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 National Bureau of Economic Research*

Since the seminal work of Stigler in 1962, economists have recognized that information is costly to acquire and leads to "search frictions." Growth in online search has lowered the cost of information acquisition. We analyze the expansion of the website Craigslist, which allows users to post job and housing ads. 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 house rental vacancy rate, and had no effect on the unemployment rate.

I. Introduction

The advent of the Internet has noticeably affected the ways that markets function. By providing an abundance of information at a very low

We thank David Card, Raj Chetty, Gordon Dahl, Stefano DellaVigna, Michael Greenstone, Ulrike Malmendier, Emmanuel Saez, 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 article. We also acknowledge the financial support of the Wharton Real Estate Center. Contact

[*Journal of Labor Economics*, 2014, vol. 32, no. 2]

0734-306X/2014/3202-0003$10.00

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 suggest 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 the 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 article, 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. An 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, which empirically tests whether online search improves aggregate matching.

Our empirical analysis of matching efficiency focuses on the labor and the apartment and house rental markets. The average duration for the vacancy of a rental unit in the United States is approximately 4 months.

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

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

[2] See Pérez-Peña (2008) for another example.

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 lengthy durations. The introduction of a potentially more efficient search technology can shed light on this question.

Both the labor and the apartment and house rental markets have experienced rapid growth in online search. We illustrate this growth for the labor market using a unique data set that records monthly job post counts for the major online job 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 house 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 house rental vacancies each would have evolved similarly across MSAs absent Craigslist. We present several pieces of evidence that provide support for this assumption. First, we show that our results are robust to the inclusion of several controls, notably pre-2005 trends. Second, 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

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

[4] We do not consider the growth in online apartment and house 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.

MSA.[6] In the empirical section, we discuss how this instrument can remove potential bias from our estimates. The second IV strategy exploits the fact that Craigslist entered different MSAs at different times and also 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 house rental conditions between 2005 and 2007, the time since entry into an MSA and the decision to charge for a job post 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 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 downward, while at the same time, the JOLTS series is trending upward. 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 house 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 IV 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 2005–7 is the only 2-year period in the past decade where Craigslist growth from 2005 to 2007 caused

[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.

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

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.

Third, in contrast to our findings in the apartment and house 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 house rental market but not the labor market, and we consider what we can learn from these results. One explanation focuses on the fact that our statistical test may lack power. Unlike the apartment and house 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 house rental market more responsive to improvements in information flows. In particular, search frictions might be much more important in the apartment and house rental market than in the labor market. We view our contrasting findings across these markets as a potentially interesting area for future research.

Our article 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, uses individual-level Current Population Survey (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 IV 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 article 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, Zettelmeyer, and Silva-Risso 2003) and price dispersion in the life insurance market (Brown and Goolsbee 2002).[8]

The remainder of this article is organized as follows. Section II provides a conceptual framework for understanding how a more efficient search mechanism can affect matching efficiency. Section III 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 IV. Section V provides a brief discussion of the results and concludes.

II. 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 article, we empirically test whether and to what degree Craigslist listings crowd out traditional print ads.

It is less clear how online search affects 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 article also relates to the literature regarding whether new goods are complements or substitutes to existing products. Most related to our article are Goolsbee (2001) and Sinai and Waldfogel (2004), which identify the relationship between online and offline retailing, and also Gentzkow (2007), which 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 toward 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 apartment and house rental 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 instead 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 it is well beyond the scope of our analysis.[10]

III. Institutional Background and Empirical Strategy

Craig Newmark founded the website "Craigslist" in 1995, and it then served exclusively the Bay Area. The website serves as a platform for users to post ads that primarily focus on jobs, housing, services, personals, and 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-10-visited English websites by the end of 2006.[11] While Craigslist began in the Bay Area, it expanded in 2000 to include other

[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 that 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, December 29, 2006.

major cities.[12] Appendix table A1 provides a list of the 162 MSAs that Craigslist entered through June 2006. As the table shows, Craigslist first entered into major cities, such as Boston, New York, and Los Angeles, and this was followed in time by entry into smaller MSAs. In Section IV, 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 house 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 house 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 that used by the other major job sites, which charge anywhere from $250 to $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* on December 8, 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 Craigslist 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."

A. Craigslist Post Data

The main source of the data that we use in this article is the 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

[12] Craigslist has a separate web page for each city. We use the phrase "Craigslist entered into a city" to signify that a specific web page was created for that city. The design and features of the website are 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, Los Angeles, Washington, DC, Boston, Seattle, San Diego, Portland, Sacramento, and Chicago.

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

162 MSAs for which Craigslist had entered by June 2006.[15] Using MSA population counts from the 2000 US 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 data are presented in appendix table A1. 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, and Santa Barbara all had more than 10 job posts per thousand. At the same time, other MSAs, such as Detroit, Cleveland, Indianapolis, and Cincinnati, all had fewer 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 noneconomic factors, such as social interactions, tipping points, and publicity. This motivates our IV 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 IV 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 relisted 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

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

[16] Corr(job posts, housing posts) = .78; Corr(job posts, personal posts) = .83; Corr(housing posts, personal posts) = .79.

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 of the coefficients.

B. 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 data set 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 eliminates 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 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 past several years. In January of 2004, the number of Monster, CareerBuilder, and Yahoo/Hotjob posts per thousand ranged from 0.3 on to 0.8 (see fig. A1). By April 2007, the number of posts ranged from 1 to 2.4 per thousand. This figure also illustrates how Craigslist was nearly nonexistent until the beginning of 2005 (the Bay Area is the major reason that it is not 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 2–3 years.

C. 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,

[17] Appendix fig. A1 provides the same plot for 121 MSAs for which Corzen data for Monster.com, CareerBuilder.com, and Yahoo/Hotjobs are available. However, Craigslist data are 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 fig. 1, 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.

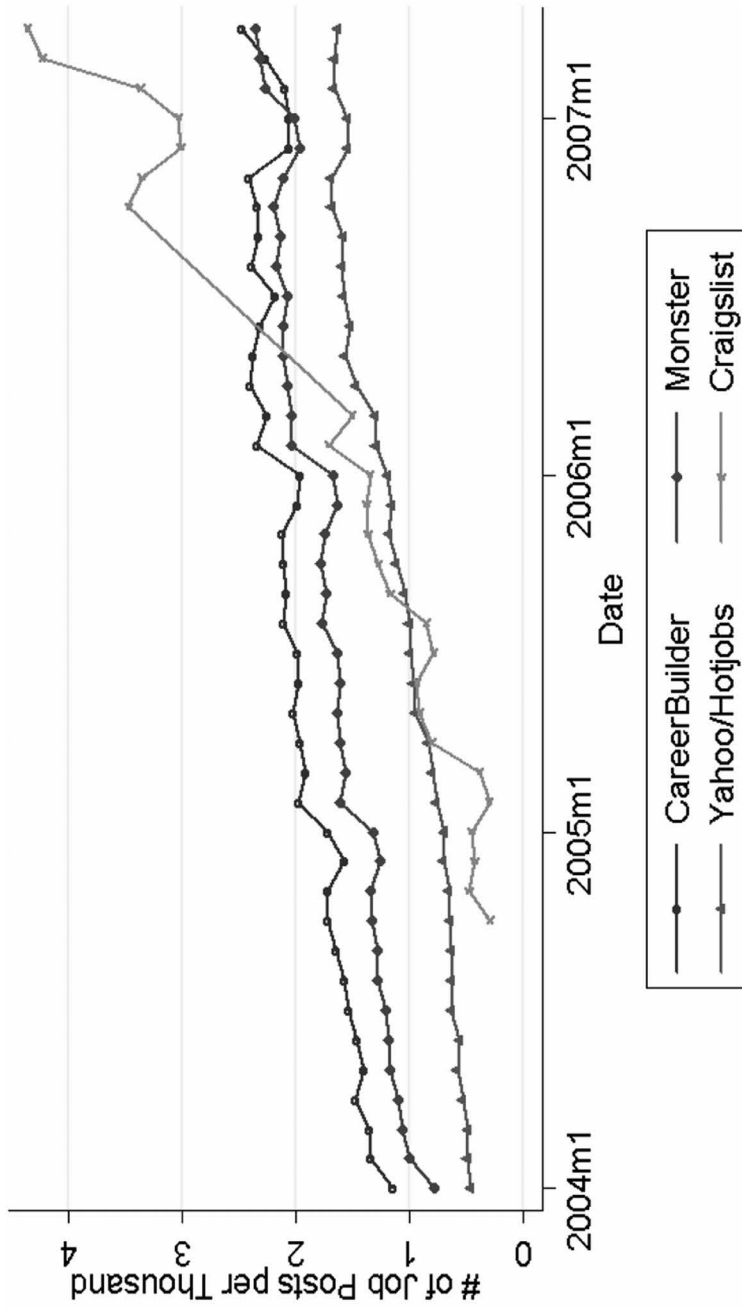

FIG. 1.—Based on 36 MSAs for which data from Corzen.com are available for Craigslist across time (36 of the first MSAs that Craigslist entered). This figure illustrates the number of new, unique job posts (per thousand population) for each month between January 2004 and April 2007. Corzen.com-collected 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. A colored version of this figure is available online.

and Yahoo/Hotjobs have all experienced increases in job posts over the past several years, but they have done so 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 A1, coupled with the information in figure 1, illustrate that over a 2-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.

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 that 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 re-

[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

gressions, we control for the change in housing prices, which can have an effect on apartment and house rentals.[20]

We also implement a feasible generalized least squares (FGLS) procedure. To do this, we estimate our baseline equation using ordinary least squares (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 regression. This potentially has the ability to improve the precision of our estimates, which is a concern given our limited sample sizes.

Finally, we explore two IV 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 apartment and house rental 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.

To address this, we propose two IV 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 cannot rule out such spillover effects, we propose a second IV 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

---

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.

[21] Craigslist charges $25 for job posts in Chicago, New York City, Los Angeles, Washington, 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.

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 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 house 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 house rental vacancy rate data to test whether there are other 2-year periods over the past decade, other than 2005-7, 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

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

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. $\Delta\text{CraigPosts}_i \equiv \text{CraigPosts}_{i,2007} - \text{CraigPosts}_{i,2005}$ measures the difference in Craigslist posts per thousand population in April of 2007 and January 2005.[25] Unfortunately, we only have

[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 this article are robust to using only time of entry as the IV that is likely to be less prone to failing to satisfy the exclusion restriction.

[23] In Sec. IV, 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 "inframarginal 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.

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

It is worth mentioning briefly two empirical approaches that we do not use in this article. First, one may wonder why we focus on a 2-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 are not available month by month. Recall that although Corzen.com records month-to-month Craigslist post data, these are only available for 36 MSAs. When we merge those MSAs with the MSAs for which outcome data are available, there are too few observations available for valid statistical inference. Furthermore, the Corzen data are 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 article 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 IV 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 adver-

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

tised on the same website. In the presence of such spillover effects, our results would overstate the impact of Craigslist and would provide a lower-bound estimate for the impact of online search overall.

D. 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 such as word of mouth. Before turning to results, anecdotal and empirical evidence can help defend this assumption and shed light on why a large variation existed in Craigslist expansion between 2005 and 2007.

According to Craigslist founder Craig Newmark when he was 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, 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 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

[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, J1. http://www.sfgate.com/cgi-bin/article.cgi?f=/c/a/2004/08/15/NEWMARK.TMP.

Table 1
Understanding Craigslist Expansion

	Dependent Variable: Growth in Job Posts per 1,000 Population (2007-2005)		Dependent Variable: Growth in Housing Posts per 1,000 Population (2007-2005)	
	(1)	(2)	(3)	(4)
Log total population	.658*	−1.416**	.875	−2.698**
	(.397)	(.612)	(.729)	(1.313)
% change in population	53.5***	33.0***	116.1***	89.5***
	(12.7)	(10.7)	(23.4)	(22.9)
Change in Internet penetration	−.041	.008	−.113	−.041
	(.089)	(.070)	(.164)	(.151)
Median age	.080	.052	.743***	.716***
	(.137)	(.112)	(.252)	(.241)
Population density (population/square mile)	.0001	.0005	.0013*	.0020***
	(.0003)	(.0003)	(.0006)	(.0006)
Pacific Coast dummy	2.73***	2.95***	1.80	.78
	(.95)	(1.15)	(1.75)	(2.46)
% in information industry	.8	4.6	−33.3	2.2
	(36.4)	(31.2)	(67.0)	(66.9)
% with college degree	.108*	.028	.352***	.206*
	(.061)	(.056)	(.111)	(.120)
Year of entry:				
2000		4.17**		9.22**
		(2.01)		(4.32)
2001		9.73***		13.18***
		(1.72)		(3.69)
2002		7.34***		13.9***
		(1.74)		(3.74)
2003		4.49***		7.5**
		(1.35)		(2.89)
2004		2.41**		5.08**
		(.949)		(2.04)
2005		1.40		3.27*
		(.88)		(1.90)
R^2	.39	.67	.41	.56

NOTE.—Number of observations = 70. Coefficient values and robust standard errors are presented from ordinary least squares 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 the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

(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.

[29] The age, population density, and education variables are from the 2000 Census. Both the population and population growth variables are from the Population

Our findings indicate that Craigslist initially entered large and growing MSAs and in MSAs with a more educated population. Consistent with 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 (near Craigslist's first site—the Bay Area) and in cities with a higher population density. Also consistent with Newmark's remarks is that there exists a lot of MSA heterogeneity that cannot be explained by demographic differences (R^2 = .39 and .41 in cols. 1 and 3, respectively).

In columns 2 and 4 of table 1, 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 that 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 (Sec. III.C), 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.

IV. 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 house rental vacancy rate and then the unemployment rate. For each of the three main outcomes investigated below, we begin by providing graphical evidence and follow this with regression estimates.

A. Help-Wanted Index Results

The Conference Board, a nonprofit 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

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 program at the Bureau of Labor Statistics, http://www.bls.gov/oes/.

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% 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 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. 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 downward, while at the same time the JOLTS series is trending upward. This evidence suggests that job advertisements migrated 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 3.0 million). One can see that overall the Help-Wanted Index decreased by nearly 25% over the period between 2005 and 2007, while at the same

[30] For further details on the construction of the index, the reader is referred to Abraham (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 detrended version of the Help-Wanted Index can serve as a very good proxy for short-term demand conditions.

[32] Five of the 51 newspapers in the Help-Wanted Index contained 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.

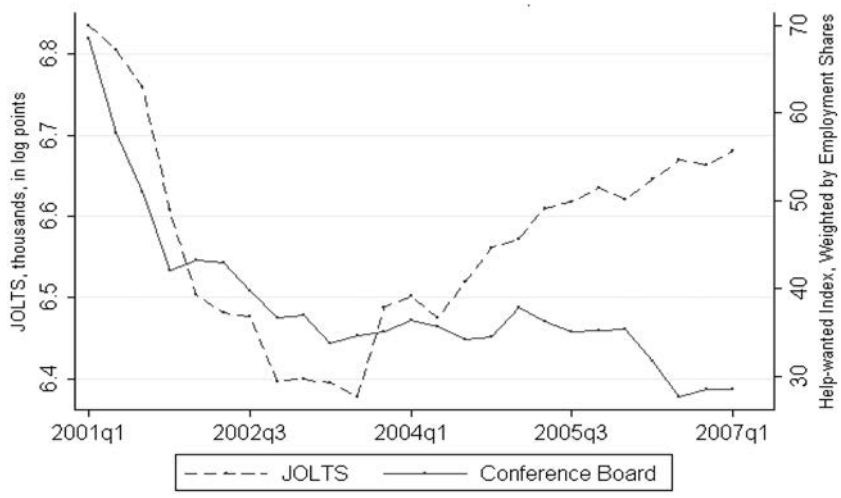

FIG. 2.—This figure plots two separate measures of job vacancies across time: the JOLTS and the Help-Wanted Index. The right axis represents the Help-Wanted Index measure at the national level, while the left axis represents the JOLTS measure of vacancies. Data are reported from the first quarter of 2001 to the first quarter of 2007. A colored version of this figure is available online.

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 and continuing through the first quarter of 2007. Figure 3 plots these 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 especially in 2006, at which point the number of job classifieds falls 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 nonparametric 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 nonlinearities in the data. Thus, we are confident that the linear model we are estimating is well specified.

Table 2
Summary Statistics

	Mean	SD
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 1,000 population (2007–2005)	2.7	2.3
Growth in Craigslist personal posts per 1,000 (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
Observations	46	
Rental vacancy rate sample:		
Rental vacancy rate (2006)	10.3	3.9
Rental vacancy rate difference (2006–2004)	−.9	3.4
Growth in Craigslist housing posts per 1,000 population (2007–2005)	6.1	5.3
Growth in Craigslist personal posts per 1, 000 population (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
Observations	57	
Unemployment rate sample:		
Unemployment rate (2006)	4.6	1.3
Unemployment rate difference (2006–2004)	−.8	.8
Growth in Craigslist job posts per 1,000 population (2007–2005)	2.3	2.2
Growth in Craigslist personal posts per 1,000 (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
Observations	162	

SOURCES.—The Help-Wanted Index data were provided by The Conference Board. Rental vacancy rate data were 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. Internet penetration data were provided by The Media Audit. Housing Index data were provided by the Office of Federal Housing Enterprise Oversight. Population estimates were obtained from http://www.census.gov.

NOTE.—Summary statistics are provided for the three key samples used in this study. 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).

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

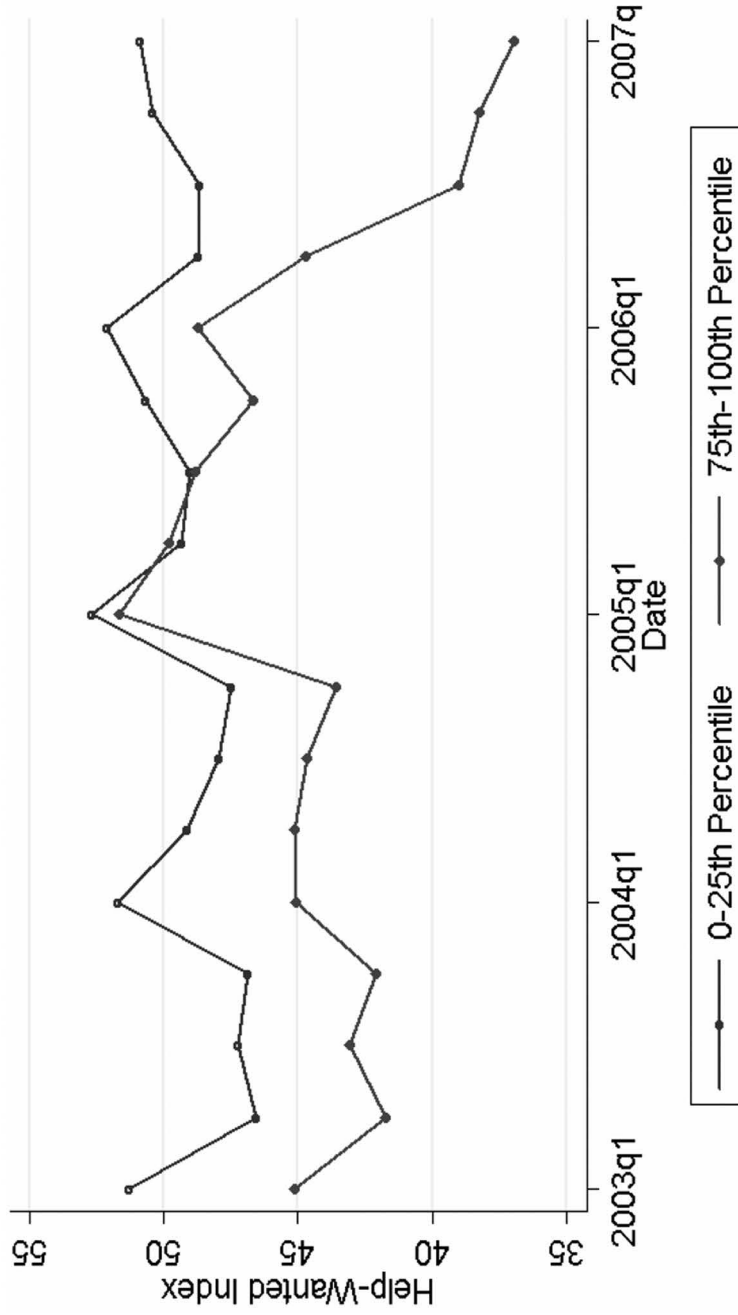

FIG. 3.—Based on 45 MSAs, this figure illustrates the value of the Help-Wanted Index each quarter from 2003 to 2007. The MSAs are divided into quartiles based on the growth in Craigslist job posts between January 2005 and April 2007. The top and bottom quartiles are plotted. A colored version of this figure is available online.

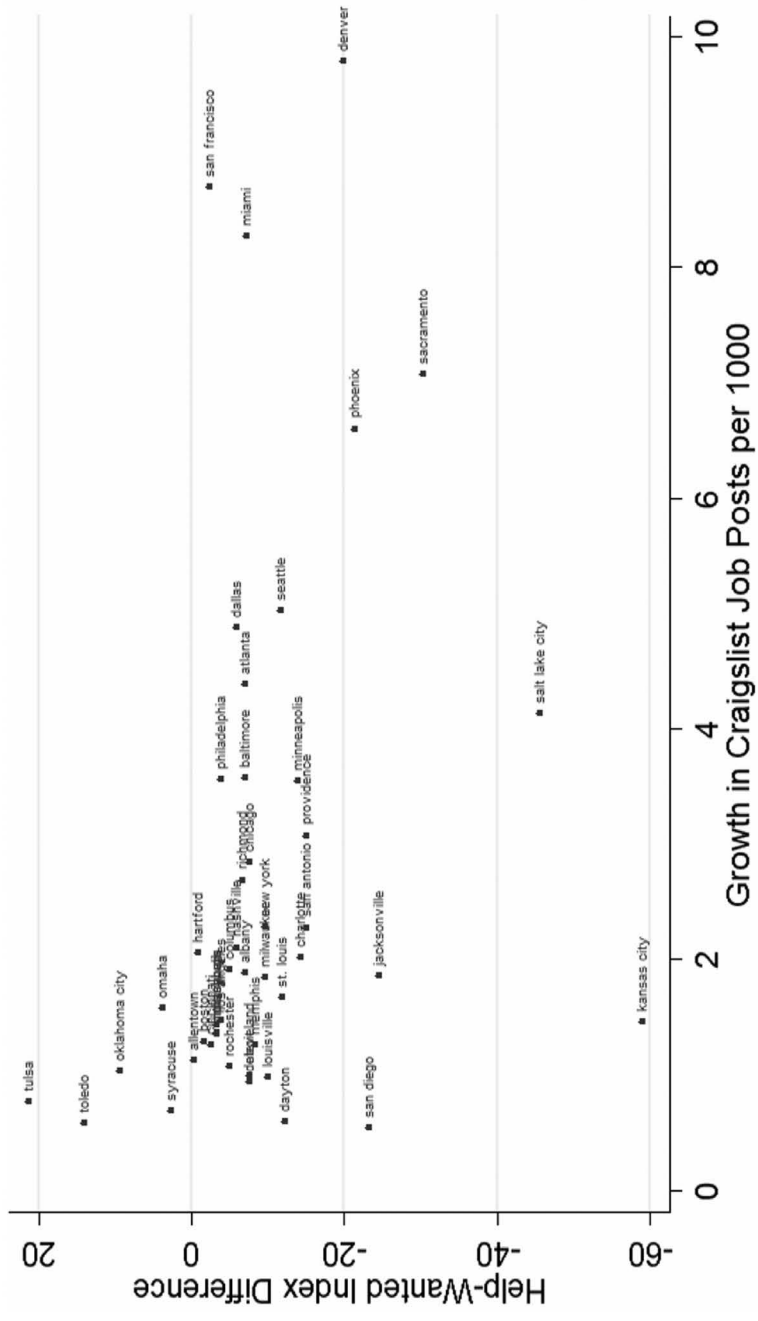

FIG. 4.—This figure provides a scatter plot of 45 MSAs. The difference in the Help-Wanted Index between the first quarter of 2007 and the first quarter of 2005 is plotted against the growth in Craigslist job posts per thousand population between January 2005 and April 2007. A colored version of this figure is available online.

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

	OLS					FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Growth in Craigslist job posts per 1,000 population	−1.56*	−1.62**	−1.45*	−1.05	−.93	−1.10	−1.27	−1.54*
	(.80)	(.80)	(.82)	(.90)	(.97)	(.77)	(.92)	(.89)
Log total population			−1.41	−1.75	−1.12			
			(1.62)	(1.62)	(1.80)			
% change in population				−105.5*	−99.2			
				(61.7)	(77.8)			
2006–2004 total Internet penetration difference					−.36			
					(.51)			
Instrument: growth in personal posts per 1,000 population							X	
Instrument: entry timing and fee								X
Pre-2005 trends		X	X	X	X	X	X	X
R^2	.07	.09	.10	.12	.12	.06	.09	.09
Observations	45	45	45	45	40	45	45	45

NOTE.—Dependent variable is Help-Wanted Index difference (Quarter 1, 2007 − Quarter 1, 2005). Coefficient values and robust standard errors are presented from ordinary least squares (OLS; cols. 1–5), feasible least squares (FGLS; col. 6), and instrumental variables (IV; cols. 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 population (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 whether Craigslist charges a fee for job posts. First-stage results are presented in table 4.

* Significant at the 10% level.

** Significant at the 5% level.

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 (col. 2) and log total population (col. 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 FGLS 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 of table 3 report the results from the IV 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 A1, the relationship is very strong 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 A1). 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 it 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

[33] The point estimate is smaller in col. 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.

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

	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 1,000 population	.23***		.46***		.27***	
	(.03)		(.05)		(.02)	
Fee is charged		−3.36**				−5.76***
		(1.39)				(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^2	.66	.72	.61	.47	.69	.38
Observations	45	45	57	57	161	162

NOTE.—Dependent variable is change in Craigslist posts (April 2007–January 2005). 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 since 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 cols. 1, 3, and 5, the growth in Craigslist personal posts is used as an instrument. In cols. 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 the 5% level.

*** Significant at the 1% level.

2.7 posts per thousand between 2005 and 2007. Let us take the lower bound of our estimates to be −1. Our results suggest that Craigslist reduced newspaper help-wanted advertisements by approximately 7% for the MSAs in our sample. This provides support for the claim that newspaper executives have been making, namely, that online job search has had a negative impact on 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 this article's final section, we discuss how we interpret this effect in terms of the crowding out of newspaper classifieds.

B. Apartment and House Rental Vacancy Rate Results

We now turn to testing whether or not Craigslist has improved matching efficiency in the apartment and house rental market and the labor market. We begin with the apartment and house rental market. There is a paucity of data available that provide even a year-to-year measure of rental vacancy rates by MSA. The data source that we are aware of that pro-

vides this information is the US Census Bureau. In its Housing Vacancies and Homeownership Survey, it provides annual rental vacancy rate data for the largest 75 MSAs in the United States.[34] Due to changes as to 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 nonmissing 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 relative to control MSAs. One potential worry in this graph (relative, e.g., 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 negative relationship 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 one 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 important, 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.

[34] A unit is defined by 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.

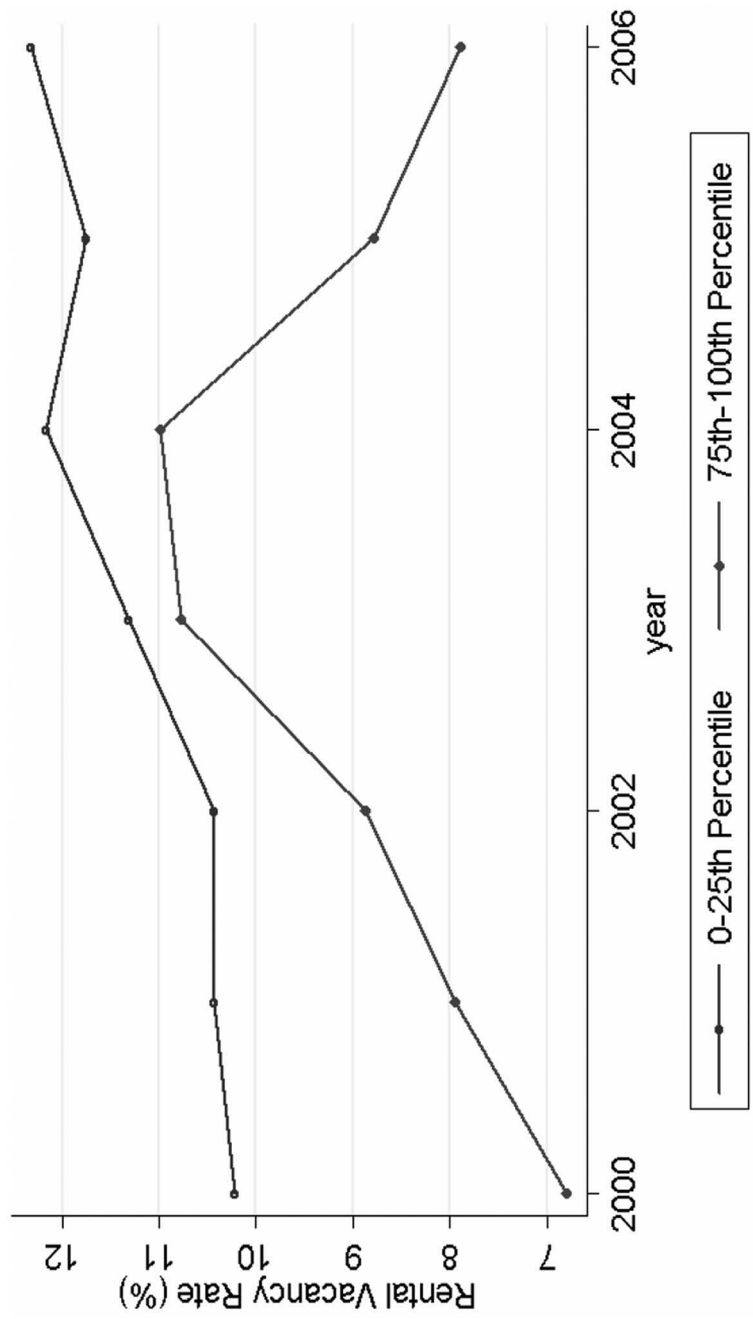

FIG. 5.—Based on 59 MSAs, this figure illustrates the value of the rental vacancy rate each year from 2000 to 2006. The MSAs are divided into quartiles based on the growth in Craigslist housing posts between January 2005 and April 2007. The top and bottom quartiles are plotted. A colored version of this figure is available online.

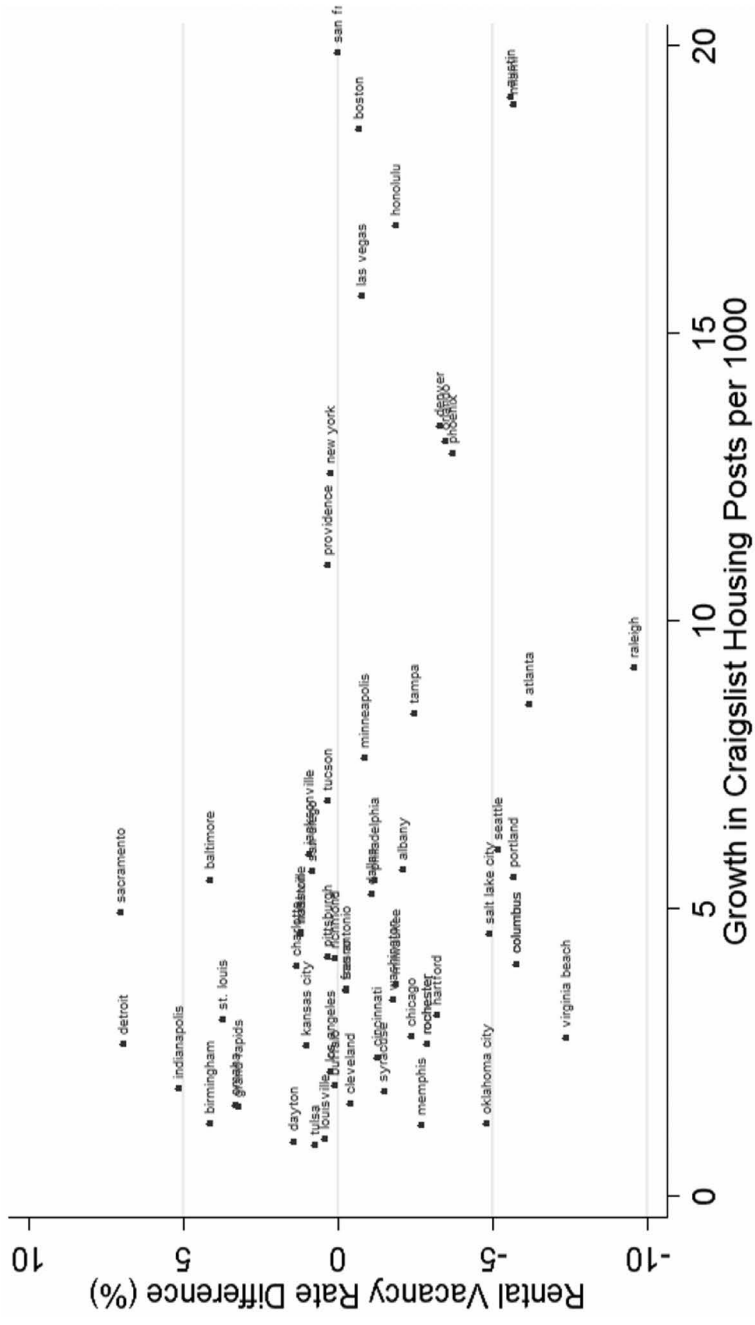

FIG. 6.—This figure provides a scatter plot of 59 MSAs. The difference in the annual rental vacancy rate between 2006 and 2004 is plotted against the growth in Craigslist housing posts per thousand population between January 2005 and April 2007. A colored version of this figure is available online.

Table 5
The Effect of Craigslist Usage on Rental Vacancy Rates

	OLS						FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Growth in Craigslist housing posts per 1,000 population	−.198***	−.209***	−.239***	−.182***	−.159***	−.153**	−.209**	−.241**	−.114
	(.069)	(.068)	(.073)	(.060)	(.060)	(.065)	(.082)	(.105)	(.112)
Log total population			.695	.543	.668	.719			
			(.554)	(.495)	(.527)	(.580)			
% change in population				−34.7**	−33.7*	−45.5**			
				(16.2)	(17.0)	(19.6)			
Housing Index difference 2006–2004					−.012	−.016			
					(.015)	(.016)			
2006–2004 total Internet penetration difference						.18			
						(.130)			
Instrument: growth in personal posts per 1,000 population								X	
Instrument: entry timing									X
Pre-2005 trends		X	X	X	X	X	X	X	X
R^2	.095	.123	.144	.209	.218	.266	.123	.121	.102
Observations	57	57	57	57	57	50	57	57	57

NOTE.—Dependent variable is percentage point difference in annual rental vacancy rates (2006 to 2004). Coefficient values and robust standard errors (in parentheses) are presented from OLS (cols. 1–6), FGLS (col. 7), and IV (cols. 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 the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

Table 4 provides the first-stage results using the sample of 57 MSAs for which rental vacancy data are 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 IV regressions in columns 8 and 9 of table 5. The coefficient on the Craigslist 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. 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 house 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,

[37] Rental vacancy rate and the home vacancy rate levels are correlated between MSAs in our data set. 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%.

Table 6
The Effect of Craigslist Usage on Home Vacancy Rates: Placebo Test Number 1

	OLS						FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Change in Craiglist housing posts per 1,000 population	.037	.035	.028	.036	.022	.038	.040	−.002	.028
	(.029)	(.029)	(.029)	(.028)	(.027)	(.027)	(.025)	(.032)	(.029)
Log total population			.158	.138	.062	−.071			
			(.170)	(.172	(.180)	(.187)			
% change in population				−4.5	−5.1	−7.5			
				(4.6)	(4.2)	(4.7)			
Housing Index difference 2006–2004					.007	.010**			
					(.004)	(.005)			
2006–2004 total Internet penetration difference						.023			
						(.036)			
Instrument: growth in personal posts per 1,000 population								X	
Instrument: entry timing									X
Pre-2005 trends		X	X	X	X	X	X	X	X
R^2	.039	.048	.060	.073	.113	.160	.056	.008	.046
Observations	57	57	57	57	57	50	57	57	57

NOTE.—Dependent variable is percentage point difference in annual home vacancy rates (2006-2004). Coefficient values and robust standard errors (in parentheses) are presented from OLS (cols. 1–6), FGLS (col. 7), and IV (cols. 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 the 5% level.

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 2-year rental vacancy differences going back to 1996, we find that the cities that aggressively took up Craigslist did not ever significantly decrease their rental vacancy rates relative to cities that did not take up Craigslist aggressively. For one 2-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 2-year period in the past decade that would have produced such effects.

C. Unemployment Rate Results

This section examines the impact of Craigslist on the labor market. Unlike apartment and house rental vacancy rates, unemployment rates are widely available. We obtained unemployment rate data from the Local Area Unemployment Statistics (LAUS) at the Bureau of Labor Statistics. 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 be-

Table 7
The Effect of Craigslist Usage on Historical Rental Vacancy Rates: Placebo Test Number 2

	Dependent Variable Is Percentage Point Difference in Annual Rental Vacancy Rates for the Years Indicated						
	2004–2002 (1)	2003–2002 (2)	2002–2000 (3)	2001–1999 (4)	2000–1998 (5)	1999–1997 (6)	1998–1996 (7)
Growth in Craigslist housing posts per 1,000 population (2007–2005)	−.007	.108	.148*	.015	−.077	−.033	−.072
	(.060)	(.080)	(.083)	(.052)	(.060)	(.062)	(.061)
R^2	.000	.027	.046	.001	.019	.004	.017

NOTE.—Observations = 57. Coefficient values and robust standard errors (in parentheses) 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 2-year period difference in annual rental vacancy rates as indicated.
* Significant at the 10% level.

tween 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]

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% 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 2,530. Thus, approximately 3.5 job posts lead to a reduction of one unemployed worker. In order to better interpret this magnitude, con-

[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 United States, and they 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.

[40] This is found by taking 2.2% of the average number of unemployed in 2004 across the 162 MSAs, 34,201.

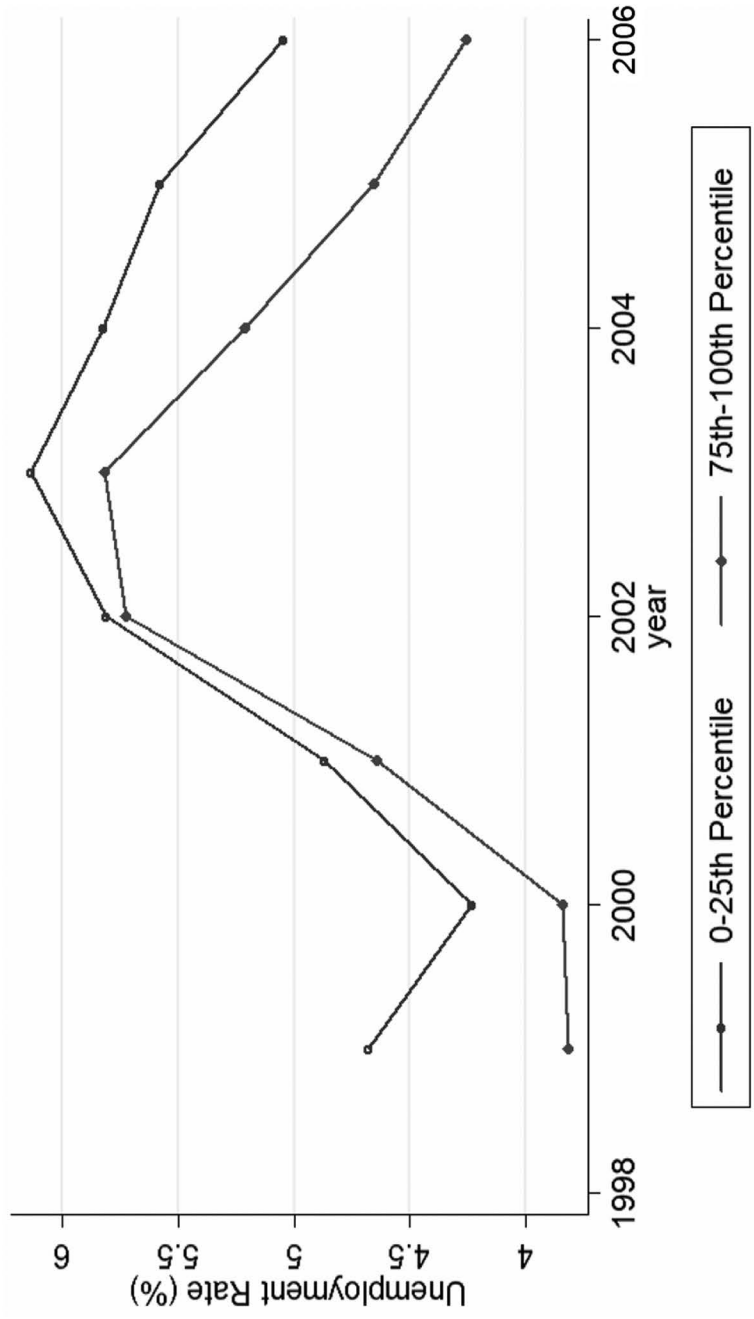

FIG. 7.—Based on 162 MSAs, this figure illustrates the value of the unemployment rate each quarter from 2003 to 2007. The MSAs are divided into quartiles based on the growth in Craigslist job posts between January 2005 and April 2007. The top and bottom quartiles are plotted. A colored version of this figure is available online.

Table 8
The Effect of Craigslist Usage on Unemployement Rates

	OLS					FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Growth in Craigslist job posts per 1,000 population	−.031	.017	.018	.038	−.024	.011	.047*	−.017
	(.023)	(.019)	(.018)	(.025)	(.021)	(.023)	(.026)	(.038)
Log total population			−.002	−.016	−.073			
			(.045)	(.053)	(.075)			
% change in population				−6.64	3.18			
				(5.91)	(2.70)			
2006–2004 total Internet penetration difference					.010			
					(.025)			
Instrument: growth in personal posts per 1,000							X	
Instrument: entry timing and fee								X
Pre-2005 trends		X	X	X	X	X	X	X
R^2	.01	.19	.19	.24	.29	.19	.17	.18
Observations	162	162	161	70	70	162	162	162

NOTE.—Dependent variable is percentage point difference in annual unemployment rates (2006-2004). Coefficient values and robust standard errors (in parentheses) are presented from OLS (cols. 1–5), FGLS (col. 6, and IV (cols. 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 the 10% level.

sider 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 affected the market for low-skilled jobs and that 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.

V. Discussion and Conclusion

Identifying the effect of the Internet on markets can be very difficult. In this article, 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 house 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 house 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 its major local newspaper. Although The Conference Board does not release raw data, we were able to obtain a rough estimate that suggests that the current number of classified advertisements per month in the 51 major newspapers

[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 United States 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.

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

	OLS					FGLS	IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Growth in Craigslist job posts per 1,000 population	−.057*	−.033	−.020	−.015	−.010	−.029	−.079*	−.060*
	(.031)	(.033)	(.034)	(.036)	(.055)	(.046)	(.043)	(.032)
Log total population			−.070	−.068	−.017			
			(.095)	(.095)	(.171)			
% change in population				−2.78	−8.46			
				(5.05)	(7.18)			
2006–2004 total Internet penetration difference					.018			
					(.041)			
Instrument: growth in personal posts per 1,000 population							X	
Instrument: entry timing and fee								X
Pre-2005 trends		X	X	X	X	X	X	X
R^2	.02	.04	.09	.10	.11	.08	.07	.08
Observations	65	65	65	65	31	65	65	65

NOTE.—Dependent variable is percentage point difference in low-wage occupation workers (2006-2004). Coefficient values and robust standard errors (in parentheses) are presented from OLS (cols. 1–5), FGLS (col. 6), and IV (cols. 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 the 10% level.

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 fewer in April of 2007 than there otherwise would have been. Our self-collected Craigslist data indicate 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 one job post in a local newspaper covered by the Help-Wanted Index, an effect that we find very plausible. Why is the relationship between online and print job posts not one-to-one? 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 also crowds out print advertisements in secondary newspapers in each market that are not reflected in our simple calibration.

Now let us 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 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 that 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 suggest 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

[42] This information was obtained through an e-mail correspondence with Ken Goldstein at The Conference Board on May 16, 2008.

[43] The figure 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.

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 to 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] Finally, 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 it 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 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 house rental market but not in the labor market is interesting and could be a fruitful area of future research.

[44] Kuhn and Skuterud'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 (http://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.

Appendix

Table A1
Craigslist Entry Dates (1995–2006) and April 2007 Posts per 1,000 Population

	Year	Month	City	MSA	State	Job Posts/ 1,000	Housing Posts/ 1,000	Personal Posts/ 1,000
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

Table A1 (*Continued*)

	Year	Month	City	MSA	State	Job Posts/ 1,000	Housing Posts/ 1,000	Personal Posts/ 1,000
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

Table A1 (*Continued*)

	Year	Month	City	MSA	State	Job Posts/ 1,000	Housing Posts/ 1,000	Personal Posts/ 1,000
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	.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

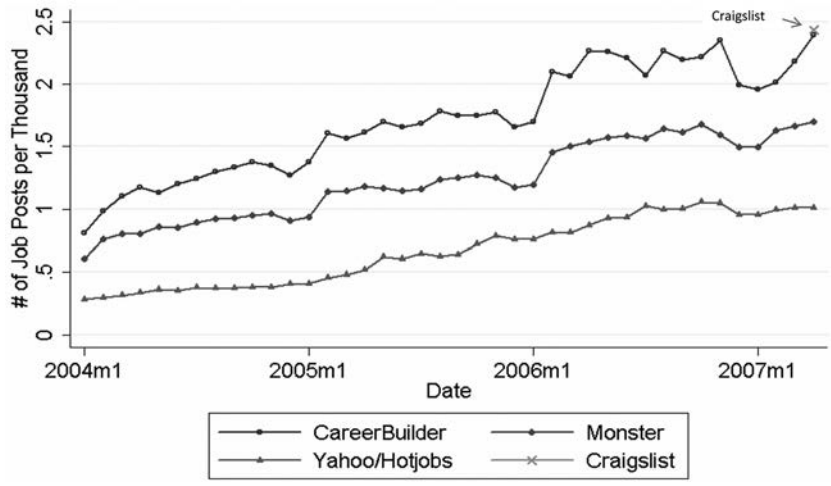

FIG. A1.—Based on 121 MSAs for which data are available, this figure illustrates the number of new, unique job posts (per thousand population) for each month between January 2004 and April 2007. Job posts are reported for Monster.com, CareerBuilder.com, and Yahoo/Hotjobs. Also, data for Craigslist for the same 121 MSAs are reported for April 2007. Craigslist data were collected by the authors using Craigslist's website. Data for Monster.com, CareerBuilder.com, and Yahoo/Hotjobs were obtained from Corzen.com. A colored version of this figure is available online.

References

Abraham, Katharine G. 1983. Structural/frictional versus deficient demand unemployment: Some new evidence. *American Economic Review* 73, no. 4:708–24.

———. 1987. Help-wanted advertising, job vacancies, and unemployment. *Brookings Papers on Economic Activity*, no. 1:207–43.

Autor, David. 2001. Wiring the labor market. *Journal of Economic Perspectives* 15, no. 1:25–40.

Bleakley, Hoyt, and Jeffrey C. Fuhrer. 1997. Shifts in the Beveridge curve, job matching, and labor market dynamics. *New England Economic Review*, September–October, 3–19.

Brown, Jeffrey R., and Austan Goolsbee. 2002. Does the Internet make markets more competitive? Evidence from the life insurance industry. *Journal of Political Economy* 110, no. 5:481–507.

Chetty, Raj. 2008. Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy* 116, no. 2:173–234.

Gentzkow, Matthew. 2007. Valuing new goods in a model with complementarities: Online newspapers. *American Economic Review* 97, no. 3: 713–44.

Goolsbee, Austan. 2001. Competition in the computer industry: Online versus retail. *Journal of Industrial Economics* 49, no. 4:487–99.
Katz, Lawrence F., and Alan B. Krueger. 1999. The high-pressure U.S. labor market of the 1990s. *Brookings Papers on Economic Activity*, no. 1, 1–65.
Kuhn, Peter, and Mikal Skuterud. 2004. Internet job search and unemployment durations. *American Economic Review* 94, no. 1:218–32.
Meyer, Bruce D. 1995. Natural and quasi-experiments in economics. *Journal of Business and Economic Statistics* 13, no. 2:151–61.
Mortensen, Dale. 2000. Panel: Modeling how search-matching technologies affect labor markets. Paper presented at the IRPP and CERF Conference on Creating Canada's Advantage in an Information Age, Ottawa, Canada, May.
Pérez-Peña, Richard. 2008. An industry imperiled by falling profits and shrinking ads. *New York Times*, February 7.
Saks, Raven E. (2005). "Job Creation and Housing Construction: Constraints on Metropolitan Area Employment Growth." Finance and Economics Discussion Series 2005-49, Board of Governors of the US Federal Reserve System, Washington, DC.
Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Risso. 2003. Consumer information and discrimination: Does the Internet affect the pricing of new cars to women and minorities? *Quantitative Marketing and Economics* 1, no. 1:65–92.
Shimer, Robert. 2005. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review* 95, no. 1:25–49.
Sinai, Todd, and Joel Waldfogel. 2004. Geography and the Internet: Is the Internet a substitute or a complement for cities? *Journal of Urban Economics* 56, no. 1:1–24.
Stevenson, Betsey. 2007. The impact of the Internet on worker flows. Unpublished manuscript, Wharton School, University of Pennsylvania.
Stigler, George. 1962. Information in the labor market. *Journal of Political Economy* 70, no. 5:94–105.