

NBER WORKING PAPER SERIES

DISABILITY INSURANCE IN THE GREAT RECESSION:
EASE OF ACCESS, PROGRAM ENROLLMENT, AND LOCAL HYSTERESIS

Melissa S. Kearney
Brendan M. Price
Riley Wilson

Working Paper 28725
<http://www.nber.org/papers/w28725>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2021

We thank Hans Ringger for excellent research assistance and Hilary Hoynes for serving as a discussant of this paper at the ASSA 2021 meetings. The views expressed in this paper are those of the authors and do not necessarily represent the views or policies of the Social Security Administration, nor those of the Board of Governors of the Federal Reserve System, nor its staff, nor the National Bureau of Economic Research. All errors are ours.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Melissa S. Kearney, Brendan M. Price, and Riley Wilson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Disability Insurance in the Great Recession: Ease of Access, Program Enrollment, and Local Hysteresis

Melissa S. Kearney, Brendan M. Price, and Riley Wilson

NBER Working Paper No. 28725

April 2021

JEL No. H53,I38,J08

ABSTRACT

Previous research has documented that Social Security Disability Insurance (SSDI) applications and awards increase during economic downturns and that expanded access to SSDI leads to a reduction in employment. We build on these insights and investigate to what extent differential access to SSDI during economic downturns leads to differential changes in SSDI enrollment and employment during the subsequent recovery. We exploit plausibly exogenous variation in SSDI appeals processing time (a measure of hassle or access) facing individuals living in ZIP codes that straddle Social Security Administration hearing office catchment borders. During the Great Recession, ZIP codes assigned to hearing offices with faster appellate processes saw a larger increase in SSDI enrollment than their cross-border neighbors. These enrollment effects are concentrated among ZIP code pairs that experienced more severe labor market downturns, and they persist as late as 2015. In the full sample, there is no clear effect of longer processing times on subsequent employment rates. However, we find some limited evidence that faster appellate processes may have weighed on the employment recovery in hard-hit ZIP codes that had high pre-recession rates of SSDI enrollment. Our findings highlight the importance of considering interaction effects between economic shocks and ease of access to the safety net.

Melissa S. Kearney
Department of Economics
University of Maryland
3105 Tydings Hall
College Park, MD 20742
and NBER
kearney@econ.umd.edu

Riley Wilson
435 CTB
Department of Economics
Brigham Young University
Provo, Utah 84602
riley_wilson@byu.edu

Brendan M. Price
Federal Reserve Board
20th Street and Constitution Ave NW
Washington, DC 20551
brendan.m.price@frb.gov

A data appendix is available at <http://www.nber.org/data-appendix/w28725>

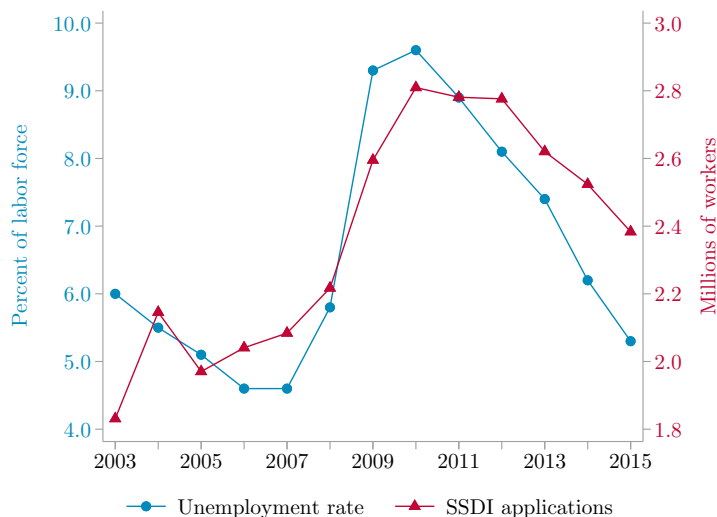
1 Introduction

The Social Security Disability Insurance (SSDI) program run by the US Social Security Administration (SSA) provides roughly \$125 billion annually in benefits to 8.4 million disabled workers, or 4.4 percent of the adult population under 65 ([Social Security Administration, 2020](#)). The program provides insurance benefits to individuals who meet work requirements and demonstrate a medical disability that is determined to preclude substantial work activity. Prior research has documented that enrollment in the program increases in response to economic downturns (e.g., [Autor and Duggan, 2003](#); [Maestas, Mullen, and Strand, 2015](#)). This is consistent with the notion of disability being somewhat fluid for a set of marginal beneficiaries who apply to the program if and when their labor market opportunities are sufficiently weak, so that enrolling in SSDI is relatively more attractive. Research has also documented that a sizable subset of SSDI beneficiaries would have worked in the years immediately following their initial SSDI application had they not been awarded benefits (e.g. [Maestas, Mullen, and Strand, 2013](#); [French and Song, 2014](#)). We build on these two insights by investigating whether an interaction effect between easier access to SSDI and an unemployment shock leads to increased program enrollment after an economic downturn and to reduced employment growth during the subsequent recovery.

We examine the context of the Great Recession in the United States. As shown in [Figure 1](#), SSDI applications rose sharply in 2009, in line with the surge in unemployment. [Maestas, Mullen, and Strand \(2018\)](#) estimate that the Great Recession induced nearly 1 million additional applicants to apply to the program, of whom 42 percent were awarded benefits (half of which were awarded on appeal). They further show that the increase was driven by applicants who were on average healthier and more equipped for the labor force. This shift in composition is consistent with the relative opportunity cost of applying to SSDI falling for less severely disabled workers when the labor market is weak. Given that SSDI is generally an absorbing state, with almost all those who enter the program staying permanently out of the workforce, the question emerges as to whether relatively easier access to the program for marginal applicants during a downturn imparts a drag to employment recovery thereafter.

This paper identifies the extent to which SSDI accessibility—as experienced through appeal wait times—amplifies or dampens the increase in SSDI enrollment in the aftermath of the Great

Figure 1: National trends in unemployment and in SSDI applications



Notes: National unemployment rate as reported by the Bureau of Labor Statistics (series LNU04000000, annual averages). SSDI applications by workers as reported by the [Social Security Administration \(2020\)](#).

Recession, as well as the extent to which it speeds or slows the subsequent recovery in employment. The general question of how SSDI accessibility affects employment rates and cyclical recovery is difficult to answer because local variation in SSDI application and enrollment rates is partly driven by factors related to local labor market conditions, some of them unobserved by the econometrician. We attempt to overcome this identification challenge by exploiting plausibly exogenous local variation in a particular element of application hassle: the length of the appellate process. Our empirical strategy takes advantage of an institutional feature of the way SSA processes SSDI appeals. Denied SSDI applicants have the option of appealing their denials to an administrative law judge (ALJ), and many do so: during the Great Recession, 28 percent of *all* SSDI applicants appealed an initial rejection to an ALJ ([Zayatz, 2015](#)). These judges hear appeals from SSA hearing offices with geographic catchment areas, to which applicants are assigned based on their ZIP code of residence. We leverage differences in average appeal processing time at the borders between hearing office jurisdictions to identify the effects of easier access to SSDI on program enrollment and employment rates in local labor markets hit more or less severely by the Great Recession.

To illustrate our empirical approach, consider two workers who live in neighboring ZIP codes that are located within the same county but in different hearing office catchment areas. Since these ZIP codes belong to the same commuting zone (by virtue of their shared county), we assume

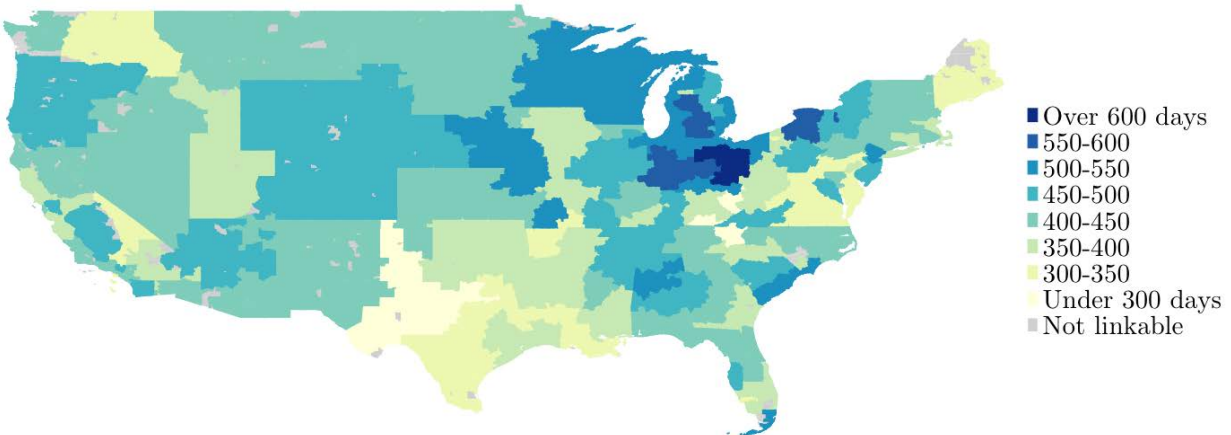
that both workers face the same labor market conditions. But if they decide to apply for SSDI, perhaps in response to the same contraction in labor demand, they will have their appeals assigned to different hearing offices and thus, in expectation, they will face different wait times. This difference in wait times amounts to differential hassle costs or ease of access and, as such, might affect a potential beneficiary’s inclination to apply, appeal, stick with the process, and ultimately be enrolled in SSDI.¹ Our strategy is to compare the evolution of SSDI enrollment and employment rates between 2003 and 2015 in neighboring ZIP codes that straddle hearing office boundaries. Our identifying assumption is that, absent cross-border differences in the ease of accessing SSDI, and conditional on fixed effects, SSDI enrollment and employment rates in these bordering ZIP codes would have responded to the Great Recession in the same way.

Key to our identification strategy, there is large variation across hearing offices in the average time elapsed from the date an SSDI appeal is filed until the date it is resolved. As shown in [Figure 2](#), some offices take an average of over 600 days to process an appeal, while others take fewer than 300. Some of this variation is driven by the composition of applications and how difficult they are to process, and some of it presumably reflects differences in administrative efficiency across offices. Our research design rests on the assumption that the variation in processing time between neighboring ZIP codes assigned to different hearing offices reflects procedural idiosyncracies uncorrelated with changes in employment outcomes except for any indirect effect coming through SSDI enrollment. Because each hearing office services many ZIP codes—150 regional hearing offices were active in 2010, compared with over 37,000 ZIP codes—any individual ZIP code will have a negligible mechanical influence on average processing time at the hearing-office level. Moreover, any difference in hearing office processing time that stems from differences in the composition of that office’s applicant pool (which might be related to latent employment outcomes) should approach zero when we look only at bordering ZIP codes.

In our first set of results, we document that one additional month of hearing office processing time—as measured in 2010, when many recession-induced appeals were being adjudicated—results in a reduction in SSDI enrollment rates following the Great Recession. In the years 2003–2008, be-

¹This notion of hassle costs affecting SSDI enrollment rates finds support in a recent paper by [Deshpande and Li \(2019\)](#), which shows that the closings of SSA field offices (where applicants receive assistance with filing disability applications) led to a persistent decline in the number of disability recipients in surrounding areas, with the decline driven primarily by increased congestion at remaining offices.

Figure 2: Average processing times for disposition of SSDI appeals at SSA hearing offices



Notes: Average processing time by hearing office as of 2010, as reported by the Social Security Administration. Processing time equals the number of days elapsed between the date a hearing is requested and the date the appeal is resolved. Colored areas represent ZIP codes linkable to a hearing office as of 2010; grey areas could not be linked.

fore the recession, there is no discernible difference in the evolution of paired ZIP codes' conditional SSDI enrollment rates as a function of average processing time at each ZIP code's assigned hearing office. In subsequent years, however, ZIP codes faced with slower appellate processes exhibit relatively slower growth in their SSDI caseloads, though the effect is only marginally significant. Our point estimates imply that an additional month of processing time leads to an SSDI enrollment rate (as a share of residents aged 30–64) that is 0.03 percentage points lower by 2015, the last year in our analysis period. With an average processing time difference of 1.8 months and a mean enrollment rate of 4.3 percent across all observations in our estimation sample, our estimates suggest that SSDI enrollment rates were persistently 0.06 percentage points (or 1.3 percent) lower in ZIP codes with longer processing times than their paired neighbors. Though our wide confidence intervals warrant caution, the long-lived relationship between processing lags and SSDI enrollment suggests that hassle costs discourage program participation following a downturn.

We next investigate how this enrollment effect varies with local labor market conditions. Here we find that the negative effects of processing time on SSDI enrollment are concentrated among ZIP code pairs in commuting zones that experienced more severe unemployment shocks during the Great Recession. In these hard-hit areas, our paired ZIP codes' SSDI enrollment rates move in parallel across hearing office borders in the lead-up to the Great Recession, then diverge sharply in its wake. The effect persists deep into the subsequent recovery: from 2008 through 2015, ZIP

codes facing longer processing times exhibit cumulative reductions in SSDI enrollment of about 0.08 percentage points (1.9 percent) relative to their paired neighbors. In other words, easier access to SSDI (through shorter processing time) leads to relative increases in program enrollment in the areas most affected by the Great Recession. In less severely hit commuting zones, by contrast, the speed of appellate review has little bearing on SSDI enrollment rates.

We then turn to the question of whether easier access to SSDI mediates the effect of the unemployment shock on subsequent ZIP-level employment. In our main sample of bordering ZIP codes, we find no clear effect of average processing time on subsequent employment rates. Though our confidence intervals preclude strong conclusions, a null relationship would be consistent with the observation that, in theory, delays in the SSDI appeal process should have offsetting effects on employment: some denied applicants are deterred from filing an appeal (and hence remain employed), while those who do so experience longer processing times (and hence incur larger scarring effects that limit their employment potential should their appeal be denied). In subsample analysis, we find qualified support for the idea that larger hassle costs may boost the employment recovery in areas that both had high SSDI enrollment rates prior to the Great Recession and experienced a more severe rise in unemployment during it. This result, though only suggestive, is consistent with the “deterrent effect” of hassle costs outweighing the “decay effect” in traditional SSDI hotspots.

This paper contributes to our understanding of how access to the safety net interacts with economic shocks of greater or less severity. We build on four existing strands of economic evidence. First, our paper is related to prior research about the link between economic downturns and SSDI enrollment (Autor and Duggan, 2003; Maestas, Mullen, and Strand, 2018). Second, our paper is related to prior research on the labor market disincentives of the SSDI program (Bound, 1989; Chen and Van der Klaauw, 2008; von Wachter, Song, and Manchester, 2011; Maestas, Mullen, and Strand, 2013; French and Song, 2014; Gelber, Moore, and Strand, 2017). Third, our paper relates to the literature on the targeting of government transfers (e.g., Nichols and Zeckhauser, 1982), especially recent work by Deshpande and Li (2019) showing that stifled SSDI access in the wake of SSA field office closures results in disproportionate screening out of potential claimants with moderately severe disabilities and low levels of education. We emphasize that, unlike Deshpande and Li (2019), our paper does not attempt a calculation of our findings’ possible welfare implications, which would depend on the relative social weights placed on encouraging work versus supporting out-of-work

individuals. We return to this point in our concluding discussion.

Fourth, our paper builds on a line of research documenting hysteresis effects in employment from local labor market shocks (e.g., [Amior and Manning, 2018](#); [Hershbein and Stuart, 2020](#)), especially recent work by [Yagan \(2019\)](#) that does so in the context of the Great Recession. Yagan estimates that, conditional on 2006 local area characteristics and fixed effects, a one percentage-point (p.p.) increase in the local unemployment shock between 2007 and 2009 led to a 0.39 p.p. reduction in the 2015 employment rate among those ages 30–49 just before the recession, a sizable persistent effect. Looking cursorily at the role of SSDI, Yagan highlights that SSDI replaced only 2 percent of lost 2015 earnings for these working-age individuals and hence played a minimal role in driving the hysteresis he finds. Our paper builds on [Yagan \(2019\)](#) in two ways. First, we look at a plausibly exogenous driver of SSDI enrollment—namely, a measure of procedural hassle. Second, we draw comparisons *within* the set of places that experienced large unemployment shocks during the Great Recession to see whether local differences in the ease of accessing SSDI muted or amplified the employment effects. Again, our main contribution is to highlight the understudied importance of potential interaction effects between economic shocks and access to the safety net.

2 Social Security Disability Insurance

We begin with a brief overview of the SSDI program, with an emphasis on the application and appeal processes as well as the program’s ramifications for disabled individuals’ incentives to participate in the labor force. For a more comprehensive overview of the SSDI program, see [Autor and Duggan \(2003, 2006\)](#); [Zayatz \(2015\)](#); [Liebman \(2015\)](#).

2.1 Program structure

The SSDI program provides earnings replacement for individuals with a medically verified disabling condition that limits their ability to engage in more than a modest amount of paid labor, referred to in program terminology as “substantial gainful activity” (SGA). In 2019, benefit outlays to disabled workers amounted to approximately \$125 billion, with an additional \$15 billion in benefits paid to dependents; the average monthly SSDI benefit for a qualifying disabled worker was about \$1,250 ([Social Security Administration, 2020](#)). For a beneficiary with median previous earnings, the

replacement rate is about 40–50 percent, with replacement rates closer to 80 percent for low-income workers (Autor and Duggan, 2003). The beneficiary population skews older, with 77 percent of recipients over the age of 50 (Social Security Administration, 2020).

SSDI caseloads have soared in recent decades. Between 1970 and 2019, the number of disabled workers receiving SSDI increased from 1.5 to 8.4 million (Social Security Administration, 2020). Growth in program enrollment has been accompanied by a change in the composition of SSDI recipients: more recipients now qualify for benefits on the basis of hard-to-verify impairments, and the program has played an increasingly important role in providing income to less-educated workers adversely affected by changes in labor demand (Liebman, 2015).

Eligibility is determined through a multi-step process in which SSA assesses whether an applicant meets both medical and non-medical program criteria. A potential beneficiary begins the process by submitting an application either online or to the local SSA field office. Applications are then sent to the field office to which the applicant’s ZIP code of residence is assigned. In an initial triage, applications are screened by field office staff to verify that the applicant is not engaging in SGA and meets a set of non-medical eligibility criteria, notably a work-history requirement.²

Applications that meet the non-medical criteria are then tested for medical eligibility. Disability examiners determine an application’s medical eligibility by verifying that the impairment is permanent or long lasting (expected to last at least 12 months or result in death) and severe enough to preclude substantial gainful activity. Disabling conditions specified in SSA’s standardized Listing of Impairments are automatically judged to satisfy this latter condition. If an applicant’s impairment does not appear on this list, the examiner must decide if it prevents the applicant from performing his or her past work-related activities, as well as any other work broadly available in the national economy to workers with the same qualifications. If an application progressing through the system does not meet these criteria, the applicant is denied a disability award. Applicants must show that they have met both the medical and non-medical criteria for a consecutive five-month waiting period before benefits are disbursed.

²In 2010 the SGA level was \$1,000 a month. Applicants must also have been employed in jobs covered by Social Security for at least five of the preceding 10 years to be eligible for SSDI. Disabled individuals without sufficient work history can file for Supplemental Security Income (SSI).

2.2 The role of hearing offices and appeals in award determination

A denied applicant can appeal the decision, first for reconsideration by a different disability examiner and then (if still denied) for a hearing through SSA’s Office of Disability Adjudication and Review (ODAR). Most applicants who meet the work history requirement are initially rejected (around 64 percent). The majority of them (56.5 percent) appeal. In 2010, over 28 percent of all applicants underwent the appeal process through ODAR and 29 percent of allowances were granted through that process (Zayatz, 2015).³ Applicants who are awarded SSDI through an appeal have less straightforward cases for receiving SSDI, and prior research has found that they are more attached to the labor force in the lead-up to application (Maestas, Mullen, and Strand, 2013).

When an individual appeals an SSDI claim denial to ODAR, the case is assigned to a hearing office based on the appellant’s ZIP code of residence. As such, individuals from two neighboring ZIP codes might have their appeals adjudicated by different hearing offices. In the 2010 fiscal year, 150 geographically demarcated hearing offices were active throughout the United States.⁴ The median hearing office served 212 ZIP codes with a median population of 1.8 million.⁵ Within a hearing office, each appeal is assigned to an administrative law judge on a rotational basis from a master docket. Judges do not specialize in types of cases, and claimants cannot choose their judge.

As seen in Figure 2, appeal processing time varies substantially across hearing offices. (We describe the underlying data in Section 4.) For appeals adjudicated in 2010, it took 424 days on average to complete the appeal process. However, some hearing offices—such as Middlesboro, Kentucky; Huntington, West Virginia; and Kingsport, Tennessee—completed the process much faster (in 270 days or less on average), while hearing offices in Lansing, Michigan; Dayton, Ohio; and Columbus, Ohio took much longer, around 600 days on average. In general, processing times were longer in the upper Midwest and shorter in Appalachia and the South Central states. In some cases, the average processing time of neighboring hearing offices might differ by as much as 50, 100,

³These numbers are right-censored as of July 2013, at which time there were still 52,245 claims pending through the appeal system. The share of awards ultimately granted on appeal was, if anything, even larger.

⁴This tally is taken from the SSA’s Hearing Office Workload Data for fiscal year 2010 and omits three hearing offices with zero dispositions as well as five National Hearing Centers (NHCs), which were opened between 2007 and 2010 to relieve appeal backlogs by hearing cases (often remotely) on behalf of local offices that are particularly congested. NHCs adjudicated only 3.1 percent of SSDI appeals in 2010, suggesting that compliance with assignment to applicants’ assigned regional hearing office was quite high during the period we analyze.

⁵Authors’ calculations using data from the US Social Security Administration and Census Bureau. These numbers are based on the subset of 148 hearing offices that we successfully match to ZIP codes with known population from the 2010 Decennial Census.

or even 200 days.

Given the frequency with which SSDI status is determined by appeals to ODAR hearing offices, potential applicants with moderate disabilities—who are more likely to face the appeals process—might take this processing time into account when deciding whether to apply, whether to appeal a denial, or whether to see a pending appeal through to its conclusion. For a potentially eligible applicant who is (by definition) not working, this additional processing time might prove costly and hence might influence the ultimate decision to apply or to instead try and find another source of income (either through work or through another transfer program). Although average processing times are probably not widely known throughout the population, potential applicants can learn about how long it takes their assigned hearing office to process appeals through consultations with field office staff, online forums, disability lawyers, or links in their social networks who have had experience with the process.⁶

2.3 Disincentives for labor force participation

After an initial trial period, SSDI beneficiaries lose program eligibility if they begin earning over the SGA level or if a continuing disability review determines that they have recovered.⁷ SSDI’s high replacement rates and strict work requirements discourage beneficiaries from taking on meaningful levels of labor force engagement. In 2014, only 9 percent of exits from SSDI were due to medical recovery or a return to employment above the SGA threshold (as opposed to death or conversion to Social Security retirement benefits), and less than one percent of SSDI recipients exited due to recovery (Zayatz, 2015). For this reason, SSDI is sometimes referred to as an “absorbing state” (Golosov and Tsyvinski, 2006).

Given the strong disincentives to work once on SSDI, there is a question of how much downward pressure SSDI exerts on employment rates. If SSDI recipients would not have worked anyway, then enrollment in the program shouldn’t have a meaningful effect on employment rates, even if recipients don’t return to work. This question has received a great deal of attention from economists. The evidence is summarized in Abraham and Kearney (2020), which reviews the relevance of a

⁶Hearing office processing times are published on the SSA website, and many disability attorney offices duplicate this information on their websites.

⁷After being awarded benefits, beneficiaries are given a nine-month trial work period during which they are allowed to work above the SGA level without losing benefits, so that they may attempt a return to the labor force.

variety of potential factors in explaining the decline in prime-age employment between 1999 and 2018.⁸ Well-identified studies of the causal link between SSDI enrollment and employment outcomes consistently find that a sizable subset of SSDI beneficiaries would have worked in the years immediately following their initial SSDI application had they not been awarded benefits.⁹

The distortionary incentives presented by SSDI might be particularly relevant during a downturn, when SSDI applications spike and less severely impaired applicants are brought into the program (Autor and Duggan, 2006; Maestas, Mullen, and Strand, 2015). For a low-income worker with a moderate impairment (such as musculoskeletal back pain), the opportunity cost of applying might fall substantially if they were laid off during a recession and have been unable to find work. Maestas, Mullen, and Strand (2015) refer to these applicants as “conditional” applicants who would prefer to work, but would apply if they lose their current job.

If economic downturns induce individuals who are on average healthier and more attached to the labor force to apply for disability insurance, this could plausibly have meaningful and persistent effects on employment after the economy has recovered. First, some of these conditional applicants may ultimately receive SSDI. Since the negative causal effect of SSDI receipt on employment is plausibly larger for conditional applicants than for severely disabled individuals who would apply for SSDI even in a tight labor market (Maestas, Mullen, and Strand, 2013), each conditional applicant who is awarded benefits represents a comparatively large reduction in employment. Second, even conditional applicants whose applications (and appeals) are denied have effectively removed themselves from the labor force while their claims are pending. This temporary withdrawal might result in human capital (or social/network capital) depreciation, making it more difficult for de-

⁸Based on the point estimates from existing studies, combined with data on age-group-specific growth in SSDI caseloads, Abraham and Kearney (2020) conclude that SSDI was a modest contributor to the decline in employment rates between 1999 and 2018, accounting for a roughly 0.09 percentage point decline in the 2018 employment-to-population ratio, as compared to a 3.8 percentage point total net reduction. This makes SSDI much less important to the aggregate employment-to-population decline than increased import competition from China and the adoption of industrial robots, which the authors estimate accounted for 0.93 and 0.43 percentage point declines, respectively.

⁹Maestas, Mullen, and Strand (2013) use administrative data to match SSDI applications to disability examiners and exploit variation in examiners’ allowance rates as an instrument for benefit receipt. Their IV approach yields the finding that, among the nearly 23 percent of applicants on the margin of program entry (meaning that their award determination depends on the leniency of the examiner), employment rates would have been 28 percentage points higher two years after initial award had they not received benefits. The estimated effect ranges from no effect for applicants with the most severe conditions to 50 percentage points for applicants with the least severe conditions. A similar finding emerges from the work of French and Song (2014), who use variation in allowance rates across administrative law judges in the second stage of the appeals process to estimate the labor supply effect of SSDI receipt. They find that the employment rate of applicants granted benefits at this stage would have been 26 percentage points higher three years after a decision had they not been granted SSDI benefits. These studies build on earlier papers by Bound (1989); Chen and Van der Klaauw (2008); von Wachter, Song, and Manchester (2011).

nied appellants to find a job once the labor market improves (Autor et al., 2015). As such, we might expect easier access to SSDI to slow the recovery in employment in areas that experienced a relatively deep unemployment shock during the Great Recession.¹⁰

3 Empirical Strategy

Because the disability program is federally administered and fairly homogeneous across the country, it is challenging to evaluate how (or if) SSDI accessibility propagates employment effects of local labor market downturns. To overcome this challenge, we compare outcomes between neighboring ZIP codes located in the same county but assigned to different hearing offices and thus subject to different processing times at the appellate stage. Because appellants must refrain from engaging in substantial gainful activity while their appeals are pending, prolongation of the appellate process can impose significant costs on SSDI applicants, especially those with greater work capacity who may be on the margin of choosing to participate in the labor market. Differences in processing time should therefore influence application and appeal decisions to the extent that potential applicants account for the costs of additional delay, which in some cases could result in up to seven additional months of waiting beyond what their cross-border neighbors experience.

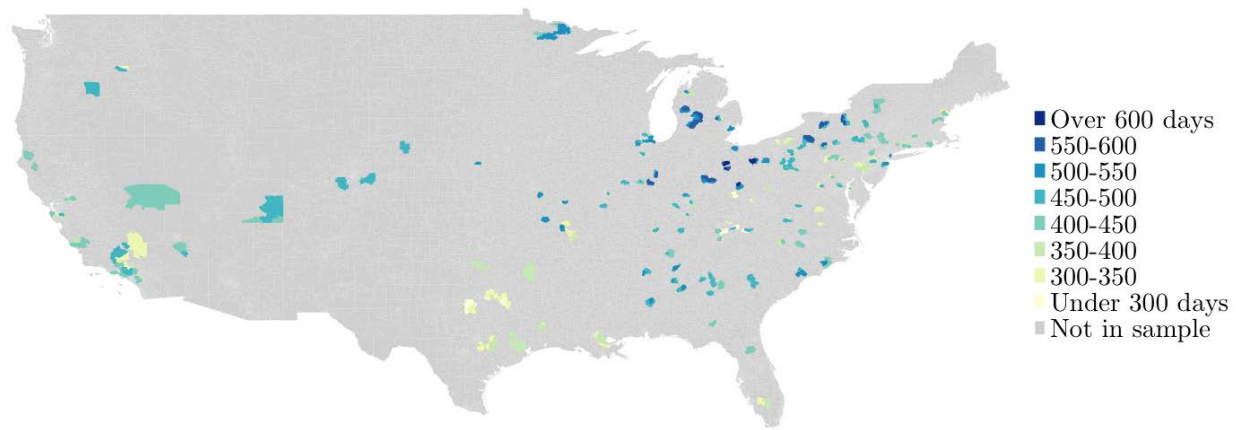
3.1 Identifying variation and sample definition

Variation in processing times across hearing offices likely stems from differences in the prevalence of disability, economic conditions, and local attitudes towards SSDI—all of which are potentially correlated with employment opportunities and labor supply decisions—as well as from idiosyncratic factors like judges’ efficiency or office staffing levels. To isolate variation that is plausibly exogenous to our outcomes of interest, we focus on neighboring ZIP codes that straddle the borders between adjacent hearing office catchment areas, so as to compare localities home to similar populations faced with similar labor market conditions but subject to differences in appeal duration.

Since hearing offices serve large geographic areas containing many ZIP codes, the average processing time of a given hearing office is unlikely to be driven by any one ZIP code, but instead by

¹⁰A couple of notable studies examine what happens to SSDI applications and enrollment when the economy improves. The general finding is that people are less likely to apply and fewer people enroll in SSDI when the local economy is doing well (Black, Daniel, and Sanders, 2002; Charles, Li, and Stephens Jr., 2018).

Figure 3: Average processing times in our estimation sample of paired cross-border ZIP codes



Notes: Average processing time is defined as in Figure 2. The colored regions are those included in our estimation sample of neighboring ZIP codes located in the same county but assigned to different hearing offices.

the characteristics of applicants throughout its catchment area. To the extent that households and firms do not sort discontinuously across hearing office boundaries, neighboring ZIP codes straddling such a boundary should be similar in population characteristics, labor market conditions, and industry composition. Because people and firms might selectively sort across these boundary lines when they also demarcate other government jurisdictions, we restrict our sample to neighboring ZIP codes that belong to the same county but are assigned to different hearing offices.¹¹ For brevity, we refer to such same-county paired ZIP codes simply as “border pairs”.

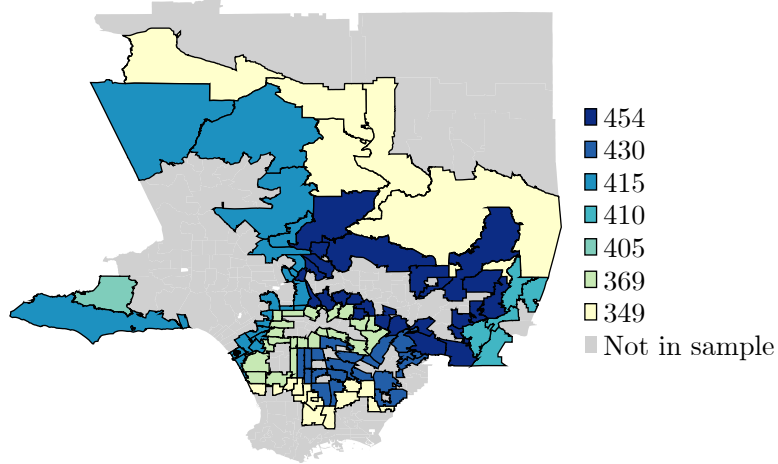
In addition to these restrictions, we also restrict the sample to a balanced panel with non-missing values for all key variables, as detailed in Appendix B. This results in a sample of 1,049 distinct ZIP codes forming 907 border pairs located in 159 counties across the country. As shown in Figure 3, our border pairs are spread throughout the country but concentrated in more densely populated areas, where counties typically contain larger numbers of ZIP codes and thus more neighboring pairs that are candidates for our analysis sample.¹² As an example of our identifying variation, Figure 4 shows the ZIP codes that meet our criteria in Los Angeles County, California.¹³ Even within Los Angeles, there is substantial variation in average processing time, with some ZIP codes facing 349 days on average, while applicants from a neighboring ZIP code might expect to

¹¹For example, if the hearing office boundary coincides with a state or county border, individuals and firms might face differences in tax policy or in the quality of publicly provided services.

¹²Data from the US Department of Agriculture’s 2003 Rural-Urban Continuum Codes show that 81.4 percent of ZIP codes in our border-pair sample are in metro areas, as compared to 56.2 percent of ZIP codes nationwide.

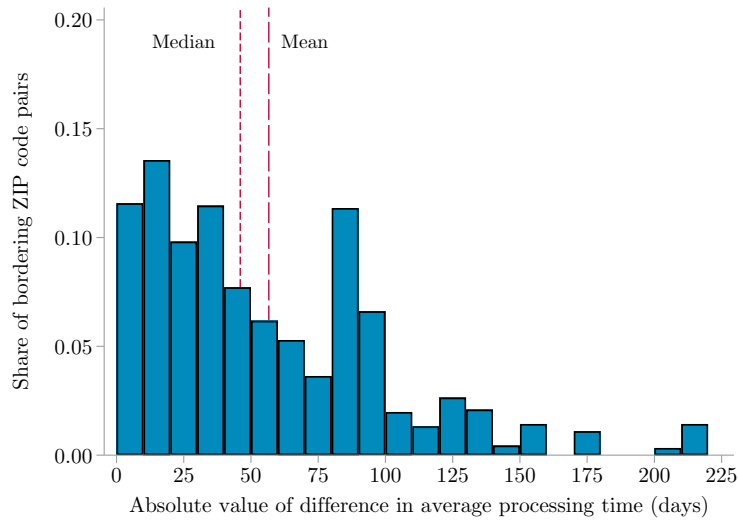
¹³Appendix Figure A1 presents a similar map for Cook County, Illinois (home to the city of Chicago).

Figure 4: Adjacent ZIP codes in Los Angeles County assigned to different hearing offices



Notes: Average processing times by hearing office as of 2010, as reported by the Social Security Administration and plotted for ZIP codes in Los Angeles, California, with one or more neighbors assigned to a different hearing office.

Figure 5: Distribution of cross-border differences in average processing time within our sample



Notes: Average processing times by hearing office as of 2010, as reported by the Social Security Administration. The sample consists of our neighboring ZIP codes located in the same county but assigned to different hearing offices.

wait over 100 days longer. Among all border pairs, the difference in average wait times varies from 0 to 219 days with an unweighted median of 46 days and a mean of 57 days (see [Figure 5](#)).

3.2 Empirical specifications

We investigate how average processing time at the appellate stage affects SSDI enrollment rates and employment rates after the Great Recession. To do so, we estimate the following equations:

$$Y_{zpt} = \sum_{\tau=2003}^{2015} \beta_{\tau} \cdot \text{Delay}_{o(z)} \cdot \mathbb{1}\{t = \tau\} + \delta_{pt} + \phi_z + \varepsilon_{zpt} \quad (1)$$

$$Y_{zpt} = \beta_{SR} \cdot \mathbb{1}\{2009-11\}_t \cdot \text{Delay}_{o(z)} + \beta_{LR} \cdot \mathbb{1}\{2012-15\}_t \cdot \text{Delay}_{o(z)} + \delta_{pt} + \phi_z + \varepsilon_{zpt} \quad (2)$$

The level of observation is the ZIP code (z) by border pair (p) by year (t), reflecting the fact that a given ZIP code may be matched with multiple neighbors.¹⁴ The subscript $o(z)$ denotes the hearing office to which ZIP code z was assigned as of 2010.

The outcome variable Y is alternately defined as the SSDI enrollment rate or employment rate for adults aged 30–64. In Equation (1), the coefficients of interest are the vector β_{τ} , which trace out the impact of average appellate processing time ($\text{Delay}_{o(z)}$)—as measured in 2010 and expressed in months—over the years preceding and following the Great Recession. The year 2008 is excluded from these interactions to serve as the reference year. Equation (2) replaces the year-by-year β_{τ} coefficients with pooled coefficients β_{SR} and β_{LR} capturing, respectively, the “short-run” years 2009–2011 and the “long-run” years 2012–2015. While we chose this particular parametrization simply as a convenient means of summarizing the patterns we observe using our more flexible year-by-year specification, the designation of 2009–2011 as a transitional phase accords well with the timing of the surge in aggregate SSDI applications and the delays typically encountered in the application and appellate processes.

In both equations, the inclusion of border pair \times year fixed effects (δ_{pt}) controls for common time shocks to a ZIP code pair; partialing out these fixed effects means that the estimated β coefficients capture differences in outcomes between two neighboring paired ZIP codes in the same year. The inclusion of ZIP code fixed effects (ϕ_z) controls for any time-invariant characteristics of a ZIP code. The identifying assumption for Equations (1) and (2) to yield causal estimates of the effect of processing time on a given outcome is that ZIP codes assigned to hearing offices with

¹⁴Our approach of admitting ZIP codes to the sample in pairs, with some ZIP codes bordering multiple neighbors assigned to other hearing offices and hence being included multiple times, mirrors the implementation of the border-county analysis used by Dube, Lester, and Reich (2010) in their study of minimum wage effects.

longer wait times would have responded to the Great Recession in the same way as their paired neighbors had they faced the same processing times. We adjust standard errors for potential two-way clustering at the hearing-office level and the border pair \times year level, and we weight our estimates by the ZIP code’s 2010 population.

Central to our investigation is the idea that the effects of SSDI hassle costs on program enrollment and employment may vary with the local severity of an economic downturn, since the marginal applicants induced to seek disability benefits in response to a recessionary shock will be more numerous in harder-hit locations. To test this hypothesis, we follow [Yagan \(2019\)](#) in computing the 2007–2009 change in the unemployment rate in each ZIP code’s commuting zone—a measure we refer to as the “Great Recession shock”.¹⁵ Since border pairs are nested within counties, which are in turn nested within commuting zones, both ZIP codes in each border pair are necessarily assigned the same shock. We then estimate [Equations \(1\) and \(2\)](#) separately for border pairs whose commuting zones experienced recession shocks above or below the sample median.

Finally, we also report specifications that further partition severely shocked border pairs into those located in counties with above-median versus below-median SSDI enrollment (as a share of the population ages 30–64) in 2007, on the eve of the Great Recession. We conjecture that marginal changes in the ease of accessing SSDI may have larger effects in localities with higher baseline receipt of SSDI, since whatever local characteristics resulted in high baseline caseloads—for example, the prevalence and complexity of physical impairment, greater knowledge or awareness of the SSDI program, or the presence of lawyers specializing in SSDI cases—are likely to amplify the responsiveness of local program enrollment and employment to a shock like the Great Recession.

4 Data and Sample Characteristics

Estimating how SSDI accessibility affects long-run SSDI enrollment and employment after the Great Recession requires ZIP code level data on program enrollment and employment rates, hearing office catchment areas, appellate processing time at each hearing office, and the geographic extent of each ZIP code. In this section, we briefly describe our data preparation and the resulting estimation sample. [Appendix B](#) provides a detailed account of how we process and combine these datasets.

¹⁵We depart slightly from [Yagan](#) in using the year-2000 (rather than year-1990) vintage of commuting zones, which we regard as better delineating local labor markets in the period we analyze.

4.1 Data sources and preparation

Our analysis relies entirely on publicly available data provided by a medley of government agencies.¹⁶ We merge these data to construct a panel at the ZIP code \times border pair \times year level spanning 2003–2015, the period preceding and following the Great Recession. The initial year of our sample period is dictated by the availability of ZIP-level data on SSDI enrollment; the final year is aligned with related work by [Yagan \(2019\)](#).

SSDI enrollment rates. The SSA website provides annual data on the number of SSDI beneficiaries at the ZIP code and county levels, together with information on the field office to which each ZIP code is assigned in a given year. These files contain separate tabulations of both disabled workers and their dependent beneficiaries, but since our interest lies in the linkages between the safety net and the labor market, we define “SSDI enrollment” as the number of disabled workers receiving SSDI benefits.¹⁷ To construct annual SSDI enrollment rates for the years 2003–2015, we divide the number of SSDI recipients in a given year by the ZIP code’s population aged 30–64 in the 2010 Decennial Census. We chose this age range because few individuals enroll in the SSDI program prior to age 30 and because SSDI beneficiaries are moved into the Social Security retirement program upon reaching retirement age. We compute ZIP-level population measures by aggregating population counts from the Census-tract level to the ZIP code level using a crosswalk provided by the US Department of Housing and Urban Development (HUD).¹⁸

¹⁶In addition to the measures described in this section, we construct population race, ethnicity, and education shares from the 2010 Decennial Census (for race and ethnicity) and the 2006–2010 pooled American Community Surveys (for education), ZIP-level industry shares from the Census LODES data, and a ZIP-level housing price index from the Federal Housing Finance Agency. We report these measures in [Table 1](#), where we assess the comparability of ZIP codes on either side of hearing office borders, but we do not otherwise use them in our analysis.

¹⁷We have also obtained annual counts of SSDI applications and allowances at the ZIP code level. We are grateful to Manasi Deshpande for facilitating the extraction of these data for us. We intended to use these data to decompose the enrollment response into applications, appeals, and approvals, but the data unfortunately did not permit a thorough analysis of these separate margins. SSA censors these data for privacy purposes: any ZIP code with fewer than 10 SSDI or SSI applications in a given year will have both application measures suppressed, and the same is true for allowances. This leads to considerable suppression in the extract we obtained from SSA: only 57 percent of our border pairs can be used to analyze applications and only 42 percent can be used to analyze allowances. Furthermore, these data are dated to the year of SSA’s initial disability determination. Since it can take years for awards to be made (especially in cases that undergo appeal), we see fewer allowances in later years, so that a larger share of cells are suppressed. Lastly, given the empirical fact that SSDI applications and allowances increase after local economic downturns ([Maestas, Mullen, and Strand, 2015](#)), the extent of data suppression varies systematically both over time and with the local severity of the Great Recession; this further complicates attempts at causal estimation. Given the sizable and systematic suppression in these data, we did not find them informative for our analysis.

¹⁸The Census Bureau reports population counts for ZIP Code Tabulation Areas rather than for the US Postal Service (USPS) ZIP codes used in our SSA data. We can therefore better match the geography of our data on SSDI recipients by starting from tract-level estimates and mapping them into USPS ZIP codes. Some Census tracts do not

Employment rates. We derive ZIP code level employment counts from the Census Bureau’s Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics, or LODES (Census Bureau, 2018). The LODES database provides annual estimates of place-of-work and place-of-residence employment, measured at the beginning of April. We use place-of-residence employment, which aligns with SSA hearing office assignment on the basis of an individual’s ZIP code of residence.¹⁹ Mirroring our construction of SSDI enrollment rates, we use the HUD crosswalk to aggregate LODES employment from tracts into ZIP codes, then compute annual employment rates by dividing year-by-year employment among workers aged 30 or older—the most appropriate age grouping available in LODES—by the 2010 ZIP code population aged 30–64.

Hearing office boundaries. We obtain hearing office catchment areas from the SSA website. Using the Internet Archive’s Wayback Machine, we pulled archived copies of the website and created a crosswalk from each field office to its assigned hearing office for each year from 2007 to 2013; after 2013, the website’s Hearing Office Locator tool became interactive and historical data cannot be retrieved. Because SSA’s tabulations of local SSDI enrollment list each ZIP code under its assigned field office, we can link ZIP codes to field offices and thereby link them to hearing offices. Using ArcGIS and an accompanying ZIP-level layer, we then identified all pairs of adjacent ZIP codes that were assigned to different hearing offices in 2010.

Appellate processing time. In 2010, SSA began producing annual reports detailing, for each hearing office, the average time elapsed from the date a denied SSDI applicant requested an appellate hearing until the date the appeal was resolved, along with judge-level tabulations of appeals granted or denied.²⁰ Since our interest is in SSDI appeals filed in the wake of the Great Recession,

perfectly align with ZIP codes, and some SSA awards will be assigned to ZIP codes that correspond to post offices or PO boxes, where there is no population. Appendix B reports match rates between our various data sources.

¹⁹We use LODES counts of “primary jobs” (each individual’s highest-earnings job during the year) to ensure that our employment measure is a count of distinct workers, rather than distinct worker-firm pairings. The Census Bureau infuses noise into the LODES data to protect worker and business confidentiality. Since we use the LODES data to construct outcomes but not regressors, this source of measurement error will reduce the precision of our employment estimates but should not introduce bias.

²⁰One reason SSA compiled these data was to better understand differences in judicial practice across administrative law judges (ALJs). In 2011, a massive ALJ reform was undertaken to retrain or remove judges who appeared to be too strict or too lenient and to promote uniformity within the appellate process. Since the composition and behavior of judges changed in 2011, a potential concern is that hearing offices with slower appellate processes in 2010 may have been more (or less) likely to experience changes in 2011 in ways that could bias our estimates. Reassuringly, however, we find that changes in hearing office award rates between 2010 and 2012 are uncorrelated with average processing time in 2010 ($\rho = 0.03$).

we rely on average processing times for appeals completed during the 2010 fiscal year (ending in September 2010), the earliest for which data are available. Given typical time lags at the application and appellate stages, most appeals resolved in 2010 would likely have begun the application process in 2007, 2008, or 2009, so that our measure should capture the administrative conditions experienced by SSDI applicants displaced in the early phases of the Great Recession.²¹

Balancing the panel. After combining these datasets, we restrict our sample to border pairs for which both ZIP codes are present throughout our sample period and have non-missing values in all years for all of the variables used in estimation. By doing so, we ensure both that the same sample is used across all of our analyses and that the effects we estimate reflect trends within ZIP codes rather than compositional shifts in the set of ZIP codes with available data. Of particular note, balancing the panel eliminates any border pairs located in Arizona, Massachusetts, Mississippi, or the District of Columbia, which enter the LODES data during our sample period and consequently lack employment-to-population rates in certain years. [Appendix Table B1](#) provides a step-by-step reckoning of the number of border pairs, ZIP codes, and ZIP \times pair \times year observations remaining in our sample after each successive sample restriction.

4.2 Sample characteristics

[Table 1](#) reports summary statistics for our estimation sample of adjacent, same-county ZIP codes assigned to different hearing offices. All statistics are weighted by 2010 ZIP code population, and all variables are measured in 2010 (with the exception of the Great Recession shock).

The average difference in 2010 average processing time is 55 days (386 versus 441 days). For the ZIP codes with longer processing times, SSDI award rates at the appellate stage are 0.7 percentage points (p.p.) higher, though the difference is not statistically significant. If anything,

²¹While individuals filing SSDI applications or appeals before 2010 could not have known the realized processing times among appeals resolved in that year, it seems likely that areas we identify as having slow appellate processes circa 2010 would have had longer wait times in earlier years as well, though we cannot test this directly. Since we observe average processing time from 2010 onwards and the assignment of field offices to hearing offices through 2013, we have a reliable ZIP-level measure of processing times for the years 2010–2013. Among paired ZIP codes belonging to our estimation sample, the population-weighted autocorrelation of average processing time is 0.70 between 2010 and 2011, 0.37 between 2010 and 2012, and 0.23 between 2010 and 2013. We suspect that the fall-off in this autocorrelation over time reflects the 2011 ALJ reform, which sought to standardize hearing practices across areas. It is therefore plausible that local processing times were more strongly correlated over time during the years when most recession-induced applicants and appellants would have been forming their expectations about processing delays.

Table 1: Characteristics of adjacent ZIP codes located in the same county but assigned to different hearing offices

	Full sample (1)	Shorter delay (2)	Longer delay (3)	Difference (4)
Explanatory variables				
Average processing time (days)	413.51	386.35	441.25	54.90***
Great Recession shock (p.p.)	4.93	4.88	4.98	0.10
SSDI appeals and enrollment				
Percent of appeals approved	65.37	65.01	65.73	0.72
Percent of pop. 30–64 on SSDI	4.60	4.65	4.54	–0.11
Employment-to-pop. rate (%)				
All, ages 30–64	59.81	59.49	60.14	0.65
Men, ages 15+	47.19	46.87	47.50	0.63
Women, ages 15+	44.92	44.50	45.36	0.85*
Population	21,839.81	22,106.90	21,717.80	–389.10
Population shares (%)				
Ages 30–64	46.05	45.71	46.39	0.67**
White	56.53	56.88	56.12	–0.75
Black or African American	18.35	18.12	18.62	0.50
Asian or Asian American	7.10	6.54	7.68	1.14
Other race + two or more races	18.02	18.46	17.58	–0.89
Hispanic or Latino ethnicity	31.91	32.76	31.11	–1.64
Less than a high school degree	21.88	22.46	21.31	–1.15
High school graduate	26.83	27.29	26.32	–0.97**
Some college education	25.46	25.26	25.66	0.40
Bachelors degree or more	25.83	24.99	26.71	1.72
Sector employment shares (%)				
Agriculture, mining, utilities	1.58	1.52	1.62	0.10
Construction	3.89	4.01	3.76	–0.25**
Manufacturing	9.79	9.81	9.76	–0.05
Trade, transport., warehousing	20.04	20.17	19.91	–0.26
Education and healthcare	22.89	22.96	22.83	–0.13
Other services	17.99	17.80	18.18	0.39*
House price index (base: 2000)	141.31	139.66	143.01	3.35
<i>p</i>-value of joint <i>F</i> test				0.03
Observations	1814.00	903.00	903.00	

Notes: Statistics are estimated means of the indicated variables; with the exception of the row for population itself, all statistics are weighted by 2010 ZIP code population. Column (1) reports statistics for our estimation sample of adjacent ZIP codes located in the same county but assigned to different hearing offices. Column (2) retains the ZIP code in each pair whose assigned hearing office had shorter average processing times as of 2010. Column (3) retains ZIP codes with longer processing times. A handful of ties are dropped from these columns. Each ZIP code may appear in our sample multiple times, once for each valid neighbor. The Great Recession shock is the change in the unemployment rate between 2007 and 2009, as measured at the commuting zone level. All other measures are 2010 values. The *F* test compares all listed variables except processing time, the Great Recession shock, and the house price index (the only variable here with missing values for some observations). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

this modest difference in award rates would bias us against finding that longer processing delays discourage SSDI enrollment and promote employment.

Our paired ZIP codes are otherwise fairly similar. Their total populations, racial and ethnic composition, sectoral composition, and 2000–2010 changes in house prices are nearly identical. ZIP codes with longer processing times had slightly more educated populations than their neighbors, as 2.1 percent more adults over the age of 25 have at least some college education, and a slightly larger share of residents aged 30–64. A joint F test of differences across all of the dimensions shown in this table—other than appellate processing time, the Great Recession shock, and the housing price index (which is missing for a significant portion of ZIP codes)—has a p -value of 0.03, suggesting that these ZIP code pairs are statistically different, even if the differences are small. We note, however, that any level differences across ZIP codes will be absorbed by our fixed effects.

5 Results


We first compare trends in SSDI enrollment across hearing office boundaries to test whether the disability caseload adjusts differently in areas that faced longer versus shorter appeal processing times during the Great Recession. We then turn to employment trends and, lastly, examine the effect of processing time on both outcomes in traditional SSDI hotspots.

5.1 Impact on SSDI enrollment

Figure 6 plots the estimated coefficients $\hat{\beta}_\tau$ from Equation (1), which estimates the effect of 2010 average processing time on SSDI enrollment rates in each of the years 2003–2015 (with the coefficient for 2008 normalized to zero). Prior to 2008, these coefficients—which capture the effect of a one-month increase in average appellate processing time—are close to zero and statistically insignificant, suggesting that neighboring ZIP codes experienced similar rates of change in SSDI enrollment in the years leading up to the Great Recession.

In the immediate wake of the recession (in 2009–2010), an additional month of average processing time reduces SSDI enrollment by approximately 0.01 percentage points, though this effect is statistically insignificant. Consistent with the time lags inherent in the appellate process, as well as the absorbing nature of SSDI, this effect then widens to 0.03 percentage points by 2012

Figure 6: Effect of appellate processing time on SSDI enrollment rates



Notes: Estimated coefficients $\hat{\beta}_\tau$ on interactions between average processing time and annual year indicators from Equation (1). Capped spikes denote 95 percent confidence intervals. The sample consists of pairs of adjacent ZIP codes located in the same county but assigned to different hearing offices as of 2010. SSDI enrollment is the ratio of the number of disabled workers receiving SSDI to the resident population ages 30–64.

and remains at that level through 2015, the end of our analysis period. Though the estimates fall narrowly short of statistical significance, with p -values between 0.1 and 0.2 from 2011 onward, they suggest that slower appellate processes may curb growth in SSDI caseloads during a downturn.


Table 2 reports, in column (1), the corresponding estimates $\hat{\beta}_{SR}$ and $\hat{\beta}_{LR}$ from Equation (2), which parametrizes these dynamics into “short-run” (2009–2011) and “long-run” (2012–2015) effects. With an average cross-border difference of 1.8 months within our border-pair sample, the long-run point estimate of -0.03 implies that the ZIP code faced with shorter processing times saw a persistent relative increase in SSDI enrollment of approximately 0.06 percentage points. This amounts to 1.3 percent of the dependent variable mean, as SSDI enrollment averages to 4.3 percent across all observations used in estimation. As with the year-by-year coefficients, we caution that the difference-in-difference estimate is at best marginally significant, with a p -value of 0.16.

We next look at how these effects vary with the severity of the local unemployment shock. To do so, we estimate Equations (1) and (2) separately for ZIP codes in commuting zones with Great Recession shocks above or below 4.5 percentage points, the median in our estimation sample. Figure 7 plots the coefficient on appellate processing time year-by-year, separately for more versus less severely shocked places. The divergent estimates for the two series reveal reductions in relative

Table 2: Effect of appellate processing time on post-recession SSDI enrollment rates

	Full sample of border pairs (1)	Less severely shocked CZs (2)	More severely shocked CZs (3)	Severely shocked and ...	
				High 2007 SSDI rate (4)	Low 2007 SSDI rate (5)
1{2009–11} × processing time	−0.015 (0.012)	0.001 (0.023)	−0.020 (0.013)	−0.033 (0.025)	−0.010 (0.010)
1{2012–15} × processing time	−0.032 (0.023)	0.006 (0.042)	−0.045* (0.024)	−0.093** (0.036)	−0.009 (0.015)
<i>p</i> -value for long-run coefficient	0.16	0.89	0.06	0.01	0.56
Mean of dependent variable	4.34	4.60	4.17	7.90	3.34
Number of border pairs	907.00	451.00	456.00	177.00	279.00
Number of ZIP codes	1,049.00	537.00	512.00	228.00	284.00
Number of observations	23,582.00	11,726.00	11,856.00	4,602.00	7,254.00

Notes: Regression results obtained from a balanced panel of paired ZIP codes located in the same county but assigned to different hearing offices as of 2010. The level of observation is ZIP code × border pair × year, stacked over the years 2003–2015. Each observation is weighted by ZIP code population in 2010. The SSDI enrollment rate is defined as the number of disabled workers receiving SSDI divided by the population ages 30–64. A ZIP code is defined as experiencing a severe shock if its commuting zone saw an above-median 2007–2009 change in the unemployment rate and as having a high 2007 SSDI rate if its county had above-median SSDI enrollment in that year. All specifications include ZIP code fixed effects and border pair × year fixed effects. Standard errors in parentheses are corrected for two-way clustering at the border pair × year and hearing-office levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 7: Effect of appellate processing time on SSDI enrollment rates, by local severity of the Great Recession

Notes: Estimated coefficients $\hat{\beta}_\tau$ on interactions between average processing time and annual year indicators from Equation (1), estimated separately for ZIP code pairs in commuting zones with 2007–2009 changes in the unemployment rate below or above the sample median. SSDI enrollment is the ratio of the number of disabled workers receiving SSDI to the resident population ages 30–64. Capped spikes denote 95 percent confidence intervals.

SSDI enrollment in commuting zones hit harder by the Great Recession. Within these hardest-hit labor markets, ZIP codes assigned to hearing offices with one month shorter processing times experienced a persistent relative increase in SSDI enrollment just shy of 0.05 percentage points relative to their neighboring ZIP code. The difference-in-difference specification reported in column (3) of [Table 2](#) confirms this impression, with a 2012–2015 point estimate of $-.045$ that just skirts 95 percent significance ($p = 0.06$). If we again scale this estimate by the typical 1.8 month disparity in processing times, these estimates suggest that SSDI enrollment rates reach 0.08 percentage points (1.9 percent) higher in the ZIP code that faced longer processing times. By contrast, ZIP code pairs in commuting zones where the recession was comparatively less severe exhibit no economically or statistically significant change in their relative SSDI enrollment trends.

In sum, then, the results we have presented thus far suggest that marginal increases in SSDI hassle costs dampen the growth in SSDI enrollment in areas subject to especially severe labor market shocks, but have no detectable effect on program enrollment in areas hit less hard. A plausible explanation of this fact pattern is that bigger local shocks result in greater numbers of dislocated “conditional applicants” for whom moderate shifts in the expected cost of obtaining SSDI benefits can potentially sway decisions to apply in the first place or to appeal a denied claim.


5.2 Impact on employment rates

As the speed of the appellate process leaves a persistent imprint on SSDI enrollment rates after a severe labor market contraction, we next explore its effect on local employment rates. All else equal, we might expect neighboring ZIP codes to experience a similar recovery in employment once the economy begins to improve. However, as SSDI enrollment appears to have been affected by processing times, we might instead observe persistent differences in employment.

[Figures 8](#) and [9](#) presents coefficients $\hat{\beta}_\tau$ from estimating [Equation \(1\)](#) for the employment-to-population rate, first in our overall sample and then, as before, splitting the sample into more versus less severely shocked local labor markets. In both figures, we find no clear indications of differences in employment trends between ZIP codes faced with longer or shorter appellate processing times. The point estimates move jaggedly in the middle of our sample period, and the confidence intervals are wide. The parametrized estimates in columns (1)–(3) of [Table 3](#) are correspondingly noisy.


While the noisiness of these estimates precludes clear conclusions, a null effect of appellate

Figure 8: Effect of appellate processing time on employment-to-population rates



Notes: Estimated coefficients $\hat{\beta}_\tau$ on interactions between average processing time and annual year indicators from Equation (1). Employment rates are the ratio of total jobs held by ZIP code residents ages 30+ to the ZIP code’s population ages 30–64. Capped spikes denote 95 percent confidence intervals. The sample consists of pairs of adjacent ZIP codes located in the same county but assigned to different hearing offices as of 2010.

Figure 9: Effect of appellate processing time on employment-to-population rates, by local severity of the Great Recession



Notes: Estimated coefficients $\hat{\beta}_\tau$ on interactions between average processing time and annual year indicators from Equation (1), estimated separately for ZIP code pairs in commuting zones with 2007–2009 changes in the unemployment rate below or above the sample median. Employment rates are the ratio of total jobs held by ZIP code residents ages 30+ to the ZIP code’s population ages 30–64. Capped spikes denote 95 percent confidence intervals.

processing time on employment rates would be consistent with the theoretical ambiguity of this effect. Though an increase in the hassle cost of appealing a denied SSDI application should have an unambiguously negative impact on SSDI enrollment, the effect on employment could go either way.

Table 3: Effect of appellate processing time on post-recession employment-to-population rates


	Severely shocked and ...				
	Full sample of border pairs	Less severely shocked CZs	More severely shocked CZs	High 2007 SSDI rate	Low 2007 SSDI rate
	(1)	(2)	(3)	(4)	(5)
1{2009–11} × processing time	0.187 (0.124)	0.305* (0.155)	0.148 (0.161)	0.473*** (0.177)	−0.091 (0.181)
1{2012–15} × processing time	0.158 (0.205)	0.098 (0.319)	0.178 (0.248)	0.833*** (0.270)	−0.303 (0.194)
<i>p</i> -value for long-run coefficient	0.44	0.76	0.48	0.00	0.13
Mean of dependent variable	60.32	62.33	59.02	56.56	59.57
Number of border pairs	907.00	451.00	456.00	177.00	279.00
Number of ZIP codes	1,049.00	537.00	512.00	228.00	284.00
Number of observations	23,582.00	11,726.00	11,856.00	4,602.00	7,254.00

Notes: Regression results obtained from a balanced panel of paired ZIP codes located in the same county but assigned to different hearing offices as of 2010. The level of observation is ZIP code × border pair × year, stacked over the years 2003–2015. Each observation is weighted by ZIP code population in 2010. The employment-to-population ratio is the number of employed residents ages 30+ divided by the population ages 30–64. A ZIP code is defined as experiencing a severe shock if its commuting zone saw an above-median 2007–2009 change in the unemployment rate and as having a high 2007 SSDI rate if its county had above-median SSDI enrollment in that year. All specifications include ZIP code fixed effects and border pair × year fixed effects. Standard errors in parentheses are corrected for two-way clustering at the border pair × year and hearing-office levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Consider a disabled worker who is deciding whether to appeal a denied claim. On the one hand, an increase in expected processing time will raise the opportunity cost of filing an appeal, since the worker must keep earnings below the SGA limit while the appeal is pending. To the extent that workers are attuned to processing times, this should discourage some from filing an appeal. We call this the *deterrence effect*.²² On the other hand, if the worker does choose to file an appeal, a slower appellate process will necessitate a longer period of below-SGA earnings before the case is resolved. Given the strong empirical evidence that joblessness erodes earnings capacity, whether through skill depreciation, discouragement, or the scarring of employer beliefs (Kroft, Lange, and Notowidigdo, 2013; Autor et al., 2015; Schmieder, von Wachter, and Bender, 2016), a lengthy period of legal limbo should depress subsequent employment among those whose claims are ultimately denied. This is the *decay effect*. The broadly similar evolution of employment rates in areas with longer versus shorter processing delays could arise if the deterrence and decay effects roughly offset each

²²Longer processing times might also deter some of those who file an appeal from seeing the process through to its conclusion. Even an appellant who knows nothing about average processing times might withdraw a pending appeal upon learning of a suitable employment opportunity, and this is more likely to occur the longer an appeal drags on.

Figure 10: Effect of processing time on SSDI enrollment and employment-to-population rates, in areas with high pre-recession SSDI enrollment and severe Great Recession shocks



Notes: Estimated coefficients $\hat{\beta}_\tau$ on interactions between average processing time and annual year indicators from Equation (1), using ZIP code pairs with above-median 2007–2009 changes in the unemployment rate (measured at the commuting zone level) and above-median 2007 SSDI enrollment rates (measured at the county level). The outcome variables are defined as in previous figures. Capped spikes denote 95 percent confidence intervals.

other, though we cannot rule out economically meaningful effects in either direction.²³

5.3 Heterogeneous effects by baseline SSDI enrollment rates

As a final exercise, we investigate whether the effects of SSDI hassle costs on program enrollment and employment rates vary with historical SSDI enrollment rates. Pre-recession SSDI caseloads span a wide range: among ZIP codes in our border panel, county-level SSDI enrollment rates at the 75th percentile of the 2007 distribution are nearly twice those at the 25th percentile (6.3 versus 3.1 percent). Among our border pairs, we see baseline SSDI enrollment in excess of 10 percent for counties in Alabama, Arkansas, California, Kentucky, Missouri, Tennessee, Virginia, and West Virginia. We categorize ZIP code pairs based on whether their (shared) county had SSDI enrollment (as a share of the population ages 30–64) in 2007, on the eve of the Great Recession, above or below the sample median of 4.1 percent. We conjecture that marginal changes in the ease of accessing SSDI (as captured through appeals processing time) may have larger effects in localities with higher baseline receipt of SSDI, since such places are likely to have stronger SSDI network and information

²³Since the LODS data are noise-infused to protect confidentiality, a more prosaic interpretation of these results is simply that measurement error is obscuring any true relationship between average processing time and employment.

effects that would likely amplify the responsiveness of local SSDI enrollment and employment to a shock like the Great Recession.

The data offer qualified support for this conjecture. [Figure 10](#) plots the year-by-year results, with the analogous difference-in-difference estimates presented in column (4) of [Tables 2](#) and [3](#). In severely shocked ZIP code pairs in counties with high 2007 SSDI enrollment rates, we find a statistically significant negative effect of processing time on SSDI enrollment and a statistically significant positive effect on subsequent employment rates. In such localities, an additional month of processing time is associated with a relative decrease in SSDI enrollment of about 0.1 percentage points—twice the magnitude we found when looking at all of the hardest-hit places. On the employment side, interpretation of our estimates is complicated by an apparent (if statistically insignificant) pretrend in the years leading up to the Great Recession. With this caveat noted, there is some indication of relatively faster employment growth in ZIP codes faced with longer appellant processing times. If we take this result at face value, it would suggest that in areas with traditionally high SSDI enrollment rates hit by deep recessionary shocks, less cumbersome access to SSDI contributed to slower employment recoveries after the Great Recession.

6 Conclusion

During economic downturns, and the Great Recession in particular, the SSDI caseload surges as labor market dislocation induces marginal, less-impaired applicants to apply for disability benefits. Since SSDI receipt is in most cases an absorbing state, the program’s distortionary labor force participation incentives have the potential to prolong the employment effects of recessions by keeping these marginal recipients—who might otherwise have remained in the labor force—from returning to work. While a long line of research has established that SSDI enrollment responds to labor market conditions, the strength of these effects is likely to depend on the ease with which potential beneficiaries can access the program. In this paper, we exploit differences across SSA hearing offices in the average processing time of SSDI appeals, coupled with residence-based assignment of potential beneficiaries to hearing offices, to estimate the effect of SSDI hassle costs on long-run program enrollment and employment rates in areas subject to labor market shocks of variable severity.

By comparing neighboring ZIP codes that straddle the borders between hearing office catch-

ment areas, we trace the evolution of SSDI enrollment and employment in areas subject to differences in the average processing time of SSDI appeals, as measured in 2010, but home to similar populations and exposed to common economic shocks. We find that neighboring ZIP codes assigned to different hearing offices (and thus differences in the speed of appellate review) experienced similar trends in SSDI enrollment prior to the Great Recession, but divergent trends thereafter. This divergence is concentrated in commuting zones that experienced particularly severe unemployment shocks, which presumably push larger numbers of displaced workers to the margin of program entry. In these hardest-hit labor markets, ZIP codes faced with longer appellate processing times exhibit persistent reductions in SSDI enrollment rates of about 0.08 percentage points (1.9 percent of the sample mean) relative to neighboring ZIP codes that face shorter processing times.

Having established that greater hassle costs lead to relatively lower SSDI enrollment after the Great Recession, especially in severely hit locations, we turn to the question of how they affected subsequent employment. In the full sample of ZIP code pairs, we find no evidence of divergent employment recovery rates after the recession in ZIP codes with longer versus shorter appeals processing time, though our confidence intervals are wide. However, in severely shocked ZIP code pairs with larger SSDI caseloads on the eve of the recession, we find that slower appellate review has large negative effects on SSDI enrollment and large positive effects on employment, consistent with larger hassle costs spurring faster employment recovery in these particular places. Though an apparent pretrend in the employment estimates warrants caution about this particular result, our findings raise the possibility that SSDI network effects—arising, for instance, through word of mouth or through the local concentration of lawyers specializing in disability cases—might amplify the importance of hassle costs in traditional SSDI hotspots.

In presenting these results, we make no statement about the welfare gains or losses associated with SSDI hassle costs, as these depend on both the social planner’s objective function and the characteristics of the complier population. In a recent paper, [Deshpande and Li \(2019\)](#) find that the SSDI applicants screened out of the program by the closure of SSA field offices disproportionately had moderately severe conditions and low education levels. If marginal SSDI applicants during the Great Recession tended to be those with greater earnings potential and lower levels of disadvantage, making SSDI more or less cumbersome to access has different welfare implications than if applicants deterred from entering the program on account of hassle costs were especially disadvantaged. We do

not attempt a welfare calculation, but merely highlight the point made by [Deshpande and Li \(2019\)](#) that the welfare implications of hassle costs depend on who is screened in or out of a program, along with the relative social weights placed on encouraging employment versus insuring workers against the earnings losses associated with disability.

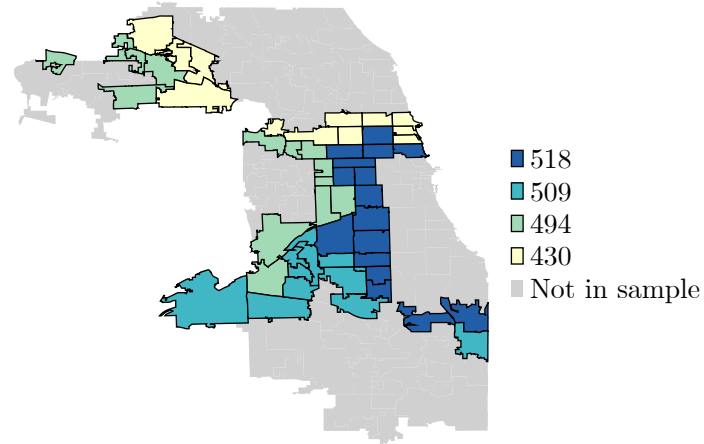
References

- Abraham, Katherine G. and Melissa S. Kearney.** 2020. “Explaining the Decline in the US Employment-to-Population Ratio: A Review of the Evidence.” *Journal of Economic Literature*, 58(3): 583–643.
- Amior, Michael and Alan Manning.** 2018. “The Persistence of Local Joblessness.” *American Economic Review*, 108(7): 1942–1970.
- Autor, David H. and Mark G. Duggan.** 2003. “The Rise in the Disability Rolls and the Decline in Unemployment.” *The Quarterly Journal of Economics*, 118(1): 157–205.
- Autor, David H. and Mark G. Duggan.** 2006. “The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding.” *Journal of Economic Perspectives*, 20(3): 71–96.
- Autor, David H., Nicole Maestas, Kathleen J. Mullen, and Alexander Strand.** 2015. “Does Delay Cause Decay? The Effect of Administrative Decision Time on the Labor Force Participation and Earnings of Disability Applicants.” *NBER Working Paper No. 20840*.
- Black, Dan, Kermit Daniel, and Seth Sanders.** 2002. “The Impact of Economic Conditions on Participation in Disability Programs: Evidence from the Coal Boom and Bust.” *American Economic Review*, 92(1): 27–50.
- Bound, John.** 1989. “The Health and Earnings of Rejected Disability Insurance Applicants.” *American Economic Review*, 79(3): 482–503.
- Census Bureau.** 2018. “LEHD Origin-Destination Employment Statistics (LODES).” Center for Economic Studies, Longitudinal-Employer Household Dynamics Program. Version 7.5.
- Charles, Kerwin Kofi, Yiming Li, and Melvin Stephens Jr.** 2018. “Disability Benefit Take-up and Local Labor Market Conditions.” *Review of Economics and Statistics*, 100(3): 416–423.
- Chen, Susan and Wilbert Van der Klaauw.** 2008. “The Work Disincentive Effects of the Disability Insurance Program in the 1990s.” *Journal of Econometrics*, 142(2): 757–784.
- Deshpande, Manasi and Yue Li.** 2019. “Who Is Screened Out? Application Costs and the Targeting of Disability Programs.” *American Economic Journal: Economic Policy*, 11(4): 213–248.
- Dube, Arindrajit, T. William Lester, and Michael Reich.** 2010. “Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties.” *Review of Economics and Statistics*, 92(4): 945–964.
- French, Eric and Jae Song.** 2014. “The Effect of Disability Insurance Receipt on Labor Supply.” *American Economic Journal: Economic Policy*, 6(2): 291–337.

- Gelber, Alexander, Timothy Moore, and Alexander Strand.** 2017. “The Effect of Disability Insurance Payments on Beneficiaries’ Earnings.” *American Economic Journal: Economic Policy*, 9(3): 229–261.
- Golosov, Mikhail and Aleh Tsyvinski.** 2006. “Designing Optimal Disability Insurance: A Case for Asset Testing.” *Journal of Political Economy*, 114(2): 257–279.
- Hershbein, Brad and Bryan A. Stuart.** 2020. “Recessions and Local Labor Market Hysteresis.” *Upjohn Institute Working Paper No. 20-325*.
- Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo.** 2013. “Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment.” *Quarterly Journal of Economics*, 128(3): 1123–1167.
- Liebman, Jeffrey B.** 2015. “Understanding the Increase in Disability Insurance Benefit Receipt in the United States.” *Journal of Economic Perspectives*, 29(2): 123–150.
- Maestas, Nicole, Kathleen J. Mullen, and Alexander Strand.** 2013. “Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt.” *American Economic Review*, 103(5): 1797–1829.
- Maestas, Nicole, Kathleen J. Mullen, and Alexander Strand.** 2015. “Disability Insurance and the Great Recession.” *American Economic Review, Papers and Proceedings*, 105(5): 177–182.
- Maestas, Nicole, Kathleen J. Mullen, and Alexander Strand.** 2018. “The Effect of Economic Conditions on the Disability Insurance Program: Evidence from the Great Recession.” *NBER Working Paper No. 25338*.
- Nichols, Albert L. and Richard J. Zeckhauser.** 1982. “Targeting Transfers through Restrictions on Recipients.” *American Economic Review, Papers and Proceedings*, 72(2): 372–377.
- Schmieder, Johannes F., Till von Wachter, and Stefan Bender.** 2016. “The Effect of Unemployment Benefits and Nonemployment Durations on Wages.” *American Economic Review*, 106(3): 739–777.
- Social Security Administration.** 2020. “Annual Statistical Report on the Social Security Disability Insurance Program, 2019.” *Social Security Administration Office of Retirement and Disability Policy*, SSA Publication No. 13-11826.
- Tolbert, Charles M. and Molly Sizer.** 1996. “US Commuting Zones and Labor Market Areas: A 1990 Update.” *Economic Research Service Staff Paper 9614*.
- von Wachter, Till, Jae Song, and Joyce Manchester.** 2011. “Trends in Employment and Earnings of Allowed and Rejected Applicants to the Social Security Disability Insurance Program.” *American Economic Review*, 101(7): 3308–3329.
- Yagan, Danny.** 2019. “Employment Hysteresis from the Great Recession.” *Journal of Political Economy*, 127(5): 2505–2558.
- Zayatz, Tim.** 2015. “Social Security Disability Insurance Program Worker Experience.” *Social Security Administration Office of the Chief Actuary*, SSA Pub. No. 11-11543.

A Additional Figures

Figure A1: Adjacent ZIP codes in Cook County (Chicago) assigned to different hearing offices

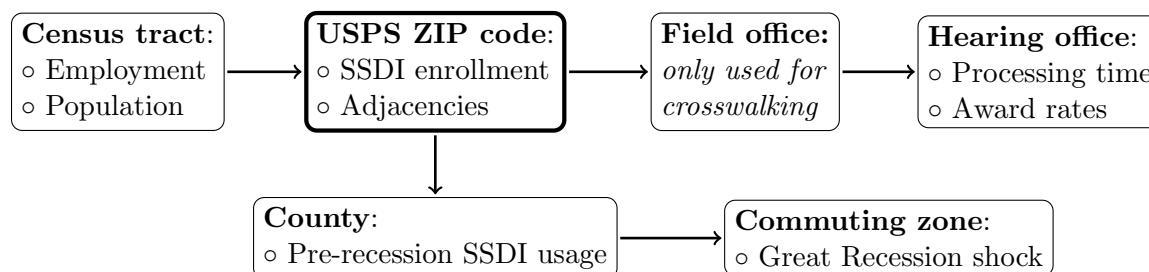


Notes: Average processing times by hearing office as of 2010, as reported by the Social Security Administration and plotted for all ZIP codes in Cook County, Illinois, with one or more neighbors assigned to a different hearing office.

B Data preparation

As summarized in [Appendix Figure B1](#), our analysis relies (entirely) on public-use data at six levels of geographic aggregation. We describe here the main steps in our data preparation. Our replication code, including detailed information on how and where we accessed the source data as well as further details about the data preparation and analysis, will be posted alongside the published version of this paper.

Figure B1: Key variables originating at each level of geographic aggregation



Notes: The bold border around USPS ZIP codes signifies that this is our unit of analysis. Arrows indicate which geographic units are typically subsumed within each other—for instance, most ZIP codes are contained within larger counties, though some ZIP codes straddle county lines.

Our data preparation proceeds as follows, with each step detailed below:

- Step 1.** Assemble ZIP \times year data on the number of disabled workers receiving SSDI benefits, together with the identity of the field office to which each ZIP code is assigned in a given year.
- Step 2.** Assemble hearing office \times year data on average appellate processing times and the percentage of appeals approved.
- Step 3.** Merge the data assembled in steps 1 and 2 by matching each ZIP code’s assigned field office, as of 2010, with its associated hearing office in that year. This yields a ZIP \times year dataset of SSDI enrollment augmented with hearing office characteristics as of 2010.
- Step 4.** Assemble ZIP-level data on population and employment by aggregating up from data reported at the Census-tract level. Merge these data into our ZIP \times year panel, and compute per capita measures of SSDI enrollment and employment.
- Step 5.** Assemble county- and commuting-zone level data, including our measure of the local severity of the Great Recession, and merge these into our ZIP \times year panel.
- Step 6.** Match each ZIP code with each of its ZIP-level neighbors, and reshape the dataset to the pair \times year level.
- Step 7.** Identify and retain paired ZIP codes that are located in the same county, assigned to different hearing offices, and not missing any of the main variables used in our analysis.

Step 1: Assemble ZIP \times year data on SSDI caseloads and assigned field office

Beginning with December 2003, the Social Security Administration (SSA) publishes [annual data files](#) reporting the number of individuals receiving Old-Age, Survivors, and Disability Insurance (OASDI) benefits at year end in each five-digit US Postal Service (USPS) ZIP code.²⁴ From these files, we extract ZIP-level counts of the number of SSDI recipients in each year during 2003–2015, together with the name and state of the field office to which each ZIP code was assigned in that year. To preserve confidentiality, the SSA files exclude ZIP \times year cells containing fewer than 15 OASDI recipients, and the beneficiary counts are rounded to multiples of five. We treat suppressed ZIP \times year cells as missing observations.

Throughout the paper, we restrict attention to the 50 US states plus the District of Columbia. After excluding non-state US territories, we observe approximately 37,000 ZIP codes per year in these SSA files, with the number fluctuating slightly as ZIP codes are created or destroyed and as OASDI enrollment crosses the threshold for data suppression.

Step 2: Assemble hearing-office data on processing times and award rates

For each fiscal year from 2009–2010 onwards, SSA reports [caseload statistics and performance metrics](#) for each of the hearing offices that adjudicate appeals of denied applications for SSDI benefits. We assemble hearing office characteristics from two sets of files:

- From the *Hearing Office Workload Data*, we extract the average processing time across all appeals concluded during each fiscal year, where processing time is defined as the number of days elapsed from the date a hearing was requested to the date the appeal was resolved.
- From the *ALJ Disposition Data*, we extract monthly data on the number of appeals decided and the number of appellants awarded benefits by each administrative law judge at each hearing office (judges may hear cases at multiple offices). We sum these counts to the hearing office \times fiscal-year level, then compute the award rate as the ratio of total awards to total decisions. Our denominator excludes dismissals, which include withdrawals by the applicant.

We refer to the 2009–2010 fiscal year as simply “2010” (and likewise for later years).

We discard data pertaining to either National Hearing Centers or the anti-fraud Special Review Cadre, which are not associated with particular geographic areas, as well as data pertaining to hearing offices located in Puerto Rico. After making these exclusions and harmonizing office names across files, we observe a total of 153 hearing offices in 2010, of which two are satellites of other, larger offices. Three hearing offices had zero appeal dispositions in 2010 and thus lack data on average processing time, and five are missing at least one of the variables described above.

A handful of hearing offices enter or exit the data over time. Net of our sample restrictions, the number of hearing offices rises to 162 by 2012 and falls to 161 in 2015, and we observe a total of 166 distinct hearing offices at some point during 2010–2015.

Step 3: Merge ZIP-level data with hearing office characteristics

To match ZIP codes with their assigned hearing offices, we must first construct a crosswalk from field office names to hearing office names. Although our crosswalk is imperfect—owing to inconsistent naming conventions across SSA files and significant changes in office assignments as offices open, close, or restructure—we believe we are able to accurately classify ZIP codes accounting for the

²⁴The SSA website notes that archival copies of some pre-2003 data are stored at the HathiTrust Digital Library. We did not attempt to access these earlier editions.

vast majority of the potential SSDI claimant population. To the extent that any slippage in these mappings takes the form of classical measurement error, it would bias us against finding effects of appeals processing time on program enrollment and employment rates.

Constructing a crosswalk from field offices into hearing offices. The SSA’s [Hearing Office Locator tool](#) provides information about which field offices are currently assigned to each hearing office, but not about how these assignments have changed over time. During 2007–2013, however, the SSA maintained and updated static webpages listing the name and state of each field office, together with its associated hearing office, in a semi-structured HTML format. We use the Internet Archive Wayback Machine to fetch one snapshot per year for each of the 10 regional pages on which this information was provided.

Next, we parse the HTML source code to construct year-by-year crosswalks from field offices into hearing offices. We use string manipulations, regular expressions, and hand-coding to standardize office naming conventions across snapshots and to align the field office names with those used in the ZIP \times year files described in step 1. In most cases, a field office is assigned to a single hearing office at a given point in time, but we make note of occasional instances in which a field office is served by multiple hearing offices, so that we can exclude such cases from our analysis.

Mapping ZIP codes to hearing offices. We use our new crosswalk to merge the ZIP \times year data assembled in step 1 with the hearing office \times year data assembled in step 2 on the basis of harmonized field-office names. Since our ultimate interest is in associating each ZIP code with the hearing office to which it was assigned in 2010, we focus our discussion on the match rate for that year.

Upon applying our crosswalk to the ZIP \times year panel on a year-by-year basis, we successfully assign 96.7 percent [respectively, 94.2 percent] of ZIP codes to at least one [exactly one] hearing office in 2010. We populate hearing office assignments for unmatched observations as follows:

- i. *Interpolate office assignments:* First, if a ZIP code was assigned to the same hearing office in years $t - 1$ and $t + 1$ but is unassigned in year t , we assign it to that hearing office in year t as well. We repeat this process successively for gaps of two, three, four, or five years in length. After this step, 97.4 percent of ZIP codes have an assigned hearing office in 2010.
- ii. *Extrapolate office assignments backwards:* Next, if a ZIP code is first assigned to a hearing office in some year $t > 2007$, we assign it to that hearing office in all years $s \in \{2007, \dots, t - 1\}$. Doing so brings our 2010 match rate up to 98.8 percent.
- iii. *Extrapolate office assignments forwards:* We extrapolate hearing office assignments forward in a similar fashion. Doing so brings our 2010 match rate above 99.9 percent.
- iv. *Handcode remaining cases.* We assign the few remaining ZIP codes manually by mimicking the hearing-office assignments of other nearby field offices.

At the conclusion of this procedure, we are left with a ZIP \times year dataset containing SSDI enrollment counts augmented with the identity of each ZIP code’s assigned hearing office—as observed in the same year, where possible, or in the closest possible year in the few cases we populate through imputation. By construction, each ZIP code that appears in our data in 2010 has a hearing office assigned (or imputed) to it in that year. Using these 2010 assignments, we merge in the 2010 values of the hearing office characteristics assembled in step 2. Although the mapping from ZIP codes to hearing offices changes somewhat over time, we posit that a locality’s office assignment

in 2010—when many recession-induced SSDI appeals were being adjudicated—is most relevant for claimants’ decisions about whether to appeal a denied SSDI claim (or whether to apply at all).

Step 4: Prepare ZIP \times year data on SSDI enrollment and employment rates

We augment our ZIP \times year panel with data on local population (measured in 2010) and employment (measured annually). With these measures in hand, we construct our two outcomes of interest: per capita SSDI enrollment and per capita employment.

Population. We obtain estimates of local population from the 2010 Decennial Census.²⁵ The US Census Bureau reports population counts not for USPS ZIP codes, but rather for a distinct set of geographic units called ZIP Code Tabulation Areas (ZCTAs). To better align the Census data with our SSA data, which are reported by USPS ZIP code, we start by loading population data at the more granular Census-tract level. We map tracts into USPS ZIP codes using a fractional, many-to-many crosswalk provided by the US Department of Housing and Urban Development (HUD).²⁶ We then compute ZIP-level population as the share-weighted sum of population in all corresponding tracts.²⁷ Our main analysis uses estimates of both overall local population and the local population aged 30–64; we also draw on population counts by educational attainment, race, and ethnicity to help assess the comparability of neighboring ZIP codes assigned to different hearing offices.²⁸

Employment. We obtain estimates of local employment from the Census Bureau’s Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics (LODES), a public-use dataset derived from the confidential LEHD. The LODES program uses matched worker-firm microdata to construct noise-infused estimates of employment at the Census-block level, with jobs grouped either by workers’ home addresses, their work addresses, or origin-destination pairs defined jointly by both addresses. Because our interest is in employment rates among the residents of each given ZIP code, we use the LODES Residence Area Characteristic (RAC) files to measure employment by place of residence; we use estimates based on workers’ primary jobs to ensure that each worker is counted only once in a given year. We sum block-level employment estimates to the Census-tract level, then aggregate to USPS ZIP codes using the same fractional HUD crosswalk we use to prepare the population data. We rely principally on measures of local employment among workers ages 30 or older, though we also draw on demographic and sectoral breakdowns of employment in our examination of sample characteristics in [Table 1](#).²⁹

The LODES database is derived from underlying state-level data, and states enter the LODES universe in different years. The employment level ascribed to ZIP code i in year t in the RAC files

²⁵While we would ideally denominate SSDI reciprocity by an annual measure of population, the only such measure we are aware of at the ZIP code level is the number of tax exemptions reported in the Internal Revenue Service’s Statistics of Income. However, this measure does not separate working-age adults from the broader population and is itself potentially impacted by changes in tax-filing and exemption-claiming rates induced by the Great Recession; furthermore, starting in 2008, ZIP codes with small populations are suppressed to preserve privacy. For these reasons, we denominate SSDI reciprocity by the static population measure observed in the Decennial Census.

²⁶Both here and in step 5, where we map ZIP codes into counties using another HUD file, we use HUD’s crosswalk file for the first quarter of 2012, the earliest to incorporate the 2010 Census geographies used in our other datasets.

²⁷Some tracts that appear in the 2010 Decennial Census are absent from HUD’s tract-to-ZIP crosswalk, but most such tracts had no residents in 2010. We successfully match tracts accounting for 99.97 percent of the 2010 US population.

²⁸Population counts by education group are taken from the pooled 2006–2010 American Community Survey. All other population counts are from the 2010 Decennial Census.

²⁹As with the population data, some tracts in the LODES data are absent from the HUD crosswalk. In every year in our sample period, we successfully match in excess of 99.9 percent of LODES employment into USPS ZIP codes.

is an estimate of total employment among residents of i at employers located in states belonging to the LODES universe in that year. As such, employment counts jump discontinuously—from a low but non-zero level reflecting only out-of-state employment—in the year when a state first enters the LODES frame. To avoid drawing erroneous inferences from these spurious leaps in employment, we exclude from our sample any ZIP codes located in Arizona (which entered LODES in 2004), Massachusetts (2011), Mississippi (2004), or the District of Columbia (2010). All other US states belong to the LODES universe throughout our sample period.

SSDI enrollment and employment rates. Armed with ZIP-level estimates of population, SSDI enrollment, and employment, we compute the local SSDI enrollment rate in a given year as the ratio of all disabled workers receiving SSDI benefits in that year to the 2010 population aged 30–64. As described in the text, we choose this denominator because SSDI reciprocity is rare at younger ages and because disabled workers transition from SSDI into Social Security retirement benefits upon reaching retirement age. We compute the local employment-to-population rate as the ratio of LODES employment among workers aged 30 or older to the 2010 population aged 30–64.³⁰

Step 5: Incorporate data at the county and commuting-zone levels

We assign each ZIP code to a single county, which is in turn nested within a commuting zone (CZ). As explained in the main text, we use county and CZ codes for three purposes. First, we restrict our analysis sample to paired ZIP codes located within the same county (and thus also the same CZ), to eliminate cases in which the border between two hearing-office catchment areas coincides with county (or state) borders. Second, we assign each border pair a “Great Recession shock” based on the 2007–2009 change in unemployment in the pair’s CZ. Third, we further split our sample based on the pre-recession SSDI enrollment rate in each border pair’s county.

Mapping ZIP codes into counties. Just as we did for our tract-to-ZIP concordance, we assign USPS ZIP codes to counties using a fractional crosswalk provided by HUD. Because USPS ZIP codes can straddle county lines, the HUD file records the share of all residential and business addresses in each ZIP code that are located within each county. Among all ZIP codes nationwide, over three-quarters are uniquely matched to a county; among those multiply matched, the median ZIP code has 93.7 percent of its addresses located in its modal county. We assign each ZIP code to its modal county.

County definitions and FIPS codes change occasionally over time, as counties are created, absorbed, or redefined. We draw on myriad internet resources, notably David Dorn’s [list of county changes](#), to standardize FIPS codes across datasets and to bridge changes in county borders.³¹

Mapping counties into commuting zones. Commuting zones (CZs) are assemblages of counties characterized both by dense commuting flows across counties within each CZ and by sparse commuting flows between CZs (Tolbert and Sizer, 1996). We map each county into a unique CZ using the US Department of Agriculture’s year-2000 CZ definitions, which partition the United States into 709 CZs.

³⁰The LODES data report employment counts among those aged 30–54 and among those aged 55 or older, but it does not subdivide the latter group into those over or under 65.

³¹Alaskan borough (county) borders change substantially during our sample period. While these changes would make it difficult for us to assign Great Recession shocks to Alaskan ZIP codes on a consistent basis, no pairs of neighboring Alaskan ZIP codes survive our other sample restrictions, so this potential issue does not arise in practice.

Constructing local recession shocks. The Bureau of Labor Statistics’ Local Area Unemployment Statistics (LAUS) database reports county-level unemployment rates at both monthly and annual frequencies. Following Yagan (2019), we define each CZ’s “Great Recession shock” as the percentage-point change in its unemployment rate between 2007 and 2009. Departing slightly from Yagan, we begin with the annually averaged data reported directly in the LAUS (rather than averaging across months ourselves), and we use county population weights to aggregate these shocks to year-2000 CZs (rather than year-1990 CZs). We favor the 2000 vintage of CZ definitions because these are closer to the start date of our analysis period, which begins in 2003.

Measuring county-level SSDI enrollment rates. To distinguish localities with higher versus lower pre-recession SSDI enrollment rates, we split ZIP pairs based on whether their (common) county had above-median or below-median per capita SSDI enrollment in the year 2007. For the numerator of this measure, we use county-level SSA data on the total number of disabled workers receiving SSDI benefits in that year. For the denominator, we use each county’s 2007 population ages 30–64, using county \times age-bin data provided by the National Cancer Institute’s Surveillance, Epidemiology, and End Results (SEER) Program.

Step 6: Construct a panel dataset of paired neighboring ZIP codes

The above steps yield a comprehensive ZIP \times year panel encompassing all US states plus the District of Columbia. To identify pairs of adjacent ZIP codes, we apply a canned ArcGIS routine to a [five-digit USPS ZIP code layer](#) provided by ESRI, the maker of ArcGIS.³² We take an expansive view of adjacency, so that two ZIP codes are deemed adjacent even if they merely touch at their corners. Among the 30,745 ZIP codes in the ArcGIS layer, we find 88,113 pairs of adjacent codes.

By merging these adjacencies onto our main panel dataset, we obtain a new dataset at the pair \times year level, with information on the characteristics of each ZIP code belonging to the pair.³³ We observe a total of 86,087 distinct pairs of adjacent ZIP codes in our data, with each pair observed an average of 12.9 times during our 13-year sample period.

Step 7: Restrict to same-county pairs straddling hearing office borders

From this nationwide set of adjacent ZIP codes, we apply a succession of sample restrictions to obtain the border-pair sample used throughout our main analysis. We summarize these restrictions in [Appendix Table B1](#), along with the sample size after each restriction is applied.

The first set of restrictions are necessitated by missing data. Out of our initial set of adjacent pairs, we lose about 100 due to missing county codes and about 3,700 due to mid-period entry into the LODES universe; we lose about 800 additional pairs due to unknown 2010 hearing offices, unknown 2010 appellate processing time at a known hearing office, unknown population, or null population aged 30–64 (such that our outcomes of interest would be undefined). Conditional on the

³²We use version 10.4 of the ArcGIS layer. ESRI has released subsequent versions since we began this project, but the version we use is closer to our sample period than those currently available on the ArcGIS website. We use the same ArcGIS layer to construct the shapefiles underlying all maps shown in the paper.

³³At the point of merge, our ZIP \times year panel contains 38,146 distinct ZIP codes, considerably more than we observe in the ArcGIS layer. In addition, 857 ZIP codes in the ArcGIS data are absent from our ZIP \times year panel. These discrepancies arise because ESRI excludes ZIP codes that do not represent polygonal delivery areas (such as those associated with PO boxes or with individual high-volume addresses) and because the set of ZIP codes active during our 2003–2015 sample period does not perfectly match those active in 2018, when our ArcGIS layer was constructed. The mismatched ZIP codes, though numerous, are sparsely populated: fully 98.6 percent of the 2010 US population resided in ZIP codes we match successfully to the GIS layer.

Table B1: Remaining sample sizes after successive sample restrictions

Sample restriction	Border pairs	ZIP codes	Observations
All pairs of adjacent ZIP codes	86,087	29,957	2,217,844
Exclusions due to missing data			
Drop if either ZIP code has a missing county code	85,982	29,926	2,215,876
Drop if either ZIP code is in a state that entered LODES post-2003	82,254	28,691	2,120,938
Drop if either ZIP code has an unknown 2010 hearing office	81,781	28,551	2,116,868
Drop if either ZIP code has unknown appellate processing time	81,472	28,465	2,108,976
Drop if either ZIP code has unknown population in 2010	81,454	28,457	2,108,508
Drop if either ZIP code has zero population ages 30–64 in 2010	81,451	28,456	2,108,430
Drop if either ZIP code has unknown SSDI enrollment	81,451	28,456	2,108,430
Drop if either ZIP code has unknown employment ages 30+	81,451	28,456	2,108,430
Drop if either ZIP code has incomplete data in any year	79,984	28,093	2,079,584
Exclusions due to our identification strategy			
Drop if ZIP codes are in different counties	49,725	27,859	1,292,850
Drop if ZIP codes are assigned to the same hearing office	1,003	1,146	26,078
Drop if either code has a non-unique assigned hearing office	907	1,049	23,582

above restrictions, we observe both SSDI enrollment and employment for all remaining observations. Finally, to ensure that our estimates are not distorted by compositional shifts in the set of ZIP codes present in our estimation sample, we then eliminate any border pairs for which any observation has been excluded in any of the preceding steps. These restrictions leave us with 79,984 surviving pairs, amounting to 92.9 percent of what we started with.

The remaining restrictions stem from our identification strategy, which harnesses differences in appeals processing time between nearby ZIP codes subject to the same state and county governments. Restricting our sample to neighboring ZIP codes located in the same county (and therefore the same state) leaves us with 49,728 pairs, and restricting further to those straddling hearing office boundaries leaves us with 1,003 pairs. Finally, we drop instances in which either member of a border pair is associated with multiple hearing offices. This leaves us with our analysis sample of 907 border pairs encompassing 1,049 distinct ZIP codes and 23,582 ZIP \times pair \times year observations.