Canadian Labour Market and Skills Researcher Network

Working Paper No. 108

Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist

Kory Kroft
University of Toronto

Devin G. Pope
University of Chicago
NBER

November 2012

CLSRN is funded by the Social Sciences and Humanities Research Council of Canada (SSHRC) under its Strategic Knowledge Clusters Program. Research activities of CLSRN are carried out with support of Human Resources and Skills Development Canada (HRSDC). All opinions are those of the authors and do not reflect the views of HRSDC or the SSHRC.

Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist*

Kory Kroft

University of Toronto

Devin G. Pope

University of Chicago and NBER

Abstract

Ever since the seminal work of Stigler (1962), economists have recognized that information in markets is costly to acquire and can lead to “search frictions”. The remarkable growth in online search has substantially lowered the cost of information acquisition. Despite this, there is little evidence concerning the extent to which this has altered the search process and raised overall matching efficiency. To address this issue, we analyze the expansion of the website “Craigslist”, which allows users to post job ads and apartment and housing rental ads at virtually no cost. Exploiting the sharp geographic and temporal variation in the availability of online search, induced by Craigslist, we produce three key findings: Craigslist significantly lowered classified job advertisements in newspapers, caused a significant reduction in the apartment and housing rental vacancy rate, and had no effect on the unemployment rate.

JEL Code: J22, J23, J24, J64, O33

Keywords: Internet, Craigslist, matching, unemployment, vacancies

Contact Kroft at kory.kroft@utoronto.ca and Pope at devin.pope@chicagobooth.edu.

*We thank David Card, Raj Chetty, Gordon Dahl, Stefano DellaVigna, Michael Greenstone, Ulrike Malmendier, Emmanuel Saez and anonymous referees and numerous seminar participants for helpful comments and discussions. We thank Robert Plummer at Corzen.com, Ken Goldstein at the Conference Board, and Karen Marvin at The Media Audit for help in obtaining access to data used in this paper. We also acknowledge the financial support of the Wharton Real Estate Center.

The advent of the Internet has noticeably affected the way markets function. By providing an abundance of information at a very low cost, Internet websites such as Facebook, eBay, Travelocity.com, and Match.com have dramatically altered the ways individuals search for jobs, apartments, friends, and consumer products. The increased availability of these new search technologies and the information they embody suggests that they may be making older search technologies obsolete and significantly reducing market frictions.

Anecdotal evidence suggests that online search has crowded out certain types of traditional search methods such as phone books and travel agents. Furthermore, several innovations (e.g. Google Maps) clearly provide more efficient features than their pre-Internet alternatives. Little is known however, about the broader impacts of improved market information. This is certainly not due to a lack of interest as there has been considerable discussion among academic economists and the public in general.¹ Rather, the key difficulty is the lack of available sources of exogenous variation in Internet usage. At the individual level, identification is difficult due to the self-selection of search method. At the aggregate level, overall Internet usage has expanded very slowly across a large number of years making it difficult to separate the effect of the Internet from a general time trend.

In this paper, we exploit a quasi-experiment created by the sharp expansion of the website “Craigslist” into cities across the United States. Craigslist is a website whose primary goal is to provide a platform to advertise jobs, apartment rentals, personals, and items for sale, at virtually no cost to the user. The first part of our empirical analysis tests whether online search on Craigslist has crowded out traditional print search. It is frequently argued in the popular press that online search is behind the recent struggles faced by major newspapers. A recent article in the December 2006 issue of *Forbes* entitled “Newspaper Killer” claims that Craigslist has “revolutionized the classified advertising market with its free listings” and that “newspaper companies, in particular, have been hit hard”.² Little is known about the overall magnitude of this effect and whether it can be interpreted as causal. However, even if online search is significantly crowding out print search, this cannot be taken as *prima facie* evidence that online search produces more matches. This motivates the second part of our analysis that empirically tests whether online search improves aggregate matching efficiency.

Our empirical analysis of matching efficiency focuses on the labor and the apartment and housing rental markets. The average duration for the vacancy of a rental unit in the United States is

¹ Katz and Krueger (1999) and Autor (2001) provide efficiency rationales for online search in labor markets.

² See Perez-Pena (2008) for another example.

approximately 4 months. Estimates of the average duration of a job vacancy are typically less than 1 month.³ An open empirical question facing both urban and labor economists is the extent to which “search frictions” underscore these vacancy measures. The introduction of a potentially more efficient search technology and the effect it has on vacancies can perhaps shed light on this question.

Both the labor and apartment and housing rental markets have experienced rapid growth in online search. We illustrate this growth for the labor market using a unique dataset that records monthly job post counts for the major online jobs boards – Monster.com, CareerBuilder.com, Yahoo/Hotjobs, and Craigslist.⁴ In particular, we show that Craigslist rose from near obscurity in 2005 to become a major contender, if not the leader, in online job posts by 2007. Our results suggest that, for the major metropolitan areas in the United States, the number of Craigslist job posts represented two-thirds the combined total of Monster, CareerBuilder, and Yahoo/Hotjobs posts in 2007. We also show that this growth varied considerably across Metropolitan Statistical Areas (MSAs). We exploit this variation to implement a differences-in-differences strategy.⁵ Specifically, we estimate the effect of the difference in Craigslist posts between 2005 and 2007 across a set of MSAs on the difference in newspaper job classifieds as represented by the Help-Wanted Index (the “crowd-out effect”), and on the difference in the apartment and housing rental vacancy rate and the unemployment rate (the “efficiency effect”).

The key advantage of our empirical strategy is that Craigslist caused a large change in the take-up of online search over a very short time period. However, this alone is not sufficient for causal identification. We also need a common trends condition to hold. Specifically, we require that job posts in newspapers, unemployment, and apartment and housing rental vacancies each would have evolved similarly across MSAs, absent Craigslist. We present several pieces of evidence that provide support for this assumption. Firstly, we show that our results are robust to the inclusion of several controls, notably pre-2005 trends. Secondly, we show that our results remain intact when using two instrumental variable (IV) regressions. The first regression instruments the growth in Craigslist job and housing posts in an MSA with the growth in Craigslist personal posts in the same MSA.⁶ In the empirical section, we

³ We computed a lower bound of 3.93 months as the average duration for rental vacancies in the US based on 2007 data from the US Census Bureau. Average job vacancy durations have often been estimated to be 2 weeks (see, for example, Abraham (1983))

⁴ We don’t consider the growth in online apartment and housing rental posts due to limited data availability.

⁵ Strictly speaking, our research design is not “differences-in-differences”, as the terminology is commonly used in the applied literature. Rather, most of the MSAs are “treated” since each receives Craigslist. The key is that the treatment intensity varies among MSAs. It is this variation that we exploit in our empirical analysis.

⁶ The majority of Craigslist housing posts are for apartment and house rentals, although they also contain posts for other categories such as vacation rentals, office and commercial space rentals, and real estate for sale. Personal posts are for individuals seeking relationships with others.

discuss how this instrument can remove potential bias from our estimates. The second instrumental variable strategy exploits the fact that Craigslist entered different MSAs at different times and the fact that Craigslist charges for job posts in some MSAs, but not others. Under the assumption that these decisions were not influenced by local labor and apartment and housing rental conditions between 2005 and 2007, the time since entry into an MSA and the decision to charge for a job posts jointly serves as a valid instrument for the increase in Craigslist posts.

We use our empirical strategy to demonstrate three key results. First, we find that Craigslist caused a significant reduction in the Help-Wanted Index, an official measure of job classifieds in print newspapers, produced by the Conference Board. For the average MSA in our sample, Craigslist reduced the number of classified job posts by approximately 7% between January 2005 and April 2007. Interestingly, this result helps us explain why the relationship between the Help-Wanted Index and a more direct measure of job vacancies obtained from the Job Openings and Labor Turnover Survey (JOLTS) is no longer stable.⁷ In particular, we show that the two series track each other very closely through the end of 2003; however, starting in the middle of 2004, we show that there is a sharp divergence between the two series – the Help-Wanted Index is trending downwards, while at the same time, the JOLTS series is trending upwards. Our results indicate that Craigslist alone can account for 15% of this divergence.

Our second key result is that Craigslist caused a significant reduction in the apartment and housing rental vacancy rate. Specifically, we find that, for the typical MSA in our sample, Craigslist reduced the rental vacancy rate by roughly 10% (1 percentage point). Using a simple calibration, we show that this implies that rental units where Craigslist was utilized for advertising take approximately 3 weeks less to rent out than they would have otherwise. In addition to showing that these results are robust to the inclusion of controls and our instrumental variable strategies, we conduct two placebo tests. First, we show that Craigslist has had no effect on for-sale home vacancy rates. This lends credibility to our results since Craigslist is typically not used to advertise home sales and as a result, we would not expect there to be a significant effect on for-sale home vacancy rates. Second, we use historical rental vacancy data to show that 2005-2007 is the only two-year period in the last decade where Craigslist growth from 2005-2007 caused a significant reduction on rental vacancy rates. Both of these false experiment tests lend support to our interpretation of a causal effect of Craigslist on rental vacancy rates.

⁷ This further explains why the Conference Board has recently developed a new index of vacancies that is based on online job posts.

Third, in contrast to our findings in the apartment and housing rental market, we find that Craigslist has had no measurable impact on labor market outcomes, as measured by local area unemployment rates. We can reject a 0.1 percentage point change in the unemployment rate due to Craigslist, which we argue is sufficient to rule out large effects of Craigslist job posts on the number of unemployed individuals. We also show that Craigslist did not have a detectable effect on low-skilled employment. We discuss several reasons for why Craigslist may have affected matching efficiency in the apartment and housing rental market, but not the labor market, and what we can learn from these results. One explanation focuses on the fact that our statistical test may lack power. Unlike the apartment and housing rental market, the labor market has many other online job boards, and so the marginal impact of Craigslist may have been muted. A second explanation is that the two markets may be fundamentally different in ways that make the apartment and housing rental market more responsive to improvements in information flows. In particular, search frictions might be much more important in the apartment and housing rental market than in the labor market. We view our contrasting findings across these markets as a potentially interesting area for future research.

Our paper contributes to the small literature that estimates the effect of the Internet on job search outcomes. Kuhn and Skuterud (2004), the first empirical paper on this topic, use individual-level CPS data, along with a supplement containing information on the method of job search, to identify the effect of online job search on subsequent unemployment duration. The main findings of this novel paper are that, after controlling for observables, online search leads to slightly *longer* unemployment durations. Kuhn and Skuterud conclude that negative selection on unobservables is the most likely explanation for their result. In a second paper, Stevenson (2007) uses similar CPS data to study the effect of online search on job-to-job flows. Stevenson finds that, after controlling for individual characteristics, Internet usage tends to increase the number of job-to-job flows. Moreover, she finds that the results are robust to using an instrumental variable strategy that relies on the way in which innovations such as the Internet typically expand across US states. Our analysis complements these studies and adds to them by looking at the effect of online search on newspaper advertisements and rental vacancy rates. Additionally, by focusing on city-level unemployment rates, we estimate the “general equilibrium” effect of online job search on labor market outcomes, which includes search externalities that may not be captured by studies that focus on individual outcomes.

Our study also contributes to the literature on Matching and the Beveridge curve (Abraham, 1987; Bleakley and Fuhrer, 1997; Katz and Krueger, 1999; Autor, 2001; Shimer, 2005). The degree of

substitution between print and online job ads is important for assessing the validity of the Help-Wanted Index, which is used to construct the Beveridge curve for the US economy. We show how our reduced-form estimates can be used to measure how growth in online search has affected the Beveridge curve, which is important for the conduct of monetary policy. Finally, our paper is related to the broader literature on how the Internet affects markets. This literature includes testing the effect of the Internet on price discrimination in the automobile market (Scott-Morton, Silva-Risso, & Zettelmeyer, 2003) and price dispersion in the life insurance market (Brown and Goolsbee, 2002).⁸

The remainder of the paper is organized as follows. Section I provides a conceptual framework for understanding how a more efficient search mechanism can affect matching efficiency. Section II discusses the institutional background of Craigslist and provides some background on Craigslist expansion, documents the changes in online search across time, and provides an explanation of our empirical strategy. The empirical results are presented in Section III. Section IV provides a brief discussion of the results and concludes.

I. Conceptual Framework

It is straightforward to conceptualize how online search affects print search. While it is possible that some job or housing posts appear both online and in print, it is likely that online posts crowd out print posts at least to some degree by providing an alternative, cheaper technology. In this paper, we empirically test whether and to what degree Craigslist crowds out traditional print ads.

It is less clear how online search impacts matching between workers and firms. Mortensen (2000) shows how a search technology like Craigslist, by reducing "search frictions", can lower frictional unemployment. There are various reasons why Craigslist provides a more efficient search method than traditional newspaper classifieds. Autor (2001) suggests that online job posts transmit much more information about a job and reach a far broader audience than their print counterparts. Autor further suggests that online job posts are more up-to-date, due to quick uploading and editing. The electronic format of job boards also allows employers and potential employees to quickly do keyword and other searches to locate potential matches.

⁸ Our paper also relates to the literature regarding whether new goods are complements or substitutes to existing products. Most related to our paper is Goolsbee (2001) and Sinai and Waldfogel (2004) who identify the relationship between online and offline retailing and also Gentzkow (2007) who tests whether online newspapers are substitutes or complements for their print counterparts.

A natural way to test for reduced search frictions is to examine the effect of Craigslist growth on the Beveridge curve, the relationship between vacancies and unemployment. When Craigslist enters a local labor market, this should shift the Beveridge curve towards the origin, reducing unemployment. Unfortunately, we are unable to implement this test due to the lack of vacancy data by MSA over time.⁹ Instead, we examine the effect of Craigslist on the unemployment rate. We also examine the effect of Craigslist on the rental vacancy rate to test for a reduction of search frictions in the rental housing market.

While Craigslist plausibly lowers search costs and increases matches, its welfare consequences are less clear. In general, if there is complete and symmetric information between workers and firms, a marginal reduction in search costs increases welfare (Autor, 2001). In markets with asymmetric information however, a reduction in search costs could exacerbate adverse selection problems. Intuitively, costly search can serve as an efficient screening device. We do not consider the welfare consequences of Craigslist. Rather, we focus on whether Craigslist reduces frictions and unemployment.

Another limitation of this study is that we do not jointly model the rental and labor markets, but rather we treat them as independent. It is natural to think that the location decision depends on the choice of job and vice-versa. Modeling such interactions would be quite interesting, but is well beyond the scope of our analysis.¹⁰

II. Institutional Background and Empirical Strategy

Craig Newmark founded the website "Craigslist" in 1995, serving exclusively the Bay area. The website serves as a platform for users to post ads that primarily focus on jobs, housing, services, personals, or for sale items. Due to tremendous growth, starting mainly around 2005, Craigslist now receives more than 8 billion page views per month making it one of the top ten visited English websites by the end of 2006.¹¹ While Craigslist began in the Bay Area, it expanded in 2000 to include other major

⁹ The typical measure of job vacancies by MSA is the Help-Wanted Index. We are unable to use this measure since our findings suggest that it has become an invalid marker of job vacancies due to online job posts. The JOLTS, a more direct measure of job vacancies that comes from surveying business establishments, does not provide such data at the MSA level.

¹⁰ Saks (2005) studies the interaction between housing and local labor markets. The key finding from this study is the impact of local labor demand shocks depend on the intensity of housing supply regulations. If housing supply is highly constrained, a boost to labor demand has a lower effect on employment, presumably since fewer potential workers can locate to the area. If Craigslist facilitates moving, this could indirectly affect employment in both the origin and destination locations.

¹¹ Alexa.com, Dec. 29, 2006

cities.¹² Appendix Table 1 provides a list of the 162 MSAs that Craigslist entered through June, 2006. As the table shows, Craigslist entered into major cities first, like Boston, New York and Los Angeles, followed in time by smaller MSAs. In Section III, we provide more background and detail on Craigslist's expansion, with a particular emphasis on relating this expansion to observable characteristics of MSAs.

For posting apartment and housing rentals, Craigslist has been a relatively unique service. While several Internet websites exist in different housing markets that allow such posts, Craigslist has been by far the largest national player for apartment and housing rental listings. For online job posts, on the other hand, Craigslist is one of several possible choices. Other major online job sites include Monster.com, CareerBuilder.com, and Yahoo/Hotjobs. Craigslist differs in several ways from these other job sites. First, Craigslist allows employers to post job ads for free in nearly all cities.¹³ This is a very different business plan than the other major job sites which charge anywhere from \$250-\$500 for a 30-60 day post. Typically, differences in price are explained by differences in quality. For the case of Craigslist, however, it is not so clear. For example, an article in the New York Times in December of 2006 entitled "Craigslist Meets the Capitalists" argues that the main objective of Craigslist is to help individuals locate jobs, cars, apartments and dates and surprisingly, not to make money doing so. According to the article, the fees that it does charge are to help defray the costs incurred by operating the site. In response to a question of how Craigslist plans to maximize revenue, CEO Jim Buckmaster replies "That is not part of the equation ... It's not part of the goal."

Craigslist post data. The main source of data that we use in this paper are monthly counts of the number of job, housing, and personal posts placed on Craigslist in the month of April, 2007.¹⁴ We self-collected these data for the 162 MSAs for which Craigslist had entered by June, 2006.¹⁵ Using MSA population counts from the 2000 Census, we normalized the data to represent the number of job, housing, and personal posts in April, 2007 per thousand individuals in each MSA. These normalized

¹² Craigslist has a separate webpage for each city. We use the phrase, "Craigslist entered into a city", to signify that a specific webpage was created for that city. The design and features of the website is uniform across all cities.

¹³ There are a handful of exceptions. Craigslist charges an employer \$75 to post an advertisement in San Francisco for 6 weeks and \$25 to post ads in New York, LA, DC, Boston, Seattle, San Diego, Portland, Sacramento, and Chicago.

¹⁴ On the website, job posts are broken down by sector and also full-time versus part-time. Housing posts contain the following categories: apartment/housing, rooms/shared, sublets/temporary, housing wanted, housing swap, vacation rentals, parking/storage, office/commercial, and real estate for sale. Finally, personal posts contain: strictly platonic, women seeking women, women seeking men, men seeking women, men seeking men, misc romance, casual encounters, missed connections, and rants and raves.

¹⁵ Strictly speaking, we collected data for the cities where Craigslist had entered and then assigned each city to the appropriate MSA.

data are presented in Appendix Table 1. The data illustrate the large amount of variation that exists in Craigslist popularity across geographic areas. In April of 2007, the Bay Area, Austin, Denver, Miami, Santa Barbara, and Medford all had more than ten job posts per thousand. At the same time, other MSAs such as Detroit, Cleveland, Indianapolis, and Cincinnati all had less than two job posts per thousand. It is also interesting to compare the job, housing, and personal posts within an MSA. For most cities, personal posts are the most popular, followed by housing posts, and then job posts. Interestingly, in cities where Craigslist is popular, it tends to be popular across all categories.¹⁶ This suggests that the popularity of Craigslist usage in an MSA may not be solely driven by market conditions (e.g. high unemployment resulting in a large number of job posts but not necessarily personal posts) but rather, may be based on non-economic factors such as social interactions, tipping points and publicity. This motivates our instrumental variable strategy, developed below, which uses personal posts as an instrument for job and housing posts.

There are several limitations to these data. While the geographic boundaries for most MSAs are reasonably clear, for large metropolitan areas such as New York and Los Angeles they are less obvious. It is possible that a Craigslist website in a certain MSA might be used by only a subset of the MSA or by more people than actually live in that MSA. This causes measurement error in our posts per population measure. Rather than try to correct this error on a case-by-case basis, we leave the posts per population numbers as they are and simply note that this error may cause attenuation bias in our estimates. The instrumental variables estimates should be purged of this bias. Another potential limitation that exists with these data is that we may be double counting posts that are re-listed on the website multiple times. We compare our job post counts with data collected by Corzen.com (described below) for a subset of cities serviced by Craigslist. Corzen uses a de-duping procedure to eliminate duplicate posts. Thus, its post measure should reflect the number of unique advertisements posted on Craigslist. We find that our self-collected Craigslist post count exceeds Corzen's Craigslist post count by 10.8%. Unfortunately, we do not have a procedure to eliminate duplicates in our self-collected Craigslist post data, since these data are only available as monthly counts for each MSA. In order to deal with repeat observations, we simply deflate all of our Craigslist post observations by the scalar 10.8% to obtain the most accurate measure of posts for each MSA. While this does nothing to the significance of our results, it slightly changes the interpretation on the coefficients.

¹⁶ $\text{Corr}(\text{job posts}, \text{housing posts}) = .78$; $\text{Corr}(\text{job posts}, \text{personal posts}) = .83$; $\text{Corr}(\text{housing posts}, \text{personal posts}) = .79$

Comparing Craigslist to other online job boards. It is important to understand how the number of Craigslist job posts compares with job post counts on other leading sites. If Craigslist is not a major player in this industry, then it is unlikely to affect city unemployment rates. In order to compare Craigslist with its competitors, we secured a dataset from a private company, Corzen.com, which has scraped data by month from Monster.com, CareerBuilder.com, and Yahoo/Hotjobs since January of 2004 for 36 large MSAs. Corzen.com uses a de-duping procedure that eliminate duplicate posts in order to get total counts for the number of new job posts each month.

Figure 1, based on the Corzen.com data, plots the average number of new job posts each month per one thousand of the population. The Figure includes all 36 MSAs for which Corzen data with monthly Craigslist counts are available.¹⁷ The Figure illustrates how online job posts have risen steadily over the last several years. In January of 2004, the number of Monster, CareerBuilder, and Yahoo/Hotjob posts per thousand ranged from 0.3 to 0.8. By April 2007, the number of posts ranged from 1 to 2.4 per thousand. This Figure also illustrates how Craigslist was nearly non-existent until the beginning of 2005 (the Bay Area is the major reason that it isn't completely zero). Starting in 2005, however, Craigslist saw dramatic increases in usage such that by April 2007, the number of Craigslist job posts per thousand far exceeded those of its competitors for these 36 large MSAs.¹⁸ This reiterates what has been reported in various newspaper articles discussing Craigslist; that Craigslist began to build momentum at the beginning of 2005 and became immensely popular over the following 2.5 years.

Empirical strategy. Figure 1 is a nice illustration of how it is difficult to identify the effect of the Internet on labor market outcomes. Monster.com, CareerBuilder.com, and Yahoo/Hotjobs have all experienced increases in job posts over the last several years, but at a very constant pace. The lack of sharp temporal changes makes it very difficult to identify the effect of these websites, or online markets in general, on matching outcomes. In contrast, the data reported in Appendix Table 1, coupled with the information in Figure 1, illustrates that over a two-year period, Craigslist experienced very rapid growth and that this growth was far from uniform across MSAs. This spatial and temporal variation motivates our empirical strategy.

¹⁷ Appendix Figure 1 provides the same plot for 121 MSAs for which Corzen data for Monster.com, CareerBuilder.com, and Yahoo/Hotjobs are available. However, Craigslist data is not available by month for these 121 MSAs. We indicate, however, the number of Craigslist job posts per thousand for these 121 MSAs from our self-collected data in April of 2007. This allows for a comparison of Craigslist to these other leading job boards for a broader sample - although only at one point in time.

¹⁸ As can be seen in the Figure, the Craigslist data are missing for all MSAs between April 2006 and August 2006. An error was made by Corzen.com in collecting and properly storing the data during those months.

In particular, we use the rise of Craigslist starting in 2005 as an event study and exploit the differences in this expansion across geographic areas. The intuition behind our empirical strategy is straightforward. When an innovative website, Craigslist, begins to receive attention, certain MSAs quickly begin to adopt the innovation. We can then test whether different outcomes resulted from cities that adopted the innovation quickly relative to cities that adopted Craigslist slowly. For ease of exposition, we refer to MSAs which experienced a large growth in Craigslist usage as “treatment MSAs” and the MSAs where Craigslist did not greatly expand as “control MSAs”.

The identification assumption necessary for consistency of our estimator is the standard “common trends” condition (Meyer, 1995), which in this case requires that the outcomes we study would have evolved similarly in treatment MSAs and control MSAs, absent Craigslist. This assumption can be violated in several ways. First, it is possible that outcomes in treatment MSAs were trending differently *prior to 2005* than outcomes in control MSAs. A natural way to evaluate this assumption is to simply plot the outcomes across time for both groups. We therefore present graphical evidence for all of our results that shed light on pre-2005 trends. In regressions, we explicitly control for pre-2005 trends.

A finding that outcomes in treatment and control MSAs were trending together pre-2005 does not necessarily imply that, *starting in 2005*, they would have evolved similarly, absent Craigslist. Although we can never know whether this assumption is satisfied, we nevertheless use a variety of methods that test the robustness of our results and that lend credibility to a causal interpretation of our findings. First, we control for observables that vary across MSAs. Motivated by evidence of how Craigslist expanded that we present below, we control for MSA population in 2004 and the growth in population from 2004 to 2006. This helps absorb economic shocks that differentially affected MSAs of different size and growth patterns. We also control for differences in the way MSAs might have evolved technologically, using as a proxy, MSA level Internet penetration data, collected from The Media Audit.¹⁹ Finally, in our rental vacancy rate regressions, we control for the change in housing prices, which can have an effect on apartment and housing rentals.²⁰

We also implement a Feasible GLS procedure. To do this, we estimate our baseline equation using OLS. We take the squared residuals from this regression and regress them on a constant and the inverse of city size. Finally, we use the inverse of the predicted values as weights in a weighted least squares

¹⁹ Each year, the Media Audit conducts a survey that asks individuals in each MSA whether or not they accessed the Internet from home in the month prior to the survey. We control for the change in the fraction of individuals in an MSA accessing the Internet from home between 2005 and 2007.

²⁰ We collected data on a housing price index from the Census and use as a control, the change in this measure between 2004 and 2006.

regression. This potentially has the ability to improve the precision of our estimates, a concern given our limited sample sizes.

Finally, we explore two instrumental variable strategies. A concern is that our estimates are biased since the number of Craigslist job or apartment posts are deterministic functions of the number of job or apartment vacancies. MSAs with a large number of apartment or job vacancies may simply take up Craigslist more aggressively to advertise these vacancies. Thus, any effects that we find could potentially be due to reverse causality. Specifically, when there are many job vacancies (and thus potentially many Craigslist job posts), the unemployment rate is likely to be low (Beveridge Curve). Thus, our OLS estimates for unemployment rate results may be negatively biased. In the rental housing market, when there are many vacancies (and thus many Craigslist rental posts), the rental vacancy rate is likely to be very high. Thus, our OLS estimates for the rental vacancy results may be positively biased downward.

To address this, we propose two instrumental variables strategies. First, we use the growth in personal posts per thousand as an instrument for growth in Craigslist job and housing posts per thousand. Under the assumption that growth in personal posts is not due to improvements in labor and apartment market conditions, using personal posts as an instrumental variable can help overcome this potential bias. Since personal posts on Craigslist are unlikely to be mechanically linked to the number of vacancies, use of this instrument can reasonably clean up a large part of the problem with reverse causality.

It is possible however, that a large number of job vacancies could result in more job posts on Craigslist, and as a result, increase familiarity with personals as well. Since we can't rule out such spillover effects, we propose a second instrumental variable strategy that exploits variation in the number of months that Craigslist has been active in an MSA, as well as variation in whether Craigslist charges money for postings in an MSA.²¹ Specifically, we use a spline in months since entry along with an indicator for whether Craigslist charges for posts as instruments to estimate the effect of Craigslist on labor and rental market outcomes.²² This instrument not only helps alleviate the concern of reverse

²¹ Craigslist charges \$25 for job posts in Chicago, NYC, LA, DC, Boston, Portland, Sacramento, Seattle, and San Diego and \$75 for job posts in San Francisco. It has been explained that it was decided that job posts should be fee based in some popular cities as a revenue source.

²² The fees that Craigslist charges for job posts in some cities may not satisfy the IV exclusion restriction since it is unclear how Craigslist decided their pricing policy. Specifically, one might worry that Craigslist's decision was influenced by city characteristics including unemployment or vacancy rates. The results that we provide in the paper are robust to using only time of entry as the IV which is likely to be less prone to failing to satisfy the exclusion restriction.

causality, but also allows us to isolate variation that is plausibly exogenous.²³ Under the assumption that Craigslist's entry decision and the decision of whether to charge for job posts is not endogenous to the labor and the apartment and housing rental market, this instrument is valid.

Our final attempt to address the potential sources of bias in our estimation is to conduct two placebo tests. We test whether the take-up of Craigslist had a significant effect on home vacancy rates. Since Craigslist had very little activity in home sale posts between 2005 and 2007, we posit that it should not have a significant effect on for-sale home vacancy rates.²⁴ We also use historical apartment and housing rental vacancy rate data to test whether there are other two-year periods over the last decade, other than 2005-2007, where rental vacancy rates in treatment MSAs were significantly different from rental vacancy rates in control MSAs. These two placebo tests provide additional evidence regarding the causality of our findings.

The precise baseline specification that we employ is

$$\Delta Y_i = \alpha + \beta \Delta \text{CraigslistPosts}_i + \gamma \Delta X_i + \delta Z_i + \theta \Delta Y_{i-1} + \varepsilon_i$$

The outcomes of interest, $\Delta Y_i \equiv Y_{i,2007} - Y_{i,2005}$, that we consider are the Help-Wanted Index (classified job advertisements in newspapers), the rental vacancy rate, and the unemployment rate. $\Delta \text{CraigslistPosts}_i \equiv \text{CraigslistPosts}_{i,2007} - \text{CraigslistPosts}_{i,2005}$ measures the difference in Craigslist posts per thousand in April of 2007 and January 2005.²⁵ Unfortunately, we only have January 2005 job post data for 36 MSAs. Thus, we assume that January 2005 posts are zero for the remaining MSAs. We consider this a minor assumption given that there were nearly zero posts in January of 2005 for every city other than a few of the larger cities (for which we have 2005 job posts).²⁶ Moreover, the MSAs for which we don't have job post data in 2005 were all cities for which Craigslist entered after October of 2004. This means that the number of posts by January of 2005 for these MSAs must have been extremely close to

²³ In section III, we present evidence on Craigslist expansion into local markets and discuss months since entry as an instrument.

²⁴ It is possible that the housing rental and sales markets are potentially linked. As Craigslist makes it easier to rent, those on the margin between selling their home and renting it out will choose the latter. We tend to think that such interactions will be minor as Craigslist is typically used to expedite the process of renting a housing unit for the "infra-marginal renters", rather than affecting the margin of whether to rent or sell a unit.

²⁵ When the dependent variable is the rental vacancy rate, this variable represents the change in Craigslist housing posts. When the dependent variable is the Help-Wanted Index or the unemployment rate, this variable represents the change in Craigslist job posts.

²⁶ 5 cities had more than 1 job post per thousand in January of 2005: Boston, Sacramento, San Diego, and San Francisco, and Seattle. All other cities had less than 1 job post per thousand (average = .27).

zero. The ΔX_i represents control variables that vary over time (e.g. Internet Penetration), Z_i represents control variables that don't vary over time (e.g. population), and ΔY_{i-1} represents the difference in outcomes from 2003 to 2004.

It is briefly worth mentioning two empirical approaches that we do not use in this paper. First, one may wonder why we focus on a two-year difference as opposed to looking at month-to-month changes in Craigslist posts and the effect of these changes on the outcomes of interest. The simple reason for this is that quality data for all MSAs is not available month-by-month. Recall that although Corzen.com records month-to-month Craigslist post data, this is only available for 36 MSAs. When we merge those MSAs with the MSAs for which outcome data is available, there are too few observations available for valid statistical inference. Furthermore, the Corzen data is missing for a 7-month period in 2006, which would erase some of the potential statistical gains from using month-by-month data to begin with.

Another empirical approach that we do not use in this paper is to test for trend breaks in outcomes when Craigslist enters an MSA. The reason we do not exploit this kind of event-study is that Craigslist did not become immediately popular when it entered an MSA. In fact, Craigslist entered 55 MSAs between 2000 and 2004 and yet there were nearly zero posts for these MSAs in January of 2005. Since no treatment took place at the Craigslist time of entry, there is simply not enough power to identify trend breaks around the date of entry. As described above however, we do use Craigslist entry dates as part of an instrumental variable strategy, based on the fact that the growth in Craigslist posts that started to occur in 2005 was correlated with the time that Craigslist entered each MSA.

Finally, it is worth noting that Craigslist growth at the MSA-level could potentially be correlated with other forms of online search. When Craigslist becomes popular in an area, there could be spillover effects onto other online search websites, such as Monster and CareerBuilder, as individuals learn about online search capabilities. We are unable to measure the local presence of these other online boards, since jobs from all cities are advertised on the same website. In the presence of such spillover effects, our results would over-state the impact of Craigslist, and would provide a lower bound estimate for the impact of online search overall.

Understanding Craigslist's Expansion. The key assumption of our empirical strategy is that the growth of Craigslist across cities is to a large extent idiosyncratic and fueled by non-measurable factors like word of mouth. Before turning to results, anecdotal and empirical evidence can help defend this assumption and shed light on why large variation existed in Craigslist expansion between 2005 and 2007.

According to Craigslist founder Craig Newmark when asked how newly added cities reach a “critical mass”, growth is “almost always a surprise. (It) might have to do with people moving from a Craigslist city to a new one, where they spread the word. That’s the only guess I have from observing rapid growth city sites, like Las Vegas and Hawaii.”²⁷

These remarks suggest that conditional on entering a new city, the growth in Craigslist is to a large extent, unpredictable. Unconditionally, and as we document below, Craigslist growth across cities between 2005 and 2007 is largely a function of when Craigslist entered a city. In an interview with the *San Francisco Chronicle* in 2004, Craig Newmark discussed the process for determining whether to enter a new city: “We put up a city based on how many people are asking us to do so. It’s also based on Jim’s perception of a city’s demographics and the city’s broadband penetration and intuition. We use word of mouth to get the word out, though sometimes the local press is kind... We just wait for things to happen. We don’t advertise, and we’re not good in promotions, so some cities grow slowly. For example, I think Memphis is growing real slow, but Las Vegas has grown quite fast.”²⁸

Motivated by this testimony, we next turn to empirically examine the effect of a city’s demographics and Internet penetration on Craigslist popularity. Our test focuses on Craigslist job posts, although similar results hold for Craigslist housing posts.

Column (1) of Table 1 displays the results from a regression of Craigslist job post penetration on several observables: (i) Log total population, (ii) percent change in total population, (iii) change in the percent of people with the Internet in their home, (iv) median age, (v) population density, (vi) an indicator for whether the MSA is located in a Pacific Coast state, (vii) Percent of employed population working in the Information industry, and (viii) fraction of the population with a Bachelors Degree.²⁹ Similarly, Column (3) of Table 1 displays the results using Craigslist housing post penetration on the same observables.

Our findings indicate that Craigslist initially entered large and growing MSAs and in MSAs with a more educated population. Consistent with Craig Newmark’s quote about the importance of word of mouth and people moving from one city to another, Craigslist grew much more rapidly on the west coast

²⁷ See <http://danzarrella.com/interview-with-craig-newmark-how-the-craigslist-meme-spread.html#>

²⁸ See Craigslist; On the Record: Craig Newmark,” *San Francisco Chronicle*, August 15, 2004, p. J1. <http://www.sfgate.com/cgi-bin/article.cgi?f=/c/a/2004/08/15/NEWMARK.TMP>

²⁹ The age, population density, and education variables are from the 2000 Census. Both the population and population growth variables are from the Population Estimates Program located at <http://www.census.gov/popest/estimates.html>. The percent of employed population working in the occupation classified as “information” is from the Occupational Employment Statistics (OES) program at the Bureau of Labor Statistics, <http://www.bls.gov/oes/>.

(near Craigslist's first site - the Bay Area) and in cities with a higher population density. Also consistent with Craig's remarks is that there exists a lot of MSA heterogeneity that can't be explained by demographic differences (R -squared = .39 and .41 in Columns (1) and (3), respectively).

In Columns (2) and (4), we include dummy variables for the entry year for each MSA. The results show that the timing of Craigslist entry is extremely important to Craigslist growth. MSAs into which Craigslist entered between 2000 and 2004 experienced significantly more growth between 2005 and 2007 (the rapid expansion years of Craigslist) than did MSAs which Craigslist entered in 2006. In fact, about 60% of the variation in Craigslist expansion seems to be explained by population and population growth, whether the MSA is on the Pacific Coast, and when Craigslist entered a city.

As discussed in the empirical strategy section, we employ a variety of tests to explore the relationship between Craigslist growth and our outcomes of interest. While we rely heavily on the parallel trend assumption since we do not have random assignment, the findings that we present here on the idiosyncrasies of Craigslist expansion help motivate our analysis.

III. Results

This section presents our main set of empirical results. We begin by analyzing the effect of Craigslist growth on the Help-Wanted Index. We then test whether the crowding out of print newspapers by Craigslist also improved matching efficiency. We do this by analyzing the relationship between Craigslist growth and the apartment and housing rental vacancy rate and then the unemployment rate. For each of the three main outcomes investigated below, we begin by providing graphical evidence followed by regression estimates.

Help-Wanted Index results. The Conference Board, a non-profit company, reports a monthly index of job classifieds found in print newspapers. Called the 'Help-Wanted Index', this measure of job vacancies is based on counts of the number of job classified advertisements appearing in 51 major metropolitan newspapers across the United States. The monthly counts from each of the newspapers are adjusted to account for seasonality as well as differences in the number of weekdays and Sundays across months. The adjusted figures are then normalized to a 1967 = 100 base and aggregated using nonagricultural payroll employment weights to form the national Help-Wanted Index. According to

Abraham (1987), in 1974, the index represented cities that accounted for roughly 50 percent of the total nonagricultural employment in the United States.³⁰

The expansion of online job search has raised questions regarding the validity of using print job advertisements as a measure of vacancies in the economy. Katz and Krueger (1999) first noticed a flattening of the index and hypothesized that this could be due to job classified advertisements in newspapers migrating to the Internet.³¹ To shed light on this issue, we plot the national Help-Wanted Index against a measure of job vacancies obtained from the Job Openings and Labor Turnover Survey (JOLTS). This new measure of vacancies comes from a survey of roughly 16,000 business establishments and is produced by the Bureau of Labor Statistics (BLS). Due to the nature of this measure, arguably it should not be affected by changes in the manner in which individuals search for jobs (e.g. online vs. print). Figure 2 shows that the two series track each other very closely through the end of 2003; however, starting in the middle of 2004, there is a sharp divergence between the two series. In particular, the Help-Wanted Index is trending downwards, while at the same time, the JOLTS series is trending upwards. This evidence suggests, although does not prove, that job advertisements are indeed migrating from newspapers to the Internet.

Using our identification approach, we can estimate how much of the observed deviation in Figure 2 can be explained by the growing popularity of Craigslist. We obtained the Help-Wanted Index for each of the 51 newspapers measured and matched each newspaper to its appropriate MSA.³² Table 2 provides summary statistics for the 46 MSAs in the Help-Wanted Index Sample, as well as the other samples that we use in our analysis. This table illustrates that the sample includes large MSAs (average population of 2.4 million). One can see that overall, the Help-Wanted Indexed decreased by nearly 25% over the period between 2005 and 2007, while at the same time, the average MSA in the sample experienced an increase in Craigslist job posts by 2.7 per thousand. We begin by graphing the Help-Wanted Index by quarter starting in the first quarter of 2003 through the first quarter of 2007, inclusive. Figure 3 plots this data separately for the MSAs that are in the top quartile of Craigslist growth (treatment MSAs) between 2005 and 2007 and for the MSAs that are in the bottom quartile of Craigslist growth (control MSAs). The Help-Wanted Index for the two sets of MSAs track each other quite well until just after 2005, and

³⁰ For further details on the construction of the Index, the reader is referred to Abraham and Wachter (1987).

³¹ Shimer (2005) makes the important point that although the unadjusted Help-Wanted Index may serve as a poor proxy for job vacancies, the de-trended version of the Help-Wanted Index can serve as a very good proxy for short-term demand conditions.

³² 5 of the 51 newspapers in the Help-Wanted Index contain missing observations for recent years and had to be dropped. We also dropped New Orleans since it fell to nearly zero after Hurricane Katrina. The 45 remaining MSAs were used in the analysis.

especially in 2006, at which point the number of job classifieds fall substantially for the treatment MSAs. This is our first piece of evidence that online search via Craigslist has crowded-out search via newspapers.

To provide more non-parametric evidence on how Craigslist has affected the Help-Wanted Index, Figure 4 provides a scatter plot of the raw 2007-2005 difference in the Help-Wanted Index against the increase in the number of Craigslist posts per thousand between 2005 and 2007. It is worth noting that the difference is negative for most of the MSAs. This most likely reflects the general growth in online job search that takes place at the national level. Reassuringly, Figure 4 shows that there do not seem to be any outliers driving the results, nor do there appear to be any non-linearities in the data. Thus, we are confident that the linear model we are estimating is well specified.

We measure the effect of Craigslist on the Help-Wanted Index more formally by estimating the baseline regression model specified in the empirical strategy section. This allows us to control for potential confounds and to quantify the magnitude of the effects. The results are reported in Table 3. Column (1) reports the results from estimating our baseline specification with no control variables. The dependent variable is the difference between the Help-Wanted Index in the first quarter of 2007 and the Help-Wanted Index in the first quarter of 2005. The key explanatory variable is the increase in the number of Craigslist job posts between January 2005 and April 2007. In Column (1), the estimate indicates that an increase in Craigslist job posts by one per thousand is associated with a drop in the Help-Wanted Index by 1.56 (3.74% off the base rate). The estimate is fairly stable across alternative specifications. In particular, it is robust to the inclusion of the difference in the Help-Wanted Index from 2003-2004 (Column (2)) and Log total population (Column (3)). In Column (4), we include the growth in population and in Column (5), we include the change in overall Internet usage.³³

We also report results using a two-step feasible GLS estimator in Column (6). Relative to Column (2), the equivalent OLS specification, the FGLS point estimate is slightly smaller and the standard error is smaller than the OLS standard error.

Columns (7) and (8) report the results from the instrumental variable regressions. We begin by presenting the first stage results in Table 4. Column (1) of Table 4 shows the first-stage relationship between growth in Craigslist personal posts between 2005 and 2007 with growth in Craigslist job posts. As would be expected given the raw data presented in Appendix Table 1, the relationship is very strong

³³ The point estimate is smaller in Columns (5) when we control for changes in overall Internet usage. However, nearly all of the change in the point estimate is due to the fact that five observations were dropped (for which we do not have Internet penetration data) and not because the results are not robust to the inclusion of Internet penetration as a control.

and significant. In fact, variation in personal post growth can explain approximately 66% of the variation in job post growth. Column (2) provides the first-stage results when using a spline with decile knots in months since Craigslist entry, as well as an indicator for whether or not Craigslist currently charges employers in a particular MSA to post a job advertisement. The number of months since entry is an important predictor of Craigslist growth (as can be seen by year of entry in Appendix Table 1). We also find a significant relationship between Craigslist job posts and the fee that Craigslist imposes on employers in certain MSAs. MSAs in which Craigslist charges a fee for employers to post a job advertisement have 3.36 fewer job posts per thousand than would be predicted given the other variables in the model. Overall, the first stage results suggest that the instrumental variables are highly correlated with the potentially endogenous independent variable.

In Columns (7) and (8) of Table 3, the second-stage results are presented for the effect of Craigslist job posts on the Help-Wanted Index. The point estimate when instrumenting with personal posts is slightly smaller than the OLS point estimates and is not statistically significant. On the other hand, the point estimate when instrumenting with months since entry and the indicator for whether Craigslist currently charges for a job post yields a point estimate slightly larger than the majority of the OLS point estimates and has a p-value of .09.

How do we interpret the results presented in Table 3? The average number of Craigslist job posts for the MSAs in this sample increased by 2.7 posts per thousand between 2005 and 2007. Let us take the lower bound of our estimates to be -1. Thus, our results suggest that Craigslist has reduced newspaper help-wanted advertisements by approximately 7% for the MSAs in our sample. This provides support to the claim that newspaper executives have been making; namely, that online job search has negatively impacted newspaper revenues. Our results further suggest that Craigslist is responsible for approximately 15% of the divergence of the Help-Wanted Index from the JOLTS data presented in Figure 2. In the last section, we discuss how we interpret this effect in terms of the crowd-out of newspaper classifieds.

Apartment and housing rental vacancy rate results. We now turn to testing whether or not Craigslist has improved matching efficiency in the apartment and housing rental market and the labor market. We begin with the apartment and housing rental market. There is a paucity of data available that provides even a year-to-year measure of rental vacancy rates by MSA. The data source that we are aware of that provides this information is the Census Bureau. In their Housing Vacancies and Homeownership Survey, they provide annual rental vacancy rate data for the largest 75 MSAs in the

U.S..³⁴ Due to changes in which MSAs were the largest each year and geographic boundary changes in 2005, only 57 MSAs consistently remained as part of the largest 75 MSAs and have non-missing data since 2000.³⁵

Using these 57 MSAs, we begin once again by graphing the annual rental vacancy rates across time. Figure 5 presents this graph separated by the MSAs that are in the top and bottom quartiles of Craigslist growth. As would be expected if Craigslist was improving matching efficiency in the rental market, we find that starting in 2005, treatment MSAs experienced a decrease in rental vacancies. One potential worry in this graph (relative, for example, to the Help-Wanted Index graph) is that prior to 2005, the paths of the two quartiles did not track each other perfectly. Thus, while this provides evidence in favor of the hypothesis that Craigslist increased the efficiency of these markets, it will be important to test for the robustness of these results using a variety of controls and other methods. Figure 6 provides a scatter plot of rental vacancy rate differences and Craigslist housing posts per thousand. Reassuringly, there do not appear to be any outliers. Rather, there appears to be a robust correlation between changes in the rental vacancy rates and Craigslist growth.

Table 5 provides the regression estimates that correspond to Figure 5. The dependent variable is the difference between the rental vacancy rates in 2006 and 2004. Column (1) finds that an increase in Craigslist housing posts by 1 per thousand results in a decrease in the rental vacancy rate by .198 percentage points. This relationship is statistically significant at conventional levels. Perhaps more importantly, this result is robust across several specifications. As in previous tables, we include the difference in outcomes from 2003 to 2004, Log total population, population growth and Internet penetration. We also control for changes in housing prices between 2005 and 2007 in each MSA.³⁶ When controlling for these factors, the estimates are all significant and range from -.153 to -.239. In Column (7), we report the FGLS estimates, which are roughly identical to the OLS estimates.

Table 4 provides the first stage results using the sample of 57 MSAs for which rental vacancy data is available. Similar to the Help-Wanted Index results, the instruments are jointly significant and the signs on the coefficients generally match intuition. Given the first stage results, we present the instrumental variable regressions in Columns (8) and (9) of Table 5. The coefficient on the Craigslist

³⁴ A unit is defined to be the Census to be “a house, an apartment, a group of rooms, or a single room occupied or intended for occupancy as separate living quarters.” A unit is defined to be vacant if no one is living in it at the time of the interview. A vacant *rental* unit represents a vacant unit that is being offered for rent (including units that are offered for rent and for sale). A vacant *home* unit represents a vacant unit that is being offered for sale only.

³⁵ We also drop New Orleans once again due to the dramatic change in vacancy rates that took place after Hurricane Katrina.

³⁶ The housing price index data at the MSA level were obtained from the Office of Federal Housing Enterprise Oversight.

growth variable for the specification that instruments with personal posts is slightly larger than the coefficient found using OLS, although the estimate is reduced in Column (9) when using the months since entry and Craigslist fee instruments.

The average MSA in our sample experienced a growth in Craigslist between 2005 and 2007 of 6.1 housing posts per thousand. Using the point estimate of -0.153 suggests that for the average MSA in our sample, Craigslist decreased the rental vacancy rate by approximately 1 percentage point, or by roughly 10% (1 percentage point off the base rate of 10.3%). We discuss the economic significance of this result in the discussion section below.

In Table 6, we present the results from the first of two placebo tests. Aside from rental vacancy rates, the Housing Vacancy and Homeowner survey also records information on for-sale home vacancy rates.³⁷ This allows for a unique robustness test. Since the majority of traffic on Craigslist is with apartment rentals as opposed to home sales, we would expect to find a much smaller effect of Craigslist on for-sale home vacancies than we did on apartment rental vacancies. Alternatively, evidence that for-sale home vacancies are significantly affected by Craigslist growth would cast doubt on the causal link between Craigslist growth and matching efficiency. Table 6 presents regression results using home vacancy rates as the dependent variable.³⁸ None of the coefficients on Craigslist growth in Columns (1)-(9) are statistically significant. In fact, nearly all of the coefficients are small suggesting very little correlation between Craigslist and home vacancy rates. The second placebo test that we conduct uses historical rental vacancy data. We find that Craigslist growth between 2005 and 2007 is correlated with rental vacancy declines during that same period. Perhaps, however, our estimates reflect a “spurious trend” phenomenon since we are comparing cities like the Bay Area, Austin, and Denver with cities like Cleveland, Detroit, and Cincinnati, whose apartment and housing rental markets might be fundamentally different. For example, it is possible that for reasons unrelated to the expansion of Craigslist, the rental markets in the Bay area, Austin and Denver may have experienced relatively strong demand starting in 2005. To investigate this more closely, Table 7 compares the historical rental vacancy rates of our treatment MSAs with our control MSAs. Looking at the effect of Craigslist on two-year rental vacancy differences going back to 1996, we find that the cities which aggressively took up Craigslist did not ever significantly decrease their rental vacancy rates relative to cities that did not take up Craigslist

³⁷ Rental vacancy rate and the home vacancy rate levels are correlated between MSAs in our dataset. Looking across the 57 MSAs in our sample, a 1% higher rental vacancy rate is associated with a .02% higher home vacancy rate ($p < .001$). We find no significant correlation between these two variables within MSA.

³⁸ As a comparison to the rental vacancy rate summary statistics provided in Table 2, the average home vacancy rate in our sample is 2.44% and the average difference between 2006 and 2004 is .62%.

aggressively. For one two-year period, we find a marginally significant positive effect. Overall, this evidence suggests that if the relationship between Craigslist growth and the decline in the rental vacancy rates is spurious, Craigslist growth happened to occur during the only two-year period in the last decade that would have produced such effects.

Unemployment rate results. This section examines the impact of Craigslist on the labor market. Unlike apartment and housing rental vacancy rates, unemployment rates are widely available. We obtained unemployment rate data from the Local Area Unemployment Statistics (LAUS) at the BLS. We begin by graphing a time series of the unemployment rate for the 162 MSAs that Craigslist had entered by June of 2006. Figure 7 graphs the annual unemployment rate for the top and bottom quartile of Craigslist growth. Unlike the figures for the Help-Wanted Index and rental vacancy rates, there does not appear to be a trend break starting in 2005 for our treatment MSAs relative to their counterparts.

Table 8 corresponds directly to Figure 7. The dependent variable is the difference in the unemployment rate between 2006 and 2004. Column (1) suggests that a one per thousand increase in the number of Craigslist job posts between 2005 and 2007 results in a .031 percentage point decrease in the unemployment rate for the average MSA in our sample. While this appears somewhat suggestive of an effect of Craigslist on unemployment rates, this effect goes away very quickly once other controls are included in the model. Specifically, including the difference in the unemployment rate from 2003 to 2004 causes the point estimate on Craigslist growth to become positive (although insignificant). The inclusion of more controls continues to influence the coefficient, but it always remains insignificantly different from zero. In Column (6), we report FGLS estimates. The FGLS estimates are essentially identical to the OLS estimates.

The first stage results for this sample of MSAs are presented in Columns (5) and (6) in Table 4. We find a strong correlation between personal posts and jobs posts as well as a strong positive relationship between months since Craigslist entry and Craigslist growth between 2005 and 2007. These instruments all have sufficient power in the first stage to make the second stage results potentially valid. Columns (7) and (8) in Table 8 provide the second stage results, which if anything suggest a positive correlation between Craigslist growth and unemployment rate changes. Overall, Table 8 provides no evidence that Craigslist has had a significant effect on matching efficiency in the labor market.³⁹

³⁹ For comparability to the housing rental results, we estimated the model after restricting the data to the 55 MSAs for which Craigslist entered prior to 2005. These MSAs are the larger MSAs in the U.S. and represent a sample that is very similar to those in our Help-Wanted Index and rental vacancy rate regressions. Once again, we find no evidence that a correlation between Craigslist growth and unemployment rates exists.

The precision of our estimates suggests that we could reject a small change in the unemployment rate due to Craigslist. We are able to reject that an increase in Craigslist job posts by one per thousand people has more than a .04 percentage point effect on the unemployment rate. Given that the average growth in Craigslist for the MSAs in the sample presented in Table 8 is 2.3, we can reject the hypothesis that the growth in Craigslist between 2005 and 2007 caused more than a 0.1 percentage point decrease in the unemployment rate. To interpret this magnitude, consider a Craigslist effect on the unemployment rate equal to the threshold value. Given a base unemployment rate of 4.6, this amounts to a 2.2 percent effect on the unemployment rate. We estimate that in April 2007, the number of unemployed in a typical MSA in our sample, was roughly 752 less due to the growth of Craigslist than otherwise would have been the case.⁴⁰ On the other hand, the number of Craigslist job posts in April of 2007 for a typical MSA was 2530. Thus, approximately 3.5 job posts lead to a reduction of 1 unemployed worker. In order to better interpret this magnitude, consider the impact on unemployment durations for those using Craigslist. The average median duration of unemployment is roughly 15 weeks (Chetty, 2008). We calculate that the ratio of Craigslist job posts to job openings in an MSA is approximately 14%.⁴¹ Using this number, we find that a 0.1 percentage point change in the unemployment rate implies a reduction of roughly 2 weeks in unemployment durations for those using Craigslist. Given this evidence, we view the upper bound for the effect of Craigslist on unemployment durations as being fairly modest and it appears that we can safely rule out a large effect.

Although Craigslist has no detectable effect on the overall level of unemployment, it is possible that Craigslist has impacted the market for low-skilled jobs and we are masking the effect that Craigslist has had on these workers. To test for this possibility, we collected data from The Occupational Employment Survey (OES), which publishes data on employment levels by occupation at the MSA level. Table 9 reports the results from a regression of Craigslist job posts on the fraction of workers employed in low-wage occupations. Consistent with our unemployment rate results, we find no impact of Craigslist on low-skilled employment.

IV. Discussion and Conclusion

⁴⁰ This is found by taking 2.2% of the average number of unemployed in 2004 across the 162 MSAs, 34201.

⁴¹ We calculate the number of job openings in our sample by taking the ratio of the number of employed in our sample to total US employment in 2004. We multiply this scalar (.70) by the number of job openings for the US in April 2007 according to the JOLTs (4,170,000) and then divide this by the number of MSAs in our sample (162). The ratio of Craigslist job posts per MSA (2,530) to this number is roughly 14%. Roughly speaking, 14% of jobs in an MSA are advertised on Craigslist.

Identifying the effect of the Internet on markets can be very difficult. In this paper, we exploit the rapid expansion of the website Craigslist in order to plausibly identify the effect of online search on traditional search methods and matching efficiency in the apartment and housing rental market and the labor market. Our results suggest that Craigslist to some degree has crowded out print newspaper advertisements. We also find evidence that Craigslist has increased matching efficiency in the apartment and housing rental market while having no effect on matching efficiency in the labor market.

How large are the effects that we find and are they plausible? Let's begin with the Help-Wanted Index results. We found that the growth in Craigslist between 2005 and 2007 caused the average MSA in our sample to experience a 7% decrease in help-wanted advertisements in their major local newspaper. Although the Conference Board does not release raw data, we were able to obtain a rough estimate suggesting that the current number of classified advertisements per month in the 51 major newspapers is approximately one-half million.⁴² Using our estimates, this implies that the growth in Craigslist caused the number of job classifieds in local newspapers to be approximately 35,000 less in April of 2007 than there otherwise would have been. Our self-collected Craigslist data indicates that in April of 2007, 350,000 Craigslist job posts were created in those same MSAs. This suggests that 10 job posts on Craigslist crowds out 1 job post in a local newspaper covered by the Help-Wanted Index, an effect we find very plausible. Why isn't the relationship between online and print job posts 1-to-1? There are many potential reasons. It is possible that many employers choose to post an advertisement on Craigslist, but continue to also post the advertisement in a local newspaper since it reaches a different audience. Another explanation is that, anecdotally, many of the jobs posted on Craigslist are for part-time, small projects that are not worth posting in the newspaper, but are worth posting on a free site such as Craigslist. Perhaps the key reason, however, is that Craigslist not only crowds out posts from the major newspaper in an MSA, but is also crowding out print advertisements in secondary newspapers in each market that are not reflected in our simple calibration.

Now let's turn our attention to the rental vacancy results. We first show that a sufficient number of new Craigslist housing posts were created in April 2007 to explain the reduction in vacant rental units that we observe in the data. This is intended to mainly serve as a "sanity-check" on our estimates, rather than to interpret the actual magnitude of our effect economically. Using the 2000 Census, we estimate that in April 2007, the number of rental vacancies in the 57 MSAs in our sample was 140,000 less due to

⁴² This information was obtained through an email correspondence with Ken Goldstein at the Conference Board on May 16, 2008.

the growth of Craigslist than otherwise would have been the case.⁴³ On the other hand, the number of Craigslist housing posts in April of 2007 for these MSAs was approximately 900,000, suggesting our estimates are not entirely unreasonable. In order to better gauge the magnitude of the effect size, consider the time on the market for the average rental unit of 4 months. Assume that 50% of rental vacancies in these MSAs in April 2007 are advertised on Craigslist (900,000 new posts and 1.4 million vacancies each month suggests that 50% is a reasonable guess). A 10% reduction in the rental vacancy rate is consistent with a decrease in the average vacancy duration for apartments advertised on Craigslist of approximately 3 weeks. Simply having Craigslist as a way to post listings daily (as opposed to the weekly real estate section that often comes out in Sunday papers), coupled with the benefit of being able to list more information and images and the easy search mechanism that Craigslist provides, could be sufficient to cause the reduction that we find.

A final question that remains is why Craigslist is having a large effect on rental vacancy rates yet no effect on unemployment rates. We speculate at least three reasons why this might be the case. First, Craigslist might not be that popular relative all other job search technologies. For example, a large number of jobs in the labor market are found via personal connections. There are also many other major online job sites in the labor market with which Craigslist competes for market share.⁴⁴ Second, aside from reducing the cost of posting a job vacancy, online posts in the labor market may not greatly improve information for job search relative to print. This could be because a lot of the relevant information for a job is communicated during an interview and the job post merely acts as a screening device. In the apartment rental market, on the other hand, online posts may be a significant innovation in terms of the search process. Most notably, the ability to post pictures allows apartment seekers to identify potentially good matches without having to visit each apartment.⁴⁵ Last, it could be that Craigslist is having an effect in the labor market that we are unable to detect. For example, Craigslist might not lower durations, but might improve the quality of matches.⁴⁶ Another example is that by reducing the time needed for workers to search, Craigslist reduced the relative costs of employed versus

⁴³ 140,000 is obtained by taking 10% (the average effect that we find across MSAs) of the number of rental vacancies identified by the 2000 Census in these 57 MSAs.

⁴⁴ Kuhn and Skuterad's (2004) finding that online search in general does not greatly affect matching efficiency suggests that this might not be the primary explanation.

⁴⁵ The Pew Internet and American Life Project Survey found that 51% of all Internet users have taken virtual tours when looking for a place to live (www.pewinternet.org/pdfs/PIP_Place_to_Live_2006.pdf). Craigslist uploads more than 10 million new images each month (this number includes images for non-apartment related posts).

⁴⁶ This can be seen most easily by considering a reservation-wage framework. Craigslist might cause an increase in the individual's reservation wage, since a given level of search will lead to a greater arrival of job offers.

unemployed search, exposing unemployed workers to more competition for jobs, and thus offsetting any reduction in the unemployment rate. Exploring the question of why a website like Craigslist can affect matching efficiency in the apartment and housing rental market, but not in the labor market, is interesting and could be a fruitful area of future research.

References

- Abraham, K. (1983) "Structural/Frictional vs. Deficient Demand Unemployment: Some New Evidence," *American Economic Review*, Vol 73, 708-724.
- Abraham, K. (1987) "Help-Wanted Advertising, Job Vacancies, and Unemployment," *Brookings Papers on Economic Activity*, Vol. 1987, Issue 1, 207-248.
- Autor, D. (2001) "Wiring the Labor Market," *Journal of Economic Perspectives*, Vol. 15, Issue 1 (Winter) 25-40.
- Bleakley, H. and J.C. Fuhrer (1997) "Shifting in the Beveridge Curve, Job Matching and Labor Market Dynamics," *New England Economic Review*, September-October, 3-19.
- Brown, J.R. and A. Goolsbee (2002) "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry," *Journal of Political Economy*, 110(5): 481-507.
- Chetty, R. (2008) "Moral Hazard vs. Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116(2): 173-234.
- Gentzkow, M. (2007) "Valuing New Goods in a Model with Complementarities: Online Newspapers," *American Economic Review*, 97(3), 713-744.
- Goolsbee, A. (2001) "Competition in the Computer Industry: Online versus Retail," *Journal of Industrial Economics*, 49(4), 487-499.
- Katz, L.F. and A.B. Krueger (1999) "The High-Pressure U.S. Labor Market of the 1990s," *Brookings Papers on Economic Activity*, Issue 1, 1-65.
- Kuhn, P. and M. Skuterud (2004) "Internet Job Search and Unemployment Durations," *American Economic Review*, 94(1): 218-232.
- Lewis, G. (2007) "Asymmetric Information, Adverse Selection and Seller Revelation on eBay Motors," Mimeo, Harvard University.
- Meyer, B.D. (1995) "Natural and Quasi-Experiments in Economics," *Journal of Business & Economic Statistics*, 13, 151-161.

Mortensen, D. (2000) "Panel: Modeling How Search-Matching Technologies Affect Labor Markets." Presentation to IRPP and CERF Conference on Creating Canada's Advantage in an Information Age, Ottawa, Canada.

Perez-Pena, Richard (2008) "An Industry Imperiled by Falling Profits and Shrinking Ads," New York Times online.

Pissarides, C. (2000) *Equilibrium Unemployment Theory*, MIT Press.

Saks, R. (2005) "Job Creation and Housing Construction: Constraints on Metropolitan Area Employment Growth," Working Paper in Finance and Economics Discussion Series (FEDS).

Scott Morton, F., Silva-Risso, J. and F. Zettelmeyer (2003) "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities," *Quantitative Marketing and Economics*, 1(1) 65-92.

Shimer, R. (2005) "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1): 25-49.

Sinai, T. and J. Waldfogel (2004) "Geography and the Internet: Is the Internet a Substitute or a Complement for Cities?" *Journal of Urban Economics*, 56(1), 1-24.

Sorkin, Andrew Ross (2006) "Craigslist Meets the Capitalists," New York Times online.

Stevenson, B. (2007) "The Impact of the Internet on Worker Flows" Mimeo, The Wharton School, University of Pennsylvania.

Stigler, G J (1962) "Information in the Labor Market", *Journal of Political Economy*, Vol. LXX, 94-105.

Figure 1. Average job posts across MSAs by company - early Craigslist entrants

Notes: Based on 36 MSAs for which data from Corzen.com is available for Craigslist across time (36 of the first MSAs that Craigslist entered). This figure illustrates the number of new, unique job posts (per thousand population) for each month between January, 2004 and April, 2007. Corzen.com collected job posts data are reported for Monster.com, CareerBuilder.com, Yahoo/Hotjobs, and Craigslist. A problem occurred in Corzen's data collection process that caused missing Craigslist data between April 2006 and August 2006 which can be seen in the figure.

Figure 2. Two measures of U.S. Job Vacancies, 2001Q1 - 2007Q1

Notes: This figure plots two separate measures of job vacancies across time: the JOLTS and the Help-Wanted Index. The right axis represents the Help-Wanted Index measure at the national level while the left axis represents the JOLTS measure of vacancies. Data is reported from the first quarter of 2001 to the first quarter of 2007.

Figure 3. Help-wanted index by Craigslist popularity

Notes: Based on 45 MSAs, this figure illustrates the value of the Help-Wanted Index each quarter from 2003 to 2007. The MSAs are divided into quartiles based on the growth in Craigslist job posts between January 2005 and April 2007. The top and bottom quartiles are plotted.

Figure 4. Help-wanted index difference by Craigslist growth (2007-2005)

Notes: This figure provides a scatter plot of 45 MSAs. The difference in the Help-Wanted Index between the first quarter of 2007 and the first quarter of 2005 is plotted against the growth in Craigslist job posts per thousand between January 2005 and April 2007.

Figure 5. Average rental vacancy rate by Craigslist popularity.

Notes: Based on 59 MSAs, this figure illustrates the value of the rental vacancy rate each year from 2000 to 2006. The MSAs are divided into quartiles based on the growth in Craigslist housing posts between January 2005 and April 2007. The top and bottom quartiles are plotted.

Notes: This figure provides a scatter plot of 59 MSAs. The difference in the annual rental vacancy rate between 2006 and 2004 is plotted against the growth in Craigslist housing posts per thousand between January 2005 and April 2007.

Figure 7. Average unemployment rate by Craigslist popularity.

Notes: Based on 162 MSAs, this figure illustrates the value of the unemployment rate each quarter from 2003 to 2007. The MSAs are divided into quartiles based on the growth in Craigslist job posts between January 2005 and April 2007. The top and bottom quartiles are plotted.

Table 1. Understanding Craigslist Expansion

	Dep Var: Growth in Job Posts per 1000 (2007-2005)		Dep Var: Growth in Housing Posts per 1000 (2007-2005)	
	(1)	(2)	(3)	(4)
Log Total Population	0.658 (.397)*	-1.416 (.612)**	0.875 (.729)	-2.698 (1.313)**
% Change in Population	53.5 (12.7)***	33.0 (10.7)***	116.1 (23.4)***	89.5 (22.9)***
Change in Internet Penetration	-0.041 (.089)	0.008 (.070)	-0.113 (.164)	-0.041 (.151)
Median Age	0.080 (.137)	0.052 (.112)	0.743 (.252)***	0.716 (.241)***
Population Density (pop/sqm)	0.0001 (.0003)	0.0005 (.0003)	0.0013 (.0006)*	0.0020 (.0006)***
Pacific Coast Dummy	2.73 (.95)***	2.95 (1.15)***	1.80 (1.75)	0.78 (2.46)
% in Information Industry	0.8 (36.4)	4.6 (31.2)	-33.3 (67.0)	2.2 (66.9)
% with College Degree	0.108 (.061)*	0.028 (.056)	0.352 (.111)***	0.206 (.120)*
Entry in 2000		4.17 (2.01)**		9.22 (4.32)**
Entry in 2001		9.73 (1.72)***		13.18 (3.69)***
Entry in 2002		7.34 (1.74)***		13.9 (3.74)***
Entry in 2003		4.49 (1.35)***		7.5 (2.89)**
Entry in 2004		2.41 (.949)**		5.08 (2.04)**
Entry in 2005		1.40 (.88)		3.27 (1.90)*
R-Squared	0.39	0.67	0.41	0.56
Observations	70	70	70	70

Notes: Coefficient values and robust standard errors are presented from OLS regressions of the growth in Craigslist job and housing posts per thousand (April, 2007 posts - January, 2005 posts) on a variety of MSA characteristics. Columns (2) and (4) include indicator variables for the year in which Craigslist entered each MSA (entry in 2006 is the omitted category).

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 2. Summary Statistics

	Mean	Standard Deviation
Help-Wanted Index Sample		
Help-Wanted Index (1st Quarter 2007)	41.7	35.5
Help-Wanted Index Difference (1st Quarter 2007- 1st Quarter 2005)	-9.2	13.1
Growth in Craigslist Job Posts per 1000 (2007 - 2005)	2.7	2.3
Growth in Craigslist Personal Posts per 1000 (2007 - 2005)	9.6	7.9
Internet Penetration Difference (2006 - 2004)	5.7	3.3
Total Population (in millions, 2004)	3.0	3.3
% Change in Population (2006-2004)	2.0	2.3
Obs	46	
Rental Vacancy Rate Sample		
Rental Vacancy Rate (2006)	10.3	3.9
Rental Vacancy Rate Difference (2006 - 2004)	-0.9	3.4
Growth in Craigslist Housing Posts per 1000 (2007 - 2005)	6.1	5.3
Growth in Craigslist Personal Posts per 1000 (2007 - 2005)	10.1	8.9
Internet Penetration Difference (2006 - 2004)	6.1	3.2
Housing Index Difference (2006 - 2004)	39.7	30.7
Total Population (in millions, 2004)	2.8	3.1
% Change in Population (2006-2004)	2.5	2.7
Obs	57	
Unemployment Rate Sample		
Unemployment Rate (2006)	4.6	1.3
Unemployment Rate Difference (2006 - 2004)	-0.8	0.8
Growth in Craigslist Job Posts per 1000 (2007 - 2005)	2.3	2.2
Growth in Craigslist Personal Posts per 1000 (2007 - 2005)	6.9	6.7
Internet Penetration Difference (2006 - 2004)	5.8	3.3
Total Population (in millions, 2004)	1.2	2.2
% Change in Population (2006-2004)	2.5	2.8
Obs	162	

Notes: Summary statistics are provided for the three key samples used in this study. The Help-Wanted Index data was provided by the Conference Board. Rental vacancy rate data was obtained from the Housing Vacancies and Homeownership Survey conducted by the Census Bureau. Unemployment rate statistics were taken from the Local Area Unemployment Statistics at the Bureau of Labor Statistics. The growth in Craigslist posts represents self-collected data in April of 2007 less the number of posts in January 2005 (when available using data from Corzen.com). Internet penetration data was provided by the Media Audit. Housing Index data was provided by the Office of Federal Housing Enterprise Oversight. Population estimates were obtained from www.census.gov.

Table 3. The Effect of Craigslist Usage on the Help-Wanted Index

	Dep Variable: Help-Wanted Index Difference (Quarter 1, 2007 - Quarter 1, 2005)							
	OLS					FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Growth in Craigslist Job Posts per 1000	-1.56 (.80)*	-1.62 (.80)**	-1.45 (.82)*	-1.05 (.90)	-0.93 (.97)	-1.10 (.77)	-1.27 (.92)	-1.54 (.89)*
Log Total Population			-1.41 (1.62)	-1.75 (1.62)	-1.12 (1.80)			
% Change in Population				-105.5 (61.7)*	-99.2 (77.8)			
2006-2004 Total Internet Penetration Difference					-0.36 (.51)			
Instrument: Growth in Personal Posts per 1000							X	
Instrument: Entry Timing and Fee								X
Pre-2005 Trends		X	X	X	X	X	X	X
R-Squared	0.07	0.09	0.10	0.12	0.12	0.06	0.09	0.09
Observations	45	45	45	45	40	45	45	45

Notes: Coefficient values and robust standard errors are presented from OLS (Columns (1)-(5)), FGLS (Column (6)), and IV (Columns (7) and (8)) regressions of the difference in the Help-Wanted Index (1st quarter 2007 - 1st quarter, 2005) on growth in Craigslist job posts per thousand (April, 2007 job posts - January, 2005 job posts). Log total population in 2004, percent change in total population (2006-2004), percent change in annual Internet penetration (2006 - 2004), and a pre-2005, 1-year linear time trend for each MSA are included as controls. Column (7) instruments Craigslist job posts growth with Craigslist personal posts growth. Column (8) instruments Craigslist job posts growth with for when Craigslist entered each MSA and for whether Craigslist charges a fee for job posts. First-stage results are presented in Table 4.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 4. First Stage: Instrumenting for Craigslist Job and Housing Posts

	Dep Variable: Change in Craigslist Posts (April 2007 - January 2005)					
	Job Posts Help-Wanted Index Sample		Housing Posts Rental Vacancy Rate Sample		Job Posts Unemployment Rate Sample	
	(1)	(2)	(3)	(4)	(5)	(6)
Growth in Craigslist Personal Posts per 1000	0.23 (.03)***		0.46 (.05)***		0.27 (.02)***	
Fee is Charged		-3.36 (1.39)**				-5.76 (1.11)***
Entry Timing (Spline with Decile Knot Points)		X		X		X
F-Value	40.5***	7.6***	41.5***	3.8***	175.5***	11.6***
R-Squared	0.66	0.72	0.61	0.47	0.69	0.38
Observations	45	45	57	57	161	162

Notes: Coefficient values and robust standard errors are presented from six first-stage regressions. Columns (1), (2), (5), and (6) provide first-stage results instrumenting for the growth in Craigslist job posts (April 2007 job posts - January 2005 job posts). Columns (1) and (2) provide results for the sample of MSAs used in Table 3. Columns (5) and (6) provide results for the sample of MSAs used in Table 8. Columns (3) and (4) provide first-stage results instrumenting for the growth in Craigslist housing posts (April 2007 housing posts - 0) and represent first stage results for Table 5. Two separate instrumental variable strategies are used. In Columns (1), (3), and (5), the growth in Craigslist personal posts is used as an instrument. In Columns (2), (4), and (6), the month and year of Craigslist entry into each MSA (a spline with knots at each decile point) along with an indicator for whether Craigslist collected job-posting fees (only relevant for job post regressions) are used as instruments.

*significant at 10%; ** significant at 5%; *** significant at 1%

Table 5. The Effect of Craigslist Usage on Rental Vacancy Rates

	Dep Variable: Percentage Point Difference in Annual Rental Vacancy Rates (2006 - 2004)								
	OLS						FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Growth in Craigslist House Posts per 1000 Pop.	-0.198 (.069)***	-0.209 (.068)***	-0.239 (.073)***	-0.182 (.060)***	-0.159 (.060)***	-0.153 (.065)**	-0.209 (.082)**	-0.241 (.105)**	-0.114 (.112)
Log Total Population			0.695 (.554)	0.543 (.495)	0.668 (.527)	0.719 (.580)			
% Change in Population				-34.7 (16.2)**	-33.7 (17.0)*	-45.5 (19.6)**			
Housing Index Difference 2006-2004					-0.012 (.015)	-0.016 (.016)			
2006-2004 Total Internet Penetration Difference						0.18 (.130)			
Instrument: Growth in Personal Posts per 1000								X	
Instrument: Entry Timing									X
Pre-2005 Trends		X	X	X	X	X	X	X	X
R-Squared	0.095	0.123	0.144	0.209	0.218	0.266	0.123	0.121	0.102
Observations	57	57	57	57	57	50	57	57	57

Notes: Coefficient values and robust standard errors are presented from OLS (Columns (1)-(6)), FGLS (Column (7)), and IV (Columns (8) and (9)) regressions of the difference in annual rental vacancy rates (2006 - 2004) on growth in Craigslist housing posts per thousand (April, 2007 housing posts - 0). Log total population in 2004, percent change in total population (2006-2004), change in annual housing index (2006-2004), percent change in annual Internet penetration (2006 - 2004), and a pre-2005, 1-year linear time trend for each MSA are included as controls. Column (8) instruments Craigslist job posts growth with Craigslist personal posts growth. Column (9) instruments Craigslist job posts growth with when Craigslist entered each MSA. First-stage results are presented in Table 4.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 6. The Effect of Craigslist Usage on Home Vacancy Rates - Placebo Test #1

Dep Variable: Percentage Point Difference in Annual Home Vacancy Rates (2006 - 2004)									
	OLS						FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Growth in Craigslist House Posts per 1000 Pop.	0.037 (.029)	0.035 (.029)	0.028 (.029)	0.036 (.028)	0.022 (.027)	0.038 (.027)	0.040 (.025)	-0.002 (.032)	0.028 (.029)
Log Total Population			0.158 (.170)	0.138 (.172)	0.062 (.180)	-0.071 (.187)			
% Change in Population				-4.5 (4.6)	-5.1 (4.2)	-7.5 (4.7)			
Housing Index Difference 2006-2004					0.007 (.004)	0.010 (.005)**			
2006-2004 Total Internet Penetration Difference						.023 (.036)			
Instrument: Growth in Personal Posts per 1000								X	
Instrument: Entry Timing									X
Pre-2005 Trends		X	X	X	X	X	X	X	X
R-Squared	0.039	0.048	0.060	0.073	0.113	0.160	0.056	0.008	0.046
Observations	57	57	57	57	57	50	57	57	57

Notes: Coefficient values and robust standard errors are presented from OLS (Columns (1)-(6)), FGLS (Column (7)), and IV (Columns (8) and (9)) regressions of the difference in annual home vacancy rates (2006 - 2004) on growth in Craigslist housing posts per thousand (April, 2007 housing posts - 0). Log total population in 2004, percent change in total population (2006-2004), change in annual housing index (2006-2004), percent change in annual Internet penetration (2006 - 2004), and a pre-2005, 1-year linear time trend for each MSA are included as controls. Column (8) instruments Craigslist job posts growth with Craigslist personal posts growth. Column (9) instruments Craigslist job posts growth with when Craigslist entered each MSA. First-stage results are presented in Table 4.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 7. The Effect of Craigslist Usage on Historical Rental Vacancy Rates - Placebo Test #2

	Dep Variable: Percentage Point Difference in Annual Rental Vacancy Rates (years indicated)						
	2004-2002	2003-2002	2002-2000	2001-1999	2000-1998	1999-1997	1998-1996
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Growth in Craigslist House Posts per 1000 Pop. (2007-2005)	-0.007 (.060)	0.108 (.080)	0.148 (.083)*	0.015 (.052)	-0.077 (.060)	-0.033 (.062)	-0.072 (.061)
R-Squared	0.000	0.027	0.046	0.001	0.019	0.004	0.017
Observations	57	57	57	57	57	57	57

Notes: Coefficient values and robust standard errors are presented from OLS regressions of the difference in annual rental vacancy rates on growth in Craigslist housing posts per thousand (April, 2007 housing posts - 0). Each column uses as the dependent variable a different two-year period difference in annual rental vacancy rates as indicated.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 8. The Effect of Craigslist Usage on Unemployment Rates

	Dep Variable: Percentage Point Difference in Annual Unemployment Rates (2006 - 2004)							
	OLS					FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Growth in Craigslist Job Posts per 1000 Pop.	-0.031 (.023)	0.017 (.019)	0.018 (.018)	0.038 (.025)	-0.024 (.021)	0.011 (.023)	0.047 (.026)*	-0.017 (.038)
Log Total Population			-0.002 (.045)	-0.016 (.053)	-0.073 (.075)			
% Change in Population				-6.64 (5.91)	3.18 (2.70)			
2006-2004 Total Internet Penetration Difference					0.010 (.025)			
Instrument: Growth in Personal Posts per 1000							X	
Instrument: Entry Timing and Fee								X
Pre-2005 Trends		X	X	X	X	X	X	X
R-Squared	0.01	0.19	0.19	0.24	0.29	0.19	0.17	0.18
Observations	162	162	161	70	70	162	162	162

Notes: Coefficient values and robust standard errors are presented from OLS (Columns (1)-(5)), FGLS (Column (6)), and IV (Columns (7) and (8)) regressions of the difference in annual unemployment rates (2006 - 2004) on growth in Craigslist job posts per thousand (April, 2007 job posts - January 2005 job posts). Log total population in 2004, percent change in total population (2006-2004), percent change in annual Internet penetration (2006 - 2004), and a pre-2005, 1-year linear time trend for each MSA are included as controls. Column (7) instruments Craigslist job posts growth with Craigslist personal posts growth. Column (8) instruments Craigslist job posts growth with when Craigslist entered each MSA and for whether Craigslist charges a fee for job posts. First-stage results are presented in Table 4.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 9. The Effect of Craigslist Usage on the Percent of the Employed Population Working in Low-Wage Occupations

	Dep Variable: Percentage Point Difference in Low-Wage Occupation Workers (2006 - 2004)							
	OLS					FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Growth in Craigslist Job Posts per 1000 Pop.	-0.057 (.031)*	-0.033 (.033)	-0.020 (.034)	-0.015 (.036)	-0.010 (.055)	-0.029 (.046)	-0.079 (.043)*	-0.060 (.032)*
Log Total Population			-0.070 (.095)	-0.068 (.095)	-0.017 (.171)			
% Change in Population				-2.78 (5.05)	-8.46 (7.18)			
2006-2004 Total Internet Penetration Difference					0.018 (.041)			
Instrument: Growth in Personal Posts per 1000							X	
Instrument: Entry Timing and Fee								X
Pre-2005 Trends		X	X	X	X	X	X	X
R-Squared	0.02	0.04	0.09	0.10	0.11	0.08	0.07	0.08
Observations	65	65	65	65	31	65	65	65

Notes: Coefficient values and robust standard errors are presented from OLS (Columns (1)-(5)), FGLS (Column (6)), and IV (Columns (7) and (8)) regressions of the difference in the percent of the employed population working in low-wage occupations (2006 - 2004) on growth in Craigslist job posts per thousand (April, 2007 job posts - January 2005 job posts). Log total population in 2004, percent change in total population (2006-2004), percent change in total population (2006-2004), percent change in annual Internet penetration (2006 - 2004), and a pre-2005, 1-year linear time trend for each MSA are included as controls. Column (7) instruments Craigslist job posts growth with Craigslist personal posts growth. Column (8) instruments Craigslist job posts growth with when Craigslist entered each MSA and for whether Craigslist charges a fee for job posts. First-stage results are presented in Table 4.

* significant at 10%; ** significant at 5%; *** significant at 1%

Appendix Figure 1. Average job posts across MSAs by company

Notes: Based on 121 MSAs for which data are available, this figure illustrates the number of new, unique job posts (per thousand population) for each month between January, 2004 and April, 2007. Job posts are reported for Monster.com, CareerBuilder.com, and Yahoo/Hotjobs. Also, data for Craigslist for the same 121 MSAs is reported for April, 2007. Craigslist data were self-collected using Craigslist's website. Data for Monster.com, CareerBuilder.com, and Yahoo/Hotjobs were obtained from Corzen.com.

Appendix Table 1. Craigslist Entry Dates (1995-2006) and April, 2007 Posts per 1000 Population

	Year	Month	City	MSA	State	Job	Housing	Personal
						Posts/1000	Posts/1000	Posts/1000
1	1995	March	SF Bay Area	San Francisco-Oakland-Fremont	CA	16.7	22.3	46.0
2	2000	June	Boston	Boston-Cambridge-Quincy	MA	4.0	20.8	8.8
3	2000	August	Los Angeles	Los Angeles-Long Beach-Santa Ana	CA	2.4	2.4	6.3
4	2000	August	San Diego	San Diego-Carlsbad-San Marcos	CA	2.1	6.3	10.7
5	2000	August	Washington DC	Washington-Arlington-Alexandria	DC		3.8	5.3
6	2000	August	Chicago	Chicago-Naperville-Joliet	IL	3.5	3.1	3.7
7	2000	August	New York	New York-Northern New Jersey-Long Island	NY	3.4	14.1	13.1
8	2000	August	Portland	Portland-Vancouver-Beaverton	OR	3.2	6.0	8.8
9	2000	August	Seattle	Seattle-Tacoma-Bellevue	WA	7.0	6.8	9.2
10	2000	October	Sacramento	Sacramento--Arden-Arcade--Roseville	CA	9.2	5.5	10.2
11	2001	April	Denver	Denver-Aurora	CO	11.6	15.0	26.8
12	2001	April	Atlanta	Atlanta-Sandy Springs-Marietta	GA	5.3	9.6	18.2
13	2001	April	Austin	Austin-Round Rock	TX	16.7	21.4	38.9
14	2002	October	Phoenix	Phoenix-Mesa-Scottsdale	AZ	7.7	14.5	23.4
15	2002	October	Miami	Miami-Fort Lauderdale-Pompano Beach	FL	10.1	21.3	17.6
16	2002	October	Minneapolis	Minneapolis-St. Paul-Bloomington	MN	4.1	8.5	17.8
17	2002	October	Philadelphia	Philadelphia-Camden-Wilmington	PA	4.4	6.2	12.8
18	2003	April	Detroit	Detroit-Warren-Livonia	MI	1.2	3.0	5.9
19	2003	April	Dallas	Dallas-Fort Worth-Arlington	TX	5.7	5.9	20.0
20	2003	April	Houston	Houston-Sugar Land-Baytown	TX	3.5	5.1	9.6
21	2003	November	Tampa Bay	Tampa-St. Petersburg-Clearwater	FL	2.9	9.4	9.5
22	2003	November	Honolulu	Honolulu	HI	5.8	18.9	14.6
23	2003	November	New Orleans	New Orleans-Metairie-Kenner	LA	2.1	4.7	1.1
24	2003	November	Baltimore	Baltimore-Towson	MD	4.0	6.2	9.7
25	2003	November	St. Louis	St. Louis	MO	2.0	3.5	7.1
26	2003	November	Raleigh	Raleigh-Cary	NC	7.4	10.3	26.2
27	2003	November	Las Vegas	Las Vegas-Paradise	NV	8.4	17.6	51.5
28	2003	November	Cleveland	Cleveland-Elyria-Mentor	OH	1.2	1.8	3.9
29	2003	November	Pittsburgh	Pittsburgh	PA	1.7	4.7	7.0
30	2004	January	Providence	Providence-Fall River-Warwick	RI	3.7	12.3	17.0
31	2004	February	Fresno	Fresno	CA	2.6	4.0	11.3
32	2004	February	Hartford	Hartford-West Hartford-East Hartford	CT	2.7	3.5	5.9
33	2004	February	Orlando	Orlando-Kissimmee	FL	4.2	14.7	16.8
34	2004	February	Indianapolis	Indianapolis-Carmel	IN	1.6	2.1	5.3
35	2004	February	Kansas City	Kansas City	MO	1.8	2.9	8.1
36	2004	February	Charlotte	Charlotte-Gastonia-Concord	NC	2.4	4.5	8.3
37	2004	February	Cincinnati	Cincinnati-Middletown	OH	1.5	2.7	4.8
38	2004	February	Columbus	Columbus	OH	2.3	4.5	8.2
39	2004	February	Nashville	Nashville-Davidson--Murfreesboro--Franklin	TN	2.5	5.2	9.7
40	2004	February	Norfolk	Virginia Beach-Norfolk-Newport News	VA	1.4	3.1	8.5
41	2004	February	Milwaukee	Milwaukee-Waukesha-West Allis	WI	1.0	1.3	2.9
42	2004	September	Anchorage	Anchorage	AK	7.5	16.9	17.1
43	2004	September	Santa Barbara	Santa Barbara-Santa Maria-Goleta	CA	10.5	15.5	17.5
44	2004	September	Boise	Boise City-Nampa	ID	4.6	10.2	15.9
45	2004	September	Albuquerque	Albuquerque	NM	2.6	4.1	4.0
46	2004	September	Buffalo	Buffalo-Niagara Falls	NY	1.0	2.2	4.5
47	2004	September	Memphis	Memphis	TN	1.4	1.4	3.9
48	2004	September	Salt Lake	Salt Lake City	UT	4.6	5.1	7.2
49	2004	November	Tucson	Tucson	AZ	3.2	7.7	7.1
50	2004	November	Omaha	Omaha-Council Bluffs	NE	1.8	1.8	4.5
51	2004	November	Reno	Reno-Sparks	NV	8.0	21.1	20.8
52	2004	November	Albany	Albany-Schenectady-Troy	NY	2.1	6.4	10.4
53	2004	November	Tulsa	Tulsa	OK	0.9	1.0	3.1
54	2004	November	Eugene	Eugene-Springfield	OR	5.5	10.8	18.8
55	2004	November	San Antonio	San Antonio	TX	2.6	4.1	6.9
56	2004	November	Spokane	Spokane	WA	5.8	10.3	14.4

	Year	Month	City	MSA	State	Job Posts/1000	Housing Posts/1000	Personal Posts/1000
57	2005	January	Bakersfield	Bakersfield	CA	2.1	2.2	5.7
58	2005	January	Stockton	Stockton	CA	2.8	5.5	14.8
59	2005	January	Jacksonville	Jacksonville	FL	2.1	6.7	6.4
60	2005	January	Des Moines	Des Moines-West Des Moines	IA	2.0	2.9	4.8
61	2005	January	Louisville	Louisville-Jefferson County	KY	1.1	1.1	2.8
62	2005	January	Richmond	Richmond	VA	3.0	4.6	9.7
63	2005	January	Burlington	Burlington-South Burlington	VT	6.6	21.2	19.8
64	2005	February	Columbia	Columbia	SC	1.3	2.7	3.7
65	2005	April	Birmingham	Birmingham-Hoover	AL	1.5	1.4	3.9
66	2005	April	Little Rock	Little Rock-North Little Rock-Conway	AR	1.2	1.1	2.9
67	2005	April	Modesto	Modesto	CA	3.5	5.8	16.3
68	2005	April	New Haven	New Haven	CT	4.8	11.0	10.3
69	2005	April	Wichita	Wichita	KS	0.8	1.0	2.5
70	2005	April	Lexington	Lexington-Fayette	KY	1.2	1.4	3.4
71	2005	April	Rochester	Rochester	NY	1.2	3.0	5.6
72	2005	April	Allentown	Allentown-Bethlehem-Easton	PA	1.3	2.8	3.8
73	2005	April	Madison	Madison	WI	1.2	1.6	2.8
74	2005	June	Chico	Chico	CA	4.1	13.0	13.9
75	2005	June	Redding	Redding	CA	3.5	8.1	12.4
76	2005	June	San Luis Obispo	San Luis Obispo-Paso Robles	CA	4.7	11.5	12.7
77	2005	June	Fort Myers	Cape Coral-Fort Myers	FL	3.5	24.5	9.8
78	2005	June	Tallahassee	Tallahassee	FL	2.0	5.5	5.0
79	2005	June	Champaign-Urbana	Champaign-Urbana	IL	2.2	8.6	12.3
80	2005	June	Western Mass.	Springfield	MA	2.4	5.6	7.7
81	2005	June	Ann Arbor	Ann Arbor	MI	2.4	8.0	6.5
82	2005	June	Grand Rapids	Grand Rapids-Wyoming	MI	0.8	1.8	3.3
83	2005	June	Jackson	Jackson	MS	0.9	1.1	2.2
84	2005	June	Asheville	Asheville	NC	3.2	7.4	10.2
85	2005	June	Syracuse	Syracuse	NY	0.8	2.0	4.5
86	2005	June	Dayton	Dayton	OH	0.7	1.1	2.0
87	2005	June	Oklahoma City	Oklahoma City	OK	1.2	1.4	4.2
88	2005	June	Harrisburg	Harrisburg-Carlisle	PA	1.5	2.0	6.2
89	2005	June	Charleston	Charleston-North Charleston	SC	2.6	6.3	4.8
90	2005	June	El Paso	El Paso	TX	0.6	0.4	3.2
91	2005	June	Charleston	Charleston	WV	1.0	0.4	1.2
92	2005	September	Mobile	Mobile	AL	0.7	1.2	1.9
93	2005	September	Montgomery	Montgomery	AL	1.5	1.6	4.4
94	2005	September	Pensacola	Pensacola-Ferry Pass-Brent	FL	1.3	6.3	7.6
95	2005	September	Savannah	Savannah	GA	2.4	6.5	10.2
96	2005	September	Baton Rouge	Baton Rouge	LA	1.2	1.8	3.4
97	2005	September	Shreveport	Shreveport-Bossier City	LA	0.8	0.6	3.5
98	2005	September	Toledo	Toledo	OH	0.7	1.2	2.7
99	2005	September	Knoxville	Knoxville	TN	0.9	1.7	3.2
100	2006	January	Colorado Springs	Colorado Springs	CO	3.4	6.4	7.3
101	2006	January	Gainesville	Gainesville	FL	2.5	9.9	7.7
102	2006	January	Lansing	Lansing-East Lansing	MI	0.8	2.7	3.0
103	2006	January	Santa Fe	Santa Fe	NM	3.0	9.0	5.4
104	2006	January	Medford	Medford	OR	11.4	6.2	8.0
105	2006	January	Chattanooga	Chattanooga	TN	0.8	1.3	3.1
106	2006	January	Bellingham	Bellingham	WA	5.5	18.4	16.7
107	2006	June	Huntsville	Huntsville	AL	1.1	0.8	2.5
108	2006	June	Fayetteville	Fayetteville-Springdale-Rogers	AR	1.1	1.3	2.5
109	2006	June	Flagstaff	Flagstaff	AZ	4.8	7.2	11.1
110	2006	June	Merced	Merced	CA	2.2	3.2	5.8
111	2006	June	Fort Collins	Fort Collins-Loveland	CO	4.9	10.2	9.9
112	2006	June	Daytona Beach	Daytona Beach-Ormond Beach	FL	1.0	4.8	3.4
113	2006	June	Sarasota	Sarasota-Bradenton-Venice	FL	1.5	7.6	4.7
114	2006	June	Athens	Athens-Clarke County	GA	3.4	9.0	15.3
115	2006	June	Augusta	Augusta-Richmond County	GA	0.6	0.9	1.9
116	2006	June	Macon	Macon	GA	0.8	0.6	3.2
117	2006	June	Peoria	Peoria	IL	0.8	0.6	1.7
118	2006	June	Rockford	Rockford	IL	1.1	1.0	3.4
119	2006	June	Springfield	Springfield	IL	1.3	0.7	3.6

	Year	Month	City	MSA	State	Job Posts/1000	Housing Posts/1000	Personal Posts/1000
120	2006	June	Bloomington	Bloomington	IN	2.4	3.8	6.8
121	2006	June	Evansville	Evansville	IN	0.5	0.5	2.2
122	2006	June	Fort Wayne	Fort Wayne	IN	0.6	0.7	2.4
123	2006	June	South Bend	South Bend-Mishawaka	IN	1.4	2.3	3.5
124	2006	June	Topeka	Topeka	KS	1.0	1.4	2.5
125	2006	June	Lafayette	Lafayette	LA	0.6	0.3	1.4
126	2006	June	Lake Charles	Lake Charles	LA	0.6	0.3	1.9
127	2006	June	Worcester	Worcester	MA	4.7	6.2	11.0
128	2006	June	Flint	Flint	MI	0.5	0.8	1.9
129	2006	June	Kalamazoo	Kalamazoo-Portage	MI	0.5	0.7	1.6
130	2006	June		Saginaw-Saginaw Township North	MI	0.3	0.5	1.5
131	2006	June	Duluth	Duluth	MN	1.2	1.8	3.4
132	2006	June	Columbia	Columbia	MO	1.9	3.8	4.9
133	2006	June	Springfield	Springfield	MO	0.9	1.5	2.3
134	2006	June	Gulfport-Biloxi	Gulfport-Biloxi	MS	0.7	0.9	2.0
135	2006	June	Fayetteville	Fayetteville	NC	0.9	2.1	4.2
136	2006	June	Wilmington	Wilmington	NC	2.2	10.9	7.2
137	2006	June	Lincoln	Lincoln	NE	1.5	0.8	3.8
138	2006	June	Binghamton	Binghamton	NY	1.0	1.5	2.4
139	2006	June	Utica	Utica-Rome	NY	0.4	0.5	1.7
140	2006	June	Akron-Canton	Akron	OH	0.8	0.8	2.2
141	2006	June	Youngstown	Youngstown-Warren-Boardman	OH	0.3	0.3	1.0
142	2006	June	Salem	Salem	OR	3.7	9.8	11.0
143	2006	June	Erie	Erie	PA	0.6	0.5	2.0
144	2006	June	Lancaster	Lancaster	PA	0.7	1.0	1.4
145	2006	June	Reading	Reading	PA	0.8	0.9	1.6
146	2006	June	Scranton	Scranton--Wilkes-Barre	PA	0.5	0.9	1.9
147	2006	June	Greenville	Greenville-Mauldin-Easley	SC	0.6	0.9	1.7
148	2006	June	Myrtle Beach	Myrtle Beach-Conway-North Myrtle Beach	SC	1.9	10.6	8.9
149	2006	June	Amarillo	Amarillo	TX	1.0	0.4	3.6
150	2006	June	Beaumont	Beaumont-Port Arthur	TX	0.6	0.2	1.7
151	2006	June	Brownsville	Brownsville-Harlingen	TX	0.4	0.3	1.8
152	2006	June	Corpus Christi	Corpus Christi	TX	1.2	1.7	3.1
153	2006	June	Laredo	Laredo	TX	0.8	0.3	3.3
154	2006	June	Lubbock	Lubbock	TX	1.0	1.0	2.7
155	2006	June	Mcallen	Mcallen-Edinburg-Mission	TX	0.3	0.3	0.9
156	2006	June	Odessa	Odessa	TX	0.7	0.3	1.7
157	2006	June	Waco	Waco	TX	1.2	1.0	2.8
158	2006	June	Provo	Provo-Orem	UT	3.3	5.6	2.2
159	2006	June	Charlottesville	Charlottesville	VA	3.5	10.5	8.7
160	2006	June	Roanoke	Roanoke	VA	1.1	1.6	4.2
161	2006	June	Yakima	Yakima	WA	1.3	1.5	3.6
162	2006	June	Eau Claire	Eau Claire	WI	3.5	11.5	14.4
163	2006	June	Green Bay	Green Bay	WI	2.2	4.1	7.4

Notes: This table lists all of the areas associated with an MSA into which Craigslist entered through June of 2006 along with information regarding the month and year of entry. The MSA and state to which the authors matched each city is also listed. The last three columns indicate the number of new Craigslist job, housing, and personal posts per 1000 population posted in April of 2007. These data were self-collected by the authors using the Craigslist website. Population statistics were obtained from the 2000 census.