The Place-Based Redistribution of Disability Insurance*

Amanda M. Michaud   Timothy J. Moore   David Wiczer

This Version: Sepetember 2024
Most Recent Version.

Abstract

There is substantial variation in Social Security Disability Insurance (DI) across the United States, with DI beneficiaries accounting from less than one percent to more than one-fifth of a county's working-age residents. We combine county-level data on DI applications and awards with measures of labor market characteristics, living costs and population health to model the geographic dispersion of DI. We find that local differences in health and income levels are about equally important in determining geographic differences, with living-cost differences also playing an important role. The DI program redistributes across space, delivering ex ante welfare gains to residents counties at the 90th percentile in terms of DI receipt that are more than twice as large as in counties at the 10th percentile. The place-based effects of DI are larger than most public policy initiatives specifically designed to support local economic activity.

*Michaud: Federal Reserve Bank of Minneapolis, Amanda.Michaud@mpls.frb.org. Moore: Purdue University and NBER, moore839@purdue.edu. Wiczer: Federal Reserve Bank of Atlanta and IZA, david.wiczer@atl.frb.org. The views in this paper are those of the authors. They do not represent an official position of the Federal Reserve Banks of Atlanta, Minneapolis, the Board of Governors, or the Federal Reserve System.

1 Introduction

Social Security Disability Insurance (DI) is a vital part of the social safety net, insuring 180 million Americans against economic hardship associated with work-limiting disabilities (Social Security Administration, or SSA 2021). Even though DI is a federal program with national policy rules and payment formulas, there is substantial geographic variation in the benefits it provides. At the tails, ten percent of the working-age population live in counties where approximately nine percent receive DI, while ten percent live in counties where that share is less than two percent (SSA, 2022).

In this paper, we document and model the county-level variation in DI applications and awards in order to determine the place-based welfare effects of DI. We combine detailed data on DI over a 19-year period with economic and health information available at the county level. We find that counties with high rates of DI applications and awards look quite different to other counties in terms of their economic, demographic and health characteristics. High mortality rates, low wage levels, high poverty rates, and low living costs are associated with high rates of both DI applications and awards.

To disentangle the relative importance of various factors, we incorporate them into a structural model of DI application choice and award determination. The model captures how place-based health and economic differences interact with program features and the application decision. We find that differences in health outcomes and income levels have large effects on both the choice to apply for DI and the probability of being deemed eligible. Finally, we relate the differences in DI outcomes to differences in the welfare value of the program, both across counties and for different types of workers.

We confront two key challenges to understand the key drivers of spatial variation in DI rates. First, we require detailed geographic data on DI outcomes. Most of the county-level variation in DI award rates occurs within states rather than across them, and sub-state variation allows us to account for state-specific factors like Medicaid and unemployment insurance. We have worked with SSA to compile a novel county-level data set that contains sex-specific information on DI applications, awards, and average payments. We merge this with county-level data on factors that are known to influence DI claiming behavior, including health (measured as mortality rates); income (measured as median wages and poverty rates); cost of living (measured as local housing costs); and healthcare costs (measured with a medical price index used by Medicare).

Second, the selection of types of workers onto DI is not apparent in data alone. For example, we document that the county-specific probability of receiving DI conditional on applying varied between 20% and 77% in 2000. This could be due to the inconsistent application of DI policies across Disability Determination Services (DDS) offices, which are state-based, or factors like county-level differences in residents' disability rates and vocational qualifications. It is difficult to distinguish the importance of each factor, as determining DI eligibility depends on detailed information and complex judgments (French and Song, 2014; Maestas et al., 2013). Furthermore, the *de facto* leniency of the determination process affects applicants' decision to apply.

We address the challenge of selection by modeling the choice between work and applying for DI explicitly. To do so, we build on established models of DI entry, such as Low and Pistaferri (2015a) and Kitao (2014), and also related to structural work on health and labor supply choices (De Nardi et al. 2024, Hosseini et al. 2022, Capatina et al. 2020). Similar structural approaches have been used to study a variety of policy and welfare questions (Seibold et al. 2022, Dal Bianco 2023). To these, we add detailed modeling of the DI administrative process following Michaud and Wiczer (2018). Finally, we add new features relevant to understanding geographic dispersion, including differences in the cost of living and prices for medical care. We also parsimoniously include alternative means of support from family-provided income (Autor et al. 2019, Lee 2020) and care or private insurance affects these decisions. The model captures meaningful variation in the data in terms of DI application and award rates. Our results provide distinct estimates of how economic and health motives interact in the application decision; how DI determination processes respond to the health status and vocational qualifications of applicants; and how the costs of applying otherwise differ across counties. Furthermore, the structure of the model allows us to elicit heterogeneous welfare valuations of the DI program and compute policy counterfactuals that include how government budget pressures affect all individuals through taxes.[1]

Our main findings speak to two areas of interest: What is it about a location that drives DI outcomes? And how does the value of the DI system vary across locations? In regard to the first question, our findings are summarized as follows. Variation in health is an important determinant of spatial differences, although variation in earnings is just as important. Aside from variation in program and residual local tastes, local health characteristics account for 45% of DI variation; local earnings differences account for 36%; and local living costs account for 19%. Several economic

[1]We abstract from how DI policy affects on labor demand and wages (Aizawa et al. 2022, Kim and Rhee 2022) or endogenous labor productivity (Millard 2022).

mechanisms account for these results. Poor health increases the disutility of work, making non-employment more attractive at lower income levels. Poor health also decreases work productivity, increasing the effective rate at which DI replaces wages. Lower local wages also increases the average replacement rate of DI, as the DI formula for converting past wages to DI payments is highly progressive. Finally, lower living costs reduce the utility loss when going onto DI.

Health and economic factors interact with both the DI application decision and the determination of eligibility. Areas that have high rates of DI applications not only have relatively poor population health, but also low wages and low living costs. For example, consider two counties in 2003. Davidson County in North Carolina had 5.2% of its working-age population of 98,000 receiving DI benefits and was near the median in terms of wages, mortality risk, and local living costs. In contrast, Etowah County in Northeastern Alabama had 7.2% of it's working age population of 63,000 receiving DI benefits, a mortality risk in the top 10%, wages close to the median, and living costs close to the 25th percentile. Our structural model allows us to predict how these factors interact with each other across their distributions. For example, the impact of an increase in the mortality rate on DI shares is higher for counties in poverty, and this relationship is non-linear. We find that equalizing health conditions alone would reduce variation in the DI award rate by 29.1%, while equalizing incomes would reduce it by 22.3%. Equalizing variation in the local price level would reduce variation in awards by 12.3%. Health is the most important factor because awards are most responsive to variation in health, whereas changing income and the local cost-of-living have approximately equal elasticity but there is more variation in income.

The second area of interest is to understand the efficiency of the DI program from a welfare perspective. Are DI payments the highest areas where they are valued most? Are the geographical differences in valuation driven by health-related needs, or due to the real replacement value of the program given local wages and cost of living? With regards to the first point, the valuation of the DI program has a range of $40,000 of present discounted value between the 90th and 10th percentile of counties. Within this top decile of program valuation, the average receipt rate is the 7th percentile, meaning a very high relationship between receipt and value of the program.

To understand the distribution, in the top 10% of counties, gets 72% more than the median county per capita. But these differences are even starker when we look at variation of lifetime welfare: there the top 10% get 91% more ex ante welfare.

We then solve for an "optimal" place-based policy by manipulating the local replacement rate.

This tool changes both the local generosity and also the composition of applications and awards. We find significant redistribution towards the low-receipt counties. These counties are underserved by the program both because they have too few recipients and because it is not generous enough. In the lowest receipt counties, they have the largest marginal gain from an additional dollar of spending.

This paper advances our understanding of spatial variation in social insurance. Spatial differences in DI have been documented since it started in the 1950s; Schmulowitz and Lynn (1966) coined the term "the disability belt" to describe the heavy concentration of DI beneficiaries in the Appalachian region, Mississippi Delta region, and nearby Southern states. Subsequent research has documented continued concentration of DI receipt (McCoy et al. 1994, McVicar 2006, Gettens et al. 2018). We show that the current geographic dispersion is historically large, and extend the analysis to application and award activity.

Information on both DI applications and awards, the non-DI data linkages and our model-based approach also allow us to characterize key drivers of DI growth at the local level, contributing to the literature on the determinants of disability insurance applications and receipt (e.g., Gruber 2000, Autor and Duggan 2006, Kreider 1999). Our focus on counties complements and extends previous research using DI geographic variation that has largely used state variation (e.g., Autor and Duggan 2003, Liebman 2015, Rupp and Stapleton 1995, Coe et al. 2011).[2] There is enormous variation within states that we incorporate into our analysis; for example, Virginia has three counties among the ten counties nationally with the highest rates of DI receipt and two of the ten counties nationally with the lowest rates. The model provides a framework for attributing the differences in DI to differences in other county characteristics, allowing the features to affect both the demand for DI and access to it.

We also quantify the large place-based differences in the welfare value from the program, building on existing studies of the welfare effects of DI (Bound et al. 2004, Chandra and Samwick 2005, Ball and Low 2014, Low and Pistaferri 2015b, Meyer and Mok 2019, Autor et al. 2019, and Deshpande and Lockwood 2022). The overall size of DI and the degree of dispersion results in estimated local welfare effects that are much larger than for specific place-based policies or state transfer programs (Colas et al., 2021). These results highlight that a federal program focused on

[2] Some studies have used county-level variation to examine specific determinants, such as local labor market demand changes induced by the coal boom and bust (Black et al., 2002) and world oil and gas prices (Charles et al., 2018).

protecting individual risks can have large, economically important effects that are amplified by geographic heterogeneity and cross-subsidies.

2 The Geography of Disability Insurance

There is substantial geographic variation in the receipt of federal disability insurance. In this section, we describe the key features of this variation, including the extent of differences across counties, the persistence of these differences over time, and the distinct roles of application behavior and award probabilities in determining these differences.

We then show that regional or state factors can only explain a minority of county-level variation, as can standard urban/rural classifications. We next outline how local differences may result from the interaction of national program rules with local labor markets, as well as from differences in residents' health characteristics, economic risks, and the value providing by disability insurance through local prices. Guided by economic theory, we look for correlations between these characteristics and DI outcomes to motivate the inclusion of these margins as potentially quantitatively significant factors in our theoretic framework.

Receipt of DI. We start with the overall number of DI primary beneficiaries, as they determine the program's size and impact.[3] We examine county-level differences using data from an annual publication produced by the Social Security Administration (SSA) that includes counts of DI beneficiaries for every county.[4] These are available for all counties from 1970 onwards, which allows us to provide a broad and long-term picture of the spatial variation in DI.

County-level DI beneficiary numbers are available for ages 18-64 years, which we scale by the resident population for that age range. The distribution of rates over the 2001-2010 period are shown in Figure 1. Although the average DI beneficiary rate is 4.65% at the county level, half of Americans live in counties where the DI rate is less than 4.25%. The mean is driven by a thick tail of counties with DI rates that are two to three standard deviations above the others. The range

[3] All of the statistics in this text are only for workers applying for and receiving DI. We do not include spousal or dependent benefits. Spatial patterns for other recipients are highly correlated with primary beneficiaries.

[4] The SSA publishes *OASDI Beneficiaries by State and County*, which we have collected, digitized and cleaned.

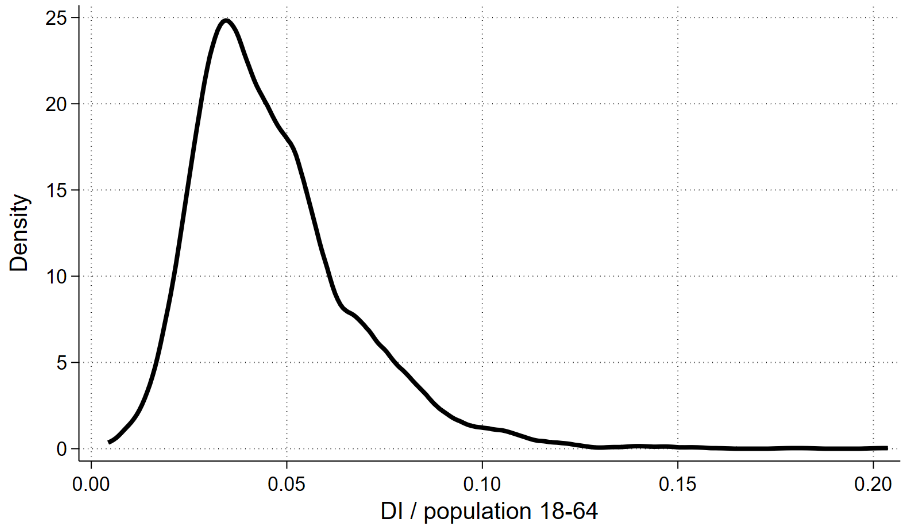

Figure 1: Distribution of DI Beneficiaries per 18-64 pop., 2001-2010.

is wide across the counties, from less than one percent to more than 20% of working-age adults receiving DI.

Figure 2 shows the map of these rates for the United States. High and low DI counties do not appear to be randomly distributed. They are also not a new phenomenon. The term ”disability belt” was coined in early research to describe the high concentration of DI beneficiaries in Appalachia and the Mississippi Delta that is evident (McCoy, Davis, and Hudson, 1994). Pockets also exist in the North, in the Great Lakes Region and the Pacific Northwest. The counties tend to be rural, but it is not the case that density has a consistent relationship to DI activity. While we do not study the time trend in spatial variation, it is noteworthy that regions with high DI have had persistently high DI almost since the inception of the program and the spatial disparity has only widened as DI rates have risen nationally. In short, geographic disparities are not random. They are not passing fads. There is something about place that is intimately close to DI outcomes.

Why are we interested in extreme outcomes, the less-populated tail of counties high in DI? From a starting point, the DI program is an important safety net. Disability payments are deeply consequential for beneficiaries and the consequences are even higher for individuals in poor health who do not utilize the program for one reason or another. Beyond welfare considerations, the tail counties are important for understanding the function of the DI system as a whole. Table 1 shows

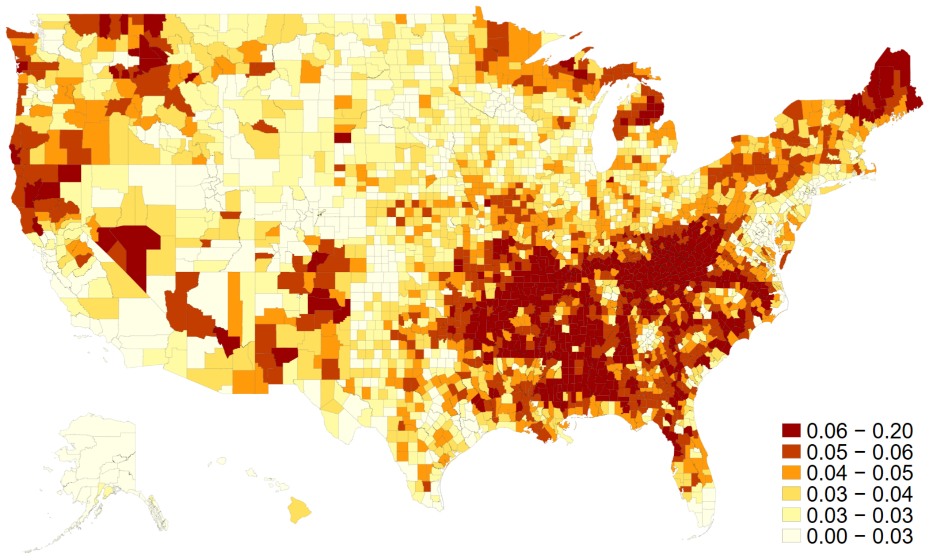

Figure 2: DI Beneficiaries as a percent of the population age 18-64, 1996-2014.

that counties with high rates of DI account for significantly more of the beneficiaries, applicants, and even benefits paid than their share of adult population. Thus, rural and other smaller areas are important for understanding national levels and trends in the program.

The use of county-level data unveils new insights into the extent of spatial disparity even within states. The state of Virginia is a striking example. Virginia contains both three of the ten counties with the highest DI beneficiary rate in the nation and two of the ten counties with the lowest beneficiary rate in the nation. This variation shows that differences in DI rates do not come from differences in state-level factors alone. There are important interactions with state-level programs, such as Medicaid, in the application decision. The determination process also has a state-level component as DDS operate mostly at the state level. Yet the maps of Virginia make bold the national trends: DI tends to be higher in poorer and rural counties. Even rural and poor is not the end of the story. There is significant additional variation that we hope to account for or to at least measure how much variation is left to be understood after the key economic and health considerations we model are accounted for.

DI applications and awards. We now turn to data on DI applications and awards. Our administrative data on DI applications and awards come from the SSA Disability Research File (DRF).

Table 1: Distribution DI metrics across counties.

Fraction of US working-age population in counties with highest DI rates	Account for what fraction of... all DI beneficiaries	all DI applications	all DI benefits paid
10%	19.7%	14.9%	19.1%
20%	34.3%	28.8%	33.5%
33%	50.1%	44.5%	49.1%
50%	67.3%	62.7%	66.4%
67%	80.8%	78.2%	80.2%
80%	89.8%	87.6%	89.4%
90%	95.1%	94.3%	94.9%

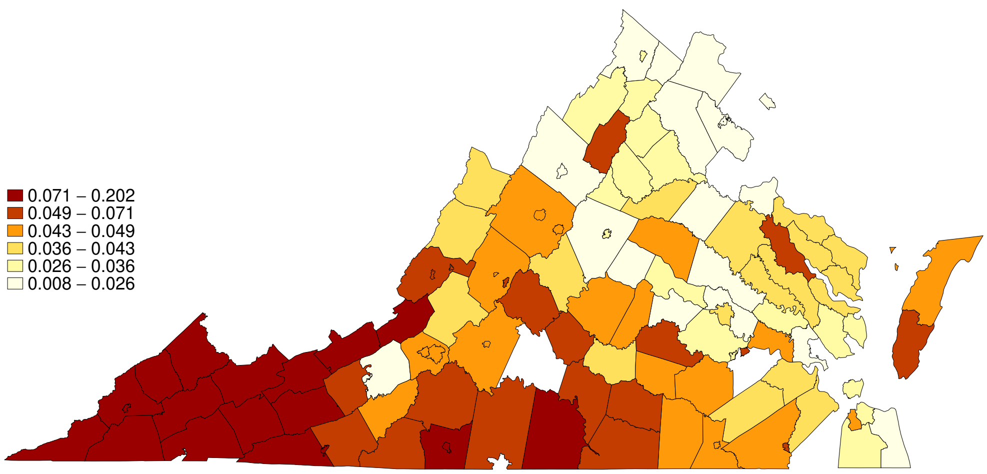

Figure 3: DI Beneficiaries per Adult age 18-64, 1995-2016.

The DRF is a data file designed to track cohorts of individuals filing for DI through the disability decision and appeal process. It is constructed by drawing on multiple administrative data sources, and updated annually. The DRF allows the status of a claim for DI to be tracked throughout the adjudicative steps, as well as providing key demographic information about the applicant, including their county and state of residence, as well as their sex and age. It has been used by other researchers to examine different aspects of DI (e.g., Costa 2017; Foote et al. 2019; Foote et al. 2019).

We obtained geographically defined counts from the DRF for claims filed from 1996 to 2014. We restrict the data to applicants aged 21 to 64 years, as 65 years was the Full Retirement Age at the beginning of the sample period (and we need consistent age ranges for linking the other data). All of the outcomes are organized in terms of the date of filing (i.e., we measure award outcomes by year of application, even if the claim is actually allowed in a subsequent year).

The role of policy rules. We now consider what may affect the spatial variation in DI, starting with policy rules. DI is a federal program with nationally consistent rules and policies. Insurance coverage and the definition of medical eligibility are set at the national level.[5] Likewise, there is a common approach to determining the current value of an applicant's earnings history, and a nationally consistent formula to map that value into a monthly cash benefit.

Such policies can interact with local labor market conditions in important ways. The DI payment formula is progressive, providing higher wage replacement rates for low-wage workers by having an initial marginal wage replacement rate of 90% for each dollar in past earnings that decreases to 32% at a certain earnings value and then 15% at a higher one.[6] In 2020, it meant that someone with average past earnings of $10,000 had 90% of their wages replaced, while replacement rates are 65% for $20,000, 45% for $50,000, and 33% for $100,000 in average past earnings.

A spatial implication of this approach is that the average replacement rates in low-wage labor markets are much higher than in high-wage ones. Consider the replacement rate at different counties' average wage levels. The replacement rate is 50% at the median, but ranges from 37% to 65% across the counties in our sample. Thus, typical workers receive very different levels of support from DI by location if they are disabled and unable to work.

Other factors. In addition to these policy features, there are also important geographic differences in the indirect determinants of the number of DI beneficiaries, including population health, living costs, and employment opportunities. Together, these differences point to multiple channels through which geographic differences can lead to markedly different DI outcomes.

[5]To be insured for DI, an individual needs needs to have Social-Security-covered earnings in at least five of the previous ten years. Medical eligibility depends on being unable to engage in any "substantial gainful activity" because of a physical or mental impairment that is expected to last at least 12 months or result in death. Substantial gainful activity is judged based on a national earnings threshold that was $1,260 per month in 2020.

[6]The monthly DI payment is known as the Primary Insurance Amount (PIA), while Average Indexed Monthly Earnings (AIME) is based on a beneficiary's past earnings. In 2020, the marginal rate at which PIA depends on AIME changed from 90% to 32% at $960 in AIME, and from 32% to 15% at $5785 in AIME. See Gelber et al. 2022 for more details.

Key candidate factors to account for the spatial variation come from knowledge of the program combined with economic principles. A first pass of the data is to construct correlations to see whether these factors appear to be quantitatively significant.

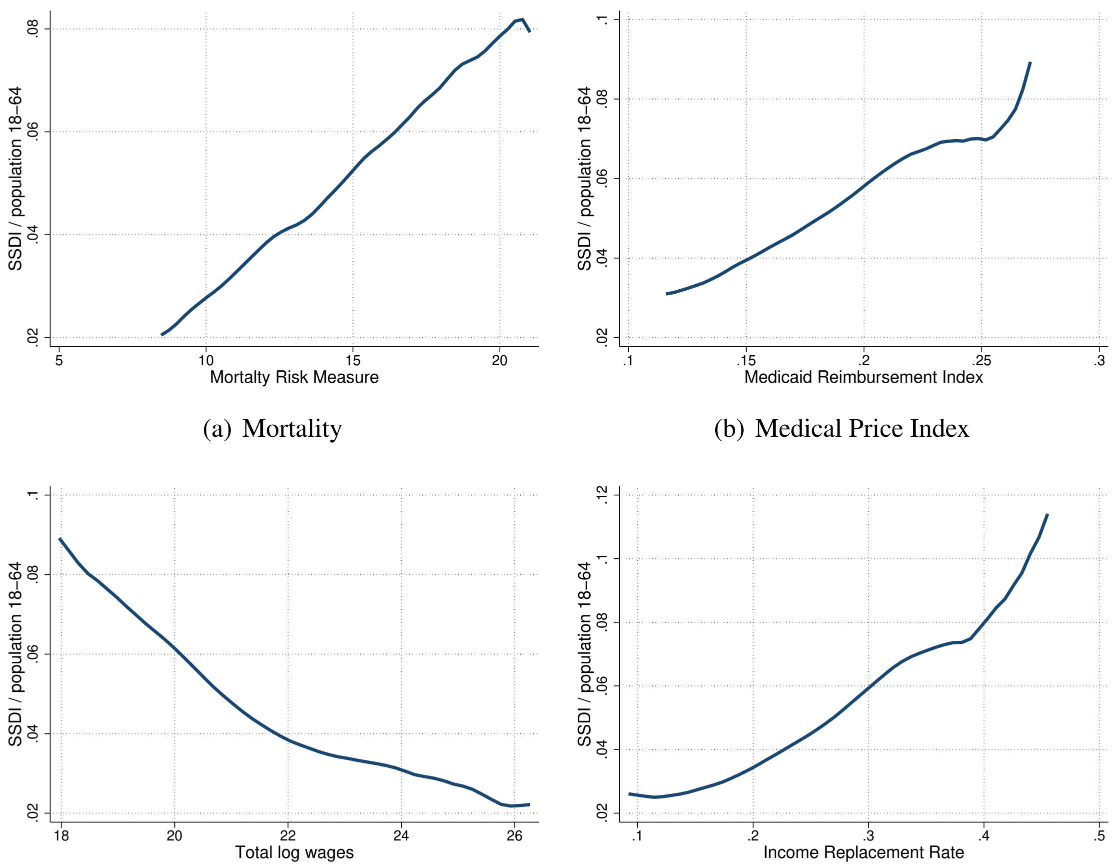

(a) Mortality

(b) Medical Price Index

(c) Log Wages

(d) Replacement Rate

Figure 4: Correlates with Beneficiary Rates over 1995-2016.

The first factor is health. The DDS determination requires that the applicant have a significant work-limiting disability expected to be permanent or result in death. While there is a subjective component to this definition, death is certainly an objective and observable outcome that fits the DDS criteria. Panel A of Figure 4 shows the relationship between county-level mortality and beneficiary rates. Mortality rates are constructed using a compilation of deindentified death records compiled from state vital statistics bureaus combined with population estimates computed using the Current Population Survey.[7] It is unsurprising that the relationship is positive but the scales and slope are striking. The range of mortality risk doubles across counties and a doubling of mortality risk is associated with a more than doubling of DI rates. We can expect variation in health to be a key part of our analysis.

The economic incentives provided by the program can be thought of in terms of income and substitution effects. The income effect is how far disability payments go in terms of consumption value. The average beneficiary receives a payment of around $1,200, monthly. Panel B of Figure 4 shows the relationship between the housing price index, a measure of the cost of living, and DI rates. In general, the relationship is negative: higher HPI correlates with lower DI rates. This correlation is weaker than our other economic factors and is non-monotone. This suggests that while DI beneficiaries could move to a lower cost-of-living county to stretch benefit dollars further, many do not. This could be for a variety of reasons such as being too financially constrained to pay the cost to move or wanting to be near friends and family for assistance with activities of daily living. For whatever reason, it suggests that understanding the mobility of DI recipients is not first-order in understanding how the program reallocates welfare across space.

The substitution effect is the amount of money benefits paid versus what an individual could earn if they continue to work despite their health burden. For a fixed benefit level, the spatial difference in the opportunity cost of going on DI is partially captured by the prevalent median wage rate in an area. Panel C of Figure 4 shows that an increase in wages is associated with a decline in the beneficiary rate. The decline is sharper for increases at the bottom of the distribution and larger in the upper half.

Variation in the substitution effect is also built into the benefit formula. The replacement rate is progressive. There is a minimum payment that increases less than linearly in an index of past earnings up to a cap. As a result, an individual with consistent earnings at $30,000 per year has a replacement rate of 82% in 2022 and one earning $100,000 has a replacement rate of 42%. Panel D of Figure 4 shows the top tail of counties has a replacement rate relative to the median earnings of around 50% while the median replacement rate is only 38.5%.

These health and economic factors are strongly correlated with DI outcomes but are also strongly correlated with each other. Take the case of Virginia shown in Figure 5. Areas that have low wages, populations with poor health outcomes, or low cost of living tend to have high

[7] Mortality data are provided by the Institute for Health Metrics and Evaluations.

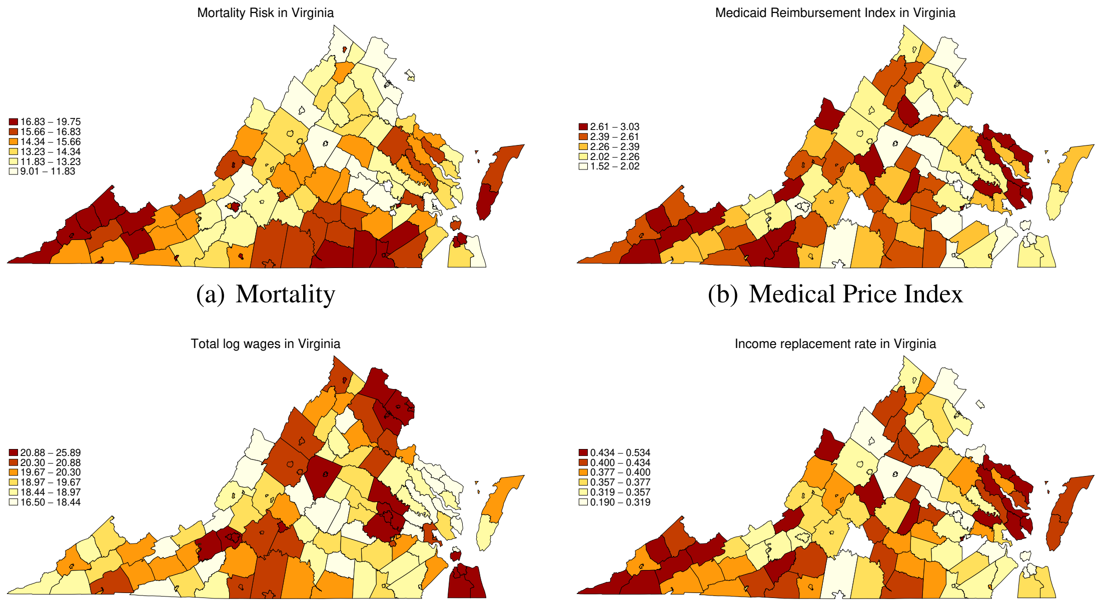

(a) Mortality

(b) Medical Price Index

(c) Log Wages

(d) Replacement Rate

Figure 5: Correlates with Beneficiary Rates over 1995-2016.

beneficiary rates. The correlations with beneficiary rates are -0.53, 0.59, and 0.50, respectively. Likewise, these factors are correlated with each other. Areas with low wages tend to have poor health outcomes and low cost of living; the correlations are -0.20 and -0.86, respectively. There are also state-level factors such as variation in processing times across DDS and interactions with other welfare programs such as Food Stamps and Medicaid. Disentangling the independent relationships of these factors and their interactions with each other in the DI application decision requires structure to be placed around how that decision is made. We provide this structure using decision theory with these margins in mind.

3 Model

Our quantitative model adds a notion of location in addition to typical factors such as health, age, and income common to structural models of DI. We focus on the DI application decision and award process and do not model trade or migration between locations. Locations instead are used to provide insight into newly explored quantitative margins. They provide rich variation in the typical factors and also cost of living and DDS award process that we will use to estimate the model and judge it's fit.

Time is infinite and discrete. The model has the following features:

Population characteristics. The model features a distribution of agents aged $t \in 0, 1, 2...T$. Each agent is characterized by their location ℓ; health status d; earnings potential y; labor force status s; and DI application history e. Their location is a fixed characteristic.[8] The other characteristics evolve according to exogenous shocks and each agent's choices. An agent's risk of death depends on their age and health, according to the probability $(1 - \phi(d, t))$.

Choices. Agents make choices about their labor force participation and whether to apply for DI benefits. A savings decision is not featured in the model. However, the flow consumption provided to disabled individuals not on DI matches the empirical findings in Meyer and Mok (2019) that disabled men experience a decline in consumption equal to approximately 90% of their decrease in net income. Employed agents choose whether to continue working or quit, and non-employed agents can choose to return to work.

An agent must be non-employed to apply for DI. Eligible agents who choose to apply for DI incur a specific cost, ψ, and receive an award with probability $\xi_\ell(d, t)$. The probability of award is modeled to reflect the DDS vocational grid. It is a function of age, health, and locality.[9] Applicants who are rejected may not apply again.[10] Finally, r is the indicator of a past application.

Income. Employed agents earn a flow income, y, that changes over time. There are finitely many values of y that evolves according to a Markov chain: $Pr[y' = y_j | y = y_i] = \mathcal{Y}_{\ell,s}(y_j, y_i)$ that depends on the agent's employment status s and his or her location ℓ.

Non-employed agents receive non-labor income either from social security disability insurance, or other means. Social Security Disability payments are location specific, denoted DI_ℓ.

[8] There is empirical support for this simplification, with research generally finding no or small migration effects of place-based incentives created by welfare and social insurance programs(Kennan and Walker 2010, Schwartz and Sommers 2014, Jia et al. 2023).

[9] Locality is a factor since local labor market conditions are considered and determinants are made by local DDS.

[10] This is a known simplification but is necessary for making data and model component consistent. This is because in data we observe an application and award count, which could not be mapped into a model in which households repeatedly applied each period.

The location specificity reflects considerably more generous replacement rate relative to average employment earnings in some locations than others. DI benefits also include the reimbursement of medical costs through medicare is provided to DI beneficiaries. Other transfers T_ℓ are paid to individuals who are neither working nor on DI. In spirit, these transfers capture other government programs such as TANF or WIC, or informal transfers from family, friends, or non-profits. These transfers are location-specific and inferred through the model to form the outside option of not-claiming DI even when non-employed.

Employed agents may choose to quit into non-employment, which happens especially as they get underlying income shocks and as their health deteriorates. Non-employed agents who do not apply to DI receive the option to work at a job-finding probability λ. Job search is free and non-employed agents who don't want to work do not need to accept this job offer.

Health and its consequences. The health, d, depreciates stochastically for each agent at a location-specific rate, a process described by π_ℓ such that $\Pr[d'|d] = \pi_\ell(d)$. As a simplification empirically and within the model, health will be indexed by mortality probability, so that the probability of death at any point is d and hence $d \in [0, 1]$. Health also affects an agent's disutility of working in the labor market. This disutility exists only for employed agents, $s = 1$

Non-employed agents who are not on DI incur flow medical expenses $m_\ell(t, d, s)$ each period.[11] This is to say, that for agents with $s = 1, 3$ $m_\ell = 0$ but with $s = 2$, agents are exposed to the full medical costs. Medical expenses, just like income, depend on an agent's location to capture the different costs of medical services in different counties. But, when working or when receiving DI, the agent does not have to pay these expenses.

The final way health enters the model is through the DI allowance probability. As discussed above, this probability is increasing in d, representing the severity of the agent's work limitation.

Preferences. Preferences are time separable and captured by a utility function $u_s(c, t, d)$ and a discount factor β. The flow utility function is increasing in consumption c, and decreasing in the extent of disability d and age t if the agent is employed. Its dependence on age t allows us to match

[11] We model medical expenses as a steady flow despite the reality that many expenses are lumpy. This is because we have no savings or borrowing in the model and do not want to overstate the welfare impact of a medical expense shock.

the age profile of labor force participation. It is also a function of labor force status s to capture potential costs associated with working in the market and these conditional costs may vary with health status d.[12]

Value Functions. We subscript value functions by age and location t, ℓ and its states are (y, d, s, r). For an employed agent ($s = 1$) the value function is given by

$$
\begin{aligned}
V_{t,\ell}(y, d, s, r) = \max_{(q \in \{0,1\})} \; & u_s(c, t, d) + \beta(1-d)E_t[V_{t+1,\ell}(y', d', s', r)] \\
& p_\ell c = y \\
& d' = \pi_\ell(d) \\
& s' = \begin{cases} 1 \text{ if } q = 0 \\ 0 \text{ if } q = 1 \end{cases}
\end{aligned}
$$

Employed agents choose q whether to continue to work or to quit to maximize their flow utility and their expected continuation utility provided by the value function $E_t[V_{t,\ell}(y', d', s', r)]$. The decision is constrained by their budget. They earn labor income $w(y, s)$ which is dependent on income shock y and their employment status s. They spend $m_\ell(t, d, y)$ of their income on medical expenses and consume whatever is left over adjusted by a local price index p_ℓ.

The problem when $s = 2$, a non-employed agent who is not a DI beneficiary, is as follows, omitting non-relevant decisions and objects.

$$
\begin{aligned}
V_{t,\ell}(y, d, s, r) = \max_{a \in \{0,1\}} \; & u_s(c, t, d) - \psi a + \beta(1-d)E_t[V_{t+1,\ell}(y', d', s', r')] \\
& p_\ell c - m_\ell(t, d, y) = T_\ell \\
& d' = \pi_\ell(d) \\
& Pr[s' = 1 | s = 2] = (1-a)\lambda \\
& Pr[s' = 3 | s = 2] = a\xi_\ell(d, t) \\
& r' = \mathbb{I}_{(a=1 \cap r=0) \cup r=1}
\end{aligned}
$$

[12] This concept of utility is often used in the structural study of disability and appears in papers.

A non-employed agent chooses whether to apply for DI if they have not already applied, $r = 0$. Their award probability depends on their state through $\xi_\ell(d,t)$. If they are rejected they remain non-employed and $r' = 1$, so they cannot apply again. In this state, their income is just the location-specific transfer T_ℓ and they also have to pay for medical expenses, m_ℓ.

An individual receiving DI benefits ($s = 3$) has no remaining choices. They cannot leave DI and so labor market state s is absorbing. They consume their income, DI_ℓ but do not pay for medical expenses, and their health depreciates exogenously. Their continuation value equals the following.

$$
\begin{aligned}
V_{t,\ell}(y,d,s,r) &= u_s(c,t,d) + \beta(1-d)E_{t,y',d'}[V_{t+1,\ell}(y',d',s,r)] \\
p_\ell c &= DI_\ell \\
d' &= \pi_\ell(d)
\end{aligned}
$$

4 Calibration and Estimation

The model is both calibrated and estimated to enable quantitative analysis. First, functional forms are chosen. Key functions include utility and the DI determination process. The length of a period is one month. The discount factor β is set to 0.996, consistent with an annual discount rate of 5%.[13]

The utility function takes the following form:

$$
u_s(c,t,d) = \frac{(c)^{1-\sigma}}{1-\sigma} + \left(\lambda_0 + \lambda_t t + \lambda_d \frac{d}{\bar{d}}\right)\mathbb{I}_{s=1}
$$

This specification includes utility over consumption c according to constant relative risk aversion parameter σ which we set to equal two. The second term in the utility function captures non-pecuniary costs incurred if an individual is employed ($\mathbb{I}_{s=1}$). The baseline cost is λ_0. The cost of working is linearly higher for older workers (λ_t) and for workers in worse health λ_d. This is a departure from Low and Pistaferri (2015a) and subsequent work where the marginal utility of

[13]Figures depicting distributions of additional parameter values and measures of model fit are contained in Appendix Section B.

consumption is affected by employment status and health.[14] This matters for welfare analysis: the higher the complementarity then the higher the welfare gains provided by DI, a cash transfer.[15] Despite the change in the utility function, our model does include a link from health to consumption through health expenditures in the budget constraint. This implies that the marginal utility of consumption is increasing in poor health. Thus, we are closer to complementarity between good health and income than is obvious from our preferences alone.

Health is a continuous variable that is represented by its mortality rate. To estimate the process at the county level, we use an exponential decay functional form so that for each location ℓ there are two parameters: a baseline rate ς_ℓ^0 and a rate at which mortality increases with age t ς_ℓ^1. The mortality hazard rate takes the form

$$\frac{\varsigma_\ell^1 e^{-\varsigma_\ell^0 + t\varsigma_\ell^1}}{1 - e^{-\varsigma_\ell^0 + t\varsigma_\ell^1}}\,.$$

We fit this to the mortality rates of 40- and 70-year-olds in each county. To further capture the location's demographics we match the fraction of the population that is between 40 and 50 and the fraction between 50-65, thus we capture the basic way in which some counties have more individuals of advanced age and of poor health.

Parameter	Value	Target
ξ_d	10.481	$corr(\xi_\ell^{\text{model}}(d,t), \xi_\ell^{\text{data}}(d,t))$
ξ_t	6.313	
λ_0	-2.610	Age profile of employment
λ_t	-2.467	
λ_d	1.0	Pr[death | worked last year] = 0.371
τ_{ss}	0.0207	balanced budget
σ	2.0	Standard IES
β	0.996	Annual discount rate of 5%
replacement rate	0.3	XX Where do we get this from? XX

Table 2: Country-wide parameters and their targets, though all are fit jointly

Medical expenditures are estimated in a two-step process. We use national data from the Medical Expenditure Panel Survey (MEPS) to estimate an average national baseline medical expenditure

[14]Whereas the goal in Low and Pistaferri (2015a) was to replicate life-cycle and cross-sectional patterns of health, consumption, and DI behavior, we match fewer life-cycle and cross-section statistics and instead introduce many more parameters to match rich spatial variation.

[15]There is, however, some disagreement in the literature as to whether health and consumption are substitutes or complements (Russo (2022)).

for people in the 40-65 age group[16]. We then adjust these values to be location-specific using Medicaid reimbursement data, which includes a local medical price deflator implied by the real value of medical services they provide.

The income process is two-state Markov process with a grid point for normal income and another for a low-income state denoted as "poverty." The process is chosen so that the average duration in poverty matches the national average of 2.4 years. The entrance into the poverty state is county-specific and is chosen such that the stationary distribution matches the poverty rate in each county from the Census' Small Area Income and Poverty Estimates.

We use two county-specific sets of parameters to exactly fit county applications and awards. The first is county-specific income from other sources (T_ℓ). The second is a DDS-specific intercept (xi_ℓ) chosen to exactly fit the probability that an application is given an allowance in that DDS region.

The estimated values of income from other sources (T_ℓ) are largely uncorrelated with other observables, just like regression residuals ought to be in a correctly specified regression model. We demonstrate this in the left panel of Figure 6 with the counties' relative health, its rank in the rate at which health depreciates, shown against the T_ℓ. The interpretation is that our model mostly characterizes the way that underlying driving factors like health and poverty differences contribute to the application choice and the set of T_ℓ captures idiosyncratic county differences orthogonal to health and economic conditions. Our choice to model these residual differences using the budget constraint is relatively innocuous given that households cannot save.

The estimated distribution of xi_ℓ, the DDS (state) specific increase in probability in an award given an application are fairly concentrated. This is shown on the right panel of Figure 6 where most of the mass is within a few percentage points of the lowest approval-rate DDS. This implies that the likelihood an application from a similarly aged and disabled individual is awarded is similar across counties, save only for some notable outliers with low award rates signified by the mass around zero.

National-level parameters must are estimated to best fit population-weighted cross-county averages of two sets of targetsWe match the age profile of labor force participation with utility function

[16]While we indexed m_ℓ by d throughout, it is inconsistent to use the realized expenditure by realized mortality from the data because so much of health spending is time-shifted and mortality is also evolving. Hence, we use a single expenditure fraction

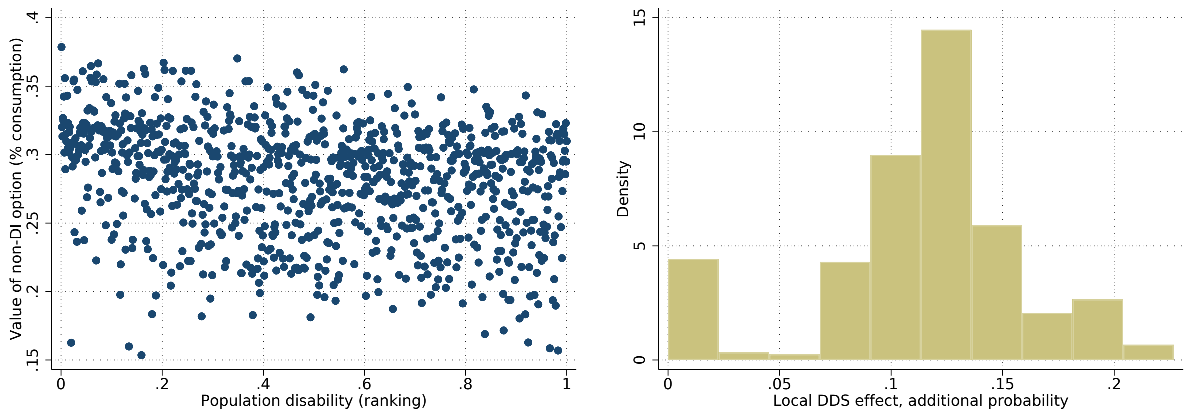

Figure 6: Non-Emp consumption value (LEFT) and distribution of DDS's ξ_ℓ (RIGHT) .

parameters, λ_0, λ_t, and the health level with νd. Simultaneously, we match the national allowance rates using the vocational grid as a guide. Specifically, we calibrate a logit-style functional form for the allowance rate

$$\xi_\ell(d,t) = \frac{\exp(\xi_d \frac{(d-\bar{d})}{\bar{d}} + \xi_t \frac{t-50}{65-40} + \xi_{DDS(\ell)})}{1 + \exp(\xi_d \frac{(d-\bar{d})}{\bar{d}} + \xi_t \frac{t-50}{65-40} + \xi_{DDS(\ell)})}$$

The $\xi_{DDS(\ell)}$ parameters give enough degrees of freedom to create the match in Figure 17 while the two other parameters ξ_t, ξ_d minimize the within-state cross-county differences in allowance rates. Of course, there are only two parameters to fit hundreds of counties' allowance rates, hence, we cannot obtain a perfect fit. The empirical correlation between county-level aging patterns, health and DI receipt bounds the goodness-of-fit that we could get from our model, given this functional form by a simple logit regression of these empirical features minus the DDS-specific fixed effects, which gives a sum of absolute allowance rate differences of 0.087. The within-model deviations approach that bound, with 0.091 with parameters $\xi_t = 2.071$ and $\xi_d = 2.924$.

To see exactly how we can use a quasi-logit structure of participation consider the utility value difference of working or not in a period which can be written as:

$$\frac{c_t(\ell_t = 1)^{1-\sigma}}{1-\sigma} + \nu d \ell_t + \phi \ell_t + \beta V_{t+1}(\cdot) - \left(\frac{c_t(\ell_t = 0)^{1-\sigma}}{1-\sigma} + \beta V_{t+1}(\cdot) \right)$$

Then we can target the fraction of employees who worked last period conditional on having died. The model equivalent is:

$$\Pr[\text{Worked}|\text{Died}] = \frac{\sum_i \frac{e^{V^{workdif}+\nu d_i}}{1+e^{V^{workdif}+\nu d_i}} \times d_i}{\sum_i d_i}$$

where the right-hand side has the model-implied difference in the value of working vs non-work, which depends on ν additively. d_i are the population-weighted probabilities of dying and we solve this non-linear equation for ν.

To see our goodness of fit, see Figure 17. We capture well the approximate halving of labor for participation during the age period we model, 40-65. In our ξ parameters we are able to capture much of the variation across counties. Some of this comes from the exact fit at the DDS-level given by $\xi_{DDS(\ell)}$ while we capture the differences cross-county by their different characteristics. The residual correlation in allowance rates is about 60%.

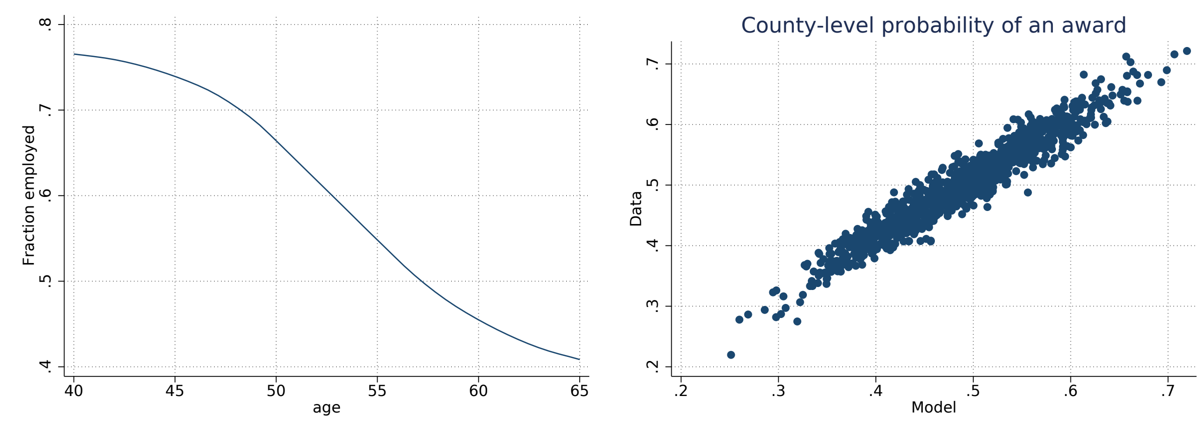

Figure 7: Fit of aggregate parameters shown by laborforce participation profile (LEFT) and county-level allowance probabilities (RIGHT) in the model and data

5 Decomposition of county differences

In this section, we present the first quantitative results of our model. We break down the causal drivers of dispersion in the DI award rates across counties and then look at the program's implications for welfare across space and across people. The largest differences in both DI utilization and welfare gain come from health differences across location. This is in part because health enters the allowance probability, but it is also because health has such a marked effect on the desirability of work.

Our first set of quantitative results decomposes the differences across counties in application and allowance rates. In our calibration, counties differed in many of the features that would drive DI differences, and in this section, we isolate the effect of each. Using the model for the decomposition allows us to isolate the particular driving forces.

Here we turn off the different fundamental differences across counties and compare the counterfactual economy to the baseline calibration. Specifically, we can set each county to have (1) the same health and demographics, (2) the same baseline income processes, (3) the same baseline set of prices or (4) the same medical costs.

Specifically, when we change the health process by making every county's health depreciation rate the same as the average county's health depreciation rate. Of course, if we are manipulating the health process, we also have to manipulate the aging process since the two are set together: health depreciation was set to match the change over ages. To keep mortality rates the same across counties, we also set every county to have the same aging process. Otherwise, we would still have health differences because of the natural differences associated with aging.

To make the income process the same, we are setting two model objects to be identical: the likelihood of a poverty shock and the nominal value of SSDI benefits. For both of these values, we set every county equal to the median poverty rate across counties. Finally, to make medical prices and goods prices the same, we simply set them to the average.

To augment this evidence, we compute the elasticity of awards and allowances to these driving forces. This is not a perfect elasticity because income and health both have two, linked dimensions over we are reducing heterogeneity. So we compute it in terms of average d for health and in terms of average y for earnings. This is a measure, somewhat akin to that which we would see in a regression using county-level differences in the characteristics, but instead we are varying these factors in the model. This is shown in the second panel of Table 3.

Notice in Table 3 the outsized role that non-health factors play in applications. However, health is the strongest determinant of an award being granted, but the effect comes more so in awards than

	Health/Demographics	Income/DI replacement rates	Local cost-of-living
Contribution to normalized IQR			
Awards	29.1	23.3	12.3
Applications	36.1	39.4	15.9
Elasticity			
Awards	6.63	2.99	3.15
Applications	5.31	4.25	4.09
Normalized IQR			
	0.049	0.28	0.14

Table 3: Effect of county-level characteristics on DI outcomes

applications. This makes sense because of the application process: Because most very unhealthy applicants are given an award so an increase in bad health and an increase in award go nearly one-for-one, whereas the other factors bring marginal applications with a lower likelihood of receiving an award.

The non-health components are also very important, especially as they contribute to the variation in applications. Looking first at the role of earnings differences, this is the largest driver of variation in applications. This is mostly because of the very large differences in earnings across space, particularly among the population at risk of going on SSDI. Recall that earnings dispersion is tied hand-in-hand with our measure of the replacement rate of SSDI because what matters to the application decision is the relative earnings compared to that which they could get in an award.

Price differences are also particularly important and again, this is because they affect the real value of the replacement rate of DI. So essentially, changes in the price-level are also changes in the replacement rate. Just as with the role of income, the local cost-of-living changes applications more than it does awards because it also changes the composition of applicants to lower-likelihood awardees. As the local cost-of-living goes down, more households choose to apply, but this set is not as old or sick than those who were applying anyway.

In the second and third panels of Table 3 we include the elasticity and normalized interquartile range of these factors to get a better sense of why each contributes. Notice that both applications and awards are very elastic to differences in health and demographics. However, there is slightly less heterogeneity across space in these features. Instead, there is a great deal of heterogeneity in income across space that then leads to large overall from this factor. Notice also that the application decision is more elastic than award decision for the income and price-level factors, but this is flipped for health. Again, that is because changes in health actually bring in applicants who are more likely to receive a DI award whereas the monetary differences bring in applicants who are less likely to get one.

5.1 Calculating welfare across counties

Given the county-level differences in exposure to DI, and the different underlying conditions, the program will have significantly different value in different counties. Of course, these differences are related to the different receipt rates, but not precisely because the same dollar has a different value in different places. The model quantifies these differences as consumption-equivalents between a household who is able to apply for the program and one who is not, i.e. $r = 0$ or $r = 1$. We then compute welfare differences by exploiting the quasi-linearity of our preferences. We solve for a percent of labor that the households would give up, a ϕ such that

$$V_t(y, d, s, r = 0) = \sum_{\tau=t}^{T} \beta^\tau c_\tau^{1-\sigma}/(1-\sigma) - \mathbb{I}_{\ell_\tau}(\lambda_0 + \lambda_1 \tau + \lambda_d \frac{d_\tau}{\bar{d}}) =$$

$$\sum_{\tau=t}^{T} \beta^\tau c_\tau^{r=1\ 1-\sigma}/(1-\sigma) - (\phi + \mathbb{I}_{\ell_\tau^{r=1}})(\lambda_0 + \lambda_1 \tau + \lambda_d \frac{d_\tau}{\bar{d}})$$

$$V_t(y, d, s, r = 0) = V_t(y, d, s, r = 1) + \frac{1-\beta^{T-t}}{1-\beta}\phi(\lambda_0 + \lambda_1(T-t) + \lambda_d \sum_{\tau=t}^{T} \frac{d_\tau}{\bar{d}})$$

We then scale the labor value of DI to dollars by multiplying by the average wage in the county and sum them over time to lifetime asset value to be consistent and comparable to Autor and Duggan (2006). In those estimates, they considered the Medicaid-inclusive value of benefits, discounted from receipt to the age at which retirement benefits take over. We first present welfare calculations akin to those, discounting the consumption-equivalent benefits for all recipients. Note that, while the median is quite similar to their figures: right under \$300,000, the tail extends quite a bit higher.

This is shown in Figure 8, where there is both a surprisingly large dispersion in welfare gains and the long-right tail of values. It is highly correlated with the fraction of workers who actually receive benefits in the county. This correlation stems both from underlying conditions, whereby if a place has an older and less healthy population, the value of receipt relative to the alternative

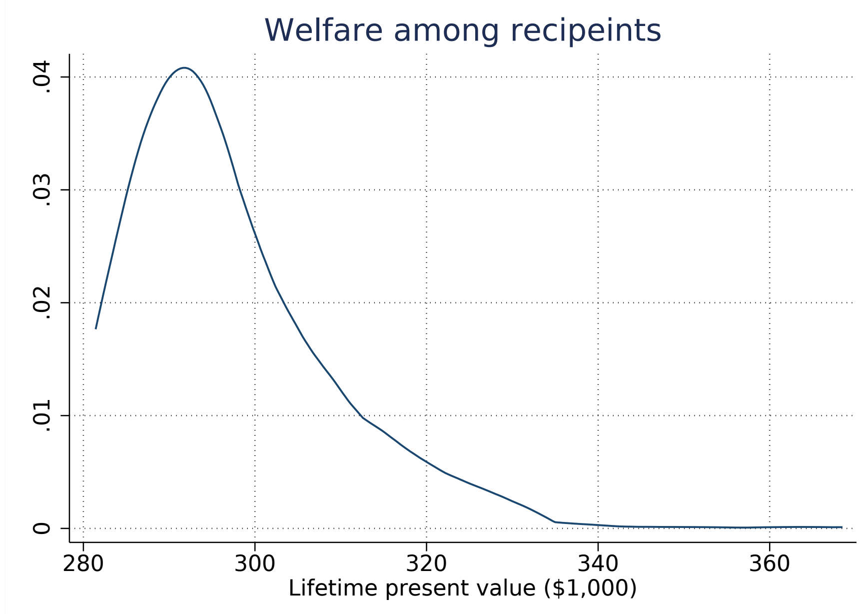

Figure 8: Cross-county Distribution of lifetime discounted welfare of DI benefits.

is better, and this also generates more applications and hence a higher fraction of the population getting awards. Overall, this rank-rank correlation between the fraction of the population receiving an award and the welfare value is 76.3.

Figure 9 demonstrates this relationship by showing the average application and award rates at each level of welfare. Notice that for the bulk of the distribution, these are rising, so that the higher award rate locations also have higher value from the program conditional on receipt. The application rate follows a roughly similar pattern, which indicates the causal direction. Higher value from the program drives demand for it, meaning more applications and many of these generate awards. Towards the top of the welfare distribution, past about $320K, the relationship becomes weaker because higher value is driven by economic motives, and that indicates a low award rate per application. At the very top, in the very tail of the distribution, there are actually lower award and application rates, indicating that these regions have a very low probability of receipt. The model must then impute very high T_ℓ values to match this data. But because there are very few counties in this area, it is difficult to draw many conclusions.

Beyond the welfare value to an individual receiver, we can also compute the value *ex ante* within a county. There, we are drawing from the ergodic distribution of all of the endogenous

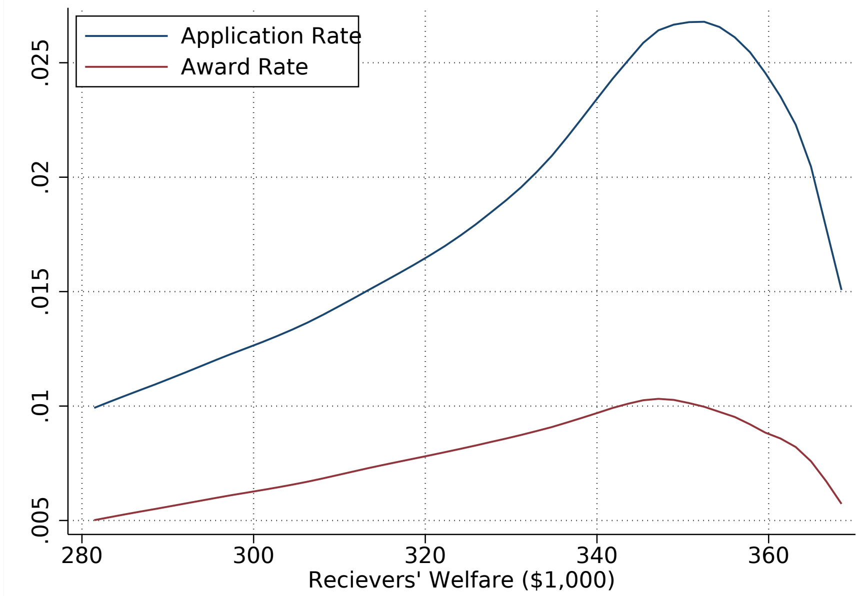

Figure 9: Value upon receipt increases, except at the top (selection)

states but excluding individuals actually receiving benefits. We then compute the lifetime present discounted value of the program for each individual. This is shown in Figure , in which the relative *ex ante* value of the program is very highly correlated with the value to a recipient, partly driven by the fact that in counties that place a higher value on the program will also have more applications and awards and then *ex ante* a higher value.

In figure 10, we show these *ex ante* welfare values against the application and award rates by county. For both of these measures, the welfare value of the program is highest among counties with the highest usage. Notably, the expected value of the program drives higher application rates: again with higher potential gain, agents choose to apply even when the probability of success is relatively low.

These counties with the highest welfare value are, of course, systematically different from those with lower gains. First and foremost, they have higher receipt rates, which is almost mechanical because being a receiver means getting more out of the program than not, so agents in counties with a higher chance of an award also get more value from it. This is potentially tempered by the implied value of the outside option, T_ℓ, of course. If this rises with receipt then that pushes against this trend since without the program they rely on this income.

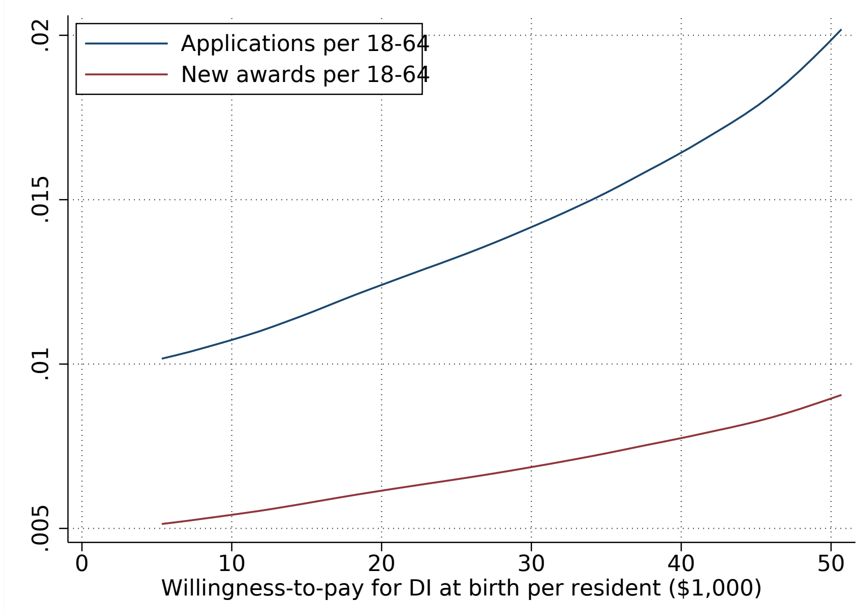

Figure 10: The most vulnerable counties gained the most

5.2 Allocative efficiency and place-based policy

With the estimated model of DI take up and its welfare implications, we can begin to consider how the program affects different places differently and how it might be perturbed. We first decompose welfare in the spirit of Benabou (2002) or, more recently Bandari et al (2024). It is easy to see that the average welfare effect $E[\Phi]$ in the economy can be decomposed into that which is associated with pure, within-individual insurance, within-county redistribution, cross-county redistribution and the overall change in efficiency. For this, we can first use the identity:

$$E[\Phi] = \int \phi_{\ell,i,t} d(\ell, i, t)$$

and

$$\phi_{\ell,i,t} = E[\phi_{\ell,i,t}] \times \frac{E_\ell[\phi_{\ell,i,t}]}{E[\phi_{\ell,i,t}]} \times \frac{E_{\ell,i}[\phi_{\ell,i,t}]}{E_\ell[\phi_{\ell,i,t}]} \times \frac{\phi_{\ell,i,t}}{E_{\ell,i}[\phi_{\ell,i,t}]}$$

Then the final term is the within-individual redistribution, i.e. insurance value, the next to last is the within-county redistribution, i.e. the risk sharing, the second term is the expected cross-county transfer and the first is the average gain. Starting with the cross-county transfer, we plot

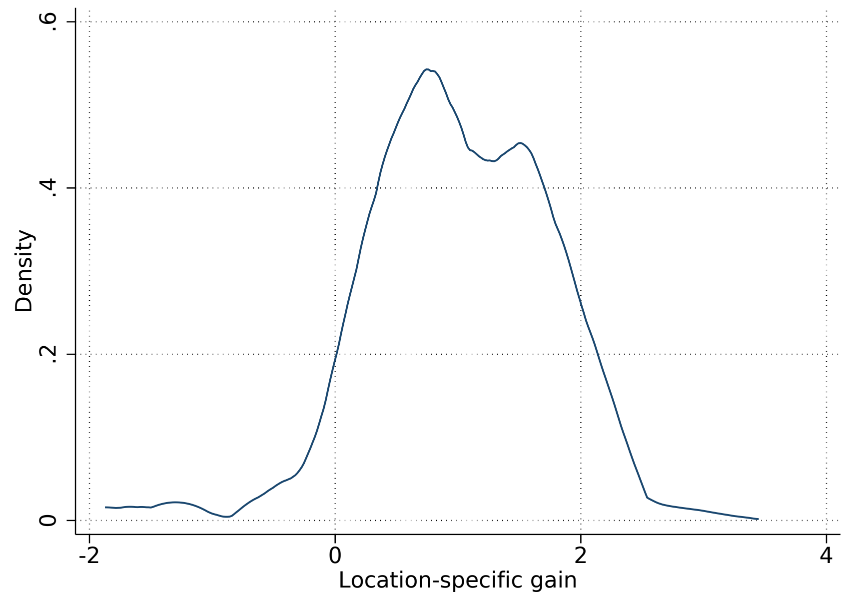

Figure 11: Location-specific gain, the cross-county transfer defined as $\frac{E_\ell[\phi_{\ell,i,t}]}{E[\phi_{\ell,i,t}]}$

this distribution in Figure 11. Notice that many of the counties have particularly large gains: about 10% of the population-weighted counties have more than twice the average welfare gain. On the other hand, about 6% get negative gains, on average.

To measure this transfer, we compute a uniform national SSDI transfer, that is to say, consider a consumption level $\hat{c}$ such that the real purchasing in every county is equalized while on DI, $\frac{\bar{p}_\ell}{p_\ell}\bar{D}I_\ell$ where $\bar{x}$ is the cross-county average. Of course, we still allow the rest of the county-specific processes to be the same and the state to be the same.

Figure 5.1 sorts counties by their welfare gain on the horizontal axis and plots the portion of this that is due to redistribution. This is to say, we compute willingness-to-pay with DI consumption equalized across counties in real terms, $\frac{\bar{p}_\ell}{p_\ell}\bar{D}I_\ell$ and compute its difference as a percent of our baseline willingness-to-pay. At the top of the distribution of willingness-to-pay, counties that gain the most from the program in general, almost half of this comes from gaining from redistribution. On the other hand, those who gain the least from the program also have the lowest real consumption stemming from it.

These very high receipt counties, those that benefit most from ex-ante transfers are not nec-

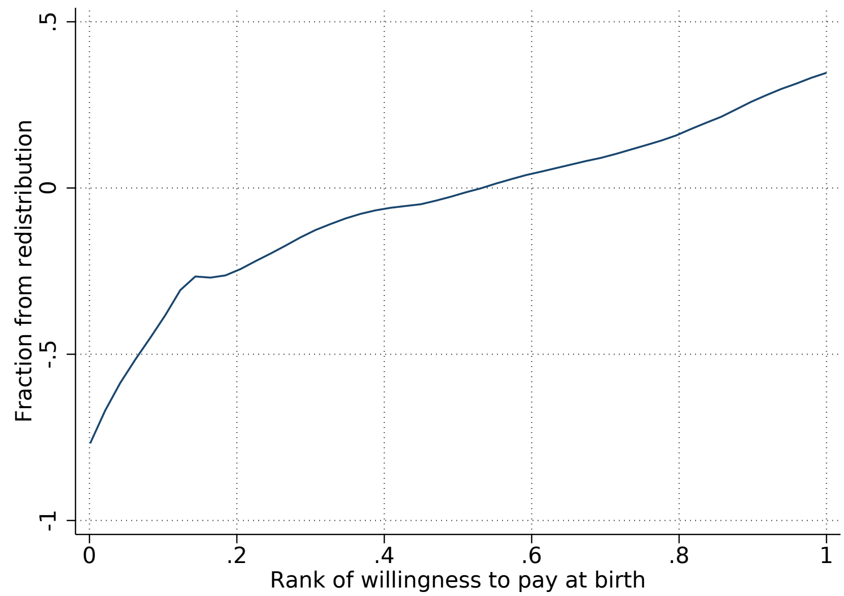

Figure 12: The redistributive effects of DI on the vertical against the total willingness-to-pay for the program

essarily getting the same benefit for the marginal dollar. Most specifically, in the highest receipt counties, the program is more generous and thus might be bringing applicants who are in somewhat less need than in the lower receipt, and lower generosity regions.

To address this, we compute an optimal distribution of DI payments holding fixed the total expenditure on the program so that it is a budget-neutral change. We are thus computing a location-based policy that can better incorporate the differences. By shifting a dollar from one county to another, we allow the marginal applicant to change, potentially bringing more people into the program in low-application counties and also reallocating payments to awardees in counties where the payment did not go far enough.

We compute these new pyaments in Figure 13. The policy tool we are using is to adjust the replacement rate, so that is plotted on the vertical axis along with the current regime labeled as "data." This replacement rate is really the number of dollars spent per recipient relative to the average wage in the county. Notice that the largest increases in payment are among the lowest receipt counties. There we are adjusting it upwards in and increasing both the number of recipients who will find it attractive to apply and the programs generosity.

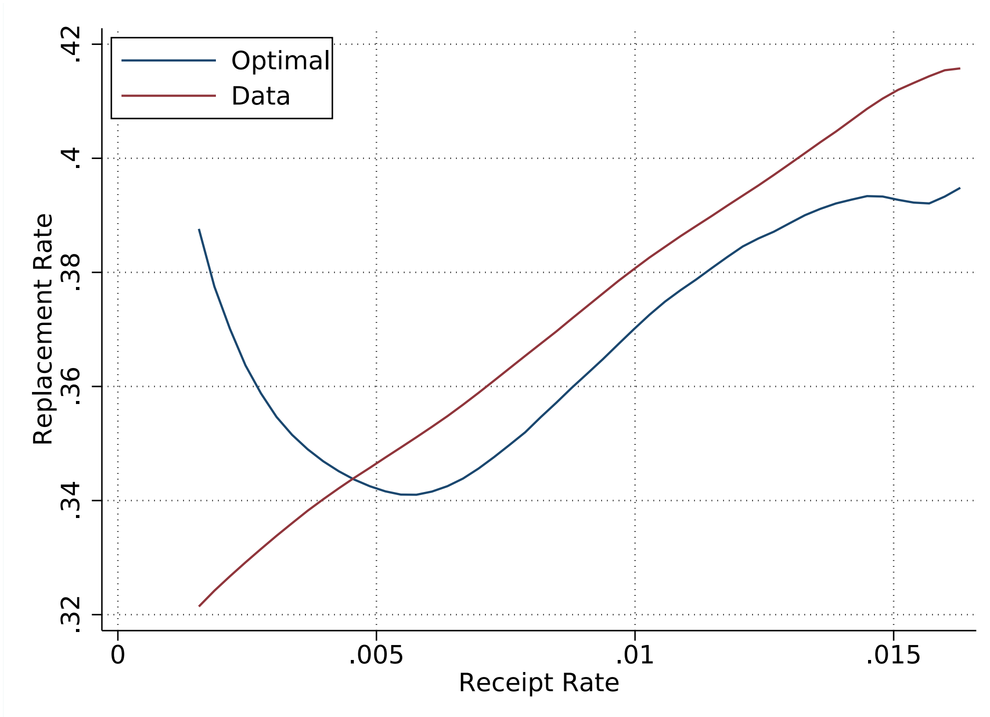

Figure 13: Optimal place-based replacement rate relative to data.

This optimal reallocation is partly a function of diminishing marginal gains from insurance. But, it also is because of the redistribution implied by the program rules. As described throughout, the real purchasing value of DI is higher in some counties than others. This redistribution also might be suboptimal, in addition to a potentially suboptimal provision of insurance. Indeed, counties receiving higher welfare from the program are also net beneficiaries of this redistribution.

5.3 Matching DI elasticities with other studies

In this section, we will compare our model's implied elasticities to those measure in other studies. When economic or policy conditions change, how do outcomes change in our model and how do they compare to other empirical studies?

This is a helpful set of exercises both for model validation but also it gives insight into empirical work: there is a tremendous amount of geographical heterogeneity in these estimated elasticities. Generally speaking, higher DI counties are also more elastic to changing conditions because places where many people are receiving DI are also places where many people are on the margin of wanting to apply.

An important set of studies look at how many DI awardees would have otherwise worked. Bound (1989) guides our analysis of this question, in which he asks whether workers rejected from the DI rolls would instead go to work. Another set of important papers then follows up with this, particularly using characteristics of the DI allowance process. For instance, Chen and van der Klaauw (2008) use the discrete change in allowance rates around rules regarding the age of the applicant. The effect that we would like to see is the labor force response of would-be applicants to a change in the allowance probability.

Within the model, we can adjust the probability of allowance by perturbing ξ_ℓ. This is full information, so agents then internalize that they are less likely to receive and adjust their behavior accordingly. There are then two interesting elasticities: on the one hand in the baseline word, we can look at those who would have applied and see whether they work. On the other hand, we can look at the overall effect. Of course, the ones who would have applied are also the sickest and least willing to work and therefore probably have the smallest response rate. We find that the elasticity of would-be applicants is -8.35% while the and the overall elasticity is -36.2%. This puts us in the middle of estimates from these studies that range from 20-30%: the lower number would be the effect of truly surprising an equivalent individual with rejection whereas the upper bound allows for all of the convoluting factors such as foresight about one's rejection and differences in who is granted an allowance.

In Figure 14 we include the distribution of elasticities among the would-be applicants. Notice that as we go out into the distribution of high recipiency regions the elasticity becomes much higher. Again, more marginal applicants here are also more marginal about employment. Hence being rejected has a large effect on their employment.

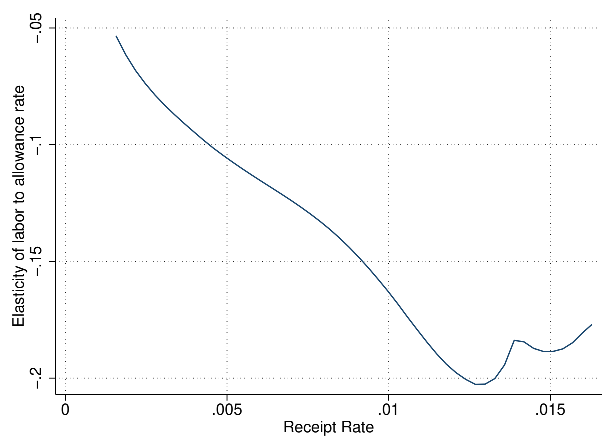

Figure 14: The higher DI counties actually have more workers that would work were DI not available.

Many studies look at employment shocks that suddenly reduce employment. The question is then how many will end up on the DI rolls. Several of the important papers in this literature are Black et al. (2002) or Charles et al. (2018). In both, an exogenous shock reduces employment in a resource-producing region. The lost employment then, to some extent is absorbed by the SSDI program and we can look at how much income and employment is replaced.

In our model, we create an analog of this experiment by increasing work disutility. To do this, we adjust λ_0 upwards, which uniformly affects ages and health status instead of moving λ_t or λ_d. It does, of course, affect the marginal workers more than the inframarginal. The key outcomes are that the change in DI receipt with respect to to employment loss is -27.6 % and the change in income is -9.4%.

Again, across space, these elasticities again vary considerably as shown in 15. This leads to an important adjustment to understand how our estimates relate to earlier studies. In many of these, the shock occurred in an already sensitive location, such as coal-producing counties of West Virginia or oil-producing ones which are again quite rural. Thus, for comparison, we look at the state of West Virginia in particular. Here, the effect on receipt is -67.4% and -24.2% on income. These compare quite well with the TSLS estimates in Charles et al. (2018) of -69.9% and -29.3% and very similar figures in Black et al. (2002).

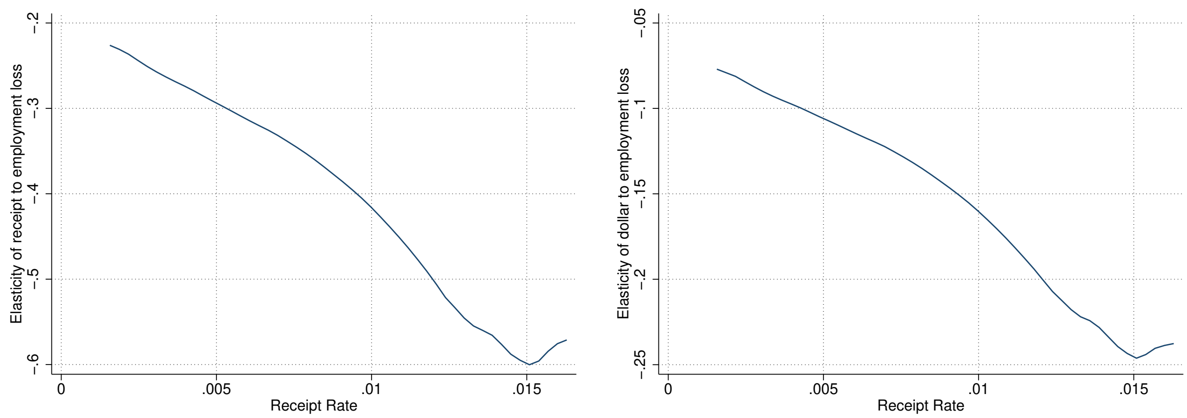

Figure 15: Cross-county elasticity of DI receipt (left) and income (right) to employment loss.

6 Conclusion

In this paper, we introduced a new set of facts about the geographic disparity in DI receipt. We showed that it is highly skewed towards high-receiver counties. These counties are generally in rural Appalachia, are poorer in less good health and have lower price levels.

We then introduced a model of DI claiming and receipt at the county level. We disciplined this model with county-level evidence on the economic environment, set of health risks and, crucially, the application and allowance rates by county. This established a national allowance function, and individual factors such as age and health were treated endogenously.

We used this model to decompose the driving factors behind DI applications and awards. The structural model is essential for this because otherwise the composition of applications cannot be observed and so correlations between e.g. health and DI receipt may not demonstrate the causal relationship.

Further, the model allowed us to study the welfare efficiency of the program. We looked at the ex ante spatial differences and found that high-recipient counties gain disproportionately. These counties also have lower marginal gains, suggesting that the program is more skewed than would be optimal.

References

Aizawa, N., C. Mommaerts, and S. Rennane (2022). Firm investment, labor supply, and the design of social insurance: Evidence from accommodations for workplace disability.". In *2022 APPAM Fall Research Conference*. APPAM.

Autor, D., A. Kostøl, M. Mogstad, and B. Setzler (2019). Disability benefits, consumption insurance, and household labor supply. *American Economic Review 109(7)*, 2613–2654.

Autor, D. H. and M. G. Duggan (2003, 02). The Rise in the Disability Rolls and the Decline in Unemployment*. *The Quarterly Journal of Economics 118*(1), 157–206.

Autor, D. H. and M. G. Duggan (2006). The growth in the social security disability rolls: a fiscal crisis unfolding. *Journal of Economic perspectives 20*(3), 71–96.

Ball, S. and H. Low (2014). Do self-insurance and disability insurance prevent consumption loss on disability? *Economica 81(323)*, 468–490.

Black, D., K. Daniel, and S. 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, J. (1989). The health and earnings of rejected disability insurance applicants. *American Economic Review 79(3)*, 482–503.

Bound, J., J. B. Cullen, A. Nichols, and L. Schmidt (2004). The welfare implications of increasing disability insurance benefit generosity. *Journal of Public Economics 88(12)*, 2487–2514.

Capatina, E., M. P. Keane, and S. Maruyama (2020). Health shocks and the evolution of earnings over the life-cycle. *ERN: Equity*.

Chandra, A. and A. Samwick (2005). Disability risk and the value of disability insurance. Technical report, NBER Working Paper No. 11605.

Charles, K. K., Y. Li, and M. Stephens Jr (2018). Disability benefit take-up and local labor market conditions. *Review of Economics and Statistics 100*(3), 416–423.

Chen, S. and W. van der Klaauw (2008). The work disincentive effects of the disability insurance program in the 1990s. *Journal of Econometrics 142*, 757–784.

Coe, N. B., K. Haverstick, A. H. Munnell, and A. Webb (2011, November). What Explains State Variation in SSDI Application Rates? Working Papers, Center for Retirement Research at Boston College wp2011-23, Center for Retirement Research.

Colas, M., R. McDonough, et al. (2021). Social transfers and spatial distortions. Technical report, Federal Reserve Bank of Minneapolis.

Costa, J. (2017). The decline in earning prior to application for disability insurance benefits. *Soc. Sec. Bull. 77*, 1.

Dal Bianco, C. (2023). Disability insurance and the effects of return-to-work policies. *Review of Economic Dynamics 49*, 351–373.

De Nardi, M., S. Pashchenko, and P. Porapakkarm (2024). The lifetime costs of bad health. *Review of Economic Studies*, rdae080.

Deshpande, M. and L. M. Lockwood (2022). Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica 90*(4), 1781–1810.

Foote, A., M. Grosz, and S. Rennane (2019). The effect of lower transaction costs on social security disability insurance application rates and participation. *Journal of Policy Analysis and Management 38*(1), 99–123.

Foote, A., M. Grosz, and A. Stevens (2019). Locate your nearest exit: Mass layoffs and local labor market response. *ILR Review 72*(1), 101–126.

French, E. and J. Song (2014). The effect of disability insurance receipt on labor supply. *American Economic Journal: Economic Policy 6(2)*, 291–337.

Gelber, A. M., T. J. Moore, Z. Pei, and A. Strand (2022). *Disability insurance income saves lives*. Stanford Institute for Economic Policy Research (SIEPR).

Gettens, J., P.-P. Lei, and A. D. Henry (2018). Accounting for geographic variation in social security disability program participation. *Soc. Sec. Bull. 78*, 29.

Gruber, J. (2000). Disability insurance benefits and labor supply. *Journal of Political Economy 108(6)*, 1162–1183.

Hosseini, R., K. A. Kopecky, and K. Zhao (2022). The evolution of health over the life cycle. *Review of Economic Dynamics 45*, 237–263.

Jia, N., R. Molloy, C. Smith, and A. Wozniak (2023). The economics of internal migration: Advances and policy questions. *Journal of economic literature 61*(1), 144–180.

Kennan, J. and J. R. Walker (2010). Wages, welfare benefits and migration. *Journal of Econometrics 156*(1), 229–238.

Kim, S. and S. Rhee (2022). Understanding the aggregate effects of disability insurance. *Review of Economic Dynamics 46*, 328–364.

Kitao, S. (2014). A life-cycle model of unemployment and disability insurance. *Journal of Monetary Economics 68*, 1–18.

Kreider, B. (1999). Social security disability insurance: Applications, awards, and lifetime income flows. *Journal of Labor Economics 17*(4), 784–827.

Lee, S. (2020). Spousal labor supply, caregiving, and the value of disability insurance. *International Economic Review*.

Liebman, J. B. (2015). Understanding the increase in disability insurance benefit receipt in the united states. *Journal of Economic Perspectives 29*(2), 123–150.

Low, H. and L. Pistaferri (2015a, October). Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review 105*(10), 2986–3029.

Low, H. and L. Pistaferri (2015b). Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review 105(10)*, 2986–3029.

Maestas, N., K. J. Mullen, and A. 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.

McCoy, J. L., M. Davis, and R. E. Hudson (1994). Geographic patterns of disability in the united staes. *Soc. Sec. Bull. 57*, 25.

McVicar, D. (2006). Why do disability benefit rolls vary between regions? a review of the evidence from the usa and the uk. *Regional studies 40*(5), 519–533.

Meyer, B. D. and W. K. Mok (2019). Disability, earnings, income and consumption. *Journal of Public Economics 171*, 51–69. Trans-Atlantic Public Economics Seminar 2016.

Michaud, A. and D. Wiczer (2018). The disability option: Labor market dynamics with macroeconomic and health risks. Technical report, CHCP Working Paper.

Millard, R. (2022). Early-onset disability, education investments, and social insurance.

Rupp, K. and D. Stapleton (1995). Determinants of the growth in the social security administration's disability programs-an overview. *Soc. Sec. Bull. 58*, 43.

Russo, N. (2022). Health-dependent preferences, consumption, and insurance. Technical report, Mimeo.

Schmulowitz, J. and H. D. Lynn (1966). *Insured and Disabled Workers Under the Social Security Disability Program: Characteristics and Benefit Payments, 1957-1963, by Jacob Schmulowitz and Henry D. Lynn*. Number 11. US Government Printing Office.

Schwartz, A. L. and B. D. Sommers (2014). Moving for medicaid? recent eligibility expansions did not induce migration from other states. *Health Affairs 33*(1), 88–94.

Seibold, A., S. Seitz, and S. Siegloch (2022). Privatizing disability insurance.

SSA (2021). Statistical supplement to the social security bulletin, 2020. Technical report, Social Security Administration.

SSA (2022). Oasdi beneficiaries by state and county, 2021. Technical report, Social Security Administration Office of Retirement and Disability Policy.

A Data

Population data. We use population and demographic data from the Census Bureau compiled by the Surveillance, Epidemiology, and End Results program of the National Cancer Institute. The data includes annual estimated population counts by sex and single years of age.

Mortality data. We use a compilation of mortality data from the Institute for Health Metrics and Evaluation. The mortality rates are created from deidentified death records from the National Center for Health Statistics, who compile data from death certificates lodged with state vital statistics bureaus. Census population data are used to create the rates. We use county-level rates by sex, and consider mortality rates for all ages, and by age ranges more focused on the working-age population (ages 25 to 64, and also split by ages 25 to 44 and 45 to 64).

Housing price index data. The Federal Housing Financing Agency constructs an index of housing prices that is available at the county level (Bogin, Doerner, and Larson, 2016). The Housing Price Index uses proprietary data held by the Agency on single family homes with roughly constant characteristics throughout the measurement period. It is constructed by regressing the change in log sale price of a home on period fixed effects and then taking the exponential of the fixed effects coefficients.

Poverty data. Poverty data come from the Small Area Income and Poverty Estimates program, which is a US Census Bureau project estimating median income and the fraction of households whose pre-tax earnings are below poverty thresholds defined by the Census Bureau. These thresholds vary by household composition and location. Thresholds are also adjusted annually by changes in the Consumer Price Index. The poverty estimates are developed using a forecasting model applying an empirical Bayesian framework to predict the aforementioned counts and American Community Survey county-level poverty counts estimates coupled with predictors coming from Census' data, including its administrative records.

Labor market data. Measures of the labor market and economic conditions come from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages (QCEW). The QCEW tabulates regional employment numbers and establishment counts among workplaces reporting to state unemployment insurance programs. An establishment is defined as a locale where goods and services are produced or provided; this means that a single business can have multiple establishments. The employment counts are the total numbers of paid jobs by the 12th of each month, irrespective of a job's characteristics. QCEW data includes roughly 97 percent of the US workforce each period, as it excludes self-employed workers as well as military personnel and a small contingency of diverse employment arrangements.

B Quantitative Model

This section presents additional figures showing the full distributions of estimated parameters across counties and model fit.

Health and economic differences across counties are captured by two gradients. Health deterioration is chosen to replicate county mortality rates of 40- and 70-year-olds given the county's age structure. Economic differences are captured by the cost of living (relative price level) and poverty rate. Figure 16 shows the exact model fit to these two features.

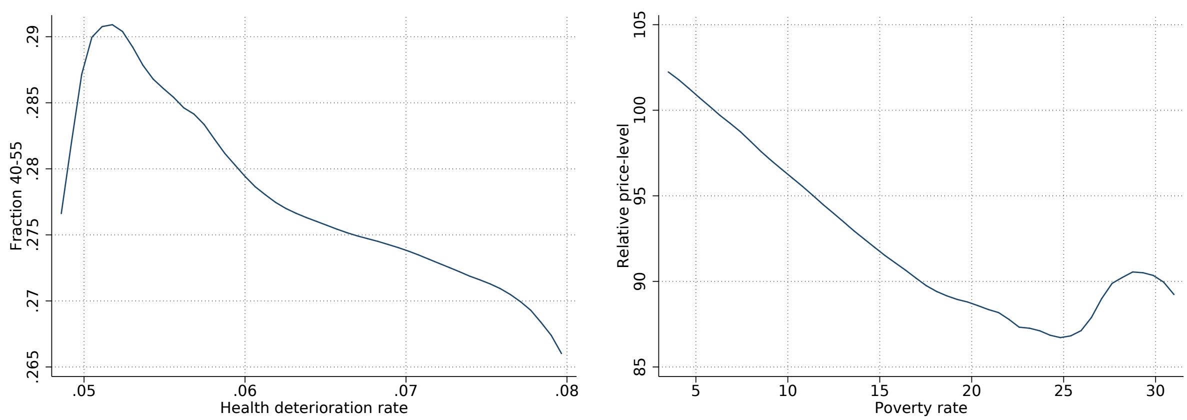

Figure 16: Exactly fitted model parameters to the health-loss rate and demographic structure (LEFT) and price level and cross-sectional poverty rate (RIGHT)

The model fit to DI applications and awards are shown in Figure 17.

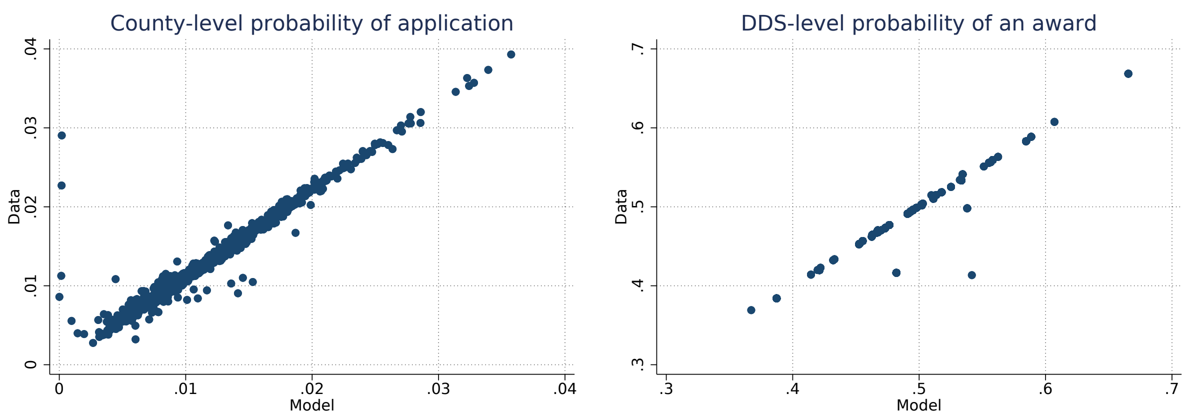

Figure 17: Application and allowance probabilities in the model and data