the binding minimum wage regards the equilibrium distribution of wages. As noted in the theory section, firms will offer the minimum wage to workers with a range of productivity levels, $m \leq \theta \leq \hat{\theta}$. As a result, the distribution of wages has a mass point at the minimum wage. In particular,

$$f(m) = \Pr[w_i = m] = \Pr\left[m \leq \theta \leq \hat{\theta} \middle| \theta \geq m\right] = \frac{G(\hat{\theta}) - G(m)}{1 - G(m)}$$

At productivity values above $\hat{\theta}$, wages are determined by (3.5), making the distribution of wages over this range continuous, provided that $G(\cdot)$ is continuous:

$$f(w) = \frac{g\left(\frac{w-(1-\alpha)\theta^*}{\alpha}\right)}{\alpha[1-G(m)]} \qquad if \;\; w > m$$

Together, these imply that the contribution to the likelihood function of an observation with wage and duration data is:

$$L_i = \left(\lambda[1-G(m)]e^{-\lambda[1-G(m)]t_i}\right) * \left(\frac{G(\hat{\theta})-G(m)}{1-G(m)}\right)^{I[w_i=m]} * \left(\frac{g\left(\frac{w-(1-\alpha)\theta^*}{\alpha}\right)}{\alpha[1-G(m)]}\right)^{I[w_i>m]}$$

$$= \left(\lambda e^{-\lambda[1-G(m)]t_i}\right) * \left(G\left(\frac{m-(1-\alpha)\theta^*}{\alpha}\right) - G(m)\right)^{I[w_i=m]}$$

$$* \left(\frac{1}{\alpha} g\left(\frac{w-(1-\alpha)\theta^*}{\alpha}\right)\right)^{I[w_i>m]} \qquad (3.6)$$

Thus, (3.6) gives the likelihood contribution of an uncensored data point. As noted above, duration and wage data are not available for all individuals because of a short 90-day follow-up period. Thus, the only information on these censored observations is that they have durations greater than 90 days. Thus, they contribute to the likelihood function this information:

$$L_i = \Pr[t \geq 90] = e^{-\lambda[1-F(m)]90} \qquad (3.7)$$

Finally, (3.6) and (3.7) are combined to construct the likelihood function.

3.3.4. Identification

As first noted by Flinn and Heckman (1982), such a model cannot be estimated without assuming a parametric distribution of $G(\cdot)$, and I will use a log-normal distribution with parameters μ and σ so that:

$$G(\theta) = \Phi\left(\frac{\ln(\theta)-\mu}{\sigma}\right)$$

where $\Phi(\cdot)$ is the standard normal distribution. To identify the model, I will also determine certain parameters ex ante. First, the minimum wage will be set at the level for District of Columbia during the time period, \$8.25 per hour. Given that a small number[33] of wage observations are below the legal minimum wage, I must also assume some measurement error in wages. I take the simplest assumption, assuming that wages measured below the minimum reflect actual wages that are at the minimum. Second, I will exogenously fix the bargaining parameter, α, to be 0.4. As Flinn (2006) notes, the bargaining parameter cannot be identified with only worker-level data, and he uses firm-level balance sheet data to estimate α with the other parameters. For simplicity, I will fix α at 0.4, which is consistent with Flinn's estimates. Having fixed α and m, I estimate the other parameters $(\lambda, \mu, \sigma, \theta^*)$ by maximum likelihood using (3.6) and (3.7) for uncensored and censored observations, respectively. Given the estimated parameters $\hat{\theta}$ and the hazard probability of leaving unemployment can be calculated easily. The value of being unemployed can still be represented by (3.2). Given the distribution of accepted wages and noting that $\theta^* = \rho U$, it can be shown that:

$$\theta^* = b + \frac{\lambda}{\rho}\left\{(m - \theta^*)\left(G(\hat{\theta}) - G(m)\right) + \alpha \int_{\hat{\theta}}^{\infty} (\theta - \theta^*) dG(\theta)\right\} \qquad (3.8)$$

The values of b and ρ are jointly identified by (3.8). Following usual practice, I will assume a value for ρ and use the values of other parameters to calculate b from (3.8).

[33] A total of twenty observations are below \$8.25, of which eight are \$8. For this reason, measurement error and "heaping" seem the most likely culprit, indicating that these wages are best interpreted as the minimum.

3.4. Results

3.4.1. Parameter Estimates

Table 3.2 provides the results of estimating the structural model of job search with Nash bargaining and a binding minimum wage. For comparison, I estimate the parameters for the whole sample and then separately for the treatment and control groups. While in the main simulation I will only use the estimated parameters from the control group, estimating parameters for the treatment group helps provide a first pass at filtering the effects of treatment through the model. The estimates for the treatment and control groups are in the second and third columns, and the parameters change dramatically between the treatment and control groups. For instance, the treatment group has a dramatically higher arrival rate of offers, 4.6578 versus 0.0070, but they come from a distribution with a much lower mean, -2.29 as opposed to 2.85. These results indicate that some caution should be taken in directly interpreting differences in structural parameters for the treatment and control groups.

However, calculated parameters that map directly to the data change between treatment and control groups in reasonable ways. For instance, the parameter estimates imply that the treatment group receives many more job offers but that a vast majority of them are obviously poor and are immediately rejected. On net, these differences cancel out. For instance, the hazard rate of leaving unemployment rises from 0.0069 in the control group to 0.0081 in the treatment group, a meaningful and plausible increase. Similarly, the shadow reservation wage θ^* increases from $2.79 to $7.15, indicating both that the assumption of a binding minimum wage is confirmed (because both are less than $8.25) and that treatment appears to generate a higher value of being unemployed, reflected in a larger shadow reservation wage. Similarly, the productivity required to earn more than minimum wage, $\hat{\theta}$, falls from nearly twice the minimum wage to less than \$2 above it at \$9.90. Since workers only workers with productivities between minimum wage and $\hat{\theta}$ receive minimum wage, this suggests that treatment leads to a large drop in the proportion of workers making minimum wage.

3.4.2. Simulating Treatment

Given these parameter estimates, we can proceed to the main goal of validating the estimated model against the experimental results. As noted above, the validation exercise involves two steps. First, the model will be estimated on the control group data only. Second, I will take measured treatment effects on search inputs from the experiment, i.e. using both control and treatment group data on search *inputs*, and use this as a guide for how to simulate treatment in the theoretical model. Importantly, this exercise will not use treatment group data on wages and unemployment durations. Thus, comparing the predictions of the model to the experimental treatment effects for wages and unemployment durations provides an external test of the validity of the model.

The first step was already completed when estimating the model on the control group data, i.e. column 3 of Table 3.2. For the second step, consider the final row of Table 3.1. Receiving the transportation subsidy treatment increased search intensity, as measured by the number of search actions taken, by approximately 15 percent. For the purposes of simulating the effects of treatment, I will assume that this is the mechanism by which treatment affects job search. In particular, I will assume that search intensity directly translates into a higher arrival rate of job offers through a constant returns technology. Since treatment increases search intensity by 15 percent, I will simulate treatment by using the control group parameters for both models but then increasing the rate at which workers meet with firms, λ, by 15 percent relative to its estimated value for the control group. The final column of Table 3.2 lists the parameters I use for this exercise. Note that the unchanged structural parameters consist of $\mu, \sigma, \rho, \alpha, m,$ and b.

Other parameters, including $\hat{\theta}, \theta^*,$ and the hazard rate, reflect endogenous choices of the job seeker and are thus potentially altered by a shock to λ. The increased arrival rate has a significant effect on labor market behavior in the model. The hazard rate of leaving unemployment only depends on the level of the minimum wage, the arrival rate of offers, and the parameters of the productivity distribution. As a result, the hazard necessarily increases due to the higher arrival rate of offers. At the same time, firms and workers do respond to the higher match rate. The worker's shadow reservation wage, θ^*, increases by more than a dollar from \$2.79 to \$3.87 and the productivity level required to earn more than minimum wage decreases from \$16.44 to \$14.82.

3.4.3. Validation

Having simulated the impact of increasing the arrival rate of offers by 15 percent on labor market behavior, I can now use the model to translate these behavioral changes into predicted treatment effects on wage and employment outcomes observed in the experiment. The first two columns of the first panel of Table 3.3 list characteristics of the treatment and control group as observed in the experiment. Given the censored data, assumptions are required for some of these characteristics. In particular, hazard rates of leaving unemployment and average unemployment durations are estimated taking into account censoring assuming an exponential distribution of unemployment durations. All other variables are directly observed in the experiment.

The final column in the first pane of Table 3.4 lists the difference between treatment and control for each variable. These are the facts that the estimated structural models must match in order to be valid for this experiment. In the experiment, treatment resulted in a large decrease in unemployment duration. On average, durations decreased by 30 days with a particularly strong effect during the first 40 days at which time the probability of being employed was 9 percentage points higher in the treatment group than in the control group. The impact of treatment on wages is smaller and statistically insignificant, though still noticeable in magnitude. The average wage rate increased by 44 cents and the proportion of workers working at or below minimum wage dropped by 12 percentage points.

The second pane of Table 3.3 provides the results of simulating treatment in the structural model. The first two columns list the outcomes predicted for the full set of control group structural parameters and then the outcome associated with shocking the matching rate of workers and firms. The final column then summarizes the net effect of this simulated treatment. The model of job search with Nash bargaining and binding minimum wages, while not perfect, does very well in matching the experimental results. It underestimates the improvement in unemployment durations somewhat, predicting a decrease of 19 rather than 30 days. This is mainly due to missing the large increase in the probability of employment within 40 days. However, by the end of the ninety day follow-up period, the model predicts the employed population almost exactly. Altogether, it fairs reasonably well in predicting hiring rates, though with some flaws.

Regarding the effect of transit subsidies on wages, this model does particularly well, predicting a 40 cent increase in average wage rates, just below the observed effect of 44 cents. It also matches the strong modal minimum wage precisely. The estimated model for the control group correctly fits the fact that over 40 percent of the sample works at or below minimum wage. Most importantly, it correctly predicts that an increase in the match rate should lead to a large reduction in the proportion of the workers at or below the minimum. It predicts a drop of 11 percentage points, just short of the actual 12 percentage points observed in the experiment.

Altogether, these results suggest that Flinn's model does a relatively good job of predicting changes in unemployment durations and an excellent job predicting changes in wages in response to the transportation subsidy experiment. In particular, the model fairs much better than a simpler search model with an exogenous, continuous distribution of posted wage offers.[34] This successful simulation helps validate using the model to answer questions beyond what can be addressed with simple observable outcomes.

3.5. Welfare Effects of Treatment

As demonstrated above, results from the transportation subsidy experiment in chapter 2 largely validate Flinn's (2006) job search model incorporating binding minimum wages and Nash bargaining. The validated model can now be used to measure outcomes not present in the data. In this section, I provide one application, using the model to calculate simple measures of welfare for the worker. Using the same simulation as above, I can then measure the effect of being assigned to the transit subsidy treatment on worker welfare. Combined with cost data from the experiment, this estimate can be used to determine whether the welfare gain of treatment

[34] See appendix B for results with the simpler model.

exceeded the cost of treatment. Assuming that externalities to job search are on net zero[35], this comparison can also be used to measure the social gains from treatment.

Measuring welfare of an unemployed individual in the search model is straightforward. It is simply the value of being unemployed, U. In the simple model with an exogenous distribution of wages, welfare can be determined by the reservation wage indifference condition:

$$U = \frac{w^*}{\rho}$$

With a binding minimum wage and Nash bargaining, the worker is not at this indifference point. However, we can make use of the shadow reservation wage, θ^*, which holds the same role. Simply rearranging the definition of θ^* we get:

$$U = \frac{\theta^*}{\rho}$$

The impact of treatment on individual welfare can then be calculated by the difference between the control group value of unemployment and that of the simulated treatment group, which can be calculated from the final two columns of Table 3.2:

$$U^S - U^C = \frac{\theta_S^*}{\rho} - \frac{\theta_C^*}{\rho} = \frac{3.87}{0.0016} - \frac{2.79}{0.0016} = \$662$$

Thus, treatment leads to an increase in welfare of \$662 for an unemployed worker who receives the transit subsidies. Relative to a treatment cost of \$140 per person, the experimental treatment resulted in a net welfare gain of \$522 for each treated individual according to the validated model. This measure provides an improvement over that in chapter 2 as it comes from a complete, validated model of job search behavior and includes the value of non-market time, b.

[35] There are at least two possible externalities. More search can help firms by making the market `thicker' but can harm other unemployed individuals competing for the same jobs.

3.6. Conclusion

Following in the example of Todd and Wolpin (2006), I use data from an experiment to validate a structural economic model. In particular, I use data from a field experiment of randomly assigned transportation subsidies in a low-wage labor market to test a model of job search with search frictions, a binding minimum wage, and Nash bargaining. I structurally estimate the model using data only from the control group. Then, I simulate the impact of treatment as an increase in the rate of any given worker matching with an employer. This simulation provides me with predicted treatment effects from both models that can be compared to the experimental results. I find that this simplified version of Flinn's (2006) model of job search fits the results of the experiment quite well, while a simpler model with an exogenous, continuous wage distribution does rather poorly.

These results underscore the importance of tailoring structural models of job search to the specific context of study. While a simple "off-the-shelf" model does quite poorly, Flinn's model fairs much better because it incorporates the main features of a low-wage labor market, a binding minimum wage and a mass point in the wage distribution at the minimum. Of course, the results do not imply that this model will prove best in all markets. Instead, it provides credence for the model of minimum wages, search frictions, and bargaining in accounting for what we observe in *low-wage* labor markets.

Finally, having validated a structural model of the labor market in question, I use the model to answer an empirical question that cannot be answered without the structure of a model: does treatment improve the welfare of an unemployed individual in a cost-effective manner? Using the validated model, I find that treatment leads to an estimated $662 increase in the welfare of someone who receives treatment, which is well in excess of program costs. Altogether, treatment leads to more than $500 in benefits in excess of program costs. This result both strengthens the conclusion of the experimental study in chapter 2, indicating that transportation subsidies can lead to large improvements in welfare for the recipient, and indicates the gains from using theory to interpret experimental results.

Table 3.1. Experimental Treatment Effects

	Observations	Control Mean	Treatment Mean	Difference
Found Job Within 40 Days	427	0.26	0.35	0.09
Found Job Within 90 Days	427	0.45	0.5	0.05
Wage Rate	202	9.59	10.03	0.44
Number of Search Actions Taken	465	2.33	2.72	0.39

Table 3.2. Estimated Structural Parameters

	All	Treatment	Control	Simulated
Assumed				
α	0.4	0.4	0.4	0.4
m	8.25	8.25	8.25	8.25
ρ	0.0016	0.0016	0.0016	0.0016
Estimated				
μ	1.79	-2.29	2.85	2.85
σ	0.72	1.51	0.35	0.35
λ	0.0225	4.6578	0.0070	0.0081
θ^*	6.49	7.15	2.79	3.87
Calculated				
b	-8.87	-7.41	-26.42	-26.42
hazard	0.0074	0.0081	0.0069	0.0079
$\widehat{\theta}$	10.89	9.90	16.44	14.82

Table 3.3. Experimental and Simulated Treatment Effects

A. Experimental Results

	Control	Treatment	Difference
Hazard Rate	0.0068	0.0085	0.0017
Average Unemployment Duration	148	118	-30
Probability Employed in 40 Days	0.26	0.35	0.09
Probability Employed in 90 Days	0.45	0.50	0.05
Average Wage	9.59	10.03	0.44
Proportion At or Below Minimum Wage	0.43	0.31	-0.12

B. Simulated Results

	Fitted Control	Simulated Treatment	Difference
Hazard Rate	0.0069	0.0079	0.0010
Average Unemployment Duration	145	126	-19
Probability Employed in 40 Days	0.24	0.27	0.03
Probability Employed in 90 Days	0.46	0.51	0.05
Average Wage	9.53	9.93	0.40
Proportion At or Below Minimum Wage	0.42	0.31	-0.11

APPENDIX A: THEORY

A.1. Search Intensity Comparative Static

As an initial step, note that substituting (2.2) into (2.1) results in:

$$w^* = b + (\alpha - 1)\int_\delta \gamma\lambda(\delta)^\alpha g(\delta)d\delta$$

Taking the derivative of this with respect to γ gives us an alternate expression for $\frac{\partial w^*}{\partial \gamma}$:

$$\frac{\partial w^*}{\partial \gamma} = (\alpha - 1)\int_\delta \lambda(\delta)^\alpha g(\delta)d\delta + (\alpha - 1)\int_\delta \gamma\lambda(\delta)^{\alpha-1} g(\delta)\frac{\partial \lambda(\delta)}{\partial \gamma}d\delta \qquad (A.1)$$

Now consider (2.2) and take the derivative with respect to γ:

$$\alpha\lambda(\delta)^{\alpha-1}g(\delta) + \alpha(\alpha-1)\gamma\lambda(\delta)^{\alpha-2}g(\delta)\frac{\partial\lambda(\delta)}{\partial\gamma} = \frac{-(1-F(w^*))}{r}\frac{\partial w^*}{\partial\gamma} \qquad (A.2)$$

Substituting (A.1) into (A.2), re-arranging terms, and changing the variable of integration provides:

$$\alpha(\alpha-1)\gamma\lambda(\delta)^{\alpha-2}g(\delta)\frac{\partial\lambda(\delta)}{\partial\gamma} + \frac{[1-F(w^*)](\alpha-1)}{r}\int_s \gamma\lambda(s)^{\alpha-1}g(s)\frac{\partial\lambda(s)}{\partial\gamma}ds$$

$$= \frac{-\left(1-F(w^*)\right)}{r}(\alpha-1)\int_s \lambda(s)^\alpha g(s)ds - \alpha\lambda(\delta)^{\alpha-1}g(\delta)$$

Define $\Gamma(\delta)$ such that:

$$\Gamma(\delta) = \frac{\left(1-F(w^*)\right)}{r}(\alpha-1)\int_s \lambda(s)^\alpha g(s)ds - \alpha\lambda(\delta)^{\alpha-1}g(\delta) > 0$$

so that

$$\alpha(\alpha-1)\gamma\lambda(\delta)^{\alpha-2}g(\delta)\frac{\partial\lambda(\delta)}{\partial\gamma} + \frac{[1-F(w^*)](\alpha-1)}{r}\int_s \gamma\lambda(s)^{\alpha-1}g(s)\frac{\partial\lambda(s)}{\partial\gamma}ds = -\Gamma(\delta) < 0 \qquad (A.3)$$

Multiplying both sides by $\frac{[1-F(w^*)]\lambda(\delta)}{r}$, integrating over δ, and some re-arranging results in:

$$\frac{[1-F(w^*)](\alpha-1)}{r}\int_{\delta} \alpha\gamma\lambda(\delta)^{\alpha-1}g(\delta)\frac{\partial\lambda(\delta)}{\partial\gamma}d\delta = \frac{-\int_{\delta} \Gamma(\delta)[1-F(w^*)]\lambda(\delta)d\delta}{r+\int_{\delta} [1-F(w^*)]\lambda(\delta)d\delta} \qquad (A.4)$$

Substituting (A.4) into (A.3) along with some re-arranging yields:

$$\alpha(\alpha-1)\gamma\lambda(\delta)^{\alpha-2}g(\delta)\frac{\partial\lambda(\delta)}{\partial\gamma} = \frac{-r\Gamma(\delta)+\int_{s} [\Gamma(s)-\Gamma(\delta)][1-F(w^*)\lambda(s)]ds}{r+\int_{s} [1-F(w^*)]\lambda(s)ds}$$

From the definition of $\Gamma(t)$:

$$\Gamma(s)-\Gamma(\delta) = \alpha\lambda(s)^{\alpha-1}g(s)-\alpha\lambda(\delta)^{\alpha-1}g(\delta) = 0$$

where the final equality follows from (2.2). Thus:

$$\alpha(\alpha-1)\gamma\lambda(\delta)^{\alpha-2}g(\delta)\frac{\partial\lambda(\delta)}{\partial\gamma} = \frac{-r\Gamma(\delta)}{r+\int_{s} [1-F(w^*)]\lambda(s)ds}$$

which is the desired result.

A.2. Reservation Wage Comparative Static

Recall equation (2.1):

$$w^* = b-\int_{\delta} \gamma\lambda(\delta)^{\alpha}g(\delta)d\delta+\int_{\delta} \lambda(\delta)\int_{w^*}^{\infty}\frac{w-w^*}{r}f(w)dwd\delta \qquad (2.1)$$

Taking the derivative of both sides with respect to γ implies:

$$\frac{\partial w^*}{\partial\gamma} = -\int_{\delta} \lambda(\delta)^{\alpha}g(\delta)d\delta+\int_{\delta} \lambda(\delta)\int_{w^*}^{\infty}\left(\frac{-1}{r}\right)\frac{\partial w^*}{\partial\gamma}f(w)dwd\delta$$

where the impact of γon $\lambda(\delta)$ has been ignored due to an envelope condition. To get the final result, simply note that $\frac{\partial w^*}{\partial\gamma}$ can be pulled out of the integral. Then, evaluating the integral and re-arranging leads to the result.

APPENDIX B: PARTIAL EQUILIBRIUM SEARCH WITH NO MINIMUM WAGE

B.1. Theory

In addition to the model of search with a binding minimum wage and Nash-bargained wages, I also estimated a simpler job search model with an exogenous, continuous distribution of wages. As before, job offers are completely characterized by a wage rate, *w*, and arrive at a Poisson rate λ. I will assume a simple partial equilibrium setting in which the wages are drawn independently from an exogenous distribution, $F(w)$. This setting matches the environment described before, leading to value of employment and unemployment characterized by (3.1) and (3.2). With no minimum wage, the reservation wage decision of the job seeker determines when a match is formed and is defined by (3.3).

B.2. Estimation

The main difference between this model and the Flinn (2006) model is that here I will simply assume an exogenous, continuous distribution of wage offers rather than determining the distribution from an endogenous bargaining process. As shown by Flinn and Heckman (1982), nearly all of the parameters of this model can be identified and structurally estimated by concentrated maximum likelihood using only data on unemployment durations and wage rates. The results of the model imply that the hazard probability of leaving unemployment, $\lambda[1 - F(w^*)]$, is constant and depends on the reservation wage. As a result, the distribution of unemployment durations is exponential:

$$\Pr[t \le t_i] = 1 - e^{-\lambda[1-F(w^*)]t_i}$$

The density of unemployment durations, $f_t(t_i)$, is:

$$f_t(t_i) = \lambda[1 - F(w^*)]e^{-\lambda[1-F(w^*)]t_i}$$

Because of the reservation wage rule, the distribution of observed, accepted wages simply becomes a truncated version of the offer distribution. So, the likelihood of being employed at a particular wage conditional on employment is:

$$\frac{f(w)}{1 - F(w^*)}$$

I will estimate the parameters of the model using concentrated maximum likelihood. As such, the contribution of an individual with wage w_i and duration t_i to the likelihood function is:

$$L_i(w_i, t_i) = \lambda[1 - F(w^*)]e^{-\lambda[1-F(w^*)]t_i} * \frac{f(w)}{1 - F(w^*)} = \lambda e^{-\lambda[1-F(w^*)]t_i} f(w) \qquad (B.1)$$

Thus, (B.1) provides the contribution to the likelihood function for an individual with observable wage and duration data. As noted above, many observations are censored due to a short follow-up period in the experiment. These observations are censored at a duration of 90 days and have no wage data. Thus, their likelihood contribution is simply:

$$L_i(w_i, t_i) = \Pr[t \geq 90] = e^{-\lambda[1-F(w^*)]90} \qquad (B.2)$$

Let c_i be an indicator variable of whether observation *i* is censored. Then, the full likelihood function is:

$$L = \prod_{i=1}^{n} \left[\lambda e^{-\lambda[1-F(w^*)]t_i} f(w)\right]^{1-c_i} \left[e^{-\lambda[1-F(w^*)]90}\right]^{c_i} \qquad (B.3)$$

To estimate the parameters by concentrated maximum likelihood, I first estimate the reservation wage using the smallest wage in the sample, which is a strongly consistent estimator of the reservation wage. Then λ and the parameters of $F(w)$ are estimated by maximum likelihood on (B.3). This process requires a parametric assumption on the distribution of wage offers. In particular, I will make the standard assumption that $F(w)$ follows a log-normal distribution so that:

$$F(w) = \Phi\left(\frac{\ln(w) - \mu}{\sigma}\right)$$

where $\Phi(\cdot)$ is the standard normal distribution. The results of the estimation identify all remaining parameters except b and ρ, which are jointly identified via (3.3). I will follow standard practice, assuming a rate of time preference in order to identify b.

B.3. Results

Table B.1 provides the results for the model with an exogenous distribution of wages and no minimum wage. Comparing the second and third columns demonstrates two important facts about these parameters. First, the estimate of the reservation wage, w^*, is quite sensitive. The minimum observed wage in the control group is \$6.50 while the minimum in the treatment group is \$7.25. In terms of the estimation, this leads my estimate of the reservation wage to increase rather dramatically. Given that both of these wages are below the legal minimum wage, measurement error could be driving this difference. So, these results should be cautiously interpreted. Second, the other parameters change dramatically between treatment and control groups. The arrival rate of offers, λ, increases from 0.0074 to 0.0370, a fivefold increase, while the wage offer distribution shifts down and spreads out. All of these changes should be interpreted cautiously given the shift in estimated reservation wages. Given these parameters, though, I can estimate the predicted hazard probability of leaving unemployment. This shows a large but reasonable increase in the hazard from 0.0069 to 0.0081. Similar to the results in chapter 3, this suggests a bit more optimism, indicating that the large swings in parameters cancel out when mapping to observable characteristics of the data.

As with the bargaining model, I simulate the effect of treatment by increasing the rate of arrival of job offers by 15 percent. The parameters are listed in the final column of Table B.1. Table B.2 then provides the results of the simulation. As can be seen, the model with an exogenous wage distribution and no minimum wage does a poor job matching the experimental results. It predicts a decrease in unemployment durations but only of 14 days, half that of the actual experimental result. Also, it fails to capture the large increases in the probability of being employed within 40 days. Even after 90 days, it only reaches half the actual effect of the experiment. Regarding wages, the model predicts smaller than observed effects with only a 12 cent increase in average wages. Not surprisingly given its continuous wage offer distribution, it does particularly poorly in predicting wages in a population with a large mode at the minimum wage. By definition, it predicts no workers making exactly the legal minimum wage and only 29 percent, relative to an actual 43 percent, of the control group earning less than or equal to minimum wage. Finally, it predicts a drop in the sub-minimum wage population of 4 percentage points, only a third of the observed effect. Altogether, the basic partial equilibrium search model with an exogenous, continuous distribution of wage offers and no legal minimum wage has a difficult time replicating the observed treatment effects. Clearly, the model considered in chapter 3 incorporating binding minimum wages and Nash bargaining performs better.

Table B.1. Estimated Parameters, Model with Exogenous Wage Distribution

	All	Treatment	Control	Simulated
Assumed				
ρ	0.0016	0.0016	0.0016	0.0016
Estimated				
μ	2.20	1.59	2.21	2.21
σ	0.26	0.51	0.23	0.23
λ	0.0083	0.0370	0.0074	0.0085
w*	6.50	7.25	6.50	6.85
Calculated				
b	-8.44	-7.23	-6.62	-6.62
hazard	0.0074	0.0081	0.0069	0.0076

Table B.2. Experimental and Simulated Treatment Effects, Model with Exogenous Wage Distribution

	Fitted Control	Predicted Treatment	Difference
Hazard Rate	0.0069	0.0076	0.0007
Average Unemployment Duration	145	131	-14
Probability Employed in 40 Days	0.24	0.26	0.02
Probability Employed in 90 Days	0.46	0.50	0.03
Average Wage	9.60	9.72	0.12
Proportion At or Below Minimum Wage	0.29	0.25	-0.04

References

Angelucci, M. (2004) ‘Aid and Migration: An Analysis of the Impact of Progresa on the Timing and Size of Labour Migration.’ *IZA Discussion Paper* No. 1187.

Angrist, J. and J.S. Pischke (2010) "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics." *Journal of Economic Perspectives*, 24(2).

Attanasio, O., C. Meghir, and A. Santiago (2011) “Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA.” *Review of Economic Studies*, 8(2).

Bhargava, A. and F. Docquier. (2007) ‘A New Panel Data Set on Physicians’ Emigration Rates (1991-2004).’

Bhargava, A. and F. Docquier. (2008) ‘HIV prevalence and migration of healthcare staff in Africa.’ *World Bank Economic Review*, 22.

Black, D., J. Smith, M. Berger, and B. Noel (2003) `Is the Threat of Reemployment Services More Effective Than the Services Themselves: Evidence from Random Assignment in the UI System.' *American Economic Review*, 93(4).

Borjas, G. (1987) ’Self-Selection and the Earnings of Immigrants.’ *American Economic Review*, 77(4).

Burdett, K. and D. Mortensen. (1998) ‘Wage Differentials, Employer Size, and Unemployment.’ *International Economic Review*, 39(2).

Card, D., S. DellaVigna, and U. Malmendier (2011) "The Role of Theory in Field Experiments." *Journal of Economic Perspectives*, 25(3)

Chen, L. and J.I. Boufford (2005) ‘Fatal Flows: Doctors on the Move.’ *New England Journal of Medicine*, 353(17).

Clampet-Lundquist, S. and D. Massey (2008) `Neighborhood Effects on Economic Self-Sufficiency: A Reconsideration of the Moving to Opportunity Experiment.' *American Journal of Sociology*, 114.

Clark, X., T.J. Hatton, and J.G. Williamson (2007) ‘Explaining US immigration, 1971-1998.’ *Review of Economics and Statistics*, 28.

Clemens, M. (2007) ‘Do Visas Kill: Health Effects of African Health Professional Emigration.’ *Center for Global Development Working Paper* No. 114.

Clemens, M. (2009) ‘Skill Flow: A Fundamental Reconsideration of Skilled-Worker Mobility and Development.’ *Center for Global Development Working Paper* No. 180.

Colson, N.E., D. Laing, and P. Wang (2001) `Spatial Mismatch in Search Equilibrium.' *Journal of Labor Economics*, 19(4).

Crepon, B., et. al. (2011) `Do Labor Market Policies Have Displacement Effect: Evidence from a Cluster Randomized Experiment.' Working Paper.

Deaton, A. (2009) “Instruments of Development: Randomization in the Tropics, and the Search for the Elusive Keys to Economic Development*,” NBER Working Paper* No. 14690.

Dovlo, D. and F. Nyonator (1999) ‘Migration of Graduates of the University of Ghana Medical School: A Preliminary Rapid Appraisal.’ *Human Resources for Health Development Journal*, 3(1).

Dolton, P. and D. O'Neill (1996) `Unemployment Duration and the Restart Effect.' *Economic Journal*, 106.

Dolton, P. and D. O'Neill (2002) `The Long-Run Effects of Unemployment Monitoring and Work-Search Programs: Experimental Evidence from the United Kingdom.' *Journal of Labor Economics*, 20(2).

Eckstein, Z. and K. Wolpin (1990) “Estimating a Market Equilibrium Search Model from Panel Data on Individuals.” *Econometrica,* 58.

Eckstein, Z., S. Ge, and B. Petrongolo (2011) “Job and Wage Mobility with Minimum Wages and Imperfect Compliance.” *Journal of Applied Econometrics*.

Flinn,C. (2006) "Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates," *Econometrica*, 74(4).

Flinn, C. and J. Heckman (1982): “New Methods for Analyzing Structural Models of Labor Force Dynamics,” *Journal of Econometrics*, 18.

Gautier, P. and Y. Zenou (2010) `Car Ownership and the Labor Market of Ethnic Minorities' *Journal of Urban Economics*, 67.

Gibson, J. and D. McKenzie (2011) ‘The Microeconomic Determinants of Emigration and Return Migration of the Best and Brightest: Evidence from the Pacific.’ *Journal of Development Economics*, 95(1).

Gobillon, L., H. Selod, and Y. Zenou (2007) `The Mechanisms of Spatial Mismatch.' *Urban Studies*, 44(12).

Hanson, G. (2008) ‘International Migration and Development.’ *Commission on Growth and Development, Working Paper* No. 42.

Holzer, H., K. Ihlanfeldt, and D. Sjoquist (1994) `Work, Search, and Travel among White and Black Youth.' *Journal of Urban Economics*, 35(3).

Holzer, H. J., J. M. Quigley, and S. Raphael (2003) `Public Transit and the Spatial Distribution of Minority Employment: Evidence from a Natural Experiment.' *Journal of Policy Analysis and Management*, 22(3).

Kain, J. (1968) `Housing Segregation, Negro Employment, and Metropolitan Decentralization.' *Quarterly Journal of Economics*, 82.

Karlan, D. and J. Zinman (2009) “Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment.” *Econometrica,* 77(6).

Kean, M. (2010) “A Structural Perspective on the Experimentalist School.” *Journal of Economic Perspectives*, 24(2).

Kling, J., J. Liebman, and L. Katz (2007) `Experimental Analysis of Neighborhood Effects.' *Econometrica*, 75.

Lise, J., S. Seitz, and J. Smith (2005) "Evaluating Search and Matching Models Using Experimental Data," *IZA Discussion Paper* No. 1717.

List, J. and I. Rasul (2010) `Field Experiments in Labor Economics.' *NBER Working Papers* No. 16062.

Lopez, R. and M. Schiff (1998) ‘Migration and the Skill Composition of the Labour Force: The Impact of Trade Liberalization in LDCs.’ *The Canadian Journal of Economics*, 31(2).

Ludwig, J., et. al. (2008) `What Can We Learn about Neighborhood Effects from the Moving to Opportunity Experiment?'" *American Journal of Sociology,* 114.

Mayda, A.M. (2010) ”International migration: A panel data analysis of the determinants of bilateral flows” *Journal of Population Economics,* 23(4).

McCall, J. (1970). "Economics of information and job search". *Quarterly Journal of Economics*, 84(1).

Meyer, B. (1995) `Lessons from the U.S. Unemployment Insurance Experiments.' *Journal of Economic Literature*, 33.

Miguel, E. and M. Kremer (2004) `Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities' *Econometrica*, 72(1).

Ministry of Health (2007) 'Human Resource Strategies and Policies for the Health Sector: 2007-2011.' Mimeo.

Okeke, E. (2009) 'An Empirical Investigation of Why Doctors Migrate and Women Fail to Go For Screening.' Unpublished dissertation, University of Michigan.

Pedersen, P. J., M. Pytlikova and N. Smith (2008) 'Selection or network effects? Migration flows into 27 OECD countries, 1990-2000.' *European Economic Review*, 52(7).

Pissarides, C. (2000): *Equilibrium Unemployment Theory* (Second Ed.). Cambridge, MA: MIT Press.

Raphael, S. and L. Rice (2002) `Car Ownership, Employment, and Earnings.' *Journal of Urban Economics*, 52(1).

Raphael, S. and M. Stoll (2000) `Racial Differences in Spatial Job Search Patterns: Exploring the Causes and Consequences.' *Economic Geography*, 76(3).

Roder, A. and S. Scrivner (2005) *Seeking a Sustainable Journey to Work: Findings from the National Bridges to Work Demonstration*. Public Private Ventures: Philadelphia.

Todd, P. and K. 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).

van den Berg, G. (1994) `The Effects of Changes of the Job Offer Arrival Rate on the Duration of Unemployment.' *Journal of Labor Economics*, 12(3).

van den Berg, G. and G. Ridder (1998) "An Empirical Equilibrium Search Model of the Labour Market." *Econometrica*, 66.

Vujicic, M., et. al. (2004) 'The Role ofWages in the Migration of Health Care Professionals From Developing Countries' *Human Resources for Health*, 2(3).

Wilson, W.J. (1997) *When Work Disappears: The World of the New Urban Poor*. Vintage Books.

World Health Organization (2010) *World Health Statistics 2010.* WHO Press, Geneva.

Yang, D. (2006) 'Why Do Migrants Return to Poor Countries? Evidence from Philippine Migrants' Responses to Exchange Rate Shocks.' *Review of Economics and Statistics,* 88(4).

Yang, D. and H. Choi (2007) 'Are Remittances Insurance: Evidence from Rainfall Shocks in the Philippines.' *World Bank Economic Review*, 21.

Zax, J. and J. Kain (1996) `Moving to the Suburbs: Do Relocating Companies Leave Their Black Employees Behind?' *Journal of Labor Economics*, 14.