THE UNIVERSITY OF CHICAGO

FROM POLICY TO PAYCHECKS: THE ECONOMIC AND POLITICAL DYNAMICS OF UI GENEROSITY

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
MASTER OF ARTS

DEPARTMENT OF ECONOMICS

BY
SHIVANT KRISHNAN

CHICAGO, ILLINOIS
SPRING 2025

TABLE OF CONTENTS

ABSTRACT

This paper investigates the distinct roles played by unemployment insurance (UI) generosity and lockdown strictness in shaping state-level labor market outcomes in the United States, particularly during mid-2021. Exploiting the cross-state variation that arose during this time, we leverage two instrumental variables – historical UI generosity and political leaning – to disentangle the supply- and demand-side effects. States at the low end of UI generosity replaced about 45% of prior wages, while the most generous states replaced about 90%, and this gap is associated with a 1.3 pp higher unemployment rate, with uncertain effects on unemployment duration. Moving from minimal to near-total lockdowns adds roughly 6 weeks to average unemployment spells and raises unemployment by about 0.4 pp. The findings contribute critical insights into effective policy evaluation during systemic shocks to supply and demand. [1]

1. This project would not have been possible without the guidance of my advisor, Robert Shimer. I am grateful to the University of Chicago's Economics Faculty for encouraging me to pursue independent research and igniting my passion for economics, especially John List, Min Sok Lee, and Pablo Peña. I also thank Harry Moroz from the World Bank Group, and Aaron Terrazas from Glassdoor, Inc. for constructive comments. Lastly, I would like to thank my family and friends for their support throughout this process.

CHAPTER 1

INTRODUCTION

Unemployment insurance (UI) serves as a cornerstone of economic policy, designed to stabilize workers' incomes during periods of unemployment and facilitate their return to the labor market. While providing immediate financial support, UI programs also raise important economic questions surrounding their optimal generosity. A longstanding debate among economists revolves around whether increased generosity primarily reduces workers' incentives to search for employment, the "moral hazard" effect, or enhances their ability to seek better job matches due to eased liquidity constraints, known as the "liquidity" or "efficient matching" effect. Resolving this debate has become particularly pressing following the significant expansions of UI programs in response to the COVID-19 pandemic.

The unprecedented nature of the pandemic-driven economic disruption led to substantial federal interventions in UI programs. The March 2020 CARES Act initially increased state unemployment benefits with a supplemental payment of $600 per week. However, by early 2021, this federal supplement was reduced to $300, prompting states to make individual decisions about whether to continue or terminate these additional benefits. This policy environment generated notable state-level heterogeneity which was strongly influenced by political considerations. Republican-led states frequently terminated enhanced UI benefits early, citing potential labor shortages and disincentives to work. In contrast, Democratic-led states typically maintained higher UI generosity, emphasizing continued economic support amidst uneven recoveries.

We leverage this natural policy divergence to address two critical questions: first, how state-level political leanings shape decisions regarding UI generosity, and second, what causal effects these politically driven differences have on labor market outcomes, particularly unemployment rate and unemployment duration. Existing empirical studies have extensively documented the moral hazard versus liquidity trade-offs, notably Chetty (2008), who found that liquidity constraints explain much of the increased unemployment duration associated with UI benefits. Yet, the political determinants of UI policy—and their implications for labor market outcomes during periods of economic recovery rather than acute crisis—remain comparatively underexplored.

To address this gap, we exploit the significant variation in state UI generosity resulting from political divergence in mid-2021, utilizing state political leanings as an instrumental variable to isolate exogenous variation in benefit generosity. This approach allows for robust causal identification, distinguishing the impacts of UI generosity from confounding factors. Specifically, political leanings, proxied by gubernatorial party affiliation and legislative compositions, provide credible instruments due to their influence on policy decisions independent of direct labor market mechanisms.

The empirical analysis relies primarily on aggregated monthly data from the Current Population Survey (CPS), including its Annual Social and Economic Supplement (ASEC), chosen for its timeliness, state-level representativeness, and comprehensive coverage of labor market indicators. This choice is motivated by the ability of the CPS to capture broad labor market trends at the state level during the critical policy transition period in mid-2021. Moreover, CPS data allow for an examination of heterogeneity across demographic groups, contributing valuable insights into how political decisions affect different segments of the labor market.

Recognizing that political decisions regarding UI policies were also influenced by factors beyond partisan leanings, the study incorporates additional data sources to control for potential confounders. To systematically represent political leanings, a Principal Component Analysis (PCA) approach is adopted, synthesizing multiple indicators into a coherent measure of state political orientation.

While this methodology is robust, several limitations must be considered. Aggregation of individual-level CPS data to state-level averages inevitably reduces granularity, which potentially masking important variations in individual behaviors and responses. Additionally, the validity of political leanings as instruments relies on the assumption that they impact labor market outcomes solely through their influence on UI generosity. We explore this assumption in more detail by assessing the institutional environment, and conducting robustness checks (placebo tests) to examine labor market trends before the policy divergence.

Overall, this paper provides an important contribution by explicitly linking political factors to UI policy decisions and subsequently evaluating their causal effects on labor market outcomes. By examining this unique period of policy divergence, the research not only clarifies the underlying mechanisms through which UI generosity affects labor markets but also informs broader policy debates regarding optimal UI design and its role in addressing fluctuations in the labor market.

In what follows: Chapter 2 provides a comprehensive literature review situating this study within existing economic research. Chapter 3 describes the institutional context, detailing the federal and state-level UI policy responses during the COVID-19 pandemic. Chapter 4 presents an extensive description of the data sources used in the analysis, including detailed subsections outlining the construction of CPS Data, political indicators, the PCA methodology, and COVID-19 lockdown measures. Additionally, it describes how both cross-sectional and panel datasets are constructed. Chapter 5 outlines the empirical strategy, emphasizing the instrumental variables (IV) framework used for causal identification and estimation via 2-stage Least Squares. Chapter 6 presents the empirical results, and Chapter 7 discusses these findings, drawing on policy implications and potential system redesigns during crisis and recovery periods.

CHAPTER 2

LITERATURE REVIEW

A consistent finding within empirical research on unemployment insurance (UI) is the positive relationship between the generosity of UI benefits and the duration of unemployment spells. However, there remains considerable debate about the underlying reasons for this relationship, as well as about its implications for the quality of subsequent employment outcomes. While most studies agree that more generous UI leads to longer unemployment durations, the magnitude and interpretation of this effect vary significantly across different contexts and methodological frameworks (Card et al., 2007; Van Ours and Vodopivec, 2008; Farber and Valletta, 2015; Nekoei and Weber, 2017).

One prominent explanation for increased unemployment durations with greater UI generosity is the moral hazard effect. According to this perspective, UI reduces the incentives for active job searching because it raises workers' reservation wages or provides financial comfort that allows them to prolong unemployment. Early studies, such as those by Moffitt (1985), Katz and Meyer (1990), and Card and Levine (2000), offered substantial empirical support for moral hazard explanations. More recent research, including Rothstein (2011) and Farber and Valletta (2015), extends these findings to the context of UI expansions during the Great Recession. These studies typically find modest but statistically significant reductions in the probability of leaving unemployment, mainly driven by declines in labor force exit rather than decreased job-finding rates. However, Kroft and Notowidigdo (2016) caution that the moral hazard costs associated with UI are highly cyclical, becoming smaller during deep recessions when labor demand is weak and the potential returns to intensive job search are limited.

In contrast to the moral hazard view, other influential literature highlights liquidity constraints as the primary driver of increased unemployment durations when UI is expanded. Acemoglu and Shimer (1999) and Marimon and Zilibotti (1999) developed theoretical models demonstrating that UI generosity can improve job matches by allowing financially constrained workers additional time to find employment that better aligns with their productivity and skill profiles. Chetty (2008) provided a critical foundation for this view, empirically distinguishing between moral hazard and liquidity effects. Using variation in state-level UI policies across the United States, Chetty finds that around 60% of the increased unemployment duration resulting from more generous UI can be attributed to liquidity constraints rather than diminished incentives to seek employment. Similarly, Card, Chetty, and Weber (2007) reinforced this conclusion by showing that actual job-seeking behavior aligns more closely with liquidity constraints than the permanent income hypothesis.

Given the theoretical ambiguity surrounding the impact of UI on labor market efficiency, subsequent studies have examined how variations in UI generosity affect job quality outcomes, such as re-employment wages. This strand of research is particularly relevant to my study, given its focus on understanding not just unemployment duration but also the broader efficiency implications of UI policies. Studies conducted in various European contexts using regression discontinuity approaches yield mixed results. Lalive (2007), Van Ours and Vodopivec (2008), and Card et al. (2007) generally report negligible impacts of UI extensions on re-employment wages. Schmieder et al. (2016), conversely, find negative impacts on wages, while Nekoei and Weber (2017) stand out for identifying a positive, significant wage increase. Nekoei and Weber interpret their findings as evidence that the liquidity effect dominates the moral hazard effect for certain segments of the population, particularly those who are financially constrained.

Farooq, Kugler, and Muratori (2020) provide a critical extension of this literature within the context of recent U.S. recessions, making their findings especially relevant for my analysis of state-level UI generosity following the COVID-19 pandemic. Using detailed U.S. Census Bureau's Longitudinal Employer-Household Dynamics data, their research demonstrates that UI extensions improve re-employment outcomes primarily through higher-quality employer-employee matches, rather than simply increasing workers' bargaining power or facilitating shifts to higher-paying firms. Significantly, these effects are strongest among more financially vulnerable populations, including women, minorities, and less-educated workers. Their methodological approach—leveraging granular firm-worker matched data—is particularly instructive for my thesis, as it closely parallels the structure and granularity of the Current Population Survey (CPS) data that I employ. By directly measuring improvements in match quality rather than relying on proxy indicators such as wages alone, Farooq et al. (2020) underscore the potential efficiency-enhancing effects of UI generosity. This finding lends empirical support to liquidity-based explanations of UI impacts and suggests that moral hazard may be less detrimental to overall labor market outcomes than previously thought, especially during severe economic downturns or subsequent recoveries.

Despite this evidence supporting liquidity-driven improvements in match quality, Meyer (1995) provides a cautionary note about the complexities of UI policy design. Meyer's review of U.S. unemployment insurance experiments, which involved policy innovations such as cash bonuses for early re-employment, highlights potential unintended consequences of attempts to reduce unemployment durations via direct incentives. While such incentives effectively shorten UI durations without significantly harming workers' subsequent earnings, Meyer notes potential moral hazard effects, including strategic timing of UI claims. His insights contribute crucially to the policy analysis component of my study, emphasizing that carefully constructed incentives—rather than broad benefit reductions—may be necessary to maintain efficiency without exacerbating moral hazard.

Finally, Landais, Michaillat, and Saez (2018) offer a macroeconomic framework that complements the micro-level insights from Chetty and Farooq et al. They propose a dynamic model in which optimal UI generosity should be countercyclical—increasing during downturns when labor market tightness is low and decreasing in expansions to preserve job search incentives. This perspective informs the broader economic context of my thesis by emphasizing the importance of controlling for cyclical labor market conditions, thus ensuring that the identified effects of UI generosity accurately reflect policy variations rather than underlying economic trends.

Overall, my study contributes to this established literature by explicitly examining how state-level political dynamics shape UI generosity and by rigorously identifying the subsequent causal effects on unemployment duration, job search intensity, and re-employment quality during the recovery period following the COVID-19 pandemic. By combining theoretical insights from the liquidity-versus-moral-hazard debate with empirical strategies informed by Farooq et al. and methodological refinements from Mogstad et al., this thesis aims to provide comprehensive evidence on the efficiency and equity implications of politically-driven UI policies.

CHAPTER 3

INSTITUTIONAL BACKGROUND

3.1 Architecture of the U.S. Unemployment Insurance Program

The modern UI system dates back to Title III of the 1935 Social Security Act, which established joint federal–state administration of cash benefits to job-losers.[1] Financing for the program is coordinated through the Federal Unemployment Tax Act (FUTA). In this 1954 act, employers remitted a uniform six-percent payroll levy on the first $7,000 of each worker's annual earnings. States that meet federal performance criteria, namely prompt payment and no-discrimination, receive a credit of up to 5.4 percent, so the effective federal rate is only 0.6 percent. The invariance of this policy across jurisdictions means that cross-state differences in generosity arise almost fully from state choices rather than from differential federal policies.[2]

Benefit design is the prerogative of the states. An individual claimant's weekly benefit amount, denoted WBA_s, is set by

$$\mathrm{WBA}_s = \min\bigl\{r_s\,\mathrm{HQW}_s,\ \mathrm{WBA}_s^{\max}\bigr\},$$

where HQW_s is the worker's highest-quarter earnings[3], r_s the statutory replacement rate[4], and $\mathrm{WBA}_s^{\max}$ an absolute dollar ceiling. Median replacement rates cluster near 0.50, but the statutory range is wide, ranging from 0.46 in Arizona to 0.65 in Massachusetts (U.S. Department of Labor, *Comparison of State UI Laws*, 2019, Table 3.1). States also determine the maximum duration of "regular" benefits. In general, most states permit up to 26 payable weeks, though statutes adopted after the Great Recession reduced the cap to as few as 12

1. Social Security Act of 1935, https://ballotpedia.org/Social_Security_Act_of_1935.

2. Congressional Research Service, *Federal Unemployment Tax Act (FUTA): Overview and Legislative Developments*, R44527, 2024.

3. The total earnings received by a worker during their highest earning quarter.

4. The percent of high quarter wages that the state will replace via UI benefits.

weeks in Florida and North Carolina. Some states tie duration to the insured unemployment rate (IUR), defined as the ratio of current benefit recipients to workers insured for UI, while others link it to an earnings ratio in the base period. Additional heterogeneity arises from waiting-week rules, and documentation of work-search activity. These provisions change the present value of benefits well beyond the baseline WBAs and are therefore crucial for interpreting interstate differences.

Historic path dependence is also a key determinant of UI generosity. Employer payroll taxes are "experience rated," so cumulative benefit amounts feed back into future tax schedules and lock in past decisions on generosity (Anderson and Meyer, 1993). Moreover, the federal government sets solvency rules for state trust funds, which prevent extremely rapid benefit expansions following recessions. In fact, the two states that exhausted their funds after 2008 enacted permanent benefit cuts in 2011 (FL) and 2013 (NC). To capture this slow-moving institutional component, the empirical analysis proxies long-run generosity with the 2000–2019 average of WBA $\times$ duration, normalized to the unit interval for computational efficiency. Simultaneously, we employ a generosity measure for July 2021 which applies the same formula, but excludes the temporary pandemic supplements discussed next, in an attempt to isolate policy choices that remained under control of the states.

3.2 Pandemic-Era Federal Expansion

The Coronavirus Aid, Relief, and Economic Security (CARES) Act, enacted on March 27, 2020 (Public Law 116-136), overlaid three temporary federal programs onto existing state UI systems. Federal Pandemic Unemployment Compensation (FPUC) provided a \$600 weekly supplement between March 29th and July 25th, 2020; the Consolidated Appropriations Act reinstated the supplement at \$300 per week on December 27, 2020, and the American Rescue Plan Act (ARPA) carried that top-up through September 5th, 2021. Pandemic Unemployment Assistance (PUA) extended benefits to gig workers and other normally ineligible groups, and set the base WBA at 50% of the state-average WBA. PUA also eventually offered eligible workers up to 86 weeks of coverage. Finally, Pandemic Emergency Unemployment Compensation (PEUC) lengthened the maximum duration of regular UI, first by 13 weeks, then 24, and finally 53 weeks under ARPA.

All three programs were reimbursed in full by the federal government, (apart from minor administrative costs). CARES Section 4105 also relaxed the trigger for Federal Extended Benefits (EB) and financed EB entirely through March 2021 (making it 75% thereafter), but EB ceased to bind in most states once unemployment fell below the 6.5% threshold later that year.

Nonetheless, states were indeed able to opt out of the federal supplements. Between May and July 2021, twenty-six[5] states withdrew from one or more of FPUC, PUA, and PEUC, according to withdrawal notices compiled in the U.S. Department of Labor UIPL 14-21. By mid-July, statutory generosity therefore varied widely once again, justifying the choice of this broader time period as the cross-section used in our baseline estimates. Figure 3.1 below reiterates the timing and duration of the federal supplements.

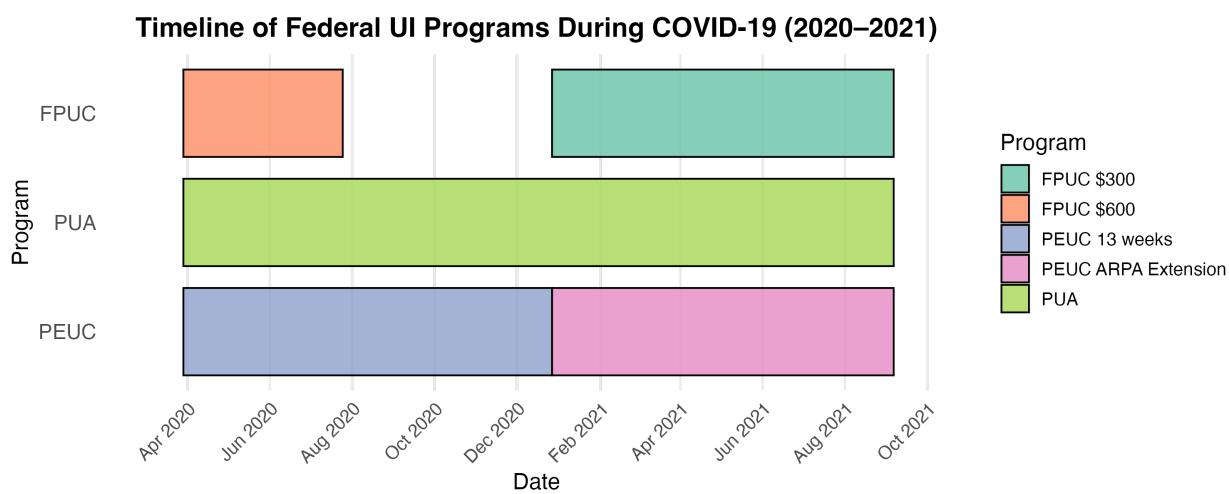

Figure 3.1: Sequencing of federal UI supplements, March 2020 – September 2021

Since the timing of federal programs was uniform nationwide and their early termination

5. Predominantly Republican-led

in some states was partisan-influenced, political orientation influences labor market outcomes in this period almost exclusively through its effect on UI policy choices. This reasoning provides logical policy evidence for the the exclusion restriction applied to the political leaning IV in what follows.

3.3 Variation in COVID-19 Policies

The widespread public health emergency left non-pharmaceutical interventions (NPIs) such as school closures and stay-at-home orders fully to the discretion of governors, which resulted in pronounced interstate heterogeneity (Adolph et al., 2022). To quantify lockdown stringency, this study aggregates eight indicators from the U.S. subset of the Oxford COVID-19 Government Response Tracker: school closures (`C1M`), workplace closures (`C2M`), public-event bans (`C3M`), gathering limits (`C4M`), transit closures (`C5M`), stay-at-home orders (`C6M`), internal mobility restrictions (`C7M`), and public-information campaigns (`H1`). Each component is rescaled to range between zero and one, summed with equal weights, and then normalized. Averaging monthly observations over 2021 yields a cross-section that produces stringency values ranging from 0.12 in South Dakota to 0.87 in California.

3.4 Political Institutions

State governors serve four-year terms and are frequently elected off the federal cycle, while legislative seats turn over every two years. In states where a single party controls both branches, both labor market and public health policy tends to align with national platforms. Thus, the political leaning index we construct below combines the governor's party affiliation, partisan seat shares in both chambers, and the 2020 presidential vote share. Evident by a manual check, no state changed party control during 2021–22, so the index is time-invariant over both the 2021 cross-section and the wider panel. Empirically, we posit that political orientation enters into the equation only through its effect on UI generosity or lockdown strictness; once those policies are controlled for, it should have no independent influence on labor outcomes.

3.5 Labor Market Context, 2020–2022

Sectoral composition also shaped both the depth of the COVID-19 shock and the pace of recovery. Tourism-dependent states like Nevada and Florida lost more than 20% of payroll employment in spring 2020, while Midwestern states experienced relatively mild declines. National unemployment jumped from 3.5% in February 2020 to 14.7% two months later, but state trajectories diverged sharply after this point.[6] Southern states that lifted restrictions sooner regained jobs faster. By contrast, prolonged school closures in parts of the Northeast delayed re-entry into the labor force, especially among women. The Delta variant which emerged in mid-2021 and the Omicron wave in late-2021 produced new disruptions, and even by December 2022, the national LFPR still trailed behind its pre-COVID peak. These persistent (yet mostly time-invariant) differences warrant the inclusion of state fixed effects in the panel specification estimated in Subsection 5.2.2.

6. Bureau of Labor Statistics, "The Employment Situation—April 2020," Table A1.

CHAPTER 4

DATA

This section describes in detail the data sources, variable construction, and data integration processes utilized in this study. The primary unified dataset is comprised of monthly state-level observations covering January to December 2021. For extensions, a supplementary panel dataset was created which covers the 50 states from January 2021 to December 2022. Both datasets integrate political indicators, COVID-19 lockdown measures, unemployment insurance claims data, and labor market outcomes derived from the Current Population Survey (CPS). Some variables, like political leaning, lockdown strictness, and UI generosity, are rescaled and normalized to [0,1] to ensure cleaner interpretations. We align the data at the state-month level to provide consistent temporal and geographical comparability, which allows for robust empirical analyses of the policy changes in question.

4.1 Current Population Survey (CPS)

The principal data source for labor market outcomes is the Current Population Survey (CPS), administered jointly by the U.S. Census Bureau and the Bureau of Labor Statistics. The CPS is a monthly, nationally representative household survey designed to measure detailed labor force characteristics, demographic information, and employment outcomes at the individual level. Specifically, the monthly Basic CPS microdata files from January through December 2021 were utilized.

Each month's CPS data were downloaded and processed individually. First, individuals' employment statuses were identified using the CPS variable EMPSTAT, which categorizes respondents as employed, unemployed, or out of the labor force. Duration of unemployment was measured via the DURUNEMP variable, capturing weeks unemployed; individuals coded as "Not in Universe" (value of 999) were recoded as missing observations. Additional demographic characteristics (such as age, sex, education, and race) were retained for subsequent analyses.

Each individual's CPS responses were mapped to states using the CPS-provided STATEFIP variable. To aggregate monthly averages, individual-level data were grouped by state and month, producing state-month-level measures of unemployment rates, average unemployment duration, and average weekly earnings. Unemployment rates were calculated as the proportion of respondents classified as unemployed relative to the total labor force (employed plus unemployed). The unemployment rates calculated through CPS were supplemented by the LAUS[1] estimates, and our results were invariant to the choice of unemployment measure. Average unemployment durations were computed conditional on respondents being unemployed. Weekly earnings measures were averaged across employed respondents, excluding those reporting zero or missing earnings.

These state-month aggregates, stored as the CPS-AGG dataset, were then merged with the other policy-oriented datasets, forming the basis for empirical analysis. The CPS aggregation provides crucial insights into monthly fluctuations of state labor market conditions, directly capturing the policies' impacts on key employment outcomes.

4.2 Construction of Labor Market Outcome Measures

Labor market outcome measures derived from CPS data were constructed in two complementary ways: as cross-sectional and panel datasets, each suited to distinct empirical purposes.

Cross-sectional measures represent state-level snapshots of labor market conditions, enabling clear cross-state comparisons at each monthly interval. Specifically, unemployment rates, average unemployment durations, and weekly earnings were calculated independently for each state-month combination, ensuring direct comparability across states. This structure provides insight into the immediate labor market impact of policy differences across

1. U.S. Bureau of Labor Statistics. Local Area Unemployment Statistics (LAUS), 2021 annual averages.

states at given points in time.

In addition, a panel structure was employed to capture longitudinal labor market dynamics within each state over the entire period from January 2021 to December 2022. This approach tracks states consistently over time, thereby allowing empirical identification strategies that leverage within-state temporal variation. Such longitudinal comparisons are particularly critical given the timing of major policy shifts, such as the federal UI supplement reductions in mid-2021. By incorporating a panel dimension, the analysis gains additional leverage in identifying causal policy impacts, since states serve as their own controls over time, mitigating concerns over time-invariant unobserved state characteristics.

4.3 Political Indicators

State-level political indicators were systematically collected from several authoritative sources, and integrated to produce a robust, comprehensive measure of political orientation. Three distinct dimensions were utilized to capture political heterogeneity across states: gubernatorial party affiliation, state legislative composition, and presidential vote shares.

Gubernatorial party affiliation data were retrieved from the National Governors Association roster, which identifies each governor's party (Democratic or Republican) throughout 2021. Legislative composition data, specifically the percentage of legislative seats held by Democrats, were sourced from the National Conference of State Legislatures (NCSL) as of February 2021. Given the typical biennial schedule for state legislative elections, this composition was assumed stable for the duration of 2021. Presidential vote shares for the most recent presidential election (2020) were obtained from a Harvard Dataverse repository, providing state-level electoral preferences for a recent national election cycle.

These three political dimensions – gubernatorial affiliation, legislative composition, and presidential vote share – were standardized to ensure comparability across scales. They were subsequently combined into a single, continuous political index. Specifically, an initial weighted average was constructed, with weights of 0.3 for gubernatorial affiliation, 0.4 for legislative composition, and 0.3 for presidential vote share. These weights were chosen to reflect a balanced representation of executive, legislative, and general voter political preferences. The resulting political index ranges from 0 (strongly Republican) to 1 (strongly Democratic), providing a nuanced scalar representation of political ideology across states.

To ensure robustness and mitigate concerns of arbitrary weighting, an alternative construction using Principal Component Analysis (PCA) was also implemented, detailed extensively below.

4.3.1 Principal Component Analysis

A central task in our analysis is to measure each state's political environment using a single scalar index. We focus on three key variables: (i) the governor's party affiliation, (ii) the legislative composition (e.g. the percentage of seats in the legislature held by Democrats), and (iii) the state's presidential vote share for the most recent election. In our baseline specifications, we aggregated these three indicators through a simple weighted average, assigning ad hoc weights to each dimension. While this approach is direct, it can overlook how strongly correlated these political variables may be. As an extension and robustness check, we now use principal component analysis (PCA) to create a more data-driven measure of political leaning.

Principal component analysis provides a natural way to extract the primary axis of variation from multiple correlated variables. Formally, let $\mathbf{A} \in \mathbb{R}^{n\times 3}$ be the data matrix for n states, where each column corresponds to one of the three political indicators and each row is a state. We center or standardize the columns of $\mathbf{A}$ to ensure that no single variable dominates the scale. The covariance matrix of these columns is then proportional to $\mathbf{A}^\top\mathbf{A}$. PCA solves an eigen-decomposition of $\mathbf{A}^\top\mathbf{A}$, yielding eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \ldots$ ordered by the magnitude of their corresponding eigenvalues. The first principal component (PC1) is associated with the largest eigenvalue and thus explains the greatest variance across these three indicators. For a given state i, we define the score on the first principal component as

$$\text{PC1}_i \;=\; \mathbf{A}_{i,\cdot}\,\mathbf{v}_1,$$

where $\mathbf{A}_{i,\cdot}$ is the row vector of the standardized political variables for state i and $\mathbf{v}_1$ is the first eigenvector. Typically, we check the sign of $\mathbf{v}_1$. In situations where a higher PC1 value represents “more conservative,” but we prefer “larger index" meaning more *liberal*, we multiply by -1. This ensures that a higher score does in fact correspond to a more liberal government, to stay consistent with our prior conventions.

From a theoretical perspective, each of the three political variables is thought to reflect a single underlying dimension of ideology. States with high presidential vote share for the Democratic candidate often have a Democratic governor and a legislature with a higher fraction of Democrats, but the exact correlation structure can be complex. PCA is attractive here because it consolidates these correlations into a single numeric index, PC1_i, that encapsulates the principal direction of variation. This empirical weighting scheme could be particularly relevant if, for example, legislative composition varies more across states than gubernatorial affiliation, or if the presidential vote share is distributed differently than we anticipated across states. By letting the data determine the relative weights on the three variables, we obtain a more empirically driven index, which we can use for secondary analyses.

It is important to verify that the loadings of each variable on the first principal component align with our conceptual idea of a left–right political dimension. In some cases, if one indicator is negatively correlated with another, PCA may assign opposite signs, potentially leading to less intuitive interpretations. Before finalizing the index, we examine the component loadings to confirm that they are consistent: for example, if a “Democratic governor” variable is positively weighted, we expect legislative “Democratic seats” or “Democratic vote share" also to load positively so that a state scoring high on all these components attains a higher PC1 score.

In practice, we treat the first principal component $PC1_i$ as our new political-leaning index and repeat it monthly for each state i whenever the underlying political variables are stable. We then substitute this PCA-based index in place of our original hand-weighted measure when we instrument for lockdown stringency in the 2SLS analysis. By comparing the resulting IV estimates (using the PCA-based index) to those obtained under the simpler weighted scheme, we can gauge how robust our findings are to alternative definitions of a state's "political environment." If the estimates are similar, it strengthens the credibility of political leaning as an instrument for lockdown strictness, irrespective of how we aggregate the underlying political variables. Conversely, if the results change dramatically, this would indicate either that the simpler weighting scheme was too crude, or that the first principal component isolates a different dimension of variation than our initial weighting had captured.

4.4 COVID-19 Lockdown Measures

Data on state-level COVID-19 lockdown measures were obtained from the Oxford COVID-19 Government Response Tracker, which compiles detailed records of governmental pandemic containment policies, including school closures, workplace restrictions, gathering limits, public information campaigns, testing initiatives, contact tracing strategies, and vaccination rollout policies. These variables are ordinal in nature, and typically take on values in $\{0,1,2\}$ or $\{0,1,2,3,4\}$.

Although these data were initially recorded at daily intervals, they were aggregated monthly to align temporally with the other state-month datasets. For each state-month, these individual lockdown indicators were first standardized, then consolidated into a single comprehensive Stringency Index. The aggregation strategy was designed to capture overall policy restrictiveness systematically, reflecting the collective intensity and scope of pandemicrelated measures implemented by each state government. The monthly aggregation allows our empirical method to consistently examine variations in policy stringency across states and over time, which is particularly important for identifying the labor market impacts associated with different policy approaches throughout the COVID-19 recovery period.

4.5 Descriptive Statistics

The final unified dataset, comprising political indicators, COVID-19 lockdown measures, UI claims data, and CPS-derived labor market outcomes, was carefully examined to produce descriptive statistics summarizing its structure and variation. These descriptive statistics offer crucial initial insights, building a foundation for the econometric modeling and interpretation that follows.

Key descriptive analyses included measures of central tendency and dispersion for unemployment rates, political index scores, lockdown stringency scores, and UI claims. In 2021, the average state unemployment rate was 2.0% (SD = 1.0 pp, range 1.0–4.0 pp), and average unemployment duration was about 25 weeks (SD = 4.9 weeks, range 15.3–35.1 weeks). Lockdown stringency scored 7.41 on our 0–12.48 scale (SD = 1.50, range 4.71–12.48), while normalized UI generosity averaged 0.10 (SD = 0.16) and historic generosity 0.13 (SD = 0.17). The political leaning index averaged 0.46 (SD = 0.23, range 0.12–0.87), with its wide range highlighting meaningful cross-state variation. Extending the lens beyond 2021 to the historical duration of UI programs, weekly initial claims averaged 6,883 (SD = 13,210), with a median of 3,428 and a maximum exceeding one million. This suggests the historical presence of not only typical amounts, but also extreme spikes – likely during downturns.

In addition, correlation matrices and heatmaps were computed to examine the interrelationships and geographic distributions of variables. Correlations between political leaning, lockdown stringency, and UI generosity provide preliminary evidence regarding the potential validity of proposed instrumental variables (Figure 4.1). Similarly, country-level heatmaps for UI generosity and political leaning (Appendix, Figures 8.3-8.4) allow us to see the geographic heterogeneity in these variables, which strongly motivates this analysis.

Beyond these larger patterns, several notable features emerge from the descriptive statistics. First, average state unemployment rates during 2021 hovered around 2.0 percent (SD = 1.0 pp) but ranged from 1.0 percent 4.0 percent, reflecting the uneven pace of reopening. Unemployment duration likewise exhibited substantial dispersion, with mean weeks unemployed spanning from about 15.3 weeks to over 35.1 weeks across states. This cross-state heterogeneity further sets up our goal of disentangling supply-side versus demand-side factors. Second, the political leaning index and the lockdown stringency score are positively correlated ($\approx$0.45), indicating that more Democratic-leaning states tended to adopt stricter NPIs. By contrast, historic UI generosity and the political index display only a modest association (0.25), suggesting that long-run institutional norms are only partially aligned with current partisan control. Finally, the cross-state maps reveal clear geographic clustering: the Northeast and West Coast combine high political-leaning scores and generous UI, while many Southern and Plains states feature both low stringency and more modest benefits.

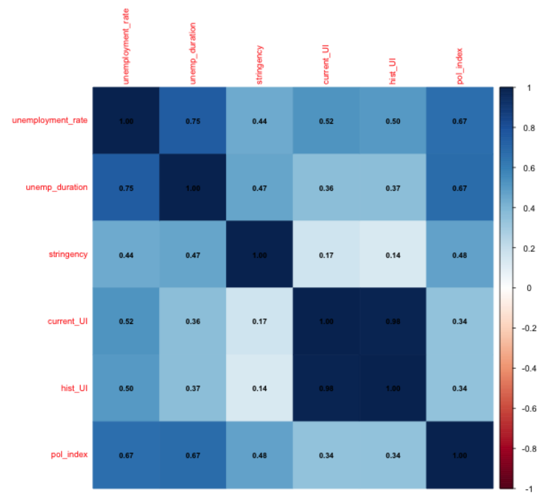

Figure 4.1: Correlation Heatmap for Key Variables

Table 4.1: Summary statistics – 2021 Cross Section

Statistic	N	Mean	St. Dev.	Min	Max
Unemployment Rate	50	0.02	0.01	0.01	0.04
Unemp. Duration (Weeks)	50	24.99	4.90	15.26	35.06
Lockdown Stringency	50	7.41	1.50	4.71	12.48
UI Generosity (Norm)	50	0.10	0.16	0.00	1.00
Historic Gen. (Norm)	50	0.13	0.17	0.00	1.00
Political Index	50	0.46	0.23	0.12	0.87

Table 4.2: Summary statistics – Weekly UI Benefits Data (1986–2025)

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Initial Claims	107,675	6,883.4	13,209.6	0	1,401	3,428	7,836	1,058,221
Continued Claims	107,675	54,529.3	96,641.3	0	11,225	26,445	59,212.5	4,805,047

CHAPTER 5

EMPIRICAL STRATEGY

This chapter formalizes the econometric framework that maps our conceptual discussions into estimable equations. In what follows, Section 5.1 establishes the conditions under which observed variations identify the causal effects of UI generosity and lockdown stringency on the labor–market outcomes in question. Section 5.2 explains the two–stage least–squares (2SLS) procedures employed in both the cross-section and panel specifications.

5.1 Identification

We specify a baseline causal model: Let Y_s denote a scalar labor-market outcome in state $s \in \{1, \ldots, S\}$ during 2021 (for example, the average unemployment rate or average unemployment duration in any month). We set up a potential outcomes framework:

$$Y_s = \beta_1 G_s + \beta_2 L_s + X_s' \gamma + \varepsilon_s, \tag{5.1}$$

where G_s is current UI generosity and L_s our lockdown stringency index, both defined in Chapter 4; X_s contains exogenous controls (demographics). The vector $\boldsymbol{\beta} = (\beta_1, \beta_2)'$ captures the supply-side and demand-side policy effects, respectively, whose signs and magnitudes constitute our first empirical question. The intuition behind these claims is that first, more generous unemployment insurance policies can reduce job search intensity or extend unemployment spells by increasing the opportunity cost of taking a job, i.e. a supply-side effect. Similarly, lockdown strictness is a "demand-side" shock because lockdowns restrict economic activity by directly reducing firms' labor demand.

However, because US states adjusted G_s and L_s in response to both political preferences and simultaneous economic conditions, OLS will generally yield biased estimates of $\boldsymbol{\beta}$. To overcome this endogeneity, we choose to exploit two sources of plausibly exogenous cross-state variation:

(i) Historical UI generosity. $Z_s^G \equiv \overline{G}_s^{2000\text{˘}2019}$, the pre-COVID mean of weekly UI claims, reflects long-run administrative norms determined decades before the pandemic. Conditional on current generosity in 2021, we argue that Z_s^G affects 2021 labor outcomes only through G_s. Chapters 2 and 3 documented that UI schedules exhibit exceptional path dependence, and thus we find it reasonable to assume that there is no causal link between a state's historical generosity and the current labor market trends once current benefits are held fixed.

(ii) Political-leaning index, denoted as Z_s^L. The baseline/naive measure averages three political indicators: gubernatorial party affiliation, partisan seat share in both legislative chambers, and the 2020 presidential vote share. As noted before, in the PCA case, we replace the equal-weighted average by the first principal component; results are invariant. In 2020 and 2021, public-health orders were administratively delegated to governors and state agencies, with more Democratic states systematically enforcing stricter non-pharmaceutical interventions (Adolph et al., 2022). We assume that – given a state's lockdown policy itself – the remaining direct influence of partisan control on labor outcomes is negligible. That is, states did not legislatively re-write unemployment law, tax codes, or subsidies during the study window.

We collect the instruments in the vector $\mathbf{Z}_s = (Z_s^G, Z_s^L)'$. Under the following conditions, $\boldsymbol{\beta}$ is identified:

1) **Relevance**

$$cov\big(G_s,\ Z_s^G\big) \ \neq\ 0 \quad \text{and} \quad cov\big(L_s,\ Z_s^L\big) \ \neq\ 0.$$

2) **Exclusion**

$$cov\big(\varepsilon_s,\ Z_s^G\big) \ =\ 0 \quad \text{and} \quad cov\big(\varepsilon_s,\ Z_s^L\big) \ =\ 0.$$

For relevance, we can observe the first-stage F–statistics in Table 6.1 and its counterparts in the Appendix. The values exceed 1,200 for the G equation and roughly 14 for the L-equation when the PCA instrument is used, which are far above conventional weak-instrument thresholds. However, the F-statistic on Z_S^L drops considerably when including state fiscal controls, so we must be cautious in interpreting those results.

In general, it is more difficult to defend the plausibility of the exclusion restriction within the context of study. We argue that for historic generosity, state benefit patterns were in place long before the pandemic. In particular, states tax employers based on how many of their former workers claimed UI benefits, so being generous today raises payroll taxes tomorrow. That creates a built-in disincentive to increase generosity quickly, so state policies tend to change slowly over time and are heavily shaped by past choices. Conditional on the 2021 benefits schedule G_s, any residual correlation between bygone statutory norms and 2021 labor outcomes would have to operate through other channels. This is further justified by a placebo regression of 2019 unemployment on historic generosity (Table 8.7, Column 2), in which we find no significance (p = 0.459). Thus, historical generosity appears to only affect labor outcomes after the 2021 cross-state changes in generosity through its impact on 2021 UI policy.

For the political leaning index, we first observe that during 2020–21 the U.S. federal government delegated virtually all lockdown/containment decisions to state executives. The index therefore shifts lockdown policy but—once L_s enters the structural equation, it has no direct channel to employment other than the minor fiscal programs already captured by G_s. To defend this empirically, we provide results for a placebo regression in Table 8.7, Column 1. We regress 2019 average unemployment on the PCA-derived political index, and find that there is no relationship between them (t = -0.175, p = 0.862). This lack of a "partisan effect" before 2021, prior to the introduction of NPIs, supports the exclusion argument: political index only matters for labor outcomes via its impact on lockdown stringency (and, by extension, UI generosity), not through any direct channel on unemployment itself.

Under assumptions 1-2, the two–stage least-squares estimator described in the following section consistently identifies β in both the cross-section and the panel extensions.

5.1.1 Caveats

While we regard Assumptions 1–2 as the most plausible empirical strategy available for this research, several concerns remain. First, the persistence of institutions may play a role. States with historically generous UI programs may also possess deeper labor market institutions, such as stronger unions, more licensing regulations for occupations, or fundamentally different industry mixes. These factors may have continued to shape employment in 2021 even after we condition on the current benefit amount. If these latent factors are insufficiently controlled for by X_s, the instrument Z_s^G could still correlate with ε_s. Additionally, the political-leaning index may influence outcomes through fiscal or public health channels that are only imperfectly captured by our variable L_s. Some examples of this: Staggered roll-out of vaccine mandates, the presence or absence of state-level stimulus checks, or political differences in the administration of financial relief. To the extent that these policies change firms' hiring decisions or job seekers' behaviors directly, the exclusion restriction for Z_s^L may be violated.

Moreover, measurement error is a non-trivial threat to this analysis. Our lockdown stringency score aggregates heterogeneous sub-indices and may misclassify short-lived executive orders, whereas UI generosity is proxied by claims data that embed administrative lags. Future study could tackle this by introducing additional lockdown measures, or perhaps assigning a method like PCA to identify the most relevant factors.

Finally, although the first-stage F–statistics are generally large, the instrument for lockdown stringency seems to be materially weaker than that for UI generosity. Weak instrument bias therefore remains a possibility, especially in the panel extension where within-state variation is much more limited. We address these issues empirically in robustness check 6.3, but we cannot rule them out entirely. The results that follow should then be interpreted as suggestive rather than definitive.

5.2 Estimation

5.2.1 Cross-section

Let $\mathbf{Y} = (Y_1, \ldots, Y_S)'$, $\mathbf{G}$ and $\mathbf{L}$ be the $S \times 1$ vectors of endogenous policies, $\mathbf{X}$ the $S \times q$ matrix of exogenous controls, and $\mathbf{Z}$ the $S \times 2$ matrix of instruments excluded from the structural model. In matrix notation, this reads $\mathbf{Y} = \mathbf{W}\theta + \varepsilon$, with $\mathbf{W} = [\,\mathbf{G} \;\; \mathbf{L} \;\; \mathbf{X}\,]$ and $\theta = (\beta_1, \beta_2, \gamma')'$. Define the full instrument matrix as $\mathbf{Q} = [\,\mathbf{Z} \;\; \mathbf{X}\,]$. Then we estimate

$$\hat{\theta}_{2\text{SLS}} = (\mathbf{W}'\mathbf{P_Q}\mathbf{W})^{-1}\mathbf{W}'\mathbf{P_Q}\mathbf{Y}, \quad \text{with } \mathbf{P_Q} = \mathbf{Q}(\mathbf{Q}'\mathbf{Q})^{-1}\mathbf{Q}'.$$

We report heteroskedasticity-consistent robust standard errors in the manner outlined by White (1980), and in the panel, we cluster at the state level.

5.2.2 Panel

For the balanced panel of states $s = 1, \ldots, S$ and months $t = 1, \ldots, T$ (January 2021–December 2022) our target equation is

$$Y_{st} = \beta_1 G_{st} + \beta_2 L_{st} + X_{st}'\gamma + \alpha_s + \lambda_t + \varepsilon_{st}, \tag{5.2}$$

where Y_{st} is either the unemployment rate or the average duration, G_{st} is the (log) current-generosity index, and L_{st} the lockdown-stringency score. The state fixed effect α_s captures all time-invariant state factors, while the month effect λ_t removes nation-wide COVID waves and other calendar shocks.

We proceed with time-invariant instruments, meaning that they are strictly cross-sectional:

- Z_s^G is historic generosity (2000–19 average).
- Z_s^L is the political-leaning index, fixed for every month of 2021–22.

Because each Z_s is constant within a state, it would be perfectly collinear with α_s once the data are de-meaned. To retain first-stage relevance we proceed by dropping the state dummies from both reduced-form regressions and keep only the month indicators φ_t. We write:

$$G_{st} = \pi_1 \, Z_s^G + \varphi_t + u_{st},$$

$$L_{st} = \pi_2 \, Z_s^L + \varphi_t + v_{st}.$$

In this situation, the month fixed effects partial out the effects of common shocks while allowing the instruments to act solely through cross-state variation.

In the second stage, we use the fitted values $\widehat{G}_{st}$ and $\widehat{L}_{st}$ from the first stage equations to estimate the baseline panel model, now with the full set of fixed effects (α_s, λ_t). Standard errors are clustered at the state level to accommodate arbitrary serial dependence within a state. This two-step arrangement avoids the mechanical collinearity that arises when instruments are time-invariant but the regressors and outcomes are time-variant. It also preserves our identifying assumption that historic generosity and political leaning of a state affect current labor outcomes only through G_{st} and L_{st}.

An important caveat is that in the panel model, identification now comes from deviations around a state's own mean. Thus, the effect we are estimating is necessarily weaker than in the cross-section. The panel estimates reported in Tables 8.5 and 8.6 can therefore be interpreted as a robustness check rather than our preferred causal result.

CHAPTER 6

RESULTS

6.1 Main IV Results

We estimate the causal effects of current UI generosity and lockdown stringency on state-level labor market outcomes in 2021. Historical UI generosity serves as an instrument for current generosity, while a naive political-leaning index instruments lockdown policy. Both instruments pass conventional relevance thresholds (first-stage F-statistics exceed 10; see Table 6.1).

Column (1) of Table 6.2 reports the 2SLS estimates for the unemployment rate, noting that the 49 observations are due to the state of Nebraska's missing weekly-claims series. Column (2) reports the corresponding estimates for average unemployment duration. We see that a full-range – least generous to most generous, since our regressor is a [0,1] index – increase in predicted UI generosity raises the unemployment rate by 1.1pp ($p<0.01$). Similarly, a full-range increase in predicted stringency increases the unemployment rate by 0.47pp ($p<0.001$), and also increases the duration of unemployment by 4.06 weeks ($p<0.001$). The coefficient plot in Figure 6.1 visualizes these effects and presents 95% confidence intervals.

The positive coefficient on generosity aligns with standard job-search theory (Katz & Meyer 1990), in which more generous benefits lengthen search and raise measured unemployment. The magnitude, $\approx$ 1pp for a full-range shock, is economically meaningful given an average state unemployment rate of 6%. By contrast, the demand-side lockdown stringency effect would seem to operate through firm-side channels, such as mandatory capacity limits and business closures which reduce vacancies, and thus increasing unemployment durations.

Table 6.1: First-Stage Regressions (Original Instrument)

	Dependent variable:	
	avg_norm_current_gen Current UI generosity	avg_stringency Lockdown stringency
	(1)	(2)
Historic generosity	0.960***	
	(0.028)	
Political-leaning index		3.042***
		(0.813)
Constant	−0.021***	6.047***
	(0.006)	(0.420)
Observations	50	49
R^2	0.962	0.230
Adjusted R^2	0.961	0.213
Residual Std. Error	0.032 (df = 48)	1.312 (df = 47)
F Statistic	1,209.171*** (df = 1; 48)	14.015*** (df = 1; 47)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 6.2: Second-Stage IV Estimates (Original Instrument)

	Dependent variable:	
	avg_unemployment_rate Unemployment rate	avg_unemp_duration Unemployment duration
	(1)	(2)
Current UI generosity	0.013**	6.217
	(0.006)	(5.517)
Lockdown stringency	0.004***	4.067***
	(0.001)	(1.242)
Constant	−0.012	−5.818
	(0.010)	(9.166)
Observations	49	49
R^2	−0.219	−0.447
Adjusted R^2	−0.272	−0.510
Residual Std. Error (df = 46)	0.007	5.915
Note:		*p<0.1; **p<0.05; ***p<0.01

6.2 Control Specification

To check the robustness of the baseline cross sectional estimates, I re-estimate the model after adding two fiscal covariates: total state spending per capita and total state tax collections per capita.[1] Since both of these covariates are defined *per-capita*, there is no need to explicitly include state population as a covariate. In other words, when we divide total spending and tax receipts by state population, any systematic scale effects that pure population size would capture are removed. What remains is a measure of how large the public sector is relative to the size of the state's economy, per-person. These controls are meant to capture differences in the sizes of state governments that might otherwise confound the relationship between labor market outcomes and the two policies of interest.

Table 8.1 reports the second-stage 2SLS estimates. Holding spending and taxes constant, pushing current-year UI generosity from its minimum to its maximum raises the unemployment rate by roughly 1.3 percentage points ($p < 0.10$). The point estimate for average unemployment duration is positive at 5.6 weeks, but the 95% confidence interval is wide enough to include zero. The fiscal controls themselves are small and statistically insignificant, and their inclusion leaves the core results unchanged. The negative R^2 seems concerning, but they simply reflect the well-known fact[2] that an IV regression can fit the in-sample variation less well than an inadequately specified OLS, so they do not undermine the causal interpretation.

The first stage diagnostics (Table 8.2) confirm that the two excluded instruments behave very differently once we include fiscal controls. Historical UI generosity is extraordinarily

1. The spending figures come from the *Fiscal Survey of States, Spring 2023* produced by the National Association of State Budget Officers (NASBO). Tax collection data are taken from the 2021 *State Government Tax Collections* release of the U.S. Census Bureau. Both series are divided by the 2021 population estimates from the Census Annual Population Survey, and expressed in thousand-dollar units.

2. One explanation of this is given in Wooldridge's *Econometric Analysis of Cross Section and Panel Data* (2010), in which he describes that in 2SLS estimation, R^2 does not have a clear interpretation; "poor in-sample fit does not call the consistency of the IV estimator into question."

powerful – the weak instrument F–statistic is 579, far above any Stock-Yogo threshold, so concerns about weak identification do not apply to the generosity equation. By contrast, the political PCA index yields an F of 3.19 with 2 and 44 degrees of freedom. This value falls well below the conventional threshold of 10 and does not exceed the Stock-Yogo 15 % critical value. This indicates that the stringency equation is weakly identified. As a result, the point estimate on stringency remains in the expected (positive) direction, but its magnitude should be viewed with skepticism.

The expanded specification adds per capita spending and tax controls, but the pattern we identified with our baseline model still persists. Specifically, higher lockdown strictness still coincides with higher unemployment rates and longer unemployment spells, while more generous UI policies still raise unemployment but leave spell length unaffected. Both fiscal controls are economically small and statistically insignificant, which implies that "state fiscal capacity" is not a confounding channel.

Finally, Wu-Hausman tests strongly reject the joint null that current UI generosity and lockdown stringency are exogenous ($p < 0.01$ in both outcome equations), so the use of instrumental variables is still warranted despite the weakness of the second instrument. Because each endogenous regressor is paired with a single excluded instrument, the model is exactly identified, and over-identification tests are not applicable.

6.3 Robustness: PCA-Based Political Instrument

The baseline instrument for lockdown stringency assigns arbitrary weights to the three political inputs. This ad-hoc weighting is ultimately arbitrary. To verify that our results are not an artefact of this choice, we have constructed a political instrument using principal-components analysis (PCA), and we now present the results of this exercise.

Looking at the first stage results, Column (2) of Table 8.3 confirms that the PCA-based index is a strong predictor of lockdown stringency. The coefficient is positive and highly significant, with an F-statistic of roughly 9.[3] Thus, replacing the equal-weighted index with PC_1 does not affect instrument relevance.

Table 8.4 replicates the IV estimates from Section 6.1 using the PCA instrument. Two patterns mainly stand out. First, the sign and magnitude of the lockdown stringency coefficient are virtually unchanged: a one-unit increase in predicted stringency raises the unemployment rate by about 0.6 percentage points and prolongs average unemployment spells by roughly 5.6 weeks—both significant at the 1% level. Second, the UI generosity coefficient remains positive but continues to be imprecisely estimated in the duration equation and marginally significant in the rate equation. Thus, the substantive conclusions of Section 6.1 survive a switch to PCA weighting.

6.4 Robustness: 2021-2022 Panel

The first stage estimates reported in Table 8.5 confirm the relevance of both instruments even after imposing a longitudinal structure (including data from 2022). Historic generosity remains an extraordinarily strong predictor of the contemporaneous UI–generosity index ($F \approx 216$), while the political-leaning index continues to explain a significant share of month-to-month variation in lockdown stringency. These estimates seem relatively large because the variables are not scaled as they were in the cross-sectional specification, but the core exclusion strategy survives when we include month fixed effects. Since our instruments are time-invariant and our regressors are time-variant, we include time fixed effects and omit state fixed effects in the first stage.

Turning to the second stage (Table 8.6), we find that, conditional on state dummies, an increase in instrument-predicted UI generosity is associated with a statistically significant, yet economically negligible, rise in the unemployment rate (a point estimate on the order

3. Historic generosity continues to explain nearly all variation in current UI generosity (column 1), in line with Table 6.1.

of 10^{-6}). Conversely, stricter lockdown policies are estimated to lower the unemployment rate and shorten average unemployment spells by roughly half a week. Both coefficients are precisely estimated, but their signs contrast with the cross-sectional results in Sections 6.1-6.2 and the PCA robustness check in Section 6.3. However, due to the reduced weakness of the lockdown/demand-side instrument, we should be careful when interpreting these coefficients as economically meaningful.

Two design features help explain this reversal. First, there is the concept of exploiting the within-state variation. Because historic generosity is time-invariant, it varies exclusively through first-stage fitted values, leaving no movement over time for a given state. Even modest measurement noise in monthly UI claims (and the fact that we only have 2 observations per state) can therefore dominate the signal and produce mechanically tiny, yet statistically significant, coefficients. Nonetheless, there is some evidence of policy co-movement taking place from the panel results. Throughout the end of 2021 and into 2022, most states simultaneously reduced both UI generosity and pandemic restrictions as national case counts declined. This policy easing could be distant from the 2021 cross-state variation which we implement in the present study.

Taken together, these estimates should be viewed purely as a sensitivity test. Their small economic magnitudes and sign instability caution against treating the within-state coefficients as precise causal effects. Rather, the exercise reveals the difficulty of leveraging monthly variation once nationwide timing and state dummies absorb most identifying power. Future work with richer, high-frequency administrative data, or truly state-specific shocks, could enhance dynamic inference. Importantly, nothing in Tables 8.5-8.6 undermines our qualitative conclusion that UI generosity and lockdown policy are empirically intertwined with labor-market outcomes. Instead, the table highlights that our primary identification stems from cross-state contrasts, not within-state month-to-month fluctuations.

CHAPTER 7

DISCUSSION AND POLICY IMPLICATIONS

The empirical results above give way to a coherent narrative regarding the distinct channels through which pandemic policies affected state labor markets. Our cross-sectional 2SLS estimates show that, conditional on the validity of our instruments, more generous UI benefits causally raise the state unemployment rate (by roughly 1.3 pp for a full-range generosity increase), while stricter lockdowns lengthen unemployment spells (by about 6 weeks) and also raise unemployment incidence ($\approx$0.4 pp). These effects are partially robust to per-capita spending and tax controls. Moreover, they are consistent with canonical search theory predictions that UI raises reservation wages and reduces search effort, whereas NPIs contract labor demand and result in slower matching.

Viewed through the lens of incomplete markets, where private credit and insurance cannot fully smooth idiosyncratic income risks, UI acts as essential public insurance. Our results illustrate the classic tradeoff in such settings: Higher UI generosity alleviates liquidity constraints and can improve match quality, but it also lengthens spells by dampening search incentives. Although we do not derive a welfare-maximizing UI model, our evidence suggests policymakers could potentially link benefit parameters to real-time labor indicators (e.g. vacancy-to-unemployment ratios or job-finding rates). Such rules would automatically raise benefits when private risk-sharing fails and phase them out as hiring recovers, therefore avoiding simultaneous supply- and demand-side disincentives.

Additionally, the use of a 24-month panel offers a valuable stress test, even if it does not sharpen our inference. Once state and time dummies absorb permanent heterogeneity, the historic generosity instrument loses within-state variation, and the remaining movement in monthly UI claims is too noisy to deliver meaningful coefficients. Moreover, most states relaxed UI supplements and containment measures simultaneously, leaving little idiosyncratic variation in stringency after month fixed effects are removed. The resulting second-stage estimates are indeed precise but mechanically small, and their signs differ from the cross-sectional benchmarks, so we are skeptical about some of the results derived here.

To refine these quantitative tradeoffs, future research could exploit matched employer–employee administrative data to measure match quality directly, and develop alternative NPI instruments (such as court-order timelines or county-level mask-mandate rollbacks). Research could also isolate the components we gathered into the lockdown stringency index as separate indicators and measure which precise policies (e.g. workplace closures) capture the primary variations in labor outcomes. Together with a direct comparison of UI policy at different points in history, these extensions would help translate our reduced-form insights into more precise guidelines for UI and public health coordination.

REFERENCES

Daron Acemoglu and Robert Shimer. Efficient unemployment insurance. *Journal of Political Economy*, 107(5):893–928, 1999. doi:10.1086/250084.

Christopher Adolph, Kenya Amano, Bryn Bang-Jensen, Nancy Fullman, and John Wilkerson. The pandemic policy u-turn: Partisanship, public health, and race in decisions to ease covid-19 social distancing policies in the united states. *Perspectives on Politics*, 20 (2):595–617, 2022. doi:10.1017/S1537592721002036.

Patricia M. Anderson and Bruce D. Meyer. Unemployment insurance in the united states: Layoff incentives and cross subsidies. *Journal of Labor Economics*, 11(1):S70–S95, 1993. URL `http://www.jstor.org/stable/2535168`.

Ballotpedia. Social security act of 1935, 1935. `https://ballotpedia.org/Social_Security_Act_of_1935`.

Timothy Besley and Anne Case. Political institutions and policy choices: Evidence from the united states. *Journal of Economic Literature*, 41(1):7–73, 2003.

David Card and Phillip B. Levine. Extended benefits and the duration of ui spells: Evidence from the new jersey extended benefit program. *Journal of Public Economics*, 78(1-2): 107–138, 2000.

David Card, Raj Chetty, and Andrea Weber. The spike at benefit exhaustion: Leaving the unemployment system or starting a new job? *American Economic Review*, 97(2):113–118, 2007. doi:10.1257/aer.97.2.113.

Raj Chetty. Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, 116(2):173–234, 2008. doi:10.1086/588585.

Congressional Research Service. Federal unemployment tax act (futa): Overview and legislative developments, 2024. `https://www.congress.gov/crs-product/R44527`.

Henry S. Farber and Robert G. Valletta. Do extended unemployment benefits lengthen unemployment spells? evidence from recent cycles in the u.s. labor market. *Journal of Human Resources*, 50(4):873–909, 2015.

Henry S. Farber, Jesse Rothstein, and Robert G. Valletta. The effect of extended unemployment insurance benefits: Evidence from the 2012-2013 phase-out. *American Economic Review*, 105(5):171–176, 2015. doi:10.1257/aer.p20151088.

Ammar Farooq, Adriana D. Kugler, and Umberto Muratori. Do unemployment insurance benefits improve match and employer quality? evidence from recent u.s. recessions. Working Paper Series. National Bureau of Economic Research, 2020.

Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry. Ipums cps: Version 12.0 [dataset], 2024. URL https://doi.org/10.18128/D030.V12.0.

Lawrence F. Katz and Bruce D. Meyer. The impact of the potential duration of unemployment benefits on the duration of unemployment. NBER Working Paper No. 2741, 1988.

Kory Kroft and Matthew J. Notowidigdo. Should unemployment insurance vary with the unemployment rate? theory and evidence. *Review of Economic Studies*, 83(3):1092–1124, 2016. doi:10.1093/restud/rdw009.

Camille Landais, Pascal Michaillat, and Emmanuel Saez. A macroeconomic approach to optimal unemployment insurance: Applications. *American Economic Journal: Economic Policy*, 10(2):182–216, 2018. doi:10.1257/pol.20160462.

Ramon Marimon and Fabrizio Zilibotti. Unemployment vs. mismatch of talents: Reconsidering unemployment benefits. *The Economic Journal*, 109(454):266–291, 1999. doi:10.1111/1468-0297.00432.

Bruce D. Meyer. Lessons from the u.s. unemployment insurance experiments. *Journal of Economic Literature*, 33(1):91–131, 1995. URL https://www.jstor.org/stable/2728911.

Robert Moffitt. Unemployment insurance and the distribution of unemployment spells. *Journal of Econometrics*, 28(1):85–101, 1985. doi:10.1016/0304-4076(85)90068-5.

Arash Nekoei and Andrea Weber. Does extending unemployment benefits improve job quality? *American Economic Review*, 107(2):527–561, 2017. doi:10.1257/aer.20150528.

Richard Rogerson, Robert Shimer, and Randall Wright. Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988, 2005. doi:10.1257/002205105775362014.

Jesse Rothstein. Unemployment insurance and job search in the great recession. *Brookings Papers on Economic Activity*, 43(2):143–213, 2011.

Robert Shimer and Iván Werning. Optimal unemployment insurance with sequential search. Manuscript, University of Chicago and MIT, 2003.

Jan C. van Ours and Milan Vodopivec. How shortening the potential duration of unemployment benefits affects the duration of unemployment: Evidence from a natural experiment. *Journal of Labor Economics*, 24(2):351–378, 2006.

Halbert White. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48(4):817–838, 1980. doi:10.2307/1912934.

CHAPTER 8

APPENDIX

Table 8.1: IV Estimates with State Controls

	Dependent variable:	
		avg_unemp_duration
	Unemp. Rate	Unemp. Duration
	(1)	(2)
Current UI generosity (norm)	0.013*	5.605
	(0.007)	(6.624)
Lockdown stringency (norm)	0.006**	6.080**
	(0.003)	(2.735)
Spending per capita	−0.00000	−0.001
	(0.00000)	(0.001)
Taxes per capita	0.00000	0.0002
	(0.00000)	(0.001)
Constant	−0.021	−16.933
	(0.018)	(18.426)
Observations	49	49
R^2	−0.957	−1.760
Note:		*p<0.1; **p<0.05; ***p<0.01
		Robust standard errors in parentheses.

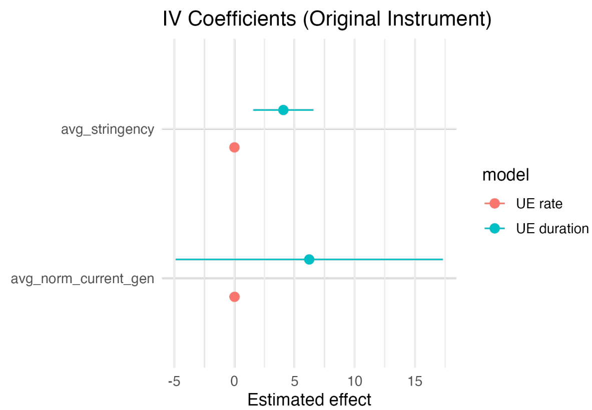

Figure 8.1: Coefficient Plot for Baseline IV Specification

Table 8.2: First-Stage Diagnostics for Cross-Sectional Specification w/Controls

Test	df1	df2	statistic	p-value	Outcome
Weak instruments (avg_norm_current_gen)	2	44	578.788	0	Unemp. Rate
Weak instruments (avg_stringency)	2	44	3.186	0.051	Unemp. Rate
Wu-Hausman	2	42	9.573	0.0004	Unemp. Rate
Sargan	0	*NA*	*NA*	*NA*	Unemp. Rate
Weak instruments (avg_norm_current_gen)	2	44	578.788	0	Unemp. Duration
Weak instruments (avg_stringency)	2	44	3.186	0.051	Unemp. Duration
Wu-Hausman	2	42	15.785	0.00001	Unemp. Duration
Sargan	0	*NA*	*NA*	*NA*	Unemp. Duration

Table 8.3: First-Stage Estimates using PCA for Political Index

	Dependent variable:	
	avg_norm_current_gen Current UI generosity	avg_stringency Lockdown stringency
	(1)	(2)
Historic generosity	0.960***	
	(0.028)	
PCA political index		2.237***
		(0.744)
Constant	−0.021***	6.370***
	(0.006)	(0.411)
Observations	50	49
R^2	0.962	0.161
Adjusted R^2	0.961	0.143
Residual Std. Error	0.032 (df = 48)	1.369 (df = 47)
F Statistic	1,209.171*** (df = 1; 48)	9.036*** (df = 1; 47)
Note:		*p<0.1; **p<0.05; ***p<0.01 Robust standard errors in parentheses.

Table 8.4: Second-Stage Estimates using PCA for Political Index

	Dependent variable:	
	Unemployment rate	Unemployment duration
	(1)	(2)
Current UI generosity	0.012	4.540
	(0.008)	(7.309)
Lockdown stringency	0.006^{***}	5.579^{***}
	(0.002)	(1.971)
Constant	-0.021	-16.911
	(0.015)	(14.507)
Observations	49	49
R^2	-0.806	-1.439
Adjusted R^2	-0.884	-1.545
Residual Std. Error (df = 46)	0.008	7.678

Note: *p<0.1; **p<0.05; ***p<0.01
Robust standard errors in parentheses.

Table 8.5: First-Stage Estimates – Panel 2021–22

	Dependent variable:	
	Current UI generosity	Lockdown stringency
	(1)	(2)
Historic generosity	$1{,}096{,}379^{***}$	
	(16,167.050)	
PCA political index		-42.989^{***}
		(4.682)
Month fixed effects	Yes	Yes
State fixed effects	No	No
Controls	Yes	Yes
Observations	1,176	1,176
Number of states	49	49
R^2	0.818	0.897

State-clustered standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table 8.6: Second-Stage Estimates – Panel 2021–22

	Dependent variable:	
	Unemployment rate	Unemployment duration
	(1)	(2)
Current UI generosity (IV)	0.00000***	0.00002***
	(0.000)	(0.00000)
Lockdown stringency (IV)	–0.0002***	–0.506***
	(0.0001)	(0.152)
State fixed effects	Yes	Yes
Month fixed effects	No	No
Controls	Yes	Yes
Observations	1,176	1,176
Number of states	49	49
R^2	0.733	0.523

State-clustered standard errors in parentheses. ***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

Table 8.7: Placebo Regressions using 2019 Unemployment Rate

	Dependent variable:	
	Unemp. (2019)	Unemp. (2019)
	(1)	(2)
Constant	3.6127***	3.5020***
	(0.2553)	(0.1542)
Political index (PC1)	−0.0809	
	(0.4628)	
Average historic UI generosity		0.5471
		(0.7319)
Observations	49	49
R^2	0.0006	0.0118
Adjusted R^2	–0.0206	–0.0093
F Statistic	0.0305	0.5586
	(df = 1; 47)	(df = 1; 47)

Robust standard errors in parentheses. ***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

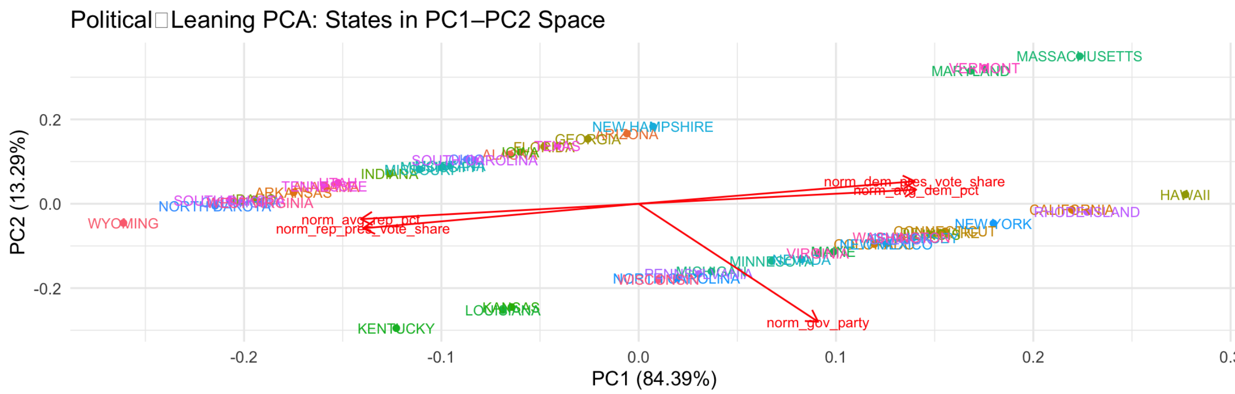

Figure 8.2: Visualization of PCA with State Labels

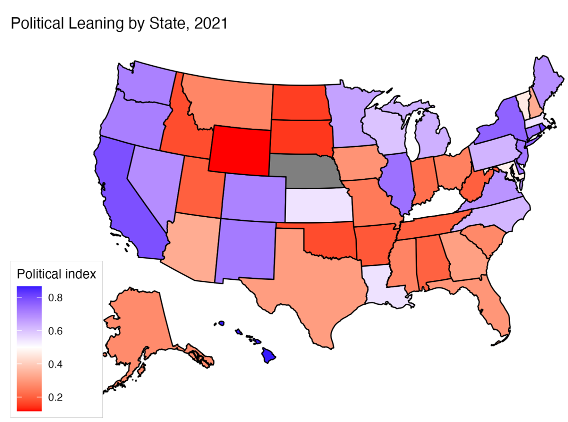

Figure 8.3: US Heatmap: Poiltical Leaning

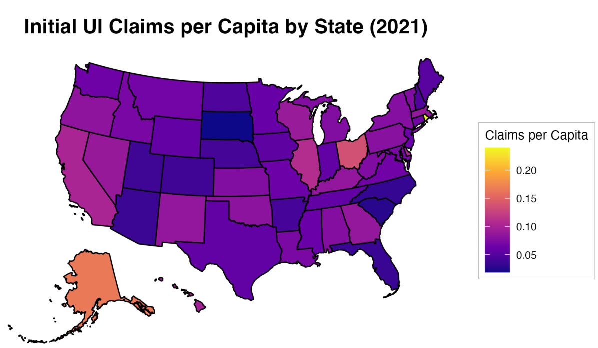

Figure 8.4: US Heatmap: 2021 UI Initial + Continued Claims

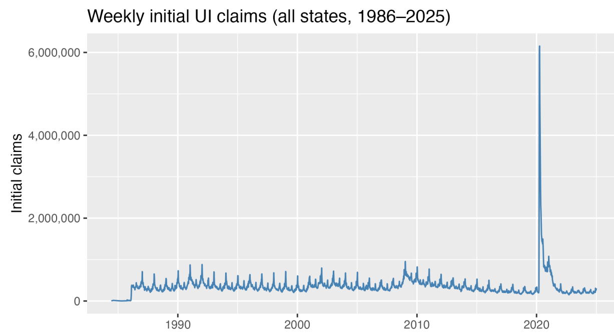

Figure 8.5: Time Series of UI Initial Claims

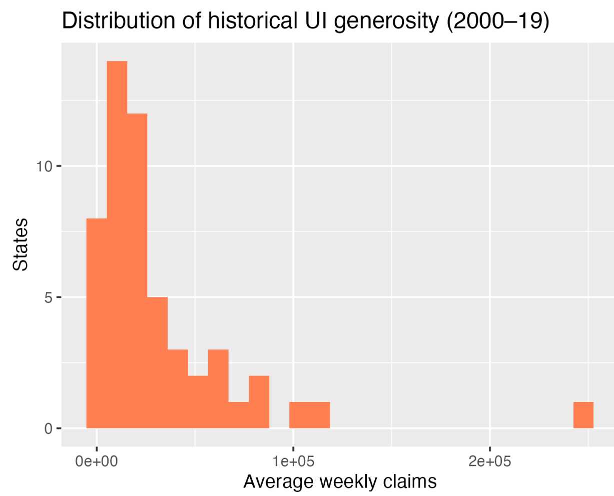

Figure 8.6: Distribution of Historical UI Generosity

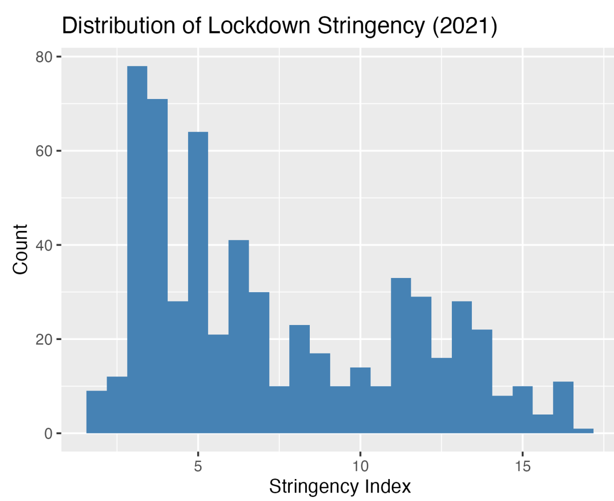

Figure 8.7: Distribution of Lockdown Stringency (2021)

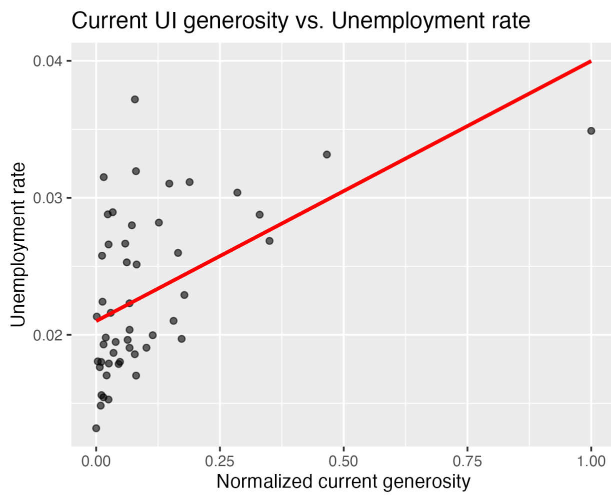

Figure 8.8: Plot of Current UI Generosity vs. Unemployment Rate, 2021