

Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data *

Peter Ganong, University of Chicago and NBER

Fiona Greig, JPMorgan Chase Institute

Pascal Noel, University of Chicago and NBER

Daniel M. Sullivan, JPMorgan Chase Institute

Joseph Vavra, University of Chicago and NBER

July 19, 2022

Abstract

We show that the largest increase in unemployment benefits in U.S. history had large spending impacts and small job-finding impacts. This finding has three implications. First, increased benefits were important for explaining aggregate spending dynamics—but not employment dynamics—during the pandemic. Second, benefit expansions allow us to study the MPC of normally low-liquidity households in a high-liquidity state. These households *still* have high MPCs. This suggests a role for persistent behavioral characteristics, rather than just current liquidity, in driving spending behavior. Third, the mechanisms driving our results imply that temporary benefit supplements are a promising countercyclical tool.

*We thank Joe Altonji, Adrien Auclert, Gabriel Chodorow-Reich, Arin Dube, Jason Furman, Jon Gruber, Greg Kaplan, Rohan Kekre, Bruce Meyer, Matt Notowidigdo, Heather Sarsons, Jesse Shapiro, Daphne Skandalis, Amir Sufi, and Rob Vishny for helpful conversations, and seminar participants at the AEA, Berkeley, BFI China, BFI Macro Finance, Boston Fed, CFPB, Chicago Booth Micro, Macro, and Finance Lunches, Clemson, Columbia, Empirical Macro Workshop-LA, Federal Reserve Board, Insper, Johns Hopkins, Kellogg, NBER Labor Studies, Public Finance and EFG, MIT, Montana, OECD, Opportunity Insights, RAND, Richmond Fed, SED, SITE, SOLE, University of Texas, the Upjohn Institute, UIUC, VMACS, Washington University St. Louis, and Yale for suggestions. We thank Maxwell Liebeskind, who was a coauthor on several of the earlier JPMCI policy briefs on pandemic UI. We thank Samantha Anderson, Timotej Cejka, Rupsha Debnath, Jonas Enders, Isaac Liu, Michael Meyer, Liam Purkey, Peter Robertson, John Spence, Nicolas Wuthenow, and Katie Zhang for excellent research assistance. We thank the Becker Friedman Institute and the Kathryn and Grant Swick Faculty Research Fund at the University of Chicago Booth School of Business for financial support. This research was made possible by a data-use agreement between three of the authors and the JPMorgan Chase Institute (JPMCI), which has created de-identified data assets that are selectively available to be used for academic research. All statistics from JPMCI data, including medians, reflect cells with multiple observations. The opinions expressed are those of the authors alone and do not represent the views of JPMorgan Chase & Co.

1 Introduction

As part of its pandemic policy response, the U.S. government implemented the largest increase in unemployment insurance (UI) benefits in history. Supplements of \$300-\$600 per week were added on top of regular benefits at various points during the pandemic. These supplements more than doubled typical benefit levels, leading most unemployed workers to receive *more* income from unemployment than they had from their prior jobs (Ganong, Noel, and Vavra 2020). In total, nearly half a trillion dollars in supplements were paid out through this program.

How did this unprecedented increase in unemployment benefits affect labor markets and household spending? This paper uses administrative bank account data covering millions of households, several causal research designs, and a dynamic structural model to estimate and interpret the responses to these supplements. We find that, gauged in several different ways, these increases in benefits had large effects on spending but small effects on job-finding.

Analyzing the impacts of this massive increase in benefits is useful for three reasons. First, this intervention was so large that measuring its impact is important for understanding both aggregate and distributional dynamics during the pandemic recession. Second, the scale and persistence of these transfers provides a unique laboratory for testing implications of benchmark structural models which have been difficult to assess before. In particular, many modern macro models explain a high marginal propensity to consume (MPC) as originating from households with temporarily low liquidity. Because the transfers we study are so large, they push otherwise low-liquidity unemployed households into a high-liquidity state. We can therefore test if the MPC falls as much as benchmark models predict. Third, countercyclical benefit levels have never before been attempted at this scale, and there is no prior evidence about their impacts. Understanding their effects is therefore important for informing future policy design.

Measuring the impacts of expanded benefits requires a dataset with information on spending, liquidity, employment transitions, and unemployment benefit receipt. We build such a dataset by using de-identified bank account transactions from the universe of Chase customers. We observe the precise week that individual households begin receiving UI supplements, and can trace out the high-frequency impact of these supplements on spending, job-finding, and the evolution of liquid balances.

We use this data set first to describe time-series patterns and then turn to causal identification strategies. On the spending side, there is a strong relationship between supplement levels and the spending of the unemployed. Spending rises when \$600 supplements begin in April 2020, falls when they end in July 2020, and then rises again when \$300 supplements begin in January 2021.

Most strikingly, we see that while the \$600 supplement is available, the spending of unemployed households *rises* after job loss, both in absolute terms and relative to the spending of employed households. This reverses the typical relationship between unemployment and spending, as spending usually falls sharply after job loss (Gruber 1997). Moreover, this increase is particularly notable since typical households substantially reduced spending during the pandemic (Chetty et al. 2020; Cox et al. 2020).

Next, we estimate the causal effect of benefits on spending using difference-in-difference research designs. First, to identify effects at the start of benefit receipt we use variation induced by processing constraints from overwhelmed state UI agencies at the onset of the pandemic. In particular, we compare spending of unemployed workers who receive benefits immediately after job loss to spending of unemployed workers who lose jobs at the same time but face delays in benefit receipt. We argue that

this timing variation is largely random and show that these workers have nearly identical spending patterns prior to benefit receipt. Second, to identify the effects of the expiration of the \$600 and start of the \$300 supplements, we compare the spending of unemployed workers to the spending of employed workers.

The main finding from these research designs is that spending is highly responsive to benefit changes, *even* when supplements are so large that recipient households have elevated liquidity. In each case, spending responds sharply in the exact week in which benefit levels change. Spending responses are quantitatively large across all specifications—with estimated one-month MPCs between 26 and 43 cents—and statistically precise with standard errors of three cents or less.

We next examine the impact of supplements on job-finding. This is particularly useful because there has been widespread uncertainty over whether more generous benefits were responsible for low re-employment rates during the pandemic. Indeed, while the supplements were in effect, one survey of economists showed that a majority were uncertain about whether they were a “major disincentive to work” (Initiative on Global Markets 2021). Furthermore, the share of economists who were uncertain about the disincentive was higher than for more than 93% of the other IGM survey questions asked since January 2020.

In contrast to the large effects of supplements on spending, we find that the supplements had only small effects on job-finding. Looking first at descriptive patterns, we see a dramatic decline in the job-finding rate at the start of the pandemic beginning a month before the start of the \$600 supplements. We then see a small increase in the job-finding rate when the \$600 supplements expire and a small decline in the job-finding rate when the \$300 supplements begin.

We estimate the precise causal effects of the supplements on job-finding using two research designs that exploit distinct sources of variation in benefits. First, we use an interrupted time-series design which relies on the overall expiration or onset of benefits. Second, we use a dose-response difference-in-difference design which compares workers with smaller and larger changes in benefits. Both research designs yield similar results. The implied duration elasticity of the \$600 supplements is 0.06-0.11 while for the \$300 supplements it is 0.10-0.21. Elasticities remain small within industry, state and age groups. Thus, they do not appear to be driven by confounding labor demand effects arising from a correlation between sector wages and sector exposure to pandemic job losses, by differences in other state-level policies, or by age-related labor supply effects. We therefore conclude that job-finding effects are small, both relative to overall fluctuations in job-finding during the pandemic and relative to pre-pandemic estimates of the effects of benefits on unemployment duration (Schmieder and von Wachter 2016).

In the second part of the paper we interpret these empirical results through the lens of a dynamic structural model. Our model framework is intentionally standard and we discipline its key features using the reduced-form causal estimates from our empirical analysis. We show that a suitably-calibrated model is able to closely match the empirical patterns we observe in the data. Furthermore, we show that the model also matches several additional spending and job-finding patterns which are not directly targeted.

A model is crucial for two primary purposes. First, the model allows us to translate the reduced-form effects we can identify in the data to the total policy effects we are most interested in measuring. Our research designs identify how supplement *changes* at a point in time affect subsequent job-finding and spending. However, we are more interested in the *total* effects of supplements. For example, we

want to know how much paying \$600 supplements from April-July decreases job-finding and increases spending from April-July. Our reduced-form evidence instead answers the slightly different question of how much ending \$600 supplements in August increases the job-finding rate and lowers spending in August. If causal effects are constant over time, then the answer to these two questions will coincide. However, in a dynamic environment with supplements which last for several months, the answers to these two questions could potentially differ. For example, if forward-looking households started searching harder for jobs in anticipation of the expiration of the \$600 supplement, then the change in the job-finding rate at expiration would underestimate the total effects of the \$600 on job-finding in earlier months.

We find that the policy's dynamic effects are limited, so our reduced-form empirical estimates are actually quite similar to the total policy effects. The limited role of dynamic effects arises because the model only matches joint patterns of spending and job-finding if households have muted responses to policy changes which will occur in the future. The lack of anticipatory behavior means that the effects of benefits on spending and job-finding are fairly constant over time.

Thus, our best-fit model delivers similar conclusions to our reduced-form estimates about the total effects of supplements: expanded UI supplements had large effects on spending and small effects on employment. This conclusion holds both relative to pre-pandemic evidence and relative to overall fluctuations in employment and spending during the pandemic. In particular, if we compare the best fit model to a pre-pandemic calibration, we find an elasticity of unemployment duration to supplements that is 80% smaller and an increase in spending while supplements are in effect that is 75% larger. When comparing to aggregate fluctuations during the pandemic, we estimate that UI supplements explain only 5% of aggregate employment shortfalls while they explain more than 25% of the recovery in aggregate spending.

The model's second main purpose is to help quantify the role of specific mechanisms driving the large spending effects and small job-finding effects of supplements. Looking backwards, understanding mechanisms is necessary for explaining this key episode in recent macroeconomic history. Looking forwards, understanding mechanisms is necessary for identifying which forces likely generalize beyond the pandemic. Such forces may provide new lessons about household behavior and future policy design.

Three forces explain why the employment distortion is small relative to prior estimates. First, supplements were temporary and implemented in a recession, when the new job-finding rate was already depressed. This limits the scope for supplements to affect employment. Second, businesses recalled many laid-off workers when they reopened, reducing the sensitivity of overall job-finding rates to benefits. Third, there were pandemic-induced reductions in the sensitivity of job-finding to benefits. The second and third forces are unlikely to generalize beyond the pandemic, but the first force explains about half of the reduced employment distortion and should generalize: short-lived increases in unemployment benefits during recessions, when the job-finding rate is depressed, are likely to induce small employment distortions. This finding has policy implications which we return to below.

Two main forces explain the large spending response to supplements. First, supplements target households who have lost their jobs. In the absence of supplements, these households are low liquidity since regular benefits do not fully replace lost earnings. This fact can help explain why unemployed households respond strongly to the *first* dollar of supplements they receive. However, in total, sup-

plements are so large that they actually drive households fully out of this liquidity-constrained state. The median unemployed worker receives more in benefits than lost earnings and sees their checking account balance more than double while supplements are in place. This means that low liquidity alone cannot explain why supplements drive the spending of unemployed *above* that of employed workers. Instead, unemployed households must also have some persistent characteristic such as greater impatience or present-bias that raises their MPCs even when they have high liquidity. Quantitatively, in order for our model to match both prior evidence on the MPC out of universal stimulus payments and our evidence on the MPC out of targeted UI supplements, we need unemployed households to have discount rates that are twice as high as those in the general population.

Overall, we take three main lessons from the results in this paper. First, benefit supplements were important drivers of aggregate and distributional dynamics during the pandemic. In particular, from an aggregate perspective we conclude that benefit supplements are important for explaining the dynamics of spending, but *not* the dynamics of employment. From a distributional perspective, the large spending responses that we document can help explain why spending rose most for low income households during the pandemic (Cox et al. 2020) even though these same households had the biggest declines in labor income (Cajner et al. 2020).

Second, permanent household characteristics—and not just current liquidity—are important for understanding household consumption patterns. Prior work has documented an empirical correlation between low liquidity and high MPCs. Drawing on this evidence, a large literature has developed models in which high MPCs arise solely from temporarily low liquidity (see Kaplan and Violante 2014 and Laibson, Maxted, and Moll 2021). Our results show that permanent heterogeneity across households is *also* important for understanding spending behavior. In most economic environments, it is hard to disentangle the role of temporary economic circumstances from that of permanent behavioral traits in driving the liquidity-MPC correlation: households facing a temporary liquidity crunch will have high MPCs, and households with permanently high propensities to spend will run down their liquidity. However, the large and sustained sequence of transfers we study provides an environment where we can observe the spending responses of normally low-liquidity households in a high-liquidity state. The high spending responses we observe provide direct evidence that temporarily low liquidity cannot be the only explanation for high MPCs.

Third, temporary benefit supplements are a promising countercyclical tool. Our finding that employment distortions hinge upon the interaction between benefit length and labor market conditions provides a new rationale in favor of countercyclical benefits. When benefits are temporary and job-finding is depressed, we show that a large wedge opens up between the per-period effect of benefits on job-finding and the total effect of benefits on unemployment durations. This force pushes up optimal benefit levels during recessions.

Large spending effects also suggest that targeted unemployment supplements are an effective way to support aggregate demand during recessions. For the last 20 years, the federal government has repeatedly used universal or near-universal tax rebate payments to combat recessions. An alternative approach is to target payments to certain households, such as the unemployed, who may have particularly high MPCs. Such targeted transfers have occasionally been implemented in the past, but never at this scale and never in a way that has allowed an econometric identification of their spending impacts. We show that such targeted payments can indeed have powerful effects. The combination of low pre-transfer liquidity among the unemployed and a higher persistent propensity to spend at

any liquidity means that even large temporary unemployment supplements may be beneficial. We find that it would be preferable from a stimulus perspective to give up to a \$2,000 one-time payment targeted to unemployed households before giving \$1 of untargeted stimulus to all households.

Our paper connects to several additional strands of past research. There is a rich empirical literature analyzing the effect of pandemic UI supplements.¹ Our paper makes four contributions relative to this literature. First, while prior work has evaluated the labor market impacts of expanded benefits, our paper is the first to study spending impacts. We show that the effects on spending were much larger than the effects on employment. Furthermore, we find that studying the joint patterns of employment and spending is crucial for disciplining household expectations about policy changes, which is necessary for interpreting reduced-form empirical results. Second, although much of this prior literature imputes likely UI reciprocity based on observable characteristics, we are able to directly study households who actually receive benefit supplements. Third, our large sample size of over one million UI recipients allows for improved precision. This is especially important for detecting non-zero but small job-finding effects. Fourth, we develop a structural model to help interpret the empirical patterns in the data. This is crucial for understanding the mechanisms driving the empirical results, enabling us to identify which lessons are likely to generalize beyond the pandemic environment.

Our paper also contributes to a literature analyzing the optimal cyclicity of unemployment benefits (Kroft and Notowidigdo 2016). One strand of this literature focuses on labor market impacts (Landais, Michaillat, and Saez 2018a,b; Mitman and Rabinovich 2015). Relative to this literature, our contribution is to document an additional force in favor of countercyclical benefits arising from the interaction between benefit length and labor market conditions. A second strand of this literature focuses on demand impacts (Kekre Forthcoming; McKay and Reis 2021). Relative to these papers, we provide direct *empirical* micro evidence that this mechanism is indeed powerful in an actual recession and also identify a mechanism (permanent heterogeneity) which amplifies the strength of countercyclical UI relative to untargeted stimulus checks.

Our paper is also connected to a growing literature on the distribution of MPCs, with a particular focus on relationships with income and liquidity. For example, Kueng (2018) finds that high income households with substantial liquidity have a high MPC out of payments from the Alaska Permanent Fund, and Fagereng, Holm, and Natvik (2021) find that even high-liquidity households still have a substantial MPC out of lottery winnings. Theoretical models by Lian (2022), Boutros (2022), and Ilut and Valchev (2020) provide psychological foundations for such high MPCs. Furthermore, several papers provide complementary evidence that permanent heterogeneity is important for explaining high MPCs. Aguiar, Bils, and Boar (2021), and Gelman (2021) combine panel data with a structural consumption model to disentangle the separate role of permanent and transitory forces. Parker (2017) shows that MPCs are correlated with persistent characteristics like impatience in surveys, and Patterson (2022) shows a relationship between unemployment risk and MPCs. The pandemic UI setting differs from prior empirical work in that we can observe that MPCs for normally low-liquidity households remain high even when these same households are moved to a high-liquidity state, and this may be a useful target for the next generation of theoretical models.

Finally, our paper contributes to a methodological literature by showing how a structural model can be useful for evaluating the validity of difference-in-difference research designs. The extent to

¹See, e.g., Bartik et al. (2020), Coombs et al. (2022), Dube (2021), Finamor and Scott (2021), Holzer, Hubbard, and Strain (2021), Marinescu, Skandalis, and Zhao (2021), and Petrosky-Nadeau and Valletta (2021).

which the effects of policy changes identified by these designs align with total policy effects will typically depend on whether these policy changes were anticipated. It is thus common to make a “no anticipation” assumption when using these designs (see Roth et al. 2022). We suggest that researchers who are concerned about bias from anticipation could instead use an economic model to quantify the potential magnitude of these identification threats. Our approach using an economic model complements statistical methods for diagnosing pre-trend violations such as those proposed by Freyaldenhoven, Hansen, and Shapiro (2019).

2 Institutions and Data

We begin with a brief discussion of the changes in unemployment insurance policies over the course of the pandemic and then describe the data that we use to analyze their impacts.

2.1 Expansion of Unemployment Benefits

The Coronavirus Aid, Relief and Economic Security (CARES) Act implemented a variety of policies in response to the emerging pandemic. One provision was a massive expansion of unemployment benefits. The CARES Act established a \$600 per week supplement from April-July 2020 on top of any amount already allotted by regular state unemployment insurance. The CARES Act also expanded eligibility for unemployment benefits to many self-employed and gig workers who would not otherwise qualify for regular benefits, through the creation of the Pandemic Unemployment Assistance (PUA) program, which was initially authorized for 39 weeks. Unemployed workers who qualified for UI through the PUA program were also eligible for the \$600 supplements. Because of data constraints, our analysis does not distinguish between regular benefits and PUA. However, most benefit recipients in the analysis sample are receiving regular benefits.² Finally, the CARES Act also established Pandemic Emergency Unemployment Compensation (PEUC), which extended benefit eligibility by an additional thirteen weeks for those who would have otherwise exhausted unemployment benefits.

The original CARES Act legislation authorized \$600 supplements through the end of July 2020. As the end of July approached, the fate of the expanded unemployment benefits remained unclear. Congressional Democrats advocated a continuation of the \$600 supplement, while some congressional Republicans advocated a \$400 supplement. Perhaps surprisingly, the two sides failed to reach any legislative compromise and the supplement fell to zero at the start of August.³

At the end of December 2020, new legislation authorized a \$300 per week supplement through mid-March 2021, and in March this supplement was extended through early September, making it available to unemployed workers for a total of eight months.

²In the two states where JPMCI can distinguish between regular benefits and PUA (Ohio and New Jersey), 74% of observed UI spells are for households receiving regular benefits. Among households who meet the account activity screens described below, the share is even higher.

³On August 8, an executive order implemented a “Lost Wages Assistance” (LWA) program to provide supplements for an additional six weeks. However, long delays meant LWA payments nearly always occurred *after* the program expired. We focus primarily on the \$600 and \$300 supplements which were paid contemporaneously with regular benefit payments for the same week, but we also briefly analyze the impacts of LWA.

2.2 Data

Our analysis sample is drawn from the 44 million households with a checking account in the JP-MorganChase Institute (JPMCI) data from January 2018 through February 2021. Our sample runs through February 2021 because this is when we can most reliably measure job-finding. Specifically, benefit eligibility extensions mean that UI exits from April 2020 through February 2021 rarely reflect benefit exhaustion and therefore usually reflect a return to work.⁴ For this reason, we generally use job-finding and UI exit interchangeably throughout the paper. The unit of observation is household-by-week. Our primary analysis sample consists of 1,493,597 unemployment benefit spells from 44 states in 2020. Figure A-1 shows a map of which states are in the sample. The exact set of states included in each part of the analysis is dependent on data availability and is described in Appendix B.1.

We measure unemployment insurance spells and labor income using information from direct deposits. We combine information on unemployment and employment spells to separate UI exits to a new job from UI exits to recall, which is when an unemployed worker returns to their prior employer. The details of how we construct these spells are described in Appendix B.2. We impose activity screens to ensure we capture workers whose primary bank account is at Chase and have stable employment prior to the pandemic (see Appendix B.3 for details).

We construct two main measures of spending. Our preferred *total spending* measure sums spending on Chase credit cards, Chase debit cards, cash withdrawals, paper checks, and various electronic payments. This measure excludes debt payments on mortgages, cars, and credit cards, as well as transfers to other accounts. While this is the most comprehensive measure of spending we can observe in Chase accounts, we note that it excludes most durable purchases and so most closely corresponds to broad non-durable spending.⁵ Nevertheless, two concerns arise from this measure. First, it includes some payments where the payee cannot be identified.⁶ It is thus possible that some of these transactions may not actually be spending. Second, it includes spending with potential for timing-related measurement error if there is a delay between when a paper check is written and when it is deposited.

We therefore also report results for a more narrow *card and cash* spending measure which excludes all paper checks and most electronic payments.⁷ Since it omits many recurring payments and eliminates the timing-related measurement error induced by paper checks, card and cash spending is better suited for measuring week-to-week spending changes caused by week-to-week income changes. Nevertheless, this more narrow measure is a conservative lower bound for non-durable spending because it omits all payments to unknown recipients, even though many of these really are spending.

We also measure household income and checking account balances. We define income as total inflows to Chase deposit accounts, excluding transfers. This definition captures take-home income because we only observe income after taxes and other deductions like retirement account contributions and health insurance premiums are withheld. We exclude transfers (e.g., from other bank accounts,

⁴For example, the California Policy Lab (Bell et al. 2022) calculates that fewer than 3 in 1000 recipients exhausted benefits during this time period. However, beginning in March 2021, there are a number of UI exits which do not reflect job-finding because of a technical issue with how UI systems pay claims for spells that last longer than a year (Bell et al. 2021). Extending our analysis to later periods therefore requires measuring job starts from direct deposit labor income in order to distinguish between UI exits arising from job-finding and exhaustion.

⁵Most durables like cars and houses are financed, and our spending measure will not include these purchases.

⁶We are unable to identify the recipient of payments made by paper check. Furthermore, apart from debt payments and transfers to other bank accounts, we are unable to categorize the majority of remaining electronic account outflows (e.g., those made via wire transfer, ACH, and other electronic channels).

⁷We include electronic payments for utility bills because it is a form of non-durable spending.

money market accounts, and investment accounts) to avoid double-counting income. Checking account balances are measured at a monthly frequency as the account balance on the final business day of the month. We sum balances across all accounts for households with multiple Chase checking accounts. Finally, we observe Economic Impact Payments, age, number of children, and industry of work for selected subsamples. Additional detail is provided in Appendix B.4.

A key strength of the JPMCI data is the ability to track monthly income, spending, and balances for the same households. Most household-level data sources track just one or two of these variables, making it impossible to uncover the effects of benefits on spending and balances. The survey and administrative datasets which capture balances typically do so at one point in time or at the end of each calendar year; such datasets are therefore unable to detect the rapid accumulation of balances by the unemployed at the start of the pandemic and decumulation after the \$600 supplement expires.

Table A-1 provides summary statistics on the main flow measures of interest—income, UI benefits, total spending, and card and cash spending—as well as checking account balances.

2.2.1 Comparison to External Benchmarks

The massive increase in unemployment benefits is readily apparent in the JPMCI data. We compare the number of continued claims in Department of Labor (DOL) data to the number of households receiving unemployment benefits in the JPMCI data. From early March to June, Figure A-2a shows that these series rose by a factor of 15 in DOL and a factor of 17 in JPMCI. The figure also shows that the increase in UI payments to households that meet the account activity screens in JPMCI is a bit smaller, rising by a factor of 11. It is not surprising that this subsample has a smaller increase in unemployment benefit receipt. The pandemic recession was particularly difficult for underbanked households, who are likely to be omitted because of these account activity screens. We calculate using the Current Population Survey that the unemployment rate from April to September 2020 has been 4 to 6 percentage points higher for underbanked households.

The JPMCI data also reproduce differences across states in the magnitude of the increase in UI as well as the level of weekly UI benefits. Figure A-2b shows that the states with the largest increase in UI claims (e.g., Florida, North Carolina) in DOL also have the largest increase in JPMCI. Conversely, the states with the smallest increases in DOL (e.g., Wyoming, Idaho, West Virginia) also have the smallest increase in JPMCI.

Figure A-2c shows that there is a strong cross-state correlation between benefit levels in DOL and benefit levels in JPMCI. The figure also shows that weekly benefit levels are a bit higher in JPMCI than in DOL. UI benefit levels are based on a worker’s pre-separation earnings, so higher benefit levels in JPMCI imply that UI recipients in JPMCI have slightly higher pre-separation earnings than the average UI recipient in each state. This pattern is largely explained by the effect of the account activity screen, which imposes a minimum level of pre-separation earnings that is more stringent than the eligibility requirements for UI. We conjecture that the consumption responses we estimate are a lower bound for consumption responses among the full population of UI recipients since this screen induces mild positive selection in terms of labor market attachment and financial well-being. The possible bias from this screen for the job-finding rate is less clear; we note that the exit rate from unemployment without this screen produces a similar aggregate series to our main estimates (see Figure A-3a).

Finally, the JPMCI data capture shifts in the industry composition of unemployment. Figure

A-4 shows that industries with the largest increase in unemployment in DOL such as retail and accommodation & food services also have the largest increase in JPMCI. Industries with the smallest increase in unemployment in DOL such as construction also have the smallest increase in JPMCI.

We conclude that the JPMCI data does a good job of capturing both the massive national increase in UI receipt as well as the cross-state and cross-industry heterogeneity which can be captured using statistics reported by the DOL. For additional analysis of the representativeness of unemployed households in JPMCI data, see Ganong and Noel (2019).

3 Spending Responses to Expanded Unemployment Benefits

This section explores the empirical effects of unemployment benefits on spending. We begin with descriptive analysis. We next identify one-month causal responses separately to the onset of benefits, to the expiration of the \$600 supplement, and to the onset of the \$300 supplement. Each of these empirical exercises has distinct advantages and disadvantages, but they all lead to the same bottom-line conclusion: spending consistently rises and falls with unemployment insurance benefits.

3.1 Descriptive Patterns: Spending of the Unemployed *Rises* After Job Loss When Expanded Benefits Available


Figure 1 compares changes in spending and income for households who become unemployed and receive UI from April 2020 through February 2021 to similar households who remain employed through February 2021. We reweight the employed sample so that it exactly matches two observable characteristics of the unemployed: 2019 income quintile and date of Economic Impact Payment (EIP) receipt. Matching by income is potentially important since low-income households were more likely to become unemployed during the pandemic (see Table A-1) and households at different points of the income distribution may have spending that evolves differently over the pandemic. Matching by EIP date is potentially important since these stimulus checks arrive around the time that \$600 supplements start.⁸

Figure 1 shows that income for the unemployed rises and falls with the ebb and flow of benefit supplements. Prior to the start of unemployment, month-to-month changes in income are nearly identical for the two groups. Note that the matching procedure described above generates similarity in only the *level* of income; the similarity of pre-pandemic *changes* in income between the two groups is not mechanical and suggests that unemployed and employed groups indeed face the same income process prior to the pandemic. Beginning with the start of unemployment, income of the two groups diverges substantially. Since the combination of regular UI plus the \$600 supplement results in average replacement rates above 100 percent, income actually rises substantially for the unemployed from April through July. At the end of July, the \$600 supplements expire, and so the income of the unemployed falls below that of the employed. Income rises briefly in September 2020 with the payment of Lost Wages Assistance (LWA). Finally, income rises again for the unemployed in January 2021 when the \$300 supplement begins.

After households become unemployed and receive \$600 weekly supplements, their spending rises substantially *above* pre-pandemic levels. The middle panel of Figure 1 shows the evolution of monthly

⁸In practice, all of our results are nevertheless quite similar if we instead use the raw unweighted employed as a comparison group.

Figure 1: Spending of Unemployed Versus Employed



Notes: This figure compares income, spending, and checking account balances of unemployed and employed households using JPMCI data. The blue line shows households that receive unemployment benefits from April 2020 through at least February 2021. The purple line shows employed households who are matched on 2019 income quintile as well as date of receipt of Economic Impact Payment. The \$600 supplement is first paid in the middle of April, so May is the first complete month during which households have the opportunity to spend the supplement.

spending for the two groups. Like income, spending of the unemployed evolves nearly identically to the employed prior to the point of unemployment and then rises sharply at the start of unemployment in April 2020. This is especially notable when compared to the declining spending of employed households in this early part of the pandemic. Usually unemployed households reduce spending relative to employed households (Gruber 1997), but during the period of \$600 supplements, these normal patterns are reversed. This sustained increase in relative spending occurs for the entire time the \$600 supplement is in place. When the supplement ends at the end of July, there is then an immediate decline in spending. This is followed by a temporary rebound when unemployed households receive temporary LWA supplements in September. Spending then remains depressed until the \$300 supplements begin in January 2021. These supplements lead to a median replacement rate of 100%

and the spending of unemployed and employed households is similar after they begin. Thus, we find a strong relationship between unemployment benefit levels and the spending of the unemployed throughout the pandemic.⁹

The bottom panel of Figure 1 shows that there is also a large and sustained increase in the checking account balances of unemployed households. As with spending, it shows an increase both in absolute terms and relative to employed households. Increases in income for unemployed households during this period were so large that they accumulated additional savings even as their spending increased.

Our finding of a significant increase in spending among unemployed workers is reminiscent of the striking pattern documented in Gerard and Naritomi (2021). They show that unemployed workers who receive severance pay in Brazil sharply increase spending upon unemployment. We complement this evidence by showing that such an increase can last even for many months. Furthermore, because we can also track liquid balances, we are able to show that such an increase can persist even when households have been pushed far off their liquidity constraints. This large increase in liquid balances for unemployed households will be important when we turn to model and policy implications.

Although the evidence in Figure 1 is *qualitatively* consistent with a causal relationship between benefit supplements and spending, we caution against interpreting the time-series difference in spending between the two groups as a *quantitative* measure of the causal effect of the supplements. The employed group in purple has roughly constant household income. But had there been no supplements, household income for the unemployed would have fallen quite substantially! Spending would likely have fallen substantially as well. Thus, the time-series of spending for the employed is useful for understanding contemporaneous aggregate changes in spending, but does not provide a counterfactual for the no-supplement spending of the unemployed. This limitation of the time-series evidence motivates the use of alternative research designs to estimate the causal impact of benefits on unemployment in the next section.

3.2 Causal Evidence

In this section, we estimate the one-month MPC out of income changes induced by three different changes in unemployment benefits.

3.2.1 Waiting for Benefit Receipt

The unemployment benefit system was slow to pay out benefits of some claimants at the start of the pandemic. Many state unemployment agencies were overwhelmed by the large increase in unemployment claims at the start of the pandemic, meaning that the payment of many claims was delayed. We use these delays in payment to identify the causal impact of benefits on spending. We compare the spending of a treatment group of unemployed households who receive benefits promptly after filing a claim to the spending of a control group of unemployed households who experience delays in receiving benefits. To construct these groups, we compare cohorts of unemployed households all of whom stop receiving paychecks at the end of March, but differ in the date of first benefit payment. Since households who face delays ultimately receive back pay for missed payments, the treatment in this natural

⁹The evolution of changes in income, spending, and checking account balances are similar in direction when looking at medians, and even larger in magnitude (See Figure A-5). Median balances more than double for unemployed workers. In addition, Figure A-6 shows that all of the same patterns are present when looking at the subset of spending on card and cash.

experiment is therefore the *timing* of the arrival of liquidity, analogous to how Johnson, Parker, and Souleles (2006) study variation in the timing of stimulus checks.¹⁰ Focusing on the difference between two groups of unemployed households removes any direct effects of job loss itself on spending and isolates the effect of benefits.

The top panel of Figure 2 shows weekly patterns of unemployment benefits for four different groups of unemployed households. By construction, benefits are zero for each group prior to the first benefit week and then jump in the first week of benefits. Groups which start benefits later have larger jumps in the first week because of back pay.

We assume that delays in benefits are orthogonal to other determinants of spending behavior. Two types of evidence are consistent with this assumption. First, owing to the high volume of claims, overall delays in payments to eligible claimants were much longer than usual (Bitler, Hoynes, and Schanzenbach 2020) and it is unlikely that state UI systems were able to prioritize claims in ways that would be correlated with spending behavior. Second, Figure 2 shows that not only are the trends in card and cash spending prior to benefit receipt similar by cohort (the standard parallel pre-trends test), but the *levels* of spending are similar across cohorts as well.¹¹ The similarity of levels and of pre-trends for spending suggests that the length of delays are close to as good as random.

The arrival of UI benefits causes an immediate increase in spending. Figure 2 shows that spending jumps sharply in exactly the week in which benefits start. It then remains at an elevated level in subsequent weeks.

We use a difference-in-difference design to estimate the MPC out of UI benefits. We compare spending for a treatment group that receives benefits at the start of April to a control group that receives benefits at the start of June. Figure A-8 shows the difference in income and spending between the cohorts. Because this is a two-period two-group research design, it is not subject to the concerns raised in the recent literature on staggered implementation difference-in-differences. We estimate an instrumental variables (IV) regression with first and second stages respectively given by equations (1) and (2).

$$y_{i,t} = \alpha + \beta Post_t \times Treat_i + Treat_i + Post_t + \epsilon_{i,t} \quad (1)$$

$$c_{i,t} = \psi + MPC \times \hat{y}_{i,t} + Treat_i + Post_t + \varepsilon_{i,t}, \quad (2)$$


where $t = \{March, May\}$, $Treat = 1$ for households who become unemployed at the end of March and start benefits in April and $Treat = 0$ for households who become unemployed at the end of March but start benefits in June, $Post_t = 1$ if $t = May$ and $Post_t = 0$ if $t = March$.

We focus on a monthly MPC (rather than weekly) because it is less subject to misclassification in the exact timing of spending and it eases comparison to the prior literature. The \$600 supplement is first paid in the middle of April, so *May* is the first complete month during which households

¹⁰If the waiting group expects to receive benefits in the future, then they should engage in intertemporal substitution and smooth consumption while waiting for benefit receipt. In that case, the waiting MPC that we estimate will only capture liquidity effects of benefit receipt and not any income effects of benefits on spending. Thus, the waiting MPC that we estimate should be a weak lower bound on the true MPC out of benefits.

¹¹For this analysis of weekly spending, we focus on the narrower spending (card and cash) definition which is more stable from week to week. The spending (total) measure—which is our preferred measure for monthly MPCs—includes more monthly recurring expenses, making it more difficult to see the immediate impact of changes in cash flow from one week to the next. Nevertheless, Figure A-7 shows similar trends and levels are also similar for the broader spending (total) measure.

Figure 2: Impact of Delays in Unemployment Benefits on Spending



Notes: This figure shows mean benefits and spending (card and cash) for several cohorts of unemployed households using JPMCI data. All households stop receiving paychecks at the end of March, but differ in the date of first benefit payment.

have the opportunity to spend the supplement. The identifying assumption is that absent the start of unemployment benefits, the change in spending between March and May for the treatment group would be the same as the change in spending for the control group ($cov(Post_t \times Treat_i, \varepsilon_{i,t}) = 0$).

Table 1: Marginal Propensity to Consume out of Unemployment Benefits

Research design	Total spending MPC
Waiting for benefits	0.43*** (0.03)
Expiration of \$600 supplement	0.30*** (0.01)
Onset of \$300 supplement	0.26*** (0.01)

Notes: This table shows estimated one-month total spending MPCs for three different unemployment benefit changes using equations (1) and (2). The waiting for benefits design compares unemployed households receiving benefits to those who face benefit delays ($N = 58,635$). The expiration and onset designs compare unemployed households to a sample of employed households matched on pre-separation income ($N = 355,548$ and $N = 287,761$ respectively). Standard errors are clustered by household.

We estimate a one-month MPC of 0.43 out of UI benefits in Table 1. This estimate implies that nearly half of unemployment benefits are spent in the first month after receipt. The spending response to benefit receipt will depend on household liquidity, the expected persistence of benefit changes and

whether these changes were anticipated or not. In Section 6.2, we discipline the role of these forces using a model matched to joint spending and job-finding moments and show that this MPC is large relative to prior work.

However, it is important to immediately note two limitations of this waiting design. First, it measures the MPC out of total unemployment benefits (regular unemployment plus the \$600 supplements) rather than the response to the supplements alone, since both payments start at the same time once benefits are received. Second, it studies the effect of variation in liquidity over a two-month horizon, some of which could reflect intertemporal substitution. We therefore use two additional identification strategies which examine persistent changes in income arising from changes in supplements.

3.2.2 Expiration of \$600 Supplement and Onset of \$300 Supplement

We next show that changes in benefit supplements have a substantial impact on spending. We study the expiration of the \$600 supplement at the end of July 2020 and the onset of the \$300 supplement at the start of January 2021. Like the waiting design above, each of these two episodes offers an opportunity to study a change in cash flow with no change in employment status.

Our identifying assumption is thus that spending *changes* would have been the same for the unemployed and employed at the time of the change in the supplement, if not for the policy expiration and onset. As in Section 3.1, we use a group of employed households matched on pre-pandemic income and EIP date. While this assumption is unlikely to hold for the time-series as a whole, there are no obvious economic events that should violate this assumptions at the exact time that the \$600 expires or the \$300 starts.¹² To validate this identifying assumption, we study the evolution of spending for unemployed and employed households prior to the policy changes. Figures A-9 and A-10 show similar pre-trends prior to the expiration of the \$600 supplement and the onset of the \$300 supplement respectively.

The supplements have an immediate visible impact on spending. Spending on card and cash drops sharply at the expiration of the \$600 supplement (Figure A-9) and rises sharply at the onset of the \$300 supplement (Figure A-10). Effects on total spending are also noticeable in Figures A-11 and A-12, but a bit harder to detect visually because of week-to-week fluctuations in spending for both the unemployed and employed.¹³

We estimate the one-month MPC out of supplement changes using the IV approach in equations (1) and (2). For expiration, we define $t = \{July, August\}$, $Post$ equals one in August, and $Post$ equals zero in July. The control group is the set of households with continuous employment, and the treatment group is the set of households who begin benefits by June 14 at the latest and continue receiving benefits through at least August 30. For onset, we define $t = \{December, January\}$, $Post$ equals one in January, and $Post$ equals zero in December. The treatment group is households that are unemployed from November 2020 through February 2021.

Spending responses are symmetric for positive and negative changes in supplements, which is consistent with the view that households were surprised by these benefit changes. Specifically, Table

¹²In contrast, we do not study the onset of the \$600 supplement using this identification strategy because the start of the \$600 supplement coincides with a transition into unemployment which likely lowers the spending of unemployed relative to employed and confounds the effects of supplements.

¹³Figure A-12 indicates that even prior to the onset of the \$300 supplement there is some evidence of a gradual downward trend in spending for unemployed relative to employed households. If we were to use a specification that accounted for this pre-trend in estimation we would likely find that the MPC out of the \$300 supplements is even larger than in our baseline estimate.

1 shows that we find an MPC of 0.30 at the expiration of benefits and an MPC of 0.26 at the onset of benefits. The sizable MPC at expiration which is similar to that at onset is particularly noteworthy. In the model in Section 5, similar spending responses to income decreases and increases will suggest that households were surprised by these benefit changes. The behavior of job-finding in Section 4 reinforces this surprise interpretation. The interpretation of expectations is in turn important for understanding whether the research designs in this section capture the full effect of the policy, and we therefore return to this issue in Section 5.3.

The large MPCs we estimate in this section are robust to a number of measurement choices and alternative sources of variation. Appendix C shows that our conclusions are robust to limiting the sample to households for whom we likely observe a more complete lens on spending, to alternative summary statistics, and to alternative measures of spending. In Appendix D, we analyze yet another source of policy variation in benefits (Lost Wages Assistance) and find a similarly high MPC of 32 cents (Table A-2).

It is likely that the MPCs out of benefits that we estimate in this depressed spending environment are lower bounds on MPCs out of benefits in more normal times. The pandemic environment is one in which many types of spending are depressed. Spending declines for employed households in the JPMCI data. Public aggregate data shows similar declines. While there were also changes to shopping patterns that might have increased spending on certain products like bulk purchases or home-related expenditures, it is important to note that the net effect of these pandemic forces is still that overall spending falls. Thus, it is particularly noteworthy that the spending of unemployed households *rises* in this period of time when everyone else’s spending falls.

4 Disincentive Effect of Benefit Supplements

The goal of this section is to estimate the effect of benefit supplements on the rate at which recipients exit unemployment. This can be used as an input to calculate the elasticity of unemployment duration with respect to benefit levels, a key measure of the disincentive effect of benefits analyzed in the prior literature. After briefly describing aggregate patterns, we describe results from two research designs. The first design in Section 4.1 uses an interrupted time-series approach. The second design in Section 4.2 uses a difference-in-difference approach. We apply both designs to the expiration of the \$600 supplement and the onset of the \$300 supplement. In Section 5.2, we show that all four estimates (two designs by two policy changes) imply similar small effects on job-finding after supplements change.

We focus primarily on the exit rate to new jobs rather the exit rate to recalls. This focus is motivated by the fact that UI eligibility requires recipients to accept any offer of “suitable work”, which should reduce the sensitivity of recalls to benefits. Nevertheless, in Section 4.4 we explore the effects of supplements on recall and show that even the upper bound of plausible causal impacts arising through recalls still implies small aggregate employment effects.


Beginning with descriptive patterns, Figure 3 shows that at the start of the pandemic, the new job-finding rate plunges by four percentage points and remains depressed thereafter.¹⁴ It also shows that the job-finding rate modestly rises and falls with the expiration and onset of the supplements.¹⁵

¹⁴Figure A-13 shows patterns for the recall rate and the total exit rate. The finding that the total job-finding rate is lower in the pandemic also holds in broader samples. Figure A-3a shows that it holds when we do not require account activity screens and Figure A-3b shows that it holds when we do not require an observed job separation.

¹⁵Pandemic UI eligibility expansions (PUA and PEUC) were originally legislated to last through the end of December

These descriptive patterns suggest that the supplements may modestly reduce job-finding rates, but that these effects were dwarfed by other factors holding back job-finding during the pandemic.¹⁶ We turn next to research designs that identify the causal effect of supplements on job-finding.

Figure 3: Exit Rate from Unemployment Benefits to New Job



Notes: This figure shows the exit rate to new jobs in JPMCI data. There is a surge in exits on January 3 and 10, which reflects a lapse in federal benefits rather than true exit to new job (see Figure A-14) and we therefore omit these weeks.

4.1 Research Design 1: Interrupted Time-series

We use an interrupted time-series design to estimate the effect of the supplements on the new job-finding rate. We make the identifying assumption that the job-finding rate would have been constant in the weeks just before and after a supplement change had there been no change in the supplement. Specifically, we study the change in the average job-finding rate in the two weeks prior to the policy change relative to the first four weeks after the policy change as illustrated in Figure A-16a. While this is a strong assumption, we note that we are using high-frequency weekly data so any confounding changes that affect the aggregate job-finding rate must occur at exactly the same time as the changes in supplements.¹⁷ In the next section, we describe an alternative research design which uses an orthogonal source of variation and is not subject to concerns about confounding high frequency aggregate shocks.

2020. Although the programs were ultimately extended, there was a brief lapse at the end of December and it took some time for states to restore benefits in early 2021. This means many workers stop receiving UI benefits on January 3 and January 10, as shown in Figure A-14. However, this does not reflect a change in the new job-finding rate. We therefore drop these dates in estimation below and in the figure.

¹⁶The sharp decline in the new job-finding rate during the pandemic is not driven by changes in the composition of who is unemployed; Figure A-15 shows that very similar aggregate dynamics are apparent for the subset of workers who were unemployed before the pandemic began.

¹⁷One example of such a possible confound is seasonality. Another example of a possible confound is policy changes (e.g., Economic Impact Payments) occurring at the same time that the \$300 payments start in January. However, Economic Impact Payments would, if anything, likely reduce job-finding and lead us to overstate rather than understate the magnitude of the already small effects we measure.

Using this design, we find small impacts of supplement changes on job-finding. The job-finding rate rises by 0.76 p.p. when the \$600 supplement expires, and falls by 0.55 p.p. after the onset of the \$300 supplement. These effects are economically small, as we discuss in more detail in Section 5.2. Further, we note that the job-finding rate is trending upward prior to the expiration of the \$600 supplement; if our estimates were to instead assume that the job-finding rate was rising linearly in the absence of the policy we would estimate an even smaller effect from the expiration of the supplement.

To assess statistical significance, we rely on the fact that the legislated duration of pandemic unemployment policies in the CARES Act was based on a highly uncertain forecast in March 2020 about the duration of the pandemic. The \$600 supplement was legislated in March 2020 to last 17 weeks, while the other pandemic-era unemployment programs were set to last 39 weeks. The timing of legislation passing a \$300 supplement was driven in part by the expiration at the end of 2020 of the 39-week programs. This motivates an approach to inference which treats the exact date of the policy change as random.

We compare the change in the job-finding rate at the actual dates of policy implementation to the change in the job-finding rate at 30 placebo dates where there was no implementation of a new policy. Figure A-16b compares the distribution of the changes in the job-finding rate at the placebo dates to the changes at the actual implementation dates. The observed changes at the policy implementation dates are more extreme than any of the changes at the 30 placebo dates. Thus the p -value for the null hypothesis that the policy has no effect and the change we observe occurred at random is $.032=1/(30+1)$ if we include the own implementation date and exclude the implementation date of the other policy.

4.2 Research Design 2: Difference-in-difference

As a complement to the interrupted time-series analysis, we use a difference-in-difference design to estimate the causal impact of the supplement on job-finding. Because the supplements added a constant dollar amount to every worker’s benefit, there is heterogeneity in the change in the replacement rate (the ratio of benefits to pre-separation earnings). For example, a worker with pre-separation earnings of \$600 per week and a regular weekly benefit of \$300 would see their replacement rate rise to 150% under a \$600 supplement, while a worker with pre-separation earnings of \$1,200 per week and a regular weekly benefit of \$600 would see their replacement rate rise to 100%. This heterogeneity in the intensity of treatment motivates a dose-response difference-in-difference research design. We first provide qualitative, graphical evidence that the effects of the supplements vary with the size of the increase in replacement rates, then describe the assumptions needed for identification of the causal effects of the supplements, and finally provide quantitative estimates.

4.2.1 Graphical Evidence

To measure the intensity of treatment for each worker we compute the percent change in benefits at the expiration or onset of a supplement. Because we compare one event where a supplement expires and another event where a supplement begins, we use the average value of benefits with and without the supplement in the denominator (symmetric percent change):

$$PctChange_i = \frac{2(b_{i,post} - b_{i,pre})}{b_{i,pre} + b_{i,post}}. \quad (3)$$

We measure each worker’s exposure to treatment using their typical benefit payment. Benefit payments can fluctuate from week to week for a number of reasons and we therefore define the typical payment $b_{i,pre}$ as the median weekly payment in the two-month period before a policy change. Given $b_{i,pre}$, we then impute $b_{i,post}$ based on statutory rules.¹⁸

Figure 4 shows the evolution of exit rates, dividing workers into groups with higher-than-median $PctChange_i$ (“more treated”) and lower-than-median $PctChange_i$ (“less treated”). Two lessons emerge from the figure. First, the two groups have similar trends in the job-finding rate before the policy changes. Second, the policy changes induce differential changes in exit rates for those who are differentially treated by the policy changes: when the \$600 supplement ends, the job-finding rate rises sharply for the high replacement rate group. Conversely, when the \$300 supplement begins, the job-finding rate falls sharply for the high replacement rate group. This is consistent with a causal effect of the supplement on exit rates.

To fully exploit the variation in replacement rates in the data, Figure 5 shows a binscatter of the relationship between each worker’s change in benefits and the change in the job-finding rate for both policy changes. Figure 5a shows that a larger decrease in benefits from the first policy change is associated with a larger increase in the job-finding rate. Figure 5b shows that a larger increase in benefits from the second policy change is associated with a larger decrease in the job-finding rate. The relationships appear to be close to linear, a point to which we return below.

4.2.2 Identification and Estimation

To quantify the effect of replacement rates on the job-finding rate, we use a regression framework. Let t index periods, i index workers, and e_{it} be an indicator for exit to new job. We use data on eight weeks where the supplement is not available and eight weeks where the supplement is available as captured by the indicator $SuppAvail_t$. We estimate the model:

$$e_{it} = \gamma PctChange_i + \alpha SuppAvail_t + \beta SuppAvail_t \times PctChange_i + \varepsilon_{it} \quad (4)$$


Identification in the dose-response difference-in-difference design requires three assumptions.

First, we make the standard orthogonality assumption: $\varepsilon_{it} \perp SuppAvail_t, PctChange_i$. The economic content of this assumption is that high and low-wage workers (who differ in $PctChange_i$) would have had the same trend in job-finding absent the policy change. This assumption has a testable prediction: parallel trends prior to the policy change. Figure 4 shows that the data are consistent with this assumption for the exit rate to new jobs. While this parallel pre-trend is reassuring, one might still be concerned about differential labor market trends for high and low-wage workers due to the uneven incidence of the pandemic across industries, locations, and workers of different ages, all of which are potentially correlated with wage levels. However, in Section 4.2.3, we show that nearly identical conclusions obtain when exploiting only within state-age-industry group variation. In addition, we note that if there were a persistent difference in job-finding trends (e.g., low-wage workers have faster employment growth because of business reopenings) then the bias in $\hat{\beta}$ will have opposite signs across the two policy changes because one is a decrease in benefits while the other is an increase in benefits.


¹⁸This imputation is necessary because we do not observe $b_{i,post}$ for workers who find a job before the \$600 supplement has expired or before the \$300 supplement has been reinstated. See Appendix B.5 for details.

Figure 4: Effect of Expanded Benefits: Event Study

(a) Expiration of \$600



(b) Onset of \$300




Notes: This figure shows the exit rate to new jobs around the expiration of the \$600 weekly supplement and the onset of the \$300 weekly supplement in the JPMCI data. The light blue line shows workers with a lower-than-median replacement rate with the supplement and the dark blue line shows workers with a higher-than-median replacement rate with the supplement. Exit rates are normalized by the average exit rate during the period before the policy change (June and July for the expiration of the \$600 and November and December for the onset of the \$300). Panel (b) omits a mechanical surge in exits on January 3 and 10 arising from the lapse in expanded UI eligibility. Figures A-17a and A-17b show the same figure but without a normalization in the pre-period. Figures A-17c and A-17d show that the same patterns hold when we look at the total job-finding rate (which includes both new job-finding and recalls).

Second, we assume that the causal effect of replacement rates on job-finding is homogeneous in the treatment group and control group. This assumption implies that raising a low-wage worker's replacement rate will have the same absolute effect on job-finding as raising a high-wage worker's replacement rate by the same absolute amount, thereby implying a linear relationship between replacement rates and exit rates. De Chaisemartin and D'Haultfoeuille (2018) show that this assumption is necessary for identification in dose-response designs. However, as we discuss above, the apparent linearity of the effect of benefit changes on the job-finding rate is consistent with a homogeneous treatment effect.

Third, we assume that jobseekers did not anticipate the changes in supplements. We discuss two tests which validate this assumption in Section 5.3.1.

Figure 5: Effect of Expanded Benefits: Difference-in-Difference Binscatter



Notes: This figure shows the change in the new job-finding rate at the expiration and onset of benefit supplements separately for deciles of the change in benefits as measured using equation (3). The top panel shows the difference in the average new job-finding rate between Jun 1-Jul 31 and Aug 1-Sep 31. The bottom panel shows the difference in the average new job-finding rate between Nov 1-Dec 31 and Jan 15-Mar 15. The slope estimates correspond to the $\hat{\beta}$ coefficients reported in Table A-3.

The key coefficient of interest in equation (4) is $\hat{\beta}$, which captures how the job-finding rate changes for more-treated vs less-treated workers. Table A-3 shows that at expiration, we estimate $\hat{\beta} = -0.017$ and at onset, we find a similar coefficient of $\hat{\beta} = -0.020$. These effects are precisely estimated with a standard error of 0.001 with standard errors clustered at the household level.

4.2.3 Robustness of Difference-in-Difference Estimates

We conduct three tests to validate these estimates. First, we estimate a version of equation (4) by week:

$$e_{it} = \gamma PctChange_i + \alpha Week_t + \beta_t Week_t \times PctChange_i + \varepsilon_{it} \quad (5)$$

This enables an event study interpretation of the coefficients. Figure A-18 shows that treatment effects from the expiration of the \$600 are largest in the three weeks after the policy expires and smaller in the subsequent five weeks. This suggests that long-term effects of expiration on the weekly job-finding rate may be even smaller than the baseline estimates from equation (4) which pool the eight weeks after expiration in Table A-3. The figure also shows a stable treatment effect from the onset of the \$300.¹⁹

In a second group of checks in Tables A-4a and A-4b, we re-estimate equation (4), adding controls X_i and $X_iSuppAvail_t$ to address concerns about differential trends by group. First, we add state (and state-by-supplement-available) fixed effects, so that identification comes from comparing the job-finding rate for higher- and lower-wage workers with different benefit replacement rates in the same state. Second, we add age (and age-by-supplement-available) fixed effects, so that identification comes from comparing the job-finding rate for higher- and lower-wage workers with different replacement rates who are in the same state and are the same age. Third and finally, we add industry (and industry-by-supplement-available) fixed effects. In this richest specification, identification comes from comparing the job-finding rate for higher- and lower-wage workers with different replacement rates who are in the same state, are the same age, and worked in the same industry. Our estimates of $\hat{\beta}$ change little from incorporating these control variables.

In addition to controlling for concerns which might bias our estimates, we also later explore how disincentive effects vary with observables like age, presence of children, liquidity, and industry. However, we view this analysis as primarily speaking to mechanisms and understanding magnitudes rather than mitigating threats to identification, and so we defer this analysis to Section 6.1.

Third, we explore whether allowing for serial correlation in standard errors among workers with similar pre-separation wages changes our conclusions about statistical precision. We do this using Conley’s (1999) method for calculating standard errors. This approach is usually used for correlation among geographic units using a two-dimensional distance metric but can be adapted to allow for correlation among workers with similar wages using a one-dimensional distance metric. Table A-5 shows that we continue to find that estimates of $\hat{\beta}$ are statistically significant under this alternative method of computing standard errors.

4.3 Comparison of Estimates Across Research Designs and Episodes

Comparing the estimates from the prior two sections requires re-scaling the four estimates (two research designs and two policy changes) into common units. Comparing within a given policy change, the interrupted time-series estimates tell us the average effect of the entire supplement while the difference-in-difference estimates tell us the effect of a marginal change in $PctChange$ across workers. We convert from the difference-in-difference estimate $\hat{\beta}$ to an estimate of the average effect of the entire supplement by assuming a homogeneous treatment effect over the entire range $[0, E(PctChange_i)]$ and computing $\hat{\tau} = \hat{\beta}E(PctChange_i)$. We note that this is a stronger assumption than that in Section 4.2.2 since $\min(PctChange_i) > 0$, and so this requires linear extrapolation out of sample.²⁰

Two types of evidence bolster the plausibility of such an extrapolation. First, within the empirical

¹⁹Figure A-18b indicates that even prior to the onset of the \$300 supplement there is already a gradual trend downward in the job-finding rate for households that receive the largest increase in benefits on January 1, 2021. If we were to use a specification that accounted for this pre-trend in estimation we would likely find that the \$300 supplement has even smaller effects on the job-finding rate.

²⁰Ideally, we would like to compare a treated group receiving a supplement (and thus $PctChange > 0$) to an untreated control group with no supplement (and thus $PctChange = 0$) but we have no such untreated control group and must instead extrapolate from comparisons across groups with different positive $PctChange$.

variation available in the data, the relationship between the intensity of treatment (size of the change in benefits) and the outcome (change in the exit rate) appears to be linear (Figures 5a and 5b). Second, in analysis of a structural model of job-finding in Section 5.1.2, we show that the effect of the supplement on the job-finding rate is also close to linear.

Table 2 shows our headline estimates of how UI supplements affect the job-finding rate. We estimate that the \$600 supplement reduces the weekly job-finding rate by 0.76 p.p. using the interrupted time-series estimates and by 1.45 p.p. using the difference-in-difference estimates. The \$300 supplement reduces the job-finding rate by 0.55 p.p. using time-series estimates and by 1.14 p.p using difference-in-difference estimates. Thus, the difference-in-difference estimates are about twice the size of the time-series estimates. However, as we discuss further in Section 5.2, all four of these estimates are economically small. Because the two different research designs rely on orthogonal sources of variation in benefits, the similarity of the estimates (in terms of their economic effects) across the two designs bolsters the conclusion that the supplements had small effects on the job-finding rate.

Table 2: Effect of Supplements on Weekly Job-Finding Rate

Effect of...	Interrupted time-series		Difference-in-difference	
	Entire supplement	Per \$100	Entire supplement	Per \$100
\$600	-0.76	-0.19	-1.45	-0.36
\$300	-0.55	-0.20	-1.14	-0.44

Notes: This table shows the effect in percentage points of benefit supplements on the weekly new job-finding rate. The first row uses estimates from the \$600 expiration and the second row uses estimates from the \$300 onset. Because we are comparing supplement increases and decreases, both of which are very large in size, we use a symmetric percent change calculation (see equation 3) when re-scaling effects for the “Per \$100” columns.

It is also useful to compare the effects of the \$600 supplement to the effects of the \$300 supplement. We report comparisons in two ways. First, we convert each estimate of the full supplement effect into an implied per-week causal effect of increasing benefits by \$100 relative to a no-supplement baseline.²¹ Table 2 shows that effects per \$100 were similar for both the \$600 supplements which came earlier in the pandemic and the \$300 supplements which came later in the pandemic.

Second, since elasticities are scale invariant, we compute a duration elasticity to benefit levels. We calculate this duration elasticity from the effects in Table 2 by assuming that supplements generate a constant per-week effect on the job-finding rate while they are in place and then calculating how this changes the average duration of unemployment. Details of this calculation are in Appendix E. We explore the robustness of this constant-effects assumption in Section 5. We find that the duration elasticity is 0.06 to the \$600 and 0.10 to the \$300 supplements using estimates from the time-series design and 0.11 and 0.21 using estimates from the difference-in-difference design. While elasticities using difference-in-difference estimates are slightly larger, overall these duration elasticities are significantly smaller than estimates in the prior literature. We discuss this comparison to prior estimates and the mechanisms explaining the low duration elasticity in more detail in Section 5.2.

Although our analysis treats the two research designs as estimating the same parameter, one possibility for why time-series estimates are smaller is if the “micro” disincentive effect of unemployment

²¹The models in Section 4.2 are estimated using symmetric percent change $PctChange_i$. The average of $PctChange_i$ is 81% for the \$600 supplement and 57% for the \$300 supplement. Note that because we are using symmetric percent change in equation (3), $PctChange_i$ is not linear in the size of the supplement. Relative to a no-supplement baseline, paying a \$100 supplement has an average value of 20% for $PctChange_i$. We therefore rescale the estimates from the \$600 supplement by 20%/81% and the estimates from the \$300 supplement by 20%/57%.

benefits (the effect of giving one worker more benefits) is bigger than the “macro” disincentive effect (the effect of giving all workers more benefits).²² The macro effect includes additional equilibrium effects.²³ The difference-in-difference estimate, which compares changes in the job-finding rate for more- and less-treated workers may be closer to a micro elasticity, while the interrupted time-series estimate, which measures the change in the job-finding rate for all workers may be closer to a macro elasticity.

4.4 Recalls

We now turn from studying the effect of the supplements on the exit rate to new jobs and instead study the effect of the supplements on the exit rate to recall. There is some evidence that the expiration of the \$600 supplement might have had a small effect on recalls but the evidence is hard to interpret, and even the upper bound of plausible causal impacts on recalls still implies small aggregate employment effects. There is no evidence of any effect of the \$300 supplement on recalls.

Figure A-13a shows time-series patterns of recall. The recall rate is highest while the \$600 supplement is still in place, suggesting it did not substantially deter recall. Indeed, more than half of unemployed workers return to work before the \$600 expires. This figure also illustrates that the recall rate is falling over time (making it hard to know what the counterfactual recall rate would have been in the absence of the supplement) and volatile (making it hard to assess statistical significance). There is evidence of a short-lived increase in recalls in the three weeks after the supplement expires. However, even if we make the aggressive assumption that recalls would have trended down through these three weeks in the absence of supplement expiration, the implied effect on the average duration of unemployment is tiny because this increase in the recall rate is so short lived. The time-series evidence around the start of the \$300 also suggests it had no effect on recalls. If anything recalls *rise* after the supplement takes effect. We also note that the aggregate recall rate is already low even before the onset of the \$300 supplement, meaning there is little scope for a further reduction from the \$300.

Difference-in-difference results cast further doubt on the possibility of substantial effects of supplements on recalls. Figure A-17e shows that recall is higher for the high replacement rate group after the \$600 supplement expires, but *not* in the three weeks when the aggregate recall rate is elevated. Instead, the recall rate for the high replacement group only rises differentially in the subsequent six weeks. Table A-6a finds a $\hat{\beta}$ coefficient for recall that is about two-thirds of the size of the exit to new job coefficient. However, the figure illustrates that this effect is again short-lived, implying that even if this delayed differential recall response is causal, it has a small effect on aggregate employment patterns.

The difference-in-difference evidence that the \$300 had tiny effects on recall is even more clear cut. Figure A-17f shows that there is little difference in recalls between workers with above- and below-median replacement rates. Table A-6b re-estimates equation (4) and reports a coefficient (0.002) that is one-tenth of the already small effect on new job-finding and economically indistinguishable from

²²Many theoretical papers on UI (e.g., Hagedorn et al. 2013 and Landais, Michaillat, and Saez 2018b) argue that the micro disincentive effect is insufficient to determine optimal benefit levels; one must also know the macro disincentive effect.

²³First, it captures the response of new vacancies to more generous benefits. More generous benefits could decrease vacancy creation by reducing match surplus (Hagedorn et al. 2013) or increase it by increasing aggregate demand (Kekre Forthcoming). Second, it captures the “rat-race” effects in Michaillat (2012), where discouraging one worker from taking a job may simply lead to another worker taking the job instead.

zero.

Overall, the supplement may have changed the timing of some recalls, but there is no evidence of *substantial* recall effects which would change the conclusions we describe in Section 5.2 that effects of supplements on employment were small. However, alternative data and/or research designs are needed to precisely quantify the effect of the \$600 supplement on recalls.

5 Model

Our reduced-form empirical results tell us how job-finding and spending responded to supplement changes. However, the causal responses to policy *changes* could potentially differ from the *overall* causal effects of these policies if the effects of supplements on job-finding are not constant over time. For example, if jobseekers anticipate that the supplements are going to expire and start searching more as the expiration date approaches, then the change in job-finding observed at the date of expiration will understate the effect of supplements in earlier periods. These types of dynamic forces would violate the assumption of constant per-week effects discussed in Section 4.3. Similarly, spending effects of supplements might vary over the pandemic as households accumulate liquidity.

We are more interested in overall effects of the supplements than in their effects at the moment in time that they changed. Thus, in order to translate the latter causal effects, which we can directly measure in the data, into the former causal effects that we most want to know, we develop a dynamic structural model of household job search and spending. We discipline this model’s key features using our causal empirical estimates and then use this model to quantify the overall effects of the supplements. We then use our model to understand the underlying mechanisms behind these effects, what insights generalize beyond the pandemic, and whether alternative policies would have had different effects.

5.1 Model Description

Our model combines an incomplete markets consumption-savings problem with a model of costly job search. Each of these elements is intentionally standard because part of the goal of the model is to understand what a standard model calibrated to pre-pandemic evidence predicts about the effect of large supplements on spending and job-finding. We describe the main elements here and provide additional details in Appendix F.1. Time is discrete and a model period is one month. Households choose consumption c , savings a with return r and a no-borrowing constraint $a \geq 0$, and job search effort when unemployed to maximize expected discounted utility $E \sum_{t=0}^{\infty} \beta^t [U(c) - \psi(\text{search})]$.²⁴

When employed, household i has a constant wage w_i . We allow the wage to differ across households to capture replacement rate heterogeneity in the data, but w_i is constant over time for household i . Employed households separate from jobs and become unemployed with constant probability π . When unemployed, a household finds a job and becomes re-employed at wage w_i with probability $f_{i,t} = \text{recall}_t + \text{search}_{i,t}$ where recall_t is a common recall rate and $\text{search}_{i,t} \in [0, 1 - \text{recall}_t]$ is household i ’s endogenous choice of search effort. Search effort induces disutility $\psi(\text{search}_t)$ while job-finding through recall requires no search effort or disutility. As in Katz (1986), we make the assumption that the probability of recall is exogenous and that the household accepts recall when it

²⁴Two-asset models like Kaplan and Violante (2014) help explain high MPCs for households with high *wealth* but low liquidity. They will not explain the high MPC for high *liquidity* households we see in our data. Thus, we use a one asset model for simplicity.

is offered. This assumption is motivated by the small empirical effects of supplements on recall that we find in Section 4.4, as well as by institutional rules.²⁵

Households who are unemployed during the pandemic are eligible for 12 months of regular unemployment benefits which are proportional to w_i , plus an additional flat weekly supplement of \$300 or \$600 in the periods when supplements are in place. In addition, all unemployed households receive some additional secondary income proportional to w_i even if they are not currently receiving unemployment benefits to capture the fact that income of the unemployed in the data is typically greater than zero even when not receiving benefits.²⁶ Finally, so that we can compare model results to those in our empirical waiting for benefits research design, we assume that some households face delays in receiving benefits but then later receive back pay once these benefits start. See Appendix F for additional discussion of the details of this benefit process and household expectations when facing unemployment insurance delays.

How do we model the evolution of supplements and the effects of the pandemic itself? Beginning from an initial steady state, the economy is hit by a sequence of aggregate policy changes that capture changes in UI policy over the pandemic: weekly supplements of \$600 are added to unemployment benefits from April-July 2020 and then \$300 supplements are added from January-August 2021. Our model accounts for the effects of the pandemic itself in two ways. First, we focus on the evolution of unemployed households relative to employed households in both the model and the data. This essentially removes pandemic effects that affect all households equally from both the model and data. Second, we introduce a one-month discount factor shock to all households in April 2020, calibrated to match the observed decline in spending for employed households during the pandemic, and we include one-time unanticipated transfers to replicate stimulus checks in April 2020 and January 2021 as well as LWA payments in September 2020. Although they add realism, none of these pandemic features have important implications for our conclusions. Figure F-1 helps illustrate the basic environment by showing the path of income for the unemployed relative to employed in the model and data for a newly unemployed worker with and without a benefit delay in our calibrated model.

In addition to the actual evolution of UI policy just described, we must specify household expectations about these policies. We assume that the initial start of the \$600 supplements is unanticipated. Once the \$600 supplement starts, households know for sure that it will last *at least* through July, since this duration was implemented in the initial legislation. However, expectations about renewal are more difficult to discipline. In our main results, we consider two specifications for expectations about renewal. In the first “perfect foresight” version, households correctly expect that supplements will revert from \$600 to \$0 in August. In the alternative “surprise expiration” specification, households instead expect that the \$600 supplement will continue through the end of their benefit spell, and are then surprised in August when it instead expires.²⁷ Once supplements expire in August, households expect them to remain at \$0 forever. For the \$300 supplements in January, we again consider one specification in which the start is a surprise and another in which it is anticipated. Once the \$300 supplement begins, households correctly anticipate that the supplements will expire in September

²⁵The institutional motivation comes from the eligibility requirement that UI recipients accept any offer of “suitable work”. Employers of a worker who refuses a recall offer have a financial incentive to challenge that worker’s eligibility for UI, since continued benefits claims translate into higher future payroll taxes for the employer.

²⁶One common source of such income is spousal income.

²⁷We have also explored versions of the model in which households expect supplement renewal with some probability $\in (0, 1)$ rather than the extremes of 0 or 1. This model is modestly more complicated to solve, but unsurprisingly it generates results that are between the two exercises we report.

2021. Finally, we assume that households anticipate a constant recall rate throughout the pandemic, although results are similar if we instead assume perfect foresight over the actual recall rate in each specific month.

5.1.1 Calibration

We now briefly describe the calibration. See Appendix F.1 for more detail on these choices. The model is monthly and we set the implied annual interest rate $r = .04$. We assume that the utility function is given by $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$ and set $\gamma = 2$. Following Krueger, Mitman, and Perri (2016) we set the exogenous separation probability $\pi = 0.028$. We set the expected recall rate constant at its average value of 0.08 but then use its actual evolution over the pandemic where relevant. We calibrate the benefits process to match the observed household income series for the waiting and receiving UI groups over the pandemic. We solve the model for five different w_i groups and we choose the variation across groups in w_i to match household income in five quintiles of the replacement rate in JPMCI data.

We assume that $\psi(\text{search}_t) = k_0 \frac{(\text{search}_t)^{1+\phi}}{1+\phi} + k_1$ and pick the search cost parameters in one of two ways. In a “pre-pandemic” calibration, we calibrate search costs to generate a pre-pandemic monthly new job-finding rate of 0.28 and an elasticity of average unemployment duration to a small six month change in benefits of 0.5. This is the median estimate from a recent meta-analysis by Schmieder and von Wachter (2016). In our “best-fit” calibration, we instead calibrate search costs to target the time-series of new job-finding over the course of the pandemic.

We similarly calibrate the discount factor in one of two ways. In the pre-pandemic calibration, we set $\beta = .99$ monthly (an 11 percent annual discount rate) to generate a 3-month MPC of 0.25 in response to a \$500 stimulus check sent to all households. Kaplan and Violante (Forthcoming) argues that macro models should target a value for this MPC of 0.15-0.25.²⁸ Havranek and Sokolova (2020) and Orchard, Ramey, and Wieland (2022) argue that even this empirical range may be overstated due to publication bias and issues with two-way fixed effect identification strategies. Our choice of an MPC target at the upper end of estimates in the data is therefore conservative, because even targeting an MPC at the upper end of these estimates generates spending responses to unemployment benefits that are too small. In our alternative best-fit calibration, we instead pick $\beta = 0.98$ (a 22 percent annual discount rate) to target the 1-month MPC of 0.43 in our waiting for benefit receipt design.


5.1.2 Model Fit Comparisons

What features of the model are necessary to match the empirical job-finding and spending patterns we observe during the pandemic? We begin with a discussion of the response to the \$600 weekly supplements paid from April through July 2020 and then discuss the behavior of the model in response to the \$300 payments started in January 2021.

Figures 6a and 6b show how job-finding and spending compare to the data in the pre-pandemic model calibration. This model is calibrated to match a pre-pandemic unemployment duration elasticity of 0.5 and a quarterly MPC of 0.25 to \$500 stimulus checks distributed to all households. We also

²⁸This range applies to the MPC on non-durable goods and services. We interpret our spending measure as most closely capturing the non-durable spending response as measured in the prior literature. For example, Parker et al. (2013) finds that nearly the entirety of the total spending response beyond non-durables is due to the purchase of vehicles, which are not included in our spending measure.

Figure 6: Job-Finding and Spending Responses to \$600 Supplement: Models vs. Data



Notes: This figure shows time-series of various models vs. the data in response to the \$600 supplements. The left panels show the monthly new job-finding rate and the right panels show the spending of unemployed (relative to employed). The pre-pandemic model calibrates to pre-pandemic evidence on duration elasticities and MPCs out of stimulus checks. The pandemic models all calibrate search costs to match the level of job-finding in the data and then the three models vary in the expectations about the renewal of the \$600 and in what MPC is targeted.

assume that households have correct expectations in advance that the \$600 supplements will expire in August 2020 (“perfect foresight”). Comparing the model lines in purple to the data lines in dashed blue, it is clear that the model misses the data dramatically. The job-finding rate responds too much and spending responds too little to the \$600 benefits in this pre-pandemic model.

Comparing results from this pre-pandemic calibration to the data already shows a sense in which empirical job-finding responses to supplements are unexpectedly small and spending responses are unexpectedly large. A model calibrated to prior evidence on the response of job-finding to increased benefits predicts that almost no workers take a new job while the generous supplements are in place, and that job-finding rises sharply when they expire. But the data shows much more muted patterns. In contrast, the opposite is true with respect to spending. A model calibrated to prior evidence on

spending responses to income changes implies spending swings which are much more muted than the data.

Panels 6c and 6d progressively introduce three changes to this model to better fit the data.²⁹ We first change our calibration strategy for search costs: instead of targeting pre-pandemic estimates of the duration elasticity, we target the time-series of job-finding rates actually observed in the pandemic. That is, we pick search costs to try to match the monthly job-finding rate in the data from April 2020 to February 2021 as closely as possible.³⁰ Because the *level* of the job-finding rate even with no supplements is so much lower than during normal times, and the *change* in the job-finding rate in response to large changes in supplements is so modest, the calibration that best matches these patterns implies that job search was much more costly and less responsive to monetary incentives during the pandemic. We discuss the potential sources of this change in behavior in Section 6.1.

The model shown in red makes this change to search costs but continues to assume perfect foresight and leaves the discount factor the same as the pre-pandemic model. This model is unsurprisingly a better fit for the job-finding rate than the pre-pandemic model, but it implies that the job-finding rate evolves too smoothly relative to the data: when households anticipate that the \$600 will expire in August, they begin searching more in the months before that. In addition, this perfect foresight version of the model does not generate a sharp spending decline at the expiration of the \$600.

The second change is to household expectations. The sharp spending response to the expiration of the \$600 in the data is a key object for disciplining the information structure in the model and helps demonstrate the benefit of our empirical analysis, which studies multiple benefit changes of various sizes and signs. In standard consumption models, it is difficult to get sharp drops in spending in response to predictable declines in income. It is possible to generate such drops by increasing the degree of present bias, as in Ganong and Noel (2019), but in this model households will still respond much more strongly to unanticipated increases in income than to anticipated declines in income.³¹ This means that a model with present-biased households in which the start of \$600 supplements is a surprise and the end of the \$600 is anticipated will generate asymmetric spending responses to the start and end of supplements, in contrast to the patterns in the data. Thus, present-biased preferences *without* a departure from the assumption of perfect foresight are not enough to explain the large and symmetric changes in spending we observe at the onset and expiration of supplements.³²

What can generate a roughly symmetric increase and decrease in spending when households are not hand-to-mouth? If both the increase and the subsequent decrease are a surprise, then it helps resolve this tension. Thus, rather than assuming perfect foresight, in the second change to the model we assume that households are surprised by the expiration of the \$600.³³ The yellow line shows this model, again with search costs targeted to match the observed job-finding series. This model is a much better fit to the data. It now implies a sharp increase in job-finding as well as a sharp decline

²⁹Note that we zoom in the y-axis in 6c relative to 6a to better highlight model deviations which are obscured by the large miss in the pre-pandemic model.

³⁰We assume constant search costs over the pandemic, so we have 3 parameters and 11 targets for monthly job-finding.

³¹Responses only become symmetric in the extreme situation where households are fully hand-to-mouth. However, this then implies an MPC of 1, which is much larger than the data. As soon as one departs from this extreme situation, spending responses to unanticipated increases and anticipated decreases become asymmetric.

³²This should not be construed as evidence *against* present-bias. Indeed, present-biased preferences *in conjunction* with a surprise change in benefits *would* be helpful in explaining the size of the spending changes. As we describe below, the third necessary change to the model is some force pushing up MPCs of unemployed households. Heterogeneity in the degree of present-bias (rather than simply in impatience) would be an alternative way to generate these patterns.

³³A more complicated model in which households think supplements will continue at \$400 with 50% probability and at \$600 with 50% probability and then unexpectedly decline to \$0 produces similar results.

in spending when the \$600 expires, since households are now surprised by this income decline.

How do we interpret these alternative household expectations? The key condition needed to better fit the data is that households respond strongly today to policy changes today but respond little today to policy changes in the future. There are at least two ways of rationalizing this. The most straightforward rationalization is simply that these changes were actually policy surprises, even for those paying close attention to policy making. For example, before the \$600 supplements expired, the main plans proposed by both Democrats and Republicans called for their extension. The debate was primarily about whether they should do so at \$600 or \$400. If the “compromise” supplement level of \$0 was a surprise to policymakers, it seems reasonable that it was also a surprise to unemployed households. There was similar uncertainty around the start of the \$300 supplements in January, with President Trump originally announcing that he would not sign the policy into law but then doing so anyway. However, status quo bias in belief formation could generate the same expectations: if naive households simply assume that current benefit levels will continue unchanged and only update expectations when benefits actually change, this would yield identical behavior in our best-fit model.

Nevertheless, even when households are surprised by supplement changes, spending responses in this calibration are still too small. For example, Figure 6d shows that both the increase in spending while the supplements are in effect in the summer of 2020 and the decrease in spending when they expire in August are about twice as large in the data as in this version of the model (shown in yellow). Furthermore, the MPC in a simulation of the empirical waiting for benefits design in this model is 0.30 relative to the empirical value of 0.43. Recall that the discount factor in this version of the model is calibrated to match a quarterly MPC of 0.25 out of a \$500 stimulus check sent to everyone.

We therefore need some force in the model that causes unemployed households’ spending to be more sensitive to income changes than the general population that is studied in the prior literature on stimulus check MPCs. Moreover, this larger spending sensitivity of the unemployed cannot be driven simply by their temporary low liquidity and income, because the model already accounts for the effect of these temporary economic circumstances on MPCs. Instead, it must reflect some persistent characteristic of the unemployed that makes them different from the general population.

Hence, in the third and final change to the model we recalibrate the discount factor to target the reduced form MPC of 0.43 from the waiting for benefits design in our data. This calibration implies that the unemployed households in our sample have discount rates twice as high as those of the general population needed to fit stimulus check MPCs (22% annually versus 11%).³⁴ We discuss further justification for, and implications of, this type of permanent heterogeneity in detail in Section 6.2.

Figures 6c and 6d show that this “best fit” model closely matches time-series patterns for both job-finding and spending.³⁵ Figure F-3 shows that similar conclusions extend to the \$300 supplements in January. In particular, the best fit calibration for the \$600 supplements also generates a good fit to time-series patterns around the start of the \$300 supplements. An alternative model in which the \$300 supplement is anticipated in advance misses both job-finding and spending patterns, and an alternative model calibrated to a 0.25 stimulus check MPC implies spending responses which are too small.

³⁴We note that it is difficult to distinguish empirically whether the type of households likely to become unemployed have permanently higher discount rates or whether unemployment exposure itself changes discount factors, but this distinction makes little practical difference for our conclusion that the spending of unemployed households is very sensitive to transfers and remains so even when they have high liquidity.

³⁵The model overstates spending in March and April 2020 because of time-aggregation issues: in the data there is 2-3 week gap between job loss and the start of benefits even for households with no delay. These high frequency patterns are not captured in our baseline monthly model in order to keep it parsimonious, but Figure F-2 shows that our conclusions are unchanged in a more complicated model that is fit to high frequency patterns immediately after job loss.

Furthermore, an additional untargeted job-finding moment is consistent with the best-fit model. Although we pick search costs in the best fit model to target the time-series for the average job-finding rate, Figure F-4 shows that this best-fit model also aligns well with the results from the difference-in-difference research design for job-finding. This means that calibrating to instead match that evidence yields similar conclusions. This comparison also provides additional validation for the linearity assumption in this difference-in-difference design.

Overall, the key elements necessary to jointly fit empirical job-finding and spending patterns are households who act as if they are surprised by changes in benefits, a persistent characteristic that pushes up the spending response to income changes of unemployed households above those of typical households, and more costly job search during the pandemic. With these three elements, the model can account for the main patterns observed in the data.

5.2 Magnitudes

With this best fit model in hand, we can now answer the question of how unemployment benefit supplements affected job-finding and spending during the pandemic and what explains these patterns.

Table 3: Effects of Supplements on Unemployment Duration and Spending

	\$600 supplement	\$300 supplement
Unemployment Duration Elasticity		
Best fit model	0.07	0.10
Statistical model: time-series	0.06	0.10
Statistical model: cross-section	0.11	0.21
MPCs (Best Fit Model)		
1-month MPC out of 1st month of supplements	0.30	0.31
3-month MPC out of 1st 3 months of supplements	0.39	0.34
Total MPC through month supplement ends	0.41	0.45
Total MPC through 3 months after supplement ends	0.53	0.58

Notes: The first panel shows how supplements affect unemployment duration in our best fit model, which includes dynamics, as well as in alternative statistical models which assume constant per-week effects when benefits are in place. The second panel shows how supplements affect spending over various horizons.


The first section of Table 3 reports duration elasticities to both the \$600 and \$300 supplements. The statistical model results are calculated under the assumption that the per-week effects of supplements estimated in Section 4 are constant for every week supplements are in place and zero after they expire. The best fit model elasticities allow for dynamic time-varying effects of supplements on the job-finding rate.³⁶

Figure 7 compares all of these elasticity estimates to pre-pandemic estimates from the literature review of Schmieder and von Wachter (2016). Three observations are of note. First, Figure 7 shows that our duration elasticity estimates are quite low relative to pre-pandemic estimates. Second, Table 3 shows that duration elasticities in the best fit model, which accounts for dynamics, are similar to the simpler statistical calculations which assume constant per-week effects. Third, Table 3 shows that

³⁶See Appendix E for more details on duration elasticity calculations.

duration elasticities are moderately larger for the \$300 supplement than for the \$600 supplement. We return to all of these observations in detail in Section 6.1.

Figure 7: Pandemic Elasticity Estimates Compared to Prior Literature



Notes: The pre-pandemic estimates are summarized from the literature by Schmieder and von Wachter (2016).

The next section of Table 3 reports spending effects over various horizons in the best fit model. Defining Δc_t and Δy_t as the difference in consumption and income for a household with and without supplements in month t , we compute $\frac{\sum_{t=1}^T \Delta c_t}{\sum_{t=1}^T \Delta y_t}$ for various values of T . In particular, for each supplement we show results for T equal to 1 month, 3 months, the length of the supplement period, and the length of the supplement period plus 3 months. It is these longer horizon MPCs which most directly answer policy questions about the overall effects of supplements on spending of the unemployed, but these cannot be reliably measured in the data.

We note that because the supplements themselves are persistent, this MPC is conceptually distinct from a standard MPC out of a temporary income shock at different horizons. Unlike in the case of temporary income shocks, both the numerator *and* the denominator of this MPC grow with T as long as supplements are in place. This means there is no mechanical reason MPCs need grow with T , as long as T is less than the supplement length. This is important because it means that the one month MPCs we estimate in the data are not necessarily lower bounds on MPCs over the entire supplement period.

We find one month MPCs of 0.30 and 0.31 to the \$600 and \$300 supplements respectively. We discuss the determinants of these MPCs in more detail in Section 6.2, but for now note that these one-month MPCs are large relative to *quarterly* MPCs out of stimulus checks of 0.15-0.25 (see the discussion in Section 5.1.1). The share of all supplements spent in the period while the supplements are still in place is even higher, with an MPC of 0.41 and 0.45 to the \$600 and \$300 supplements. By three months after supplement expiration, MPCs rise to 0.53 and 0.58.

These MPCs imply sizeable effects in terms of dollar spending. Of the \$2,400 in supplements received in April 2020, \$700 is spent by the end of that month. A worker unemployed from April through July 2020 receives \$10,200 in \$600 weekly supplements. \$4,200 is spent by the end of July


and \$5,400 is spent by the end of October.

Another way to gauge the magnitude of these effects is to compare them to overall aggregate fluctuations in employment and spending during the pandemic. In Figure 8 we use the best fit model scaled by the total number of workers receiving unemployment benefits in the data to generate a simple partial equilibrium counterfactual for aggregate employment and spending over the course of the pandemic had there been no \$300 or \$600 supplements.³⁷

Figure 8 shows that although the supplements had some negative effects on employment, these effects were small relative to overall employment changes during the pandemic. The \$600 weekly supplements from April 2020 to July 2020 reduced employment by an average of 0.6% while the \$300 supplements reduced employment by an average of 0.4%. Overall this amounts to only around 5% of the overall employment gap generated by the pandemic, so these supplements played a small role in explaining aggregate employment dynamics during this time period.

The effects of supplements on spending were three to five times larger than their effects on employment. Specifically, the \$600 and \$300 supplements boosted aggregate spending by an average of 2.9% and 1.3%, respectively. This means that supplements helped to close a large fraction of the aggregate spending gap during the pandemic. Before they expired at the end of July 2020, the supplements were responsible for 25% of the spending recovery during the pandemic.

Figure 8: Aggregate Implications



Notes: This figure shows aggregate employment and spending dynamics implied by our best-fit model. We first estimate individual effects and then scale these to partial equilibrium aggregate effects using the number of workers receiving benefits at each date. The left panel compares to total employment from the BLS establishment survey while the right panel compares to PCE spending.

While we focus on average effects since this is what is relevant for thinking about the overall effects of the policy, it is important to note that the effective treatment size was larger for low wage workers than for high wage workers. Thus, we would expect supplements to have larger effects on job-finding for low wage than for high wage workers. Indeed, this is exactly the pattern that we see in our difference-in-difference design in Section 4.2. Since our model includes wage heterogeneity, we can also use it to assess these distributional effects over the wage distribution.

Appendix Table F-1 shows that the employment effects for the lowest wage quintile of unemployed workers were about twice as large as the employment effects for the highest wage quintile. Thus, even

³⁷Since the payroll survey excludes self-employed workers we reduce the counts of benefit recipients when scaling employment effects by the number of self-employed PUA UI recipients in DOL data.

though supplements played a small role in employment dynamics for the typical unemployed worker, they played a somewhat more important role for the lowest wage workers amongst the unemployed. This heterogeneity might help explain anecdotal evidence that some businesses faced difficulty hiring at very low wages even if the typical unemployed worker’s search behavior responded little to supplements, so that aggregate effects were small.

5.3 Limited Dynamics

One of the main motivations for the model is to assess whether anticipatory search behavior or liquidity accumulation lead to non-constant search dynamics and thus misleading conclusions from reduced-form exercises which assume constant effects. In this section, we illustrate that neither anticipatory search behavior nor liquidity appear to have substantial impacts on the job-finding rate during this time period, and thus the reduced-form analysis approximately captures the full effect of supplements.

5.3.1 Anticipation Effects

The fact that there are limited dynamics from anticipatory search follows from the information structure necessary to fit the time-series patterns in the data. Figure 6 shows that the model only fits spending and job-finding time-series if households are surprised by policy changes. The lack of anticipatory dynamics is thus explained simply by the fact that households do not adjust behavior in advance of an unexpected policy change.


A second, complementary source of information further bolsters the case that anticipation induces little dynamic bias. Specifically, the weekly difference-in-difference coefficients prior to expiration shown in Figure A-18 are stable from week to week. We can compare these coefficients to effects in the model under the same two assumptions about expectations—perfect foresight and surprise expiration—that we study in Section 5.1.2. Figure 9 shows that the perfect foresight model is at odds with these weekly difference-in-difference regression estimates: the perfect foresight model implies that the effect of supplements on job-finding two months prior to expiration is seven times larger than the effect in the week prior to expiration. The surprise expiration model instead implies stable effects in the weeks leading up to expiration, consistent with the data.

While we focus on a particular application, this analysis illustrates a broader methodological point: an economic model can be useful for addressing questions about anticipation bias in a difference-in-difference research design. Anticipation bias is a common concern for estimates from these designs, so a standard identification assumption is that the treated group does not modify its behavior in advance of treatment. This is often called the “no anticipation” assumption (Roth et al. 2022).³⁸ This is typically justified either by using the institutional record or by a statistical test for differential pre-trends. The latter is an omnibus test which captures both anticipation effects and other confounding differences between treatment and control groups. However, with a statistical test alone, it is difficult to know how big an economically plausible anticipation effect would be if the no anticipation assumption were violated.

We suggest that researchers who are concerned about bias from anticipation—either because of

³⁸The no anticipation assumption can be relaxed in some cases. If a policy is implemented at time 0 and a researcher has *a priori* knowledge that anticipation effects are only plausible in the G periods before implementation, then they can use periods from G to 0 to measure anticipation effects and periods before G to diagnose pre-trend violations. However, in some cases, the time horizon for anticipation effects is not clear *a priori* or the researcher does not have data prior to G .

Figure 9: Dynamic Treatment Effects Under Different Informational Assumptions



Notes: This figure compares dynamic treatment effects in the model to the difference-in-difference (DiD) estimates in the weeks prior to the expiration of the \$600 supplement. The model series in red and green use the same informational and preference assumptions as Figure 6. The data series uses empirical estimates from Figure A-18. The dependent variable is the effect of the supplement or the DiD estimate normalized by its own value in the final week before expiration.

a pre-trend or on *a priori* grounds—use an economic model to quantify anticipation effects. In a setting where the researcher *does* observe a statistically significant pre-trend, they can ask whether the magnitude is consistent with anticipation effects. If it is, they would use that same model to quantify the total effect of the policy, which is likely greater than the observed change in the outcome at the time of policy implementation. Alternatively, if the pre-trend cannot be explained by anticipation bias, the researcher might choose to abandon that research design.

Economic models of anticipation are also useful in the scenario from Figure 9 where there is no statistically significant pre-trend. In this figure, we conclude that meaningful anticipatory behavior can be ruled out by the confidence intervals for the point estimates. However, it is possible in other applications that the confidence interval for the pre-trends will contain both the predictions from a model with anticipation and a model without anticipation. Because the research design is not powered to detect anticipation, the researcher may want to account for this additional source of uncertainty in their estimates. Thus, our approach using an economic model complements statistical methods for diagnosing pre-trend violations such as those proposed by Freyaldenhoven, Hansen, and Shapiro (2019).

5.3.2 Liquidity Effects

Another potential concern about dynamics arises from liquidity effects on job search. Prior research finds that unemployment benefits reduce job search in part by relaxing liquidity constraints (Card, Chetty, and Weber 2007, Chetty 2008). Figure 1 shows that the supplements were associated with a large increase in liquidity for the unemployed. If the job-finding rate remains depressed *after* the \$600 supplement expires because of this elevated liquidity, then the research designs in Section 4 will understate the full effect of the supplements. We use two complementary approaches to address this concern.

First, we show that this bias is small in the model. We begin by noting that the model includes borrowing constraints, and reassuringly it replicates the finding from the prior literature that liquidity

has important effects on job search. Specifically, it replicates untargeted results from Card, Chetty, and Weber (2007) that two months of severance pay reduces the subsequent log job-finding hazard by 0.076-0.109. Performing this same exercise in our model delivers a value of 0.079, so our model yields credible liquidity effects. Nevertheless, the model implies that bias induced by liquidity accumulation is tiny. For example, Figure F-5 shows that liquidity accumulation indeed leads to a depressed job-finding rate after expiration but that this effect is negligible: supplements reduce the monthly job finding rate by 3.3% just prior to expiration (capturing both the effect of incentives and the effect of liquidity) and by only 0.13% after expiration (capturing only the effect of liquidity).³⁹ The effect of liquidity accumulation is small in the model because liquidity constraints only affect search when they bind. Liquidity is already elevated by April 2020 when supplements start and then grows further from there, but this additional liquidity accumulation has little additional effect on search.

Second, we study empirically how the estimates in Section 4 vary with liquidity.⁴⁰ Using a triple-difference design, Tables A-7 and A-8 show that a one-standard deviation increase in liquidity is associated with a decline in the disincentive effect of the \$600 from -.0164 to -.0142, and with a decline in the disincentive effect of the \$300 from -.0191 to -.0184. A simple extrapolation from this cross-sectional heterogeneity in treatment effects to the time-series of liquidity for the unemployed during the pandemic implies that liquidity accumulation reduces the disincentive coefficient $\hat{\beta}$ by only 5%, from -.017 in April 2020 to -.0164 in July 2020.⁴¹ Figure A-19 shows similar conclusions visually. This figure recomputes the binscatter difference-in-difference regressions from Figure 5 but stratified by terciles of household checking account balances. The highest liquidity terciles exhibit more muted relationships between benefit changes and exit rates but this effect is quantitatively small.

6 Understanding Mechanisms

What forces drive the small employment effects and large spending effects to the supplements implemented during the pandemic? Answering this question is important for understanding which findings are more likely to generalize to other environments. We begin with a discussion of employment effects and then turn to a discussion of spending effects.

6.1 Understanding Employment Magnitudes

Three forces explain why the unemployment duration elasticity we estimate is lower than estimates in the prior literature. First, there is a mechanical effect due to the fact that the supplements we study are temporary. This mechanical effect is amplified by the fact that the supplements are implemented in a labor market with an already depressed job-finding rate. Second, the large share of recalls further dampens the effect of the supplements. Third, the per-week behavioral response to the supplements while they are in effect is also lower than in prior times. This sub-section discusses each of these three forces in turn.

³⁹It is worth emphasizing that the model-based duration elasticities in Table 3 include these effects of liquidity accumulation and so avoid even this small bias.

⁴⁰UI supplements raise balances by more for households with lower values of pre-supplement balances than those with high pre-supplement balances. We focus on pre-supplement balances to eliminate two endogeneity issues with the change in balances over time: 1) staying on unemployment for longer mechanically leads to more supplements and thus larger balance increases and 2) households who expect to have trouble finding a job are likely to save more.

⁴¹Mean checking account balances increase by \$1,388 for unemployed relative to employed households from April to July 2020, and the standard deviation of balances in our regressions is \$4,897. Thus $-.0164 - (1388/4897) * .0022 = -.017$.

To understand the mechanical force, it is useful to note that the duration elasticity to increased benefits is a function of (a) how much the job-finding rate changes while a benefit increase is in place (known as the per-week “hazard” elasticity) and (b) how long this increase in benefits lasts. Holding constant the per-week hazard elasticity, it is mechanical that a supplement that is in place for a shorter amount of time will have a smaller effect on average unemployment duration.

Intuitively, employment distortions will be small from supplements that only distort behavior for a short amount of time. For example, even if the hazard elasticity were so large that job-finding dropped to zero when supplements were in place, the effect on unemployment duration would be negligible if the benefit increase (and therefore the change in the job-finding hazard) lasted only for a single day. Thus, there is a wedge between the per-week hazard elasticity and the full duration elasticity when benefit increases are temporary.

The mechanical wedge between duration and hazard elasticities disappears as benefit changes become longer lasting. For example, if the hazard elasticity is one, then permanently doubling benefits will double the average duration of unemployment. Thus, the duration elasticity will also be one. The past literature typically studies the effect of relatively long-lived benefit changes where the mechanical wedge is small and there is little distinction between the hazard and duration elasticities. Indeed, some papers actually estimate the hazard elasticity and then interpret this estimate as the duration elasticity. In our context, this distinction matters.

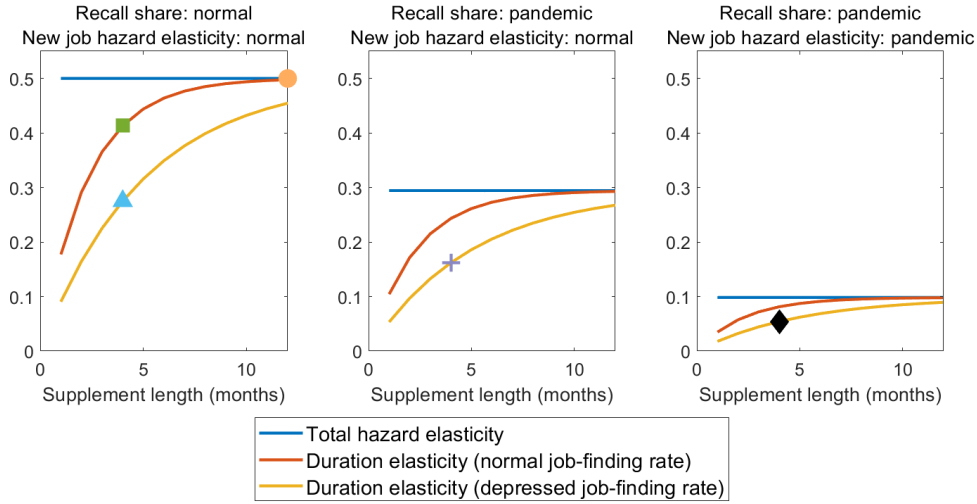
Furthermore, the mechanical wedge between the duration elasticity and the hazard elasticity grows when the level of baseline job-finding rate is depressed, as it was during the pandemic. For example, reducing the job-finding rate from 50% to 25% for one month will lead to larger growth in the duration of unemployment than will reducing the job-finding rate from 5% to 2.5% for that month.

The left panel of Figure 10 quantifies the importance of the mechanical channel. This panel shows hazard and duration elasticities when we choose a total (i.e., recall + new) job-finding hazard elasticity of 0.5, to match pre-pandemic estimates from the past literature.⁴² Given this hazard in blue, the red line in this figure shows the implied duration elasticity for supplements of various lengths if the baseline total job-finding rate is set to a “normal” level equal to its pre-pandemic average. The orange circle in this figure corresponds to the pre-pandemic duration elasticity of 0.5 in response to long-lived benefit changes, which we take from the past literature. For shorter supplement lengths, a mechanical wedge between the duration and hazard elasticities emerges since a change in the hazard translates less than one-for-one into a change in duration if that change is short-lived. The green square at four months corresponds to the length of the \$600 supplement and shows that the short length of the supplement alone can explain a reduction in the duration elasticity from 0.5 to about 0.4.

The yellow line in this figure repeats the same exercise but with a level of the baseline job-finding rate chosen to match the depressed job-finding rate during the pandemic (in fall 2020 when no supplements were in place). When the baseline job-finding rate is low, the hazard elasticity converges more slowly to the duration elasticity and so the mechanical wedge between the two is larger for any supplement length. The blue triangle shows that the duration elasticity in response to a four month supplement is then around 0.25, even though the behavioral hazard elasticity remains at 0.5. This means that the combination of short supplements and a depressed job-finding rate explains about

⁴²To simplify intuition, we analyze these effects under the assumption of a constant hazard elasticity, both over time and as a function of supplement length. However, in models of optimal search, the hazard elasticity itself would decline as the length of supplements declines, amplifying the conclusion that short supplements lead to lower duration elasticities.

Figure 10: Forces for the Low Unemployment Duration Elasticity



Notes: The figure computes the duration elasticity to supplements of different lengths. Within each panel we compute the duration elasticity under a normal pre-pandemic level of the total job-finding rate as well as for a depressed job-finding rate. Within a panel we do not change the recall share as we move from the normal to the depressed scenario, meaning both recall and new job-finding rates decline proportionately. Moving from the left panel to the middle panel, we lower the new job hazard elasticity from its pre-pandemic value to its pandemic value, and we leave the recall share at its pre-pandemic value. Moving from the middle panel to the right panel we raise the recall share from its pre-pandemic to its pandemic value. This means that the yellow line in the right panel corresponds to the full set of pandemic forces while the red line in the left panel corresponds to the full set of normal conditions.

half of the difference between the pre-pandemic duration elasticity of 0.5 and the pandemic duration elasticity of 0.07. This shows that employment distortions induced by unemployment benefits hinge upon interactions between benefit length and the broader labor market environment.

In addition to this mechanical effect, there are two pandemic-specific effects which are also important for generating the low duration elasticity. First, the recall share of total exits is elevated during the pandemic. We assume that recalls are insensitive to changes in benefit levels based on both institutional constraints as well as empirical evidence in Section 4.4. This is important because if recalls are insensitive to benefits, then raising the recall share of exits while holding constant the hazard elasticity of *new* job-finding leads to a decline in the *total* job-finding hazard elasticity. Intuitively, if a larger fraction of exits from unemployment are insensitive to benefit increases, then these increases will generate smaller employment distortions. The purple cross in the middle panel of Figure 10 shows that the elevated recall share in the pandemic combined with the depressed pandemic level of job-finding implies a duration elasticity of about 0.18 in response to a four month supplement.

Second, the new job-finding effects that we estimate in Table 2 imply a lower per-week behavioral hazard elasticity than implied by pre-pandemic estimates. The right panel of Figure 10 recomputes the results with this lower hazard elasticity and shows that this reduced sensitivity of search to benefits during the pandemic lowers the duration elasticity of a four month supplement from 0.18 (the purple plus) to 0.07 (the black diamond). This black diamond corresponds to the \$600 supplement duration elasticity in our best fit model and completes the steps moving us from the pre-pandemic yellow dot to our pandemic estimate. By looking at this final panel on the right hand side, we can also see that a majority of the slightly larger duration elasticity that we estimate in response to the \$300 supplements can be explained by the fact that this supplement is eight months rather than four months.

The above discussion shows that a substantial share of the reduced duration elasticity during the

pandemic is explained by the interaction between short supplements and a low job-finding rate on the one hand, and an increase in the share of exits to recall on the other. However, it also shows that a third force, the behavioral reduction in the hazard elasticity, is also important.

We empirically explore whether any worker-level observable characteristics can help account for the reduction in the hazard elasticity but ultimately find little evidence of heterogeneity. Table A-9 shows the role of observables.⁴³ There is some evidence that age (as a potential proxy for health risk during the pandemic) and the presence of kids (as a proxy for childcare constraints) reduce sensitivity to the \$600 supplements. However, these patterns are reversed in response to the \$300 supplements and what differences exist are small in absolute terms. To explore whether recall expectations might be important, we compute duration elasticities for households who were already unemployed pre-pandemic, and thus unlikely to expect recall. If anything, these low recall propensity households have lower duration elasticities, suggesting that this channel is unimportant. In unreported results, we explored the role of sectoral heterogeneity and found no consistent cross-sector patterns. Finally, we again note the small effects of liquidity discussed in Section 5.3.2.

Since the small behavioral response to supplements does not appear to be driven importantly by any of these obvious channels, we think that the most likely remaining explanation is a shift induced by the pandemic’s effect on working conditions, including exposure to health risk (which may be only imperfectly proxied by age), discomfort from masking, and more challenging customer interactions.

Summarizing the results from this section, the small employment effects of supplements arise because: 1) supplements were short-lived and the job-finding rate was depressed even with no supplements, 2) the recall share was elevated, and 3) there were pandemic-induced reductions in search behavior. While the second and third forces are unlikely to generalize beyond the pandemic, the first effect arises solely from *mechanical* forces and so should generalize: short-lived increases in benefits during recessions, when the job-finding rate is depressed, are likely to induce smaller distortions than permanent increases in benefits which are not tied to economic conditions. We return to this point when we discuss policy counterfactuals in Section 7.

6.2 Understanding Spending Magnitudes

What explains the large spending responses to unemployment supplements? To understand this, it is useful to compare them to more commonly studied one-time stimulus payments. In particular, recall that our pre-pandemic model is calibrated to match a 0.25 quarterly MPC out of a one-time \$500 payment. Relative to these one-time payments, UI supplements differ in their targeting, size, and persistence. Furthermore, these supplements were implemented in a pandemic environment with depressed overall spending and elevated liquidity.

Table 4 shows that each of these forces is important for understanding the spending responses to supplements. Each row varies one element at a time to illustrate the forces shaping spending responses to supplements during the pandemic. The first row shows the one month MPC out of combined regular benefits plus the \$600 supplement (\$2,400 monthly) for an unemployed household who is currently receiving no benefits during the pandemic in our best fit model. This combination of elements replicates the empirical environment for the waiting design in Section 3.2.1. Since our

⁴³The table shows duration elasticity by group, assuming a constant effect of supplements on job-finding which we estimate using a group-specific interrupted time-series approach as in Appendix E. Conclusions were similar when exploring difference-in-difference designs interacting with group-specific observables.

Table 4: Interpreting MPC Differences Between Waiting for Benefits and \$500 Stimulus

	Nature of transfer	Who receives	Calibration	Environment	MPC		
					Horizon	Model	Data
(1)	\$2400 Persistent+Reg UI	Unemp no UI	Waiting	Pandemic	Month	0.43	0.43
(2)	\$2400 Persistent	Unemp w/ reg UI	Waiting	Pandemic	Month	0.30	
(3)	\$2400 Persistent	Unemp w/ reg UI	Waiting	Normal	Month	0.45	
(4)	\$2400 One Time	Unemp w/ reg UI	Waiting	Normal	Month	0.28	
(5)	\$500 One Time	Unemp w/ reg UI	Waiting	Normal	Month	0.45	
(6)	\$500 One Time	Everyone	Waiting	Normal	Month	0.23	
(7)	\$500 One Time	Everyone	\$500	Normal	Month	0.09	
(8)	\$500 One Time	Everyone	\$500	Normal	Quarter	0.25	0.25

Notes: This table compares MPCs across various model specifications and shows empirical counterparts, where available. The “Nature of transfer” column shows the particular transfer for which we compute an MPC. The \$2,400 monthly transfer corresponds to \$600 weekly supplements while \$500 transfers correspond to past stimulus checks. “Who receives” shows which households are receiving that transfer. “Calibration” shows the target calibration used in that model. Models either target the MPC from the waiting design or target a 0.25 quarterly MPC out of \$500 stimulus checks sent to everyone. “Environment” is either a pre-pandemic environment which includes discount factor shocks and stimulus checks or a normal environment which does not. MPCs are calculated primarily at the monthly horizon but we show also a quarterly result for stimulus checks to everyone to ease comparison with empirical estimates. Note that even though we show effects one element at a time, these interactions are non-linear and so this is not an additive decomposition.

best-fit model is calibrated to match this MPC, the model MPC equals the empirical MPC of 0.43 by construction. This MPC mixes the spending responses to regular benefits and supplements. Thus, in row 2 we compute the model MPC to supplements alone, which are targeted at a household already receiving regular unemployment benefits. This MPC of 0.30 corresponds to the supplement MPC in Table 3. We note that this value is essentially the same as the untargeted MPC of 0.30 and 0.26 to the end of the \$600 and start of \$300 supplements estimated in Section 3.

In row 3, we compute this same MPC but in a “normal” economic environment which eliminates pandemic stimulus checks and discount factor shocks. Since turning off these pandemic forces decreases liquidity, the MPC to UI supplements rises to 0.45. Row 4 then shows that the MPC out of a one-time \$2,400 payment falls to 0.28, implying that some of the spending response to supplements comes from the fact that they are persistent rather than transitory transfers. In row 5, we decrease the size of the transfer from \$2,400 to \$500, which corresponds to the stimulus check sizes that we target in pre-pandemic model calibrations. This increases the MPC substantially due to the concavity of the consumption function. Unemployed households’ liquidity constraints and MPCs are relaxed more in response to a large transfer than in response to a small transfer.

Next, we modify who receives the the transfer. Row 6 shows that providing a \$500 transfer to everyone instead of just to unemployed households reduces the MPC from 0.45 to 0.23 since unemployed households have higher MPCs. Finally, row 7 and row 8 show the effect of changing the calibration of the discount factor so that the model hits a quarterly MPC of 0.25 to this \$500 transfer. Row 8 shows that the model hits this quarterly MPC of 0.25 by construction, and row 7 shows that the monthly MPC to this same shock is 0.09 so it can be compared to MPCs from other rows which are calculated at monthly horizons.

Putting all these forces together leads to a one month MPC out of supplements that is more than three times larger than the one month MPC out of \$500 stimulus checks (0.30 vs 0.09).

What are the implications of these spending results for consumption modeling? A large literature has developed models in which liquidity constraints play a key role in generating high MPCs. For

example, modern two-asset precautionary savings models are able to simultaneously match the distribution of wealth and the high MPC out of stimulus checks in the data, because many households have high wealth but little liquidity (see Kaplan and Violante Forthcoming for a review of this literature). A key feature of these models is that the MPC is high only when households have low liquidity.

However, unemployed households do *not* have low liquidity during the pandemic, so the standard channel based on low current liquidity cannot be the only mechanism driving high MPCs out of supplements. In particular, Figure 1 shows that not only do unemployed households have greater than normal liquidity during the pandemic, they have even greater liquidity than pre-pandemic employed households. This is because stimulus checks, general pandemic-driven spending declines, and the large UI supplements themselves all increase liquidity. Indeed, between January and April 2020 the median household who became unemployed during the pandemic moved *up* an entire tercile of the pre-pandemic liquidity distribution (from the 35th percentile to the 65th percentile).

Our evidence that formerly low-liquidity households *still* have large spending responses even when they are in a high-liquidity state suggests that there must be some *persistent* behavioral characteristic that helps explain MPCs. If current liquidity was the only driver of high MPCs, the spending responses to supplements for high-liquidity unemployed households in the pandemic would be smaller than what we see in the data.

Perhaps the clearest way to see visually why some sort of persistent characteristic is needed to match the data is to revisit the spending patterns in Figure 6d. The spending response to supplements is so large that it dramatically *reverses* typical spending patterns during unemployment. Instead of falling, the spending of unemployed households rises, both in absolute terms and relative to the spending of employed households. A model calibrated to match prior evidence of MPCs but without permanent heterogeneity can only explain a fraction of the sustained rise in spending while the \$600 supplement is in place. Households in this environment have high liquidity and know their incomes must eventually fall significantly (either when supplements expire or when they return to employment). Generating a large and sustained spending increase for households in this environment requires a force that pushes up spending sensitivity to current income even when households have high liquidity. We capture this in the model via permanent heterogeneity in discount rates.⁴⁴

The fact that we can observe spending responses for households receiving such a large and sustained sequence of transfers is crucial for disentangling the role of permanent heterogeneity from that of current low liquidity in driving MPCs. There is a well-documented correlation between liquidity and MPCs. However, it is difficult to determine the causal mechanisms underlying this correlation. Permanent heterogeneity affects MPCs by changing both the distribution of wealth and the MPC at a given level of wealth. Thus, it is difficult in general to determine whether the empirical correlation between liquidity and MPCs that is documented in many papers is explained solely through transitory economic shocks interacting with borrowing constraints or whether there is also an important role for more permanent heterogeneity. This matters for evaluating the consequences of many policy interventions which affect the level of liquidity in the economy. The pandemic environment provides a unique context where we can observe spending responses of normally low-liquidity households in a high-liquidity state. The high MPCs we observe even in this environment show that permanent heterogeneity matters for understanding consumer behavior.⁴⁵

⁴⁴We note that alternative modeling devices could also generate this pattern. The key is that there must be some persistent dimension that leads some households to have high MPCs even when they have high liquidity.

⁴⁵Note that liquidity still also matters for MPCs in our model. With no discount factor heterogeneity and in a normal

The evidence for permanent heterogeneity as a driver of MPCs in this section complements two other types of recent evidence. First, many papers correlate MPCs with household observables to infer the separate role of permanent and transitory characteristics. For example, Parker (2017) uses Nielsen data to show that high MPCs are correlated with survey measures of impatience and lack of financial planning. Aguiar, Bils, and Boar (2021) shows that in standard consumption models, permanent heterogeneity is necessary to match the relationship between MPCs and future spending growth in PSID data. Gelman (2021) uses panel data on the behavior of average cash buffers, current cash on hand and MPCs to infer an important role for permanent heterogeneity in a buffer-stock spending model. In contrast, our approach directly shows that MPCs remain large even after large increases in liquidity. That MPCs are high even in a high liquidity environment shows that transitory low liquidity cannot be the only force driving high MPCs.

Second, Patterson (2022) documents a strong correlation between ex-ante unemployment risk and MPCs. This is consistent with our argument that households with high unemployment exposure are also those with high spending sensitivity. Our contribution is to document that this correlation between high MPCs and unemployment risk is not driven solely by the liquidity-reducing impact of unemployment itself, but is true even when unemployed households have high liquidity. This suggests that part of the correlation is due to persistent characteristics. This conclusion is also supported by other moments in our data. In particular, even when employed, households who later become unemployed have about 25% less liquidity conditional on income than households who remain employed. This lower persistent liquidity, even conditional on income, and even when employed, is consistent with a persistent characteristic that increases spending.

The finding of high sensitivity of spending to unemployment benefits also buttresses a crucial assumption for leading behavioral models of job search. In models of job search with reference dependent preferences, the job-finding rate is highest when there is a large gap between current consumption and recent consumption (DellaVigna et al. 2017, DellaVigna et al. 2022). However, most datasets on unemployment benefits measure only income and the job-finding rate. The analyst therefore must assume a set of preferences which leads consumption to move suddenly around changes in UI benefits. The evidence in this paper provides evidence for the missing link that consumption does indeed fluctuate in response to benefit changes.

Summarizing the results from this section, persistence and targeting towards households with high propensities to spend are key forces for high MPCs out of supplements. Their large size decreases the MPC since it relaxes liquidity constraints, but on net the first two forces (i.e. persistence and targeting) dominate and so spending responses to UI supplements are much larger than spending responses to stimulus checks. This conclusion should generalize beyond the pandemic, since these forces are not pandemic-specific.⁴⁶

environment (no supplements and no pandemic-induced liquidity increases), the MPC of unemployed households is 0.3 larger than that of employed households, reflecting the role of current income and liquidity. Adding discount factor heterogeneity raises the gap to 0.44, implying that two-thirds of the MPC variation with unemployment status in the best-fit model is explained by current economic circumstances, while one-third is due to permanent characteristics.

⁴⁶Indeed, this conclusion is likely to be even stronger in a more normal recession when aggregate spending is not reduced due to pandemic-specific reasons (i.e., Table 4 row 3 is greater than row 2).

7 Policy Implications

The mechanisms underlying the small employment effects and large spending effects of temporary UI supplements during the pandemic have several policy implications that are likely to extend beyond the pandemic. This section discusses these policy implications. Rather than make specific normative recommendations about optimal policy, our goal in this section is to highlight new channels for consideration and to quantify relevant trade-offs.

7.1 A Countercyclical Motive for Temporary Supplements

Our finding that employment distortions hinge upon the interaction between benefit length and labor market conditions provides a new rationale in favor of countercyclical unemployment benefits. It is already well established theoretically that optimal benefit levels rise when unemployment duration elasticities fall (Kroft and Notowidigdo 2016; Landais, Michailat, and Saez 2018b). This is because reduced duration elasticities imply smaller distortions from UI, and it is optimal to provide greater insurance when the distortion is smaller. Kroft and Notowidigdo (2016) evaluate this channel by studying how the *hazard* elasticity changes with the unemployment rate (they then make the common assumption that this hazard elasticity is the same as the *duration* elasticity). They find that the hazard elasticity falls modestly as the unemployment rate rises, suggesting that UI benefits should be slightly more generous during recessions.

What we add is that even if the hazard elasticity is constant, the fact that job-finding rates typically fall during recessions provides an additional rationale for temporary countercyclical benefits. When the no-supplement job-finding rate falls, the wedge between duration and hazard elasticities grows. This means that the welfare-relevant distortion (the total duration elasticity) is lower for any given level of the hazard elasticity. Within the standard framework, this force pushes up optimal benefit levels during recessions, even if hazard elasticities are constant.⁴⁷

This rationale should be added on top of the forces considered in the prior literature in any full accounting of the optimal cyclicity of unemployment benefits. In particular, the prior literature has pointed out that the optimal cyclicity of benefits also depends on labor market externalities (Mitman and Rabinovich 2015; Landais, Michailat, and Saez 2018a,b) and on aggregate demand externalities (Kekre Forthcoming). A full accounting of all of the costs and benefits of countercyclical UI is outside the scope of this paper. Our point here is simply to highlight one new channel pushing in favor of more generous benefits in recessions. We next turn to a discussion of the *size* of a potential benefit increase.


7.2 Aggregate Demand Management: Size and Targeting

If temporary UI supplements are beneficial during recessions, how large should supplements be? In this section we quantify several of the key trade-offs in two steps, under the assumption that the government wants to use fiscal policy to stabilize aggregate demand. First, we look backwards to the pandemic environment and perform counterfactual simulations in our model to analyze the impacts of supplements with the same short four-month duration actually implemented, but with different

⁴⁷Schmieder, von Wachter, and Bender (2012) make a related point about optimal potential duration (as opposed to benefit levels) over the business cycle. They find empirically that the disincentive effect of UI extensions is close to constant over the business cycle. However, because the benefit exhaustion rate is higher in recessions, the welfare cost of extensions (which they capture as the ratio of the disincentive effect to the mechanical cost) is lower in recessions.

sizes. Second, we consider the impact of literal “severance” payments that last only one month. We consider these impacts in the case of a general recession without pandemic-specific effects. In this environment, we evaluate the relative spending impacts of providing targeted one-time stimulus to unemployed households (“severance”) versus providing universal one-time payments to all households (“stimulus checks”).

Figure 11: Effects of Different Supplement Sizes on Spending and Unemployment



Notes: This figure shows how spending and unemployment duration would have responded to supplements of alternative sizes implemented in April 2020. As with the actual \$600 supplements, we assume that the supplement lasts for four months, and we assume that expiration is unexpected as in our best fit model. The left panel computes the MPC out of supplements of various sizes. The solid line computes the MPC out of the entire supplement while the dashed line computes the MPC out of the last \$50 of the supplement. The right panel computes the change in unemployment duration (in weeks) to the supplement. The solid line shows the effect of the entire supplement while the dashed line shows the effect of the last \$50 of the supplement.

As discussed above, in the summer of 2020 there was substantial debate about the costs and benefits of supplements of various sizes, with some concerned about disincentive effects on labor supply from large replacement rates and others concerned about supporting the spending of the unemployed. In Figure 11 we explore what effects supplements of alternative sizes would have had on spending and employment during this time period. We summarize spending effects with the one-month MPC and we summarize disincentive effects with the increase in the length of unemployment induced by the supplement.⁴⁸ The solid lines in this figure show the total effects of supplements of various sizes while the dashed line shows marginal effects of the last dollar of supplements on spending and unemployment.

From this figure, it is clear that spending effects are large when supplements are small but then drop off rapidly for supplements around \$300, which induce a median replacement rate of 100%. This is because of the interactions between supplement size and MPCs discussed above: the high MPC out of small unemployment supplements is driven in part by the fact that the first dollar from supplements targets households who are currently liquidity constrained. However, if the supplements are large enough, then this is no longer the case on the margin. When the total supplement is large, the marginal dollars of the supplement are flowing to unconstrained rather than constrained households. The marginal spending impact therefore falls because households are at a flatter region

⁴⁸Conclusions are similar when measuring MPCs over longer horizons rather than just on impact. See Figure F-9.

of their consumption function.

In contrast, the marginal effects of supplements on employment are nearly constant, and there is nothing particularly special about larger supplements. This is because in the dynamic problem characterizing optimal search, households compare the future value of unemployment to the future value of employment. Just because current benefits are above the wage does not mean the dynamic value of unemployment exceeds that of employment. Thus, although many commentators were concerned about large supplements with greater than 100 percent replacement rates because of their potential labor market impacts, our results suggest that the bigger concern at that particular point in time should perhaps have come from declining spending impacts.

Thus, the supplement size counterfactuals in Figure 11 show two main lessons. First, there are large spending effects and small job-finding distortions from small, temporary supplements. Second, the ratio of spending benefits to job-finding costs declines as supplement size increases. This is because the spending impacts decline non-linearly while the job-finding distortions grow linearly. Every additional dollar of supplement induces less additional spending, but has the same incremental distortion to unemployment durations.

This analysis suggests that bang-for-the-buck is maximized at small supplements but it does not necessarily mean that small supplements are optimal. Indeed, even at \$600 per week the job-finding distortions from supplements are relatively small, increasing average unemployment durations by only 2.5 *weeks*. However, the longer the duration of supplements of a given size, the greater the distortion they will induce on aggregate employment. Indeed, short-lived “severance-like” UI supplements which trigger on during recessions could provide boosts to aggregate demand without much distortion to job-finding.


To explore the effects of actual severance payments, we analyze the effects of one-month supplements. We do this in a model environment which removes any pandemic-specific effects in order to more closely approximate a typical recession. Since employment effects from severance are minimal, we focus on the spending impacts. We focus specifically on how the spending impacts of severance (which is targeted to the unemployed) compares to the spending impacts of alternative universal stimulus payments (which go to the population as a whole), as a way to evaluate their potential as tools for aggregate demand management. This analysis is shown in Figure 12. The solid lines compare the average quarterly MPC out of a severance to to the MPC out of an equal-sized stimulus check, for an individual receiving each transfer. The dashed lines compare the marginal effect of the last \$50 of transfers.

We find that the spending impact of severance is larger than the impact of universal stimulus checks, although this difference declines with the size of the transfer. The impact of severance is larger than that of universal stimulus checks because severance targets individuals with a high propensity to spend while universal stimulus checks do not.⁴⁹ The difference is largest for small transfers, because when transfers are small the unemployed are temporarily liquidity constrained *and* have a high propensity to spend for any level of liquidity due to the permanent heterogeneity discussed in Section 6.2. For large transfers, unemployed households are no longer liquidity constrained and only the effect of targeting those with permanently high propensities to spend remains.

The combination of low pre-transfer liquidity and high persistent propensity to spend at any liquidity means that even large severance payments targeted specifically to the unemployed may be

⁴⁹Note that since we are now studying the effect of severance, there is no additional amplifying effect from persistence.

Figure 12: Spending Impacts of Severance vs. Untargeted Stimulus



Notes: This figure shows the quarterly spending responses to severance as well as to one-time stimulus checks of various sizes in a non-pandemic environment with normal liquidity levels. The x-axis compares severance payments and untargeted stimulus of equal size and the y-axis reports the quarterly MPC. We compute these responses in the best fit model with discount factor heterogeneity. Note that we calibrate the degree of heterogeneity in this model so that it still produces a quarterly MPC out of \$500 stimulus checks of 0.25. Solid lines compute MPCs out of the entire transfer while dashed lines compute MPCs out of the last \$50 of transfers.

beneficial relative to untargeted stimulus. Figure 12 shows that the spending impact from the last dollar of a \$2,000 one-time payment to the unemployed is larger than the spending impact of the first dollar of untargeted stimulus.⁵⁰

Severance payments can be interpreted as a means of front-loading UI benefits, and in this sense our results relate to several other papers which argue for front-loading (Shavell and Weiss 1979; Hopenhayn and Nicolini 1997; Mitman and Rabinovich 2021). Lindner and Reizer (2020) provide empirical evidence showing front-loading benefits led to shorter durations. Our analysis of severance adds a novel motive for front-loading, which is that it is an effective way to stimulate aggregate demand.

The analysis in this section is relatively simple, but it demonstrates the potential power of temporary UI supplements which are triggered on by recessions. This is potentially a much cheaper way of providing fiscal stimulus than unconditional stimulus checks. However, there are three caveats worth emphasizing. First, severance pay is not targeted to the long-term unemployed who are most in need of insurance (Ganong and Noel 2019; Gerard and Naritomi 2021). A full analysis of optimal policy would incorporate both the aggregate demand motive we study and the insurance motive studied in much of the prior literature. Second, the larger the “severance-like” payment, the larger incentive there might be for employees and employers to collude, generating false terminations to claim this benefit. Such a policy would therefore need to be carefully designed to mitigate this risk. Finally, as payments become large, the unemployed are no longer low-liquidity and so the spending benefit relative to universal payments comes only from targeting payments to households that have persistently higher MPCs. The relevance of this force in future recessions depends on the extent to which households who become unemployed in the future share the same permanent characteristic that boosts

⁵⁰The conclusion that the spending impact of UI supplements is larger than the spending impact of untargeted stimulus for a wide range of intervention sizes also holds for equal *cost* interventions instead of equal *size* payments. See Appendix F.3.

the spending sensitivity of the households we study.

8 Conclusion

We use administrative bank account data to estimate the causal effects of the largest unemployment benefit expansion in U.S. history on spending and job-finding. Our reduced-form research designs and dynamic structural model deliver consistent conclusions: expanded benefits had large effects on spending but small effects on job-finding. The small job-finding effects were driven in part by the fact that supplements were temporary and implemented in an environment with an already depressed job-finding rate. The large spending effects were driven in part by the fact that they were targeted towards households with high spending propensities.

These conclusions have lessons for future policy design: countercyclical severance-like payments should be considered alongside stimulus checks as an additional instrument for fiscal stimulus. The conclusions also have lessons for future modeling: the large spending responses we find even in a high liquidity environment imply that incorporating permanent heterogeneity in behavioral characteristics into models is important for generating realistic empirical predictions, especially in situations with dramatic changes in liquidity.

Our results suggest several avenues for future research. Studying the effects of supplements in the summer and fall of 2021 would be useful since the labor market changed rapidly over this period. Because households start to exhaust benefits in the summer of 2021, studying this period requires an alternative empirical methodology measuring job-finding directly, instead of inferring it from UI exits.

It would also be interesting to explore in more detail whether additional observable characteristics of the unemployed predict their large spending responses. Is unemployment *per se* an important tag for high MPCs or is it merely correlated with other features like income levels or education that might also affect MPCs? Answering this question is important for determining whether other forms of targeted stimulus might achieve the same ends without working through the UI system.

Finally, it is interesting to explore the implications of supplement-induced spending for other outcomes: what categories of spending and particular firms are most affected? What are the implications for cross-sector relative price movements and aggregate inflation? Understanding the possible inflationary effects of supplement-induced spending is particularly important given the rapidly changing macro environment in 2020 and 2021. Early in the pandemic, there was a strong case for fiscal stimulus (see, e.g., Guerrieri et al. 2022), but any output gap was likely much smaller by mid 2021 when UI supplements eventually expired. This means that additional spending at that point may have contributed to inflation rather than helping the economy. Future work linking spending effects to price data might shed light on these mechanisms.

References

- Aguiar, Mark A., Mark Bils, and Corina Boar. 2021. “Who are the Hand-to-Mouth?” Working paper.
- Bartik, Alexander, Marianne Bertrand, Feng Lin, Jesse Rothstein, and Matthew Unrath. 2020. “Measuring the labor market at the onset of the COVID-19 crisis.” Working paper.
- Bell, Alex, Thomas J. Hedin, Peter Mannino, Roozbeh Moghadam, Carl Romer, Geoffrey Schnorr, and Till von Wachter. 2022. “Increasing Equity and Improving Measurement in the U.S. Unemployment System: 10 Key Insights from the COVID-19 Pandemic.” Working paper, California Policy Lab.

- Bell, Alex, Thomas J. Hedin, Roozbeh Moghadam, Geoffrey Schnorr, and Till von Wachter. 2021. "An Analysis of Unemployment Insurance Claims in California During the COVID-19 Pandemic." Policy Brief, California Policy Lab.
- Bitler, Marianne P., Hilary W. Hoynes, and Diane Whitmore Schanzenbach. 2020. "The Social Safety Net in the Wake of COVID-19." Brookings Papers on Economic Activity.
- Boutros, Michael. 2022. "Windfall Income Shocks with Finite Planning Horizons." Working paper.
- Cajner, Tomaz, Leland Crane, Ryan Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz. 2020. "The U.S. Labor Market during the Beginning of the Pandemic Recession." WP 27159, NBER.
- Card, David, Raj Chetty, and Andrea Weber. 2007. "Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market." *The Quarterly Journal of Economics* 122 (4):1511–1560.
- Chetty, Raj. 2008. "Moral Hazard versus Liquidity and Optimal Unemployment Insurance." *Journal of Political Economy* 116 (2):173–234.
- Chetty, Raj, John Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team . 2020. "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." WP 27431, NBER.
- Conley, T.G. 1999. "GMM estimation with cross sectional dependence." *Journal of Econometrics* 92 (1):1–45.
- Coombs, Kyle, Arindrajit Dube, Calvin Jahnke, Raymond Kluender, Suresh Naidu, and Michael Stepner. 2022. "Early Withdrawal of Pandemic Unemployment Insurance: Effects on Employment and Earnings." *AEA Papers and Proceedings* 112:85–90.
- Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, and Fiona Greig. 2020. "Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data." Brookings Papers on Economic Activity.
- de Chaisemartin, C. and X. D'Haultfoeulle. 2018. "Fuzzy Differences-in-Differences." *Review of Economic Studies* 85 (2):999–1028.
- DellaVigna, Stefano, Jörg Heining, Johannes F Schmieder, and Simon Trenkle. 2022. "Evidence on Job Search Models from a Survey of Unemployed Workers in Germany." *The Quarterly Journal of Economics* 137 (2):1181–1232.
- DellaVigna, Stefano, Attila Lindner, Balázs Reizer, and Johannes F. Schmieder. 2017. "Reference-Dependent Job Search: Evidence from Hungary*." *The Quarterly Journal of Economics* 132 (4):1969–2018.
- Dube, Arindrajit. 2021. "Aggregate Employment Effects of Unemployment Benefits During Deep Downturns: Evidence from the Expiration of the Federal Pandemic Unemployment Compensation." WP 28470, NBER.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics* 13 (4):1–54.
- Finamor, Lucas and Dana Scott. 2021. "Labor market trends and unemployment insurance generosity during the pandemic." *Economics Letters* 199:109722.
- Freyaldenhoven, Simon, Christian Hansen, and Jesse M. Shapiro. 2019. "Pre-event Trends in the Panel Event-Study Design." *American Economic Review* 109 (9):3307–38.
- Ganong, Peter and Pascal Noel. 2019. "Consumer Spending during Unemployment: Positive and Normative Implications." *American Economic Review* 109 (7):2383–2424.
- Ganong, Peter, Pascal Noel, and Joseph Vavra. 2020. "US unemployment insurance replacement rates during the pandemic." *Journal of Public Economics* 191.
- Gelman, Michael. 2021. "What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics." *Journal of Monetary Economics* 117:521–542.
- Gerard, François and Joana Naritomi. 2021. "Job Displacement Insurance and (the Lack of) Consumption-Smoothing." *American Economic Review* 111 (3):899–942.
- Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *The American Economic Review* 87 (1):192–205.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2022. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" *American Economic Review* 112 (5):1437–74.
- Hagedorn, Marcus, Fatih Karahan, Iourii Manovskii, and Kurt Mitman. 2013. "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects." WP 19499, NBER.

- Havranek, Tomas and Anna Sokolova. 2020. “Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say “probably not”.” *Review of Economic Dynamics* 35:97–122.
- Holzer, Harry J., R. Glenn Hubbard, and Michael R. Strain. 2021. “Did Pandemic Unemployment Benefits Reduce Employment? Evidence from Early State-Level Expirations in June 2021.” WP 29575, NBER.
- Hopenhayn, Hugo A. and Juan Pablo Nicolini. 1997. “Optimal unemployment insurance.” *Journal of Political Economy* 105 (2):412–438.
- Ilut, Cosmin L. and Rosen Valchev. 2020. “Economic Agents as Imperfect Problem Solvers.” WP 27820, NBER.
- Initiative on Global Markets. 2021. “Unemployment Benefits.” URL <https://www.igmchicago.org/surveys/unemployment-benefits/>. Retrieved on 2022-07-13.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. “Household Expenditure and the Income Tax Rebates of 2001.” *American Economic Review* 96 (5):1589–1610.
- Kaplan, Greg and Giovanni L. Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments.” *Econometrica* 82 (4):1199–1239.
- . Forthcoming. “The Marginal Propensity to Consume in Heterogeneous Agent Models.” *Annual Reviews of Economics* .
- Katz, Lawrence F. 1986. “Layoffs, Recalls and the Duration of Unemployment.” WP 1825, NBER.
- Kekre, Rohan. Forthcoming. “Unemployment Insurance in Macroeconomic Stabilization.” *Review of Economic Studies* .
- Kroft, Kory and Matthew J. Notowidigdo. 2016. “Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence.” *The Review of Economic Studies* 83 (3):1092–1124.
- Krueger, D., K. Mitman, and F. Perri. 2016. “Macroeconomics and Household Heterogeneity.” In *Handbook of Macroeconomics*, vol. 2. Elsevier, 843–921.
- Kueng, Lorenz. 2018. “Excess Sensitivity of High-Income Consumers*.” *The Quarterly Journal of Economics* 133 (4):1693–1751.
- Laibson, David, Peter Maxted, and Benjamin Moll. 2021. “Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy.” WP 29094, NBER.
- Landais, Camille, Pascal Michaillat, and Emmanuel Saez. 2018a. “A macroeconomic approach to optimal unemployment insurance: Applications.” *American Economic Journal: Economic Policy* 10 (2):182–216.
- . 2018b. “A Macroeconomic Approach to Optimal Unemployment Insurance: Theory.” *American Economic Journal: Economic Policy* 10 (2):152–181.
- Lian, Chen. 2022. “Mistakes in Future Consumption, High MPCs Now.” Working paper.
- Lindner, Attila and Balázs Reizer. 2020. “Front-Loading the Unemployment Benefit: An Empirical Assessment.” *American Economic Journal: Applied Economics* 12 (3):140–74.
- Marinescu, Ioana, Daphné Skandalis, and Daniel Zhao. 2021. “The impact of the Federal Pandemic Unemployment Compensation on job search and vacancy creation.” *Journal of Public Economics* 200:104471.
- McKay, Alisdair and Ricardo Reis. 2021. “Optimal Automatic Stabilizers.” *The Review of Economic Studies* 88 (5):2375–2406.
- Michaillat, Pascal. 2012. “Do Matching Frictions Explain Unemployment? Not in Bad Times.” *The American Economic Review* 102 (4):1721–1750.
- Mitman, Kurt and Stanislav Rabinovich. 2015. “Optimal unemployment insurance in an equilibrium business-cycle model.” *Journal of Monetary Economics* 71:99–118.
- . 2021. “Whether, when and how to extend unemployment benefits: Theory and application to COVID-19.” *Journal of Public Economics* 200:104447.
- Orchard, Jacob, Valerie A Ramey, and Johannes Wieland. 2022. “Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates.” Working paper.
- Parker, Jonathan A. 2017. “Why Don’t Households Smooth Consumption? Evidence from a \$25 Million Experiment.” *American Economic Journal: Macroeconomics* 9 (4):153–83.

- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6):2530–53.
- Patterson, Christina. 2022. "The Matching Multiplier and the Amplification of Recessions." Working paper.
- Petrosky-Nadeau, Nicolas and Robert G. Valletta. 2021. "UI Generosity and Job Acceptance: Effects of the 2020 CARES Act." IZA Discussion Papers 14454, Institute of Labor Economics (IZA).
- Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe. 2022. "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature." Working paper.
- Schmieder, J. F., T. von Wachter, and S. Bender. 2012. "The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years." *The Quarterly Journal of Economics* 127 (2):701–752.
- Schmieder, Johannes F. and Till von Wachter. 2016. "The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation." *Annual Review of Economics* 8 (1):547–581.
- Shavell, Steven and Laurence Weiss. 1979. "The optimal payment of unemployment insurance benefits over time." *Journal of Political Economy* 87 (6):1347–1362.

Online Appendix to “Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data”


Peter Ganong, Fiona Greig, Pascal Noel, Daniel M. Sullivan, Joseph
Vavra

Contents

A Additional Figures and Tables	1
B Additional Data Description	26
B.1 States Included	26
B.2 Unemployment and Employment Spells	27
B.3 Sample Restrictions	27
B.4 Additional Variable Detail	28
B.5 Measuring Weekly Benefit Amount	30
C MPC Robustness	32
D Lost Wages Assistance Spending Responses	33
E Calculating Duration Elasticities	34
F Additional Model Details and Figures	35
F.1 Additional Model Details	35
F.2 Model Robustness and Additional Model Results	38
F.3 Policy Counterfactuals – Additional Results	44

A Additional Figures and Tables


Figure A-1: States included in the JPMCI sample




Notes: This figure shows the subset of states which are included in various analyses. See Appendix B.1 for details.

Figure A-2: UI Claims in JPMCI versus DOL


(a) Time-series



(b) State-level Change at Pandemic Onset




(c) Weekly Benefit Amount Pre-Pandemic




Notes: This figure compares total claims and benefits levels in JPMCI data to Department of Labor ETA Form 539 and Form 5159. Panel (a) compares aggregate time-series patterns, panel (b) compares state-by-state changes, and panel (c) compares benefit levels by state. Panel (b) depicts the ratio of the number of payments in May 2020 to the number of payments in 2020 prior to the declaration of national emergency. Panels (b) and (c) include a 45-degree line and drop states with less than 300 observations.

Figure A-3: Exit Rate Robustness


(a) Drop Account Activity Screens



(b) Job Separation Observed at Start of Spell




(c) Vary Minimum Weeks Without Receipt to Define Exit



Notes: This figure shows that our main job-finding patterns are robust to various alternative sample screening choices. Panel (a) shows patterns for the total job-finding rate after dropping account activity screens. Panel (b) shows the total job-finding rate for those with and without observed job separation (note that we can only separate the total job-finding rate into recalls and new job-finding for the sample with observed separations.) Panel (c) shows robustness of the total job-finding rate to alternate thresholds for defining an exit from unemployment.

Figure A-4: Changes in Industry Composition of Unemployment, JPMCI vs DOL




Notes: This figure compares the change in unemployment composition in DOL to JPMCI. The diagonal line is a 45-degree line. Making these data sets comparable requires two adjustments. Letting i denote industry, t denote period and s state, we start by defining total JPMCI claims as c_{it}^{JPMCI} and state-specific claims from DOL ETA Form 203 as c_{ist}^{DOL} . From 203 excludes recipients of federal programs for the long-term unemployed, so we drop recipients in JPMCI data with spells > 26 weeks. We also re-weight the DOL data to account for the fact that Chase has more customers in some states than others using weight $w_{st} = \frac{c_{st}^{JPMCI}}{\sum_{s'} c_{s't}^{JPMCI}}$. We then measure industry share in JPMCI as

$p_{it}^{JPMCI} = \frac{c_{it}^{JPMCI}}{\sum_{i'} c_{i't}^{JPMCI}}$ and industry share in DOL as $p_{it}^{DOL} = \frac{\sum_s (w_{st} c_{ist}^{DOL})}{\sum_{i'} \sum_s (w_{st} c_{i'st}^{DOL})}$. Unsurprisingly, UI claim shares by industry differ between the two datasets (“Construction” and “Agriculture” are most under-represented in JPMCI and “Administrative and Support Services” and “Finance and Insurance” are most over-represented) but we are primarily interested in the extent to which the *changes* in the composition of unemployment during the pandemic appear in the JPMCI data. To quantify this, we analyze the shift in the composition of UI claims from pre-covid (January 2019 to March 2020) to the height of the pandemic (April 2020 to December 2020). Because the increases in UI claims are so large (and therefore the changes in proportions are highly skewed), we measure composition changes with the symmetric percent change: $2 \times \frac{p_{it} - p_{i,t-1}}{p_{it} + p_{i,t-1}}$.

Figure A-5: Spending of Unemployed versus Employed (Median)


Percent difference from January 2020 (median)



- Unemployed (get benefits from April 2020 through February 2021)
- Employed


Notes: This figure replicates Figure 1 using sample medians instead of means.

Figure A-6: Spending of Unemployed Versus Employed (card and cash)




Notes: This figure shows that the spending (total) patterns in Figure 1 also hold for spending (card and cash).

Figure A-7: Impact of Delays in Unemployment Benefits on Spending (Total, Levels)




Notes: This figure shows that the spending (card and cash) patterns in Figure 2 also hold for spending (total), although this broader spending measure has noisier week-to-week fluctuations due to recurring payments and potential measurement error in the exact timing of when paper checks are deposited.

Figure A-8: Impact of Delays in Unemployment Benefits on Spending (Total, Differences)



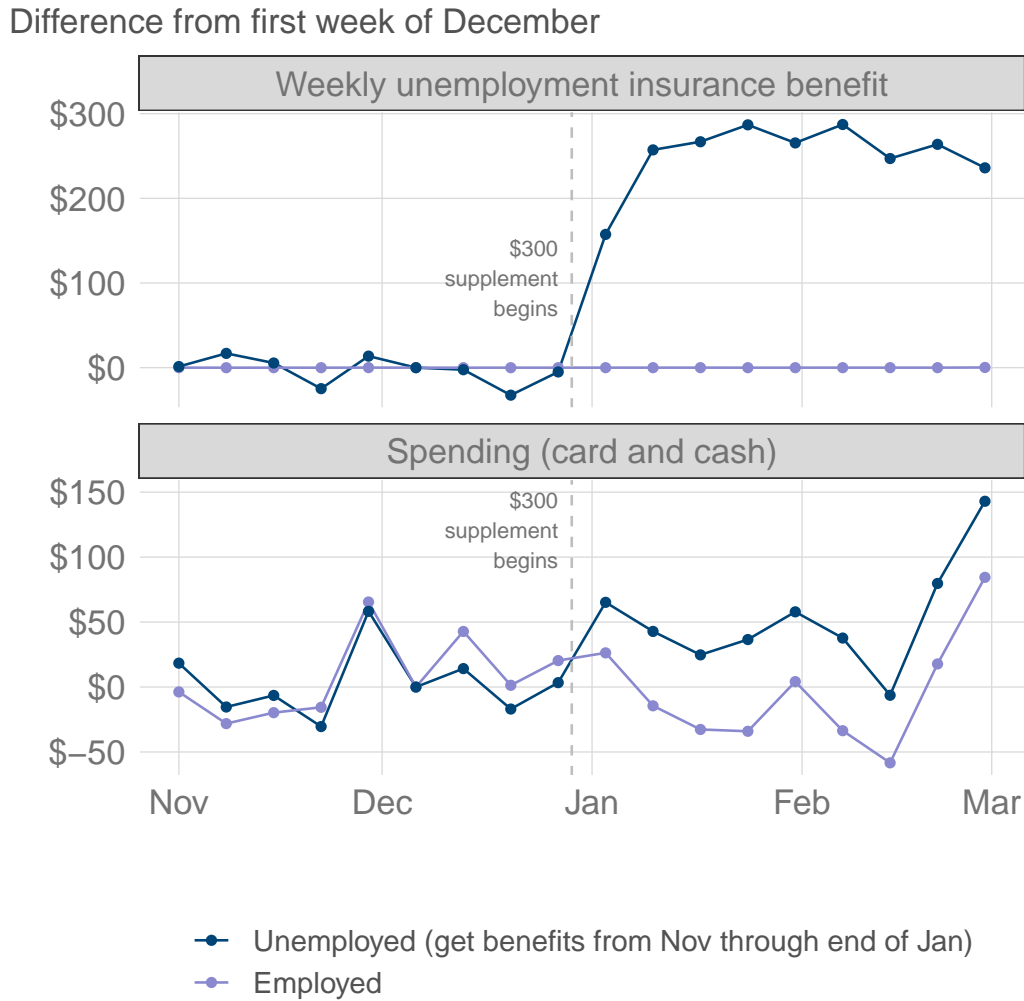
Notes: This figure shows mean income and spending (total) differences from the May 31 control group for various cohorts in the waiting for benefit receipt research design. Our MPC is based on the April 5th treatment group, which has no benefit delay.

Figure A-9: Impact of Expiration of the \$600 Supplement on Spending (Card and Cash)




Notes: This figure measures the causal impact of the expiration of the \$600 supplement on spending. The benefit amount declines over two weeks in August (rather than one week) because some states pay benefits once ever two weeks and therefore paid out the supplement for the last week of July during the first week of August. The benefit amount rises in September because states begin to pay the temporary \$300 supplement. The control group is employed workers who are matched on 2019 income levels as well as date of receipt of Economic Impact Payment. The dependent variables are mean benefits and mean spending, measured as a change relative to the first week of July. See Section 3.2.2 for details.

Figure A-10: Impact of Onset of the \$300 Supplement on Spending (Card and Cash)



Notes: This figure measures the causal impact of the onset of the \$300 supplement on spending. The benefit amount rises over two weeks in January (rather than one week) because some states pay benefits once every two weeks. The control group is employed workers who are matched on 2019 income levels as well as date of receipt of Economic Impact Payment. The dependent variables are mean benefits and mean spending, measured as a change relative to the last week of December. The figure depicts November 2020 through March 2021. See Section 3.2.2 for details.

Figure A-11: Impact of Expiration of the \$600 Supplement on Spending (Total)



Notes: This figure measures the causal impact of the expiration of the \$600 supplement on spending. The benefit amount declines over two weeks in August (rather than one week) because some states pay benefits once every two weeks and therefore paid out the supplement for the last week of July during the first week of August. The benefit amount rises in September because states begin to pay the temporary \$300 supplement. The control group is employed workers who are matched on 2019 income levels as well as date of receipt of Economic Impact Payment. The dependent variables are mean benefits and mean spending, measured as a change relative to the first week of July. Spending is noticeably higher in weeks which contain the first of the month, likely because many households pay bills at this time. See Section 3.2.2 for details.


Figure A-12: Impact of Onset of the \$300 Supplement on Spending (Total)




Notes: This figure repeats Figure A-10 but for total income and spending. The control group is employed workers who are matched on 2019 income levels as well as date of receipt of Economic Impact Payment. The dependent variables are mean benefits and mean spending, measured as a change relative to the last week of December. The figure depicts November 2020 through March 2021. Spending is noticeably higher in weeks which contain the first of the month, likely because many households pay bills at this time, so the third panel differences out these high-frequency fluctuations. Income for employed households similarly exhibits spikes at the end of the month for the same reason.

Figure A-13: Exit Rate from Unemployment Benefits

(a) Recall

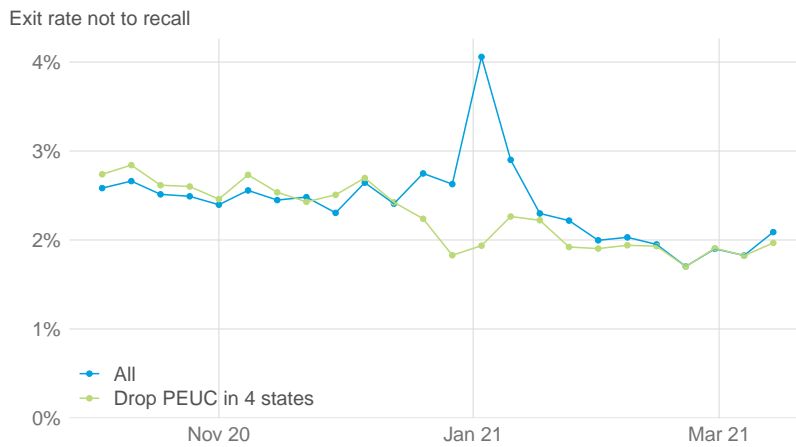


(b) Total (New Job + Recall)




Notes: This figure shows the exit rate to recall and the total exit rate in the JPMCI data. UI exit is defined as three contiguous weeks without receipt of UI benefits. Recall is measured using receipt of labor income from a prior employer. Exit rate to new job is from Figure 3.

Figure A-14: Exit Rate at Expiration of PEUC




Notes: This figure shows the evolution of the exit rate not to recall from October 2020 through February 2021. The blue series is the same as the one shown in Figure 3, except that here the series includes January 3 and January 10. The y-axis title is the “exit rate not to recall” instead of the “exit rate to new jobs” because some of the exits arise from a policy seam. The green series drops the 71,000 households that have received at least 20 weeks of benefits in 2019 and 2020 in Indiana, California, New Jersey, and Ohio. These households are likely to be recipients of Pandemic Emergency Unemployment Compensation, which temporarily lapsed at the end of December and these four states were slow to restore benefits after the lapse. The difference between the blue series and the green series reveals that the lapse triggered a surge in *measured* exits in four states. In additional unreported results, we find that the measured exits in the blue series do not show evidence of starting a new job via direct deposit of payroll from a new employer. We therefore omit January 3 and January 10 from the plot in Figure 3.

Figure A-15: Exit Rate by Start Date of Unemployment Benefit Spell



Notes: We define the pre-pandemic group as those who started receiving unemployment insurance benefits during or before the week of March 15, 2020.


Figure A-16: Effect of Expanded Benefits on Job-Finding: Interrupted Time-series Design




Notes: The top panel of this figure shows the exit rate to a new job in the JPMCI data from April 2020 through February 2021. The red horizontal bars indicate the average exit rate in the two weeks prior to and four weeks following a change in the supplement amount. We form a test statistic for the impact of the supplement using the difference between the red horizontal bars. We omit January 3 and 10 because they show a mechanical surge in exits arising from a policy lapse. We recompute the test statistic for every placebo date shown in the top panel, where we define placebo windows as those with no policy change. The bottom panel of this figure shows the distribution of the test statistic using blue bars. The changes at the actual supplement changes are more extreme than the changes at any of the placebo dates. If we assume that the date of the supplement change is random, this implies that we reject the null hypothesis of no effect of the supplement with $p \leq 1/31$.

Figure A-17: Exit Rate by Replacement Rate


(a) Expiration – New Job




(b) Onset – New Job




(c) Expiration – Total




(d) Onset – Total



(e) Expiration – Recall




(f) Onset – Recall




Notes: This figure shows several alternative specifications for Figure 4. Panels (a) and (b) report raw (unnormalized) exit rates to new job. Panels (c) and (d) report “total” exit rates including exit to new job and exit to recall. Panels (e) and (f) report exit rates to recall.

Figure A-18: Weekly Event Study Coefficients (Continuous Specification)

(a) Expiration of \$600

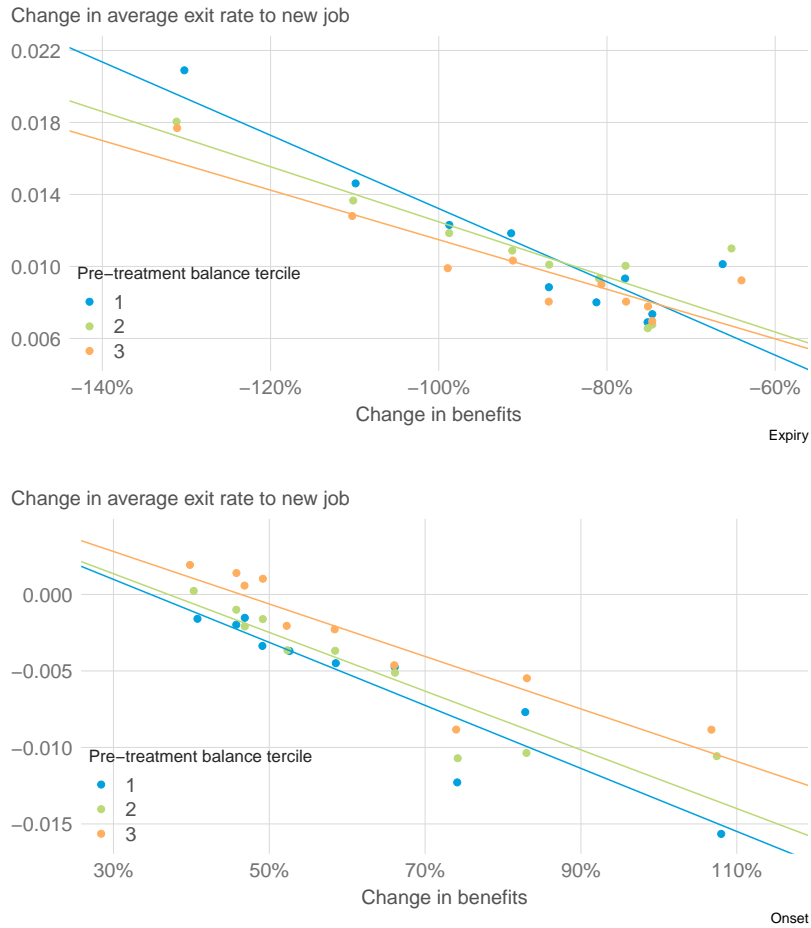


(b) Onset of \$300



Notes: This figure shows the results from the weekly difference-in-difference specification defined in Section 4.2.3 equation (5). This specification captures the effect of supplements on job-finding in each week around the supplement change. Panel (b) indicates that even prior to the onset of the \$300 supplement there is already a gradual trend downward in the job-finding rate for households that receive the largest increase in benefits on January 1, 2021. If we were to use a specification that accounted for this pre-trend in estimation we would likely find that the \$300 supplement has even smaller effects on the job-finding rate.

Figure A-19: Binscatter by Pre-treatment Liquidity



Notes: This figure repeats the binscatter between benefit changes and the new job-finding rate from Figure 5 but separately for three tertiles of the checking account balance distribution. For the \$600 expiration in panel (a), checking account balances are measured in March 2020 and for the \$300 expiration in panel (b), checking account balances are measured in October 2020. We measure balances at these dates to capture liquidity before the household receives any supplement payments. See footnote 40 for details.

Table A-1: Monthly Income and Spending in Employed and Unemployed Households

Group (months)	Income	Benefits	Spending		Account balance
			Card and cash	Total	
Mean					
Employed (Jan-Feb 2020)	\$6850	\$0	\$2470	\$4669	\$5262
Employed (Apr-Oct 2020)	\$6839	\$0	\$2322	\$4538	\$5884
Pandemic unemployed (Jan-Feb 2020)	\$5854	\$16	\$2506	\$4248	\$3488
Pandemic unemployed (Apr-Oct 2020)	\$7036	\$3947	\$2780	\$4638	\$5249
Median					
Employed (Jan-Feb 2020)	\$5353	\$0	\$2064	\$3834	\$2815
Employed (Apr-Oct 2020)	\$5466	\$0	\$1925	\$3739	\$3389
Pandemic unemployed (Jan-Feb 2020)	\$4549	\$0	\$2109	\$3495	\$1624
Pandemic unemployed (Apr-Oct 2020)	\$5784	\$3834	\$2477	\$4044	\$3242

Notes: This table shows monthly values of income, unemployment benefits, spending (card and cash), spending (total), and checking account balances for employed and unemployed households, before and during the start of pandemic. “Employed” households do not receive benefits or have a job separation from January 2020 through February 2021. “Pandemic unemployed” households begin an unemployment spell in April 2020. A very small number of these households also received benefits in a separate spell which ended prior to April 2020 in January and February 2020, which is why the pre-pandemic mean benefits for this group is \$16.

Table A-2: MPC Robustness to Summary Statistic and Sample

Episode	(1)	(2)	(3)	(4)
Waiting for benefits	0.43	0.49	0.47	0.43
Expiration of \$600 supplement	0.3	0.37	0.33	0.27
Onset of \$300 supplement	0.26	0.29	0.3	0.23
Lost wages assistance	0.32	0.34	0.35	0.28
Sample	All	All	No non-Chase credit card	Make ACH debt payments
Summary statistic	Mean	Median	Mean	Mean

Notes: This table re-computes our MPC results for a number of alternative samples. The first three rows compute MPCs to the three identification strategies show in the main text and the fourth row computes MPCs to an additional temporary “Lost wages assistance” program. Column 1 repeats the main specification from Table 1. Column 2 computes MPCs using median instead of mean spending. Column 3 excludes households who make debt payments to non-Chase credit cards (for whom we are potentially missing some spending). Column 4 restricts to households who make debt payments via ACH (whom we are more confident are not making mis-classified debt payments via paper check). See Appendix C for additional details.

Table A-3: Regression Estimates for Effect of Expanded Benefits on Job-Finding

	<i>Dependent variable:</i>	
	Exit to new job	
	Expiration of \$600	Onset of \$300
	(1)	(2)
PctChange	0.022*** (0.001)	0.021*** (0.001)
SuppAvail	0.004*** (0.001)	0.007*** (0.001)
PctChange:SuppAvail	-0.017*** (0.001)	-0.020*** (0.001)
Constant	0.002*** (0.001)	0.009*** (0.001)
Observations	2,390,259	2,031,053

Notes: This table estimates the difference-in-difference model $e_{it} = \gamma PctChange_i + \alpha SuppAvail_t + \beta SuppAvail_t \times PctChange_i + \varepsilon_{it}$ from equation (4) using a window of eight weeks prior to and eight weeks after the two policy changes (expiration of the \$600 supplement and onset of the \$300 supplement). For expiration, the supplement available period is June and July 2020 and the no-supplement period is August and September 2020. For onset, the supplement available period is January 15-March 15 2021 and the no-supplement period is November and December 2020. Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01.

Table A-4: Micro Effect of Expanded Benefits: Robustness to Controls

(a) Expiration of \$600

	<i>Dependent variable:</i>			
	Exit to New Job			
	(1)	(2)	(3)	(4)
PctChange*SuppAvail	−0.0166*** (0.0011)	−0.0146*** (0.0011)	−0.0152*** (0.0012)	−0.0163*** (0.0022)
PctChange	X	X	X	X
SuppAvail	X	X	X	X
State*SuppAvail FE		X	X	X
Age*SuppAvail FE			X	X
Industry*SuppAvail FE				X
Observations	2,390,259	2,390,259	2,111,698	573,632

(b) Onset of \$300

	<i>Dependent variable:</i>			
	Exit to New Job			
	(1)	(2)	(3)	(4)
PctChange*SuppAvail	−0.0196*** (0.0012)	−0.0174*** (0.0012)	−0.0201*** (0.0013)	−0.0200*** (0.0024)
PctChange	X	X	X	X
SuppAvail	X	X	X	X
State*SuppAvail FE		X	X	X
Age*SuppAvail FE			X	X
Industry*SuppAvail FE				X
Observations	2,031,053	2,031,053	1,811,428	505,300

Notes: This table reports estimates of $\hat{\beta}$ from equation (4), adding increasingly stringent control variables. The first column is the same as in Table A-3. Column (2) adds state by time fixed effects. Column (3) adds age bin by time fixed effects. Column (4) adds prior industry by time fixed effects. Prior industry is available only for workers who worked at the 1000 largest firms in the data and therefore uses a smaller sample than the other columns. Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01.

Table A-5: Disincentive Estimates with Conley Standard Errors

	<i>Dependent variable:</i>	
	Exit to new job	
	Expiration of \$600	Onset of \$300
	(1)	(2)
PctChange	0.0211*** (0.0038)	0.0192*** (0.0057)
SuppAvail	0.0027 (0.0020)	0.0102*** (0.0038)
PctChange:SuppAvail	-0.0144*** (0.0024)	-0.0229*** (0.0063)
Constant	0.0014 (0.0036)	0.0086** (0.0035)
Observations	59756	50775

Notes: This table re-estimates Table A-3 using the Conley (1999) method of constructing standard errors. For computational reasons, we use a 2.5% subsample. As a result, the point estimates are different than the full sample estimates in Table A-3. We cluster using the weekly benefit, which is a proxy for pre-separation wages. The distance metric is the range of benefits divided by 40 and we use a uniform kernel. *p<0.1; **p<0.05; ***p<0.01.

Table A-6: Micro Effect of Expanded Benefits: Alternative Measures of Exit

(a) Expiration of \$600

	<i>Dependent variable:</i>		
	New job	Recall	Total
	(1)	(2)	(3)
SuppAvail*PctChange	-0.0166*** (0.0011)	-0.0117*** (0.0012)	-0.0283*** (0.0016)
Observations	2,390,259	2,390,259	2,390,259

(b) Onset of \$300

	<i>Dependent variable:</i>		
	New job	Recall	Total
	(1)	(2)	(3)
SuppAvail*PctChange	-0.0196*** (0.0012)	-0.0021*** (0.0008)	-0.0217*** (0.0014)
Observations	2,031,053	2,031,053	2,031,053

Notes: This table reports estimates of $\hat{\beta}$ from equation (4) specified for four different outcome variables. The first column is the same as in Table A-3. Column (2) is exit to recall and column (3) is total exit (new job or recall). Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01.

Table A-7: Disincentive by Liquidity – Expiration of \$600

	<i>Dependent variable:</i>		
	Base (1)	Control for liquidity (2)	Triple difference (3)
PctChange	0.0216*** (0.0009)	0.0219*** (0.0009)	0.0218*** (0.0009)
SuppAvail	0.0040*** (0.0009)	0.0040*** (0.0009)	0.0038*** (0.0009)
StdBalance		0.0006*** (0.0001)	0.0025*** (0.0008)
PctChange*SuppAvail	-0.0166*** (0.0011)	-0.0166*** (0.0011)	-0.0164*** (0.0011)
SuppAvail*StdBalance			-0.0016* (0.0009)
PctChange*StdBalance			-0.0023** (0.0009)
PctChange*SuppAvail*StdBalance			0.0022** (0.0011)
Constant	0.0022*** (0.0008)	0.0019** (0.0008)	0.0019** (0.0008)
PctChange*SuppAvail if balance 1 sd above mean			-0.0142
PctChange*SuppAvail if balance 1 sd below mean			-0.0186
Observations	2,390,259	2,390,157	2,390,157

Note: “Std Balance” measures liquidity using checking account balance at the end of March 2020. Balances are winsorized at the 90th percentile and then standardized to have mean 0 and standard deviation 1. By measuring liquidity in March, we capture liquidity before the household has received any supplement payments (see footnote 40 for details). Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01.

Table A-8: Disincentive by Liquidity – Onset of \$300

	<i>Dependent variable:</i>		
	Base (1)	Control for liquidity (2)	Triple difference (3)
PctChange	0.0215*** (0.0009)	0.0215*** (0.0009)	0.0210*** (0.0009)
SuppAvail	0.0069*** (0.0007)	0.0069*** (0.0007)	0.0066*** (0.0007)
StdBalance		-0.00002 (0.0001)	0.0021*** (0.0005)
PctChange*SuppAvail	-0.0196*** (0.0012)	-0.0196*** (0.0012)	-0.0191*** (0.0012)
SuppAvail*StdBalance			0.0004 (0.0007)
PctChange*StdBalance			-0.0045*** (0.0009)
PctChange*SuppAvail*StdBalance			0.0007 (0.0012)
Constant	0.0090*** (0.0005)	0.0090*** (0.0005)	0.0092*** (0.0005)
PctChange*SuppAvail if balance 1 sd above mean			-0.0184
PctChange*SuppAvail if balance 1 sd below mean			-0.0198
Observations	2,031,053	2,030,681	2,030,681

Note: “Std Balance” measures liquidity using checking account balance at the end of October 2020. Balances are winsorized at the 90th percentile and then standardized to have mean 0 and standard deviation 1. By measuring liquidity in October, we capture liquidity before the household has received any supplement payments (see footnote 40 for details). Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01.

Table A-9: Duration Elasticity by Various Household Characteristics

Category	Duration elasticity	
	\$600	\$300
Full sample	0.06	0.10
Age: [18,24]	0.07	0.11
Age: (24,34]	0.07	0.09
Age: (34,44]	0.07	0.08
Age: (44,54]	0.05	0.10
Age: (54,107]	0.04	0.12
No kids	0.07	0.1
Kids	0.05	0.12
Pre-covid unemployed	0.02	

Notes: This table calculates a duration elasticity separately in various groups, using the separate time-series of recall and new job-finding rates by each group and constant effects assumption described in Appendix E. For kids-related results we restrict to households under age 44 with and without kids to avoid confounding the effect of kids with the effects of age. We identify children using the size of EIP payments. We do not report an elasticity for the \$300 supplements for the pre-covid unemployed sample, since the sample sizes by January 2021 are too small to compute reliable job-finding rates for this group at \$300 onset.

B Additional Data Description

B.1 States Included

We divide states into three groups:

1. Fully included (15 states)
2. Partially included (29 states + DC)
3. Not in sample (6 states)

44 states plus DC are included in the benchmarking analysis in Section 2.2.1 and in the spending analysis in Section 3. Six states are excluded because we are unable to identify a transaction string that is unique to UI payments. In some states we are able to identify a transaction string that appears to include both UI payments but they are mixed with other transfer programs. The one exception to this is that we include California, which appears to use the same transaction string for UI, Disability Insurance, and Paid Family Leave. However, in public data, only 7% of California recipients of payments from the Employment Development Department in 2020 are receiving disability insurance (DI) or paid family leave (PFL). In addition, we drop anyone with benefits greater than the maximum weekly benefit for UI, which drops a large share of the DI and PFL recipients.

Relative to the 44 state sample, additional data cleaning is needed for the job-finding analysis in Section 4 because we need to know the state’s withholding rate for income taxes (discussed below) and we drop workers receiving more than the state’s weekly maximum benefit and less than the state’s weekly minimum benefit. We attempted to clean data for the 16 largest states by number of UI payments, which collectively account for 97% of the UI payments in the bank data. We succeeded for all of the states except Florida, which accounts for 6% of the UI payments and Colorado, which accounts for 2% of the UI payments. Florida has a high rate of false exits because of PUC overpayments in May 2020 (enough to distort the aggregate time-series for job-finding) and Colorado has a 20% weekly exit rate at the start of December 2020 in what is presumably a error. The set of states that are included in the job-finding analysis therefore accounts for 89% of the total UI payments in the bank data.

Finally, we make two further restrictions for some parts of the job-finding analysis. In the difference-in-difference expiration analysis in Section 4.2, we exclude states for which we are unable to separate LWA from regular payments (LWA payments make it impossible to measure the true date of exit): Texas, Connecticut, Louisiana, and Wisconsin. The remaining sample is CA, GA, IL, IN, MI, NJ, NY, OH, OR, and WA. In the difference-in-difference onset analysis, we exclude Oregon and Michigan. Oregon has a 25% weekly exit rate at the start of December 2020 in what is presumably a bug. Michigan sees a very high share of its low-benefit workers exit UI receipt because of the PUA/PEUC cliff. Unfortunately, the data cleaning procedure we use for other states to handle the cliff (described in footnote 15) is not effective in Michigan because there still is a sharp increase in measured exits even after data cleaning. The remaining sample is CA, CT, GA, IL, IN, LA, NJ, NY, OH, TX, WA, and WI. Finally, for analysis which requires a consistent job-finding rate throughout the entire sample period (e.g., Figure 3 and Section 4.1) we use the nine states (CA, IL, IN, MI, NJ, NY, OH, OR, and WA) that are present in both the expiration analysis and the onset analysis.

B.2 Unemployment and Employment Spells

We measure unemployment insurance spells (henceforth “unemployment spells”) using the payment of unemployment benefits. An unemployment spell starts with a worker’s first benefit payment in the sample frame, which is January 2019. In most states, a spell ends when a worker has three consecutive weeks with no benefit receipt. In states which pay benefits every other week, we instead define a spell end as four consecutive weeks without benefit receipt. Figure A-3c shows that using alternative numbers of consecutive weeks to define the end of a spell changes the level of the exit rate but not aggregate dynamics (particularly the depressed job-finding rate during the pandemic).

We measure employment outcomes using receipt of labor income paid by direct deposit. An employment spell begins with a worker’s first paycheck from an employer. We identify employers using a version of the transaction description associated with a payroll direct deposit which is purged of personal identifying information (see Ganong et al. 2020 for additional details). An employment spell ends (henceforth a “separation”) if a worker has five consecutive weeks with no paycheck from that employer. We define a separation as being associated with an unemployment spell if a worker has a separation between eight weeks before and two weeks after the start of an unemployment spell. This eight week lag allows for time for UI claims to be filed, processed, and paid, while the two week lead accounts for the fact that last paychecks can be paid after the date of last employment. 55 percent of benefit recipients have a detected separation at the time of benefit receipt.

We do not detect separations for every benefit recipient for two reasons. First, the JPMCI data do not include labor income paid via paper check or direct deposit labor income without a transaction description that mentions payroll or labor income. Second, in some cases more than eight weeks elapse between the last paycheck and first benefit payment; this scenario can arise if a state UI agency is slow to process a worker’s benefit claim or if a worker does not file for benefits immediately after separation.

We combine information on unemployment and employment spells to separate UI exits to a new job from UI exits to recall, which is when an unemployed worker returns to their prior employer. We are able to observe recalls only for unemployed workers for whom we also observe a job separation. We define a worker as having been recalled when they begin an employment spell with their prior employer between five weeks before and three weeks after the end of a benefit spell. We choose these thresholds based on the timing of job starts relative to the end of unemployment spells. The data on unemployment spells and employment spells jointly offer something comparable to the administrative datasets used to study unemployment in European countries (Kolsrud et al. 2018; Schmieder, von Wachter, and Bender 2012).

B.3 Sample Restrictions

Households in our analysis sample must meet two account activity screens: 1) at least five transactions per month and 2) annual pre-pandemic labor income of at least \$12,000. We impose these screens to focus on workers whose primary bank account is at Chase. For households that get benefits in 2020 (but not in 2019), we impose the transaction screen from Jan 2018-Mar 2020, and the labor income screen in 2018 and again in 2019. For households that get benefits in 2019, we impose the transaction screen for Jan 2018-Mar 2020 and the labor income screen in 2018. Among households that meet the activity screens, 11.6 percent receive unemployment benefits at some point during the pandemic.

This is lower than the rate for the U.S. as a whole, primarily because JPMCI only captures benefits paid by direct deposit. The Census Household Pulse Survey shows that 31 percent of households with at least one working-age person received UI benefits between March 13, 2020 and the end of October 2020. Finally, we limit the sample to customers who are present in the sample from January 2020 through March 2021 with positive inflows, positive outflows, and non-null account balance in every month.

The narrowest sample we use is a sample of customers who meet the account activity screens described in the prior paragraph and receive benefits from one of nine states (CA, IL, IN, MI, NJ, NY, OH, OR, and WA). Table B-1 shows summary statistics on the number of unemployed households in this sample. Figure B-1 shows that aggregate unemployment surges at the start of the pandemic and then declines as the economy recovers. Figure B-2 shows that this pattern is present in all of the 25 largest states in the sample. We also analyze data on a random sample of 187,000 *employed* workers who meet the transaction and labor income screens for 2018 and 2019, do not ever receive UI benefits in 2019 and 2020, and do not have a job separation in 2020.

Table B-1: Unemployment Spells During the Pandemic

Month	Average active spells	Number of spell starts	Number of spell exits	Exit rate			Share			
				All	Start UI in April	Job separation observed	Continuously unemployed (since Apr)	Continuously unemployed (since May)	Unemployed repeatedly (since Apr)	Exit to recall
Jan	25,752	11,814	10,666	49%	–	–	–	–	–	25%
Feb	25,504	9,520	8,946	38%	–	31%	–	–	–	25%
Mar	27,242	12,941	9,592	38%	–	30%	–	–	–	29%
Apr	120,360	153,175	16,864	16%	–	9%	100%	–	0%	31%
May	239,353	142,534	68,823	29%	19%	20%	70%	100%	3%	68%
Jun	260,224	77,910	73,699	29%	27%	25%	52%	81%	21%	76%
Jul	251,781	59,620	62,869	24%	19%	20%	44%	68%	34%	64%
Aug	233,503	44,297	71,389	31%	20%	23%	40%	61%	45%	52%
Sep	209,398	37,534	55,811	27%	17%	22%	37%	56%	57%	52%
Oct	194,068	36,077	49,028	25%	14%	18%	34%	52%	61%	40%

Notes: This table shows the number of unemployment spells in our data. The number of active spells is the monthly average, while the spell exits and starts are the sums for each month. Continuously unemployed are uninterrupted spells since April or May. The share of repeated unemployed workers is calculated since the beginning of the pandemic in April. Exit to recalls are workers returning to their previous employer.

B.4 Additional Variable Detail


We measure age as the age of the primary account holder (the first name listed on the bank account) at the start of an unemployment spell.

The Economic Impact Payments (EIP) authorized by the CARES Act had maximum amounts of \$1200 per single adult (and \$2400 per married filing jointly) and \$500 per child. For the subset of people who receive an EIP by direct deposit, we can infer the number of children in the household from the EIP amount. If a household receives multiple EIPs, we use the value of the first EIP. If a household does not receive an EIP by direct deposit, but deposits a paper check whose sum is a multiple of \$1200 (or \$2400) and a multiple of \$500, we infer the number of children from that.


JPMCI has hand-categorized firms into 20 industry groups based on NAICS codes for approximately 2,000 employers associated with the most bank accounts. Some households have multiple labor income streams in their bank account. We assign households to industries using the firm that paid them the most in the three months prior to UI receipt. Industry is available for about one-quarter of

Figure B-1: Patterns of Unemployment Insurance Receipt

Number of active spells



Number of benefit recipients

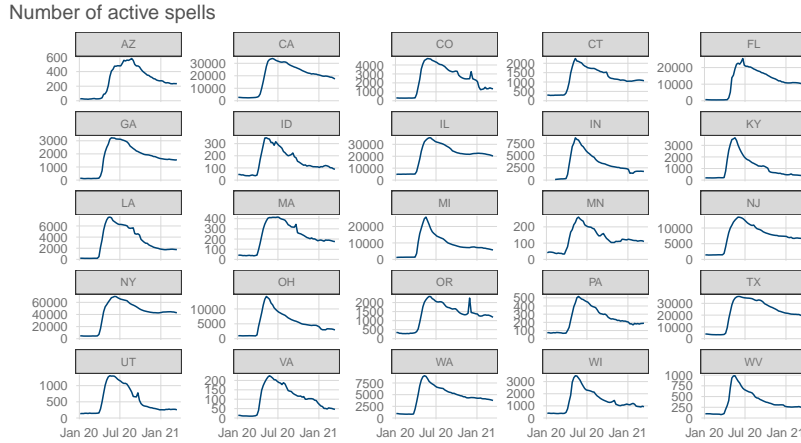


Notes: The first panel of this figure reports the number of active unemployment spells by week in JPMCI data. The second panel shows the number of households starting unemployment and leaving unemployment for new jobs and for recall (i.e., returning to their former employer).

UI spells.

To limit the influence of high income, high spending, and high asset households on means and MPC estimates, income, spending, and balances are winsorized at the 90th percentile.

Figure B-2: Patterns of Unemployment Insurance Receipt by State



Notes: This figure shows the number of active unemployment spells by week in JPMCI data for the 25 largest states in the sample.

B.5 Measuring Weekly Benefit Amount

Some individuals have UI inflow amounts that vary from week to week, for example, due to backpay. We require a single weekly benefit amount to calculate a percentage change in benefits from a supplement. We estimate the benefit amount $b_{i,pre}$ as the median benefit paid to an individual in the two-month period before the \$600 expiration or the two-month period before the \$300 onset. We then drop the first payment and the final payment and compute a median for the remaining weeks. Some states (CA, FL, MI, CO, TX, and IL) pay benefits once every two weeks and so we divide the median payment by two to capture the amount paid per week.

We measure state UI minimum and maximum benefits using the January 2020 “Most Recent Significant Provisions of State UI Laws” publication from the Department of Labor. If a state pays a dependent allowance we use the maximum benefit with dependents and the minimum benefit without dependents. We measure each state’s rate of income tax *withholding* using Whittaker and Isaacs (2022).

For the difference-in-difference analysis, we estimate workers’ regular weekly benefit amounts in the absence of any supplements. For the \$300 reinstatement, we estimate workers’ regular weekly benefit amount as $wba_i = b_{i,pre}/(1 - withholding)$. For the \$600 expiration, we estimate workers’ regular weekly benefit amount as $wba_i = \frac{b_{i,pre} - (1 - withholding)600}{1 - withholding}$. California and New Jersey did not withhold from the supplement so we instead use $wba_i = \frac{b_{i,pre} - 600}{1 - withholding}$.

We limit the sample to workers with plausible regular weekly benefit amounts wba_i . Define each state’s minimum weekly benefit as b_{min} and maximum as b_{max} . We keep workers with $wba_i \in [(1 - withholding)b_{min}, b_{max}]$. These restrictions will remove customers who have a median payment that includes substantial backpay.

Calculating wba_i requires knowing whether a worker decided to withhold, but we generally do not observe withholding at the worker level. Because more than 50% of UI recipients withhold in every state in our sample, our default assumption is that workers *are* withholding at the rates reported in Whittaker and Isaacs (2022). However, if a worker has $wba_i > (1 - withholding)b_{max}$ then the withholding assumption implies that they are receiving an invalid weekly benefit amount.

Thus, for workers with $wba_i \in ((1 - withholding)b_{max}, b_{max}]$, we recalculate wba_i assuming that $withholding = 0$.

Recall that our object of interest is the *change* in benefits from the expiration or onset of a supplement, which we construct in equation (3). For the \$300 reinstatement, we estimate equation (3) as $PctChange_i = \frac{2 \times 300}{2wba_i + 300}$. For the \$600 expiration, we estimate equation (3) as $PctChange_i = \frac{-2 \times 600}{2wba_i + 600}$.

C MPC Robustness

The MPC estimates in Section 3.2 are robust to limiting the sample to households for whom we are likely to observe a more complete lens on spending, to alternative summary statistics, and to alternative measures of spending. For example, many Chase customers have non-Chase credit cards. We will understate MPCs if these households increase spending on non-Chase credit cards when receiving supplements. Table A-2 shows that when we limit the sample to households with no ACH payments made towards non-Chase credit cards, MPCs are three to four cents higher. This suggests that the presence of non-Chase credit cards leads us to slightly understate MPCs in our baseline sample. As a second example, Chase customers may make some debt payments using paper checks, which our methodology will misattribute as spending (because we do not observe the content of account outflows where the payment method is paper check), leading us to overstate the MPC. Table A-2 therefore reports a robustness check which limits the sample to households which use their account for debt payments via ACH—thereby mitigating the concern that paper checks are mistakenly including some debt payments as spending—and finds that MPCs are indeed slightly (zero to four cents) lower than the baseline estimates but still large. Finally, although our baseline MPC estimate follows much of the prior literature by looking at means, Table A-2 shows that estimates using the median change in spending and income results in MPCs that are two to seven cents higher.

In addition, using the narrow *card and cash* measure of spending which is less subject to concerns about misclassification still delivers large spending responses. Table C-1 shows that MPCs out of this more narrow spending measure are mechanically slightly smaller but the elasticity of this narrow spending measure to benefits is actually larger. For example, Table A-1 shows that the narrow measure drops 40% of spending, but the MPC is reduced across the three research designs by only 24 to 33%. This implies that the responses observed in card and cash spending are proportionally larger than the responses in total spending and so when expressed as elasticities, the response of this subset of spending is larger than the response of total spending. This suggests that the large spending responses that we find are not driven by misclassification.

Table C-1: MPC Robustness to Spending Measure

Episode	Spend total		Spend card and cash	
	MPC	Elasticity	MPC	Elasticity
Waiting for benefits	0.43	0.60	0.32	0.74
Expiration of \$600 supplement	0.30	0.41	0.20	0.47
Onset of \$300 supplement	0.26	0.37	0.20	0.48
Lost wages assistance	0.32	0.44	0.21	0.50

Notes: This table shows the robustness of our MPC results to the measure of spending. The first columns show MPCs and elasticities for spending (total), which is our preferred specification. The third and fourth columns recompute MPCs and elasticities for spending (card and cash) instead of spending (total).


D Lost Wages Assistance Spending Responses

As a robustness exercise, we examine the spending response to the temporary \$300 supplement known as “Lost Wages Assistance” (LWA). This policy was announced in August and applied to unemployment spells in August but was largely paid out in September. To estimate the effects of these payments, we exploit state-level variation in their timing. In particular, most states paid out the \$300 payments in September, but New Jersey made these payments in October. Our empirical specification thus compares the spending of a treatment group which receives payments in September to a control group consisting of unemployed households in New Jersey, who receive payments in October.

Figure D-1 shows that spending (card and cash) of the treatment and control group track each other closely until the treatment group receives these \$300 payments, at which time their spending rises.

MPCs using the IV regression in equations (1) and (2) are shown in Table A-2, with a control group of unemployed households in New Jersey, and a treatment group of households in other states. These payments were typically made over a six week period beginning in September, so the treatment period is the 6 weeks from 9/6 to 10/11 and the pre-period is the 6 weeks from 7/26 to 8/30.

Figure D-1: Impact of Temporary LWA Supplement on Spending (Card Cash)



Notes: This figure measures the causal impact of the temporary LWA supplement on spending. The treatment group is eight states which paid the supplement in September. The control group is New Jersey, which paid the supplement at the end of October. The dependent variables are mean benefits and mean spending, measured as a change relative to the third week of August.

E Calculating Duration Elasticities

Call the total exit hazard observed in the data (which includes the effect of the supplement when it is in place) $\lambda_{t,\text{with supp}} = e_t + \text{recall}_t$, with observed new job-finding rate e_t and observed recall rate recall_t . We assume e_t and recall_t are constant at their sample averages after the end of the observed data.

We then construct a counterfactual total exit hazard with no supplement: $\lambda_{t,\text{no supp}} = \lambda_{t,\text{with supp}} + \tau_t \times I_t(\text{supp} = \text{on})$, where τ_t is an estimate of the effect of a given supplement on the job-finding rate at date t , and $I_t(\text{supp} = \text{on})$ is an indicator for whether a supplement is on or off in week t . For the statistical exercises assuming constant effects, τ_t is assumed constant at values from Table 2, while in the model τ_t is calculated from the full model dynamics. That is, the simple statistical counterfactual without supplements just shifts up the observed job-finding rate by the constant amount τ while the supplement is in effect while in the model τ_t varies with any dynamic forces. Note that the recall rate in period t is assumed to be the same with and without the supplement, so the shift in the total exit rate λ is given by just the change in the new job-finding rate τ_t .

Given $\lambda_{t,\text{with supp}}$ and $\lambda_{t,\text{no supp}}$ we can compute expected unemployment durations with and without the supplements and thus the duration elasticity by converting the job-finding hazards λ to a survival function. Specifically, let $S_{t,\text{with supp}} = \prod_{j=1}^t (1 - \lambda_{j,\text{with supp}})$ and $S_{t,\text{no supp}} = \prod_{j=1}^t (1 - \lambda_{j,\text{no supp}})$ be the cumulative survival functions with and without supplements. The expected duration with supplements is then given by

$$ED_{\text{with supp}} = \lambda_{1,\text{with supp}} + \sum_{j=1}^{\infty} (\lambda_{j+1,\text{with supp}}) (S_{j,\text{with supp}}) (j + 1),$$

and the average duration without supplements is given by

$$ED_{\text{no supp}} = \lambda_{1,\text{no supp}} + \sum_{j=1}^{\infty} (\lambda_{j+1,\text{no supp}}) (S_{j,\text{no supp}}) (j + 1).$$

Note that these expected durations will depend on both the time-series of job-finding and the length of time that the supplement is on (i.e., the number of weeks in which $I_t(\text{supp} = \text{on}) = 1$). The latter will vary for cohorts that enter unemployment at different dates since cohorts entering closer to supplement expiration will have a shorter period of time in which supplements are in effect. This means that the statistical-based duration elasticity will differ for different cohorts but will be maximized in most cases for cohorts starting unemployment in the same week that supplements start. For this reason, we conservatively report duration elasticities for an unemployed cohort starting April 1, 2020 for the \$600 supplements and January 1, 2021 for the \$300 supplements.

Given expected unemployment durations, the duration elasticity is then given by

$$\frac{\frac{ED_{\text{with supp}}}{ED_{\text{no supp}}} - 1}{\text{Supp Size/Regular Benefit Size}}.$$

F Additional Model Details and Figures

F.1 Additional Model Details

This section describes the model setup and calibration in additional detail. Each month, households choose consumption c and savings a with return r and a no-borrowing constraint $a \geq 0$ to maximize expected discounted utility $E \sum_{t=0}^{\infty} \beta^t U(c)$. When employed, household i has constant wage w_i which differs across households but is constant over time. Employed households become unemployed with constant probability π . When unemployed, a household finds a job at wage w_i with probability $f_{i,t} = recall_t + search_{i,t}$, where $recall_t$ is a common exogenous recall rate and $search_{i,t} \in [0, 1 - recall_t]$ is household i 's endogenous choice of search effort. Search effort induces disutility $\psi(search_{i,t})$. Recall requires no search effort or disutility. When households are unemployed, they receive unemployment benefits as well as additional secondary income proportional to the lost job: hw_i .

Income for an unemployed household depends on aggregate UI policy and whether they are waiting for benefits. Regular benefits last 6 months. Benefits for a household newly unemployed during the pandemic last 12 months.⁵¹ Benefit levels depend on the current aggregate UI supplement in place: $m \in \{0, 300, 600\}$. To speak to our empirical research design, we allow for the possibility that an unemployed household may face delays in receipt of UI and in turn later receive backpay. This means unemployed households can be in one of four receipt statuses: $d \in \{normal, delayed, backpay, expired\}$.

The regular benefit policy is intentionally simple: unemployed households receive benefits which replace a constant fraction b of w_i . When available, supplements add m to these baseline benefits. This means that an unemployed household getting benefits without delay receives $bw_i + m$.

Unemployed households can also be in a delayed receipt status and not currently receiving benefits if $d = delayed$. In this case, current earnings are given by hw_i . When households exit this status, they receive backpay equal to $\alpha(bw_i + m)$, where α is chosen to match ‘‘backpay’’ observed in the data. When regular benefits expire after 12 months, income again drops to hw_i . This means that total earnings y for a household with wage w_i , employment status s , supplement m , and delay status d are given by:

$$y(w_i, s, m, d) = \begin{cases} w_i & \text{if } s = e \\ bw_i + m + hw_i & \text{if } s = u \text{ and } d = normal. \\ hw_i & \text{if } s = u \text{ and } d = delayed. \\ \alpha(bw_i + m) + hw_i & \text{if } s = u \text{ and } d = backpay. \\ hw_i & \text{if } s = u \text{ and } d = expired. \end{cases} \quad (F1)$$

The economy begins in a steady-state UI policy environment with $m = 0$ and $d = normal$. Households expect UI policy will never change. Beginning from this initial steady state, the economy is hit by policy changes which mimic UI policy changes over the pandemic. In April 2020, the economy switches from $m = 0$ to $m = 600$ and remains in this state for 4 months. In August 2020, it switches to $m = 0$. In January 2021, it switches to $m = 300$, and in September 2022 it switches back to $m = 0$.

This describes the evolution of actual policy through this period, but we must also specify expectations. We assume the initial switch from $m = 0$ to $m = 600$ is unanticipated.⁵² Once the 600 is

⁵¹To simplify the computational setup, we assume that pandemic benefits only obtain for the first unemployment spell and that when a household returns to employment they return to the regular benefits policy for future UI spells.

⁵²For computational tractability, we assume employed households continue to expect regular benefits after the pandemic starts until they actually become unemployed. Since we focus on a cohort of unemployed households beginning

implemented, households know for sure that it will last at least 4 months. In our main results, we consider two different specifications for expectations about m after these 4 months.⁵³ In the perfect foresight specification, households correctly expect that m will revert from 600 to 0 in August. In the alternative specification, households instead expect that $m = 600$ for the duration of their remaining benefit spell and are then surprised in August when it expires. Once $m = 600$ expires in August, households expect that $m = 0$ forever. For the \$300 weekly supplements, we study a newly unemployed household in November of 2020. They either anticipate or are surprised that the $m = 300$ supplement begins in January, 2021.⁵⁴ Once the \$300 supplement begins, households anticipate that supplements will expire in September 2021.⁵⁵

Expectations about UI delay are simpler. Households who are in the $d = \textit{nodelay}$ state anticipate that they will remain in this state. That is, households do not anticipate delays in benefit receipt. When households are in the $d = \textit{delay}$ they always assume that they will be in $d = \textit{backpay}$ next period and that they will be in $d = \textit{nodelay}$ the period after that. That is, households always anticipate that benefit delays will be resolved next month. However, even though households always anticipate that delays will be resolved once they enter this state, the realized length of $d = \textit{delay}$ can extend for multiple periods. That is, just as households are surprised by initial delays in benefits, they can also be surprised by a longer than expected waiting period. In our main simulations, the actual benefit delay lasts two months to match what we observe in the waiting design. We discuss various other specifications for delays in Section F.2.3. Conclusions are quantitatively similar, although our baseline specification is a slightly better fit to the data.

Households anticipate a constant recall rate throughout the pandemic, although results are similar if we instead assume perfect foresight over the actual recall rate. We deal with pandemic effects in two ways: First, we focus on the evolution of unemployed households relative to employed households in both model and data. This means any effects of the pandemic which affect all households equally are effectively removed. Second, we directly model several pandemic events. We introduce a one-month discount factor shock to all households in April 2020, which we calibrate to match the decline in spending for employed households during the pandemic.⁵⁶ Since we focus on the behavior of unemployed households relative to employed households, this shock has little effect on our conclusions, but it means that we do a better job of hitting absolute spending and liquidity changes over the pandemic rather than just matching relative changes. We also introduce additional one-time unanticipated transfers to replicate stimulus checks in April 2020 and January 2021 as well as LWA payments in September 2020, but this again has little effect on our conclusions.

Letting n represent the expected number of periods until $m = 0$, the household optimization problem of a household unemployed during the pandemic can be written as:⁵⁷

in April, this choice has little practical effect beyond simplifying computation.

⁵³Intermediate versions of expectations unsurprisingly produces results between these two version.

⁵⁴We simulate a separate cohort of unemployed households becoming unemployed in November so that we do not have to also model extensions of the duration of regular benefits which happened periodically throughout the pandemic.

⁵⁵Versions of the model where the extensions to September legislated in March 2021 were a surprise produce spending and search patterns at odds with the data and mildly complicate the analysis. Similarly, since our analysis ends before the summer of 2021, we abstract from the early expiration of benefits in some states during that period.

⁵⁶Using a sequence of discount factor shocks introduces additional complication but does not change the results much since only April 2020 exhibits a very sharp swing in spending.


⁵⁷For simplicity, this notation ignores the discount factor shock. In April 2020, the model is solved using a different β for a single period before transitioning back to this specification

$$\begin{aligned}
V_u(a, s = u, m, d, n) &= \max_{c, a', search} U(c) - \psi(search) \\
&\quad + \beta E_{m', n', d'} [(search + recall)V_e(a', s = e, m', d', n') \\
&\quad + (1 - search - recall)V_u(a', s = u, m', d', n')] \\
&\quad s.t. \\
&\quad a' + c = y(w_i, s, m, d) + (1 + r)a, \\
&\quad a' \geq 0, \\
&\quad \text{Equation F1,} \\
&\quad \text{and expectations of } s', m', d', n'.
\end{aligned}$$

The value function of an unemployed household pre-pandemic is analogous except that m is always 0, d is never delayed and n lasts for 6 instead of 12 months. The value function of an employed household is also analogous except that $s = e$ so $y(w_i, s, m, d) = w_i$, they have no search decision and they transition to the regular pre-pandemic unemployment value function with exogenous separation rate π .

We set the annual interest rate $r = .04$. We assume that the utility function is given by $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$ and set $\gamma = 2$. We set the exogenous separation probability $\pi = 0.028$ to match pre-pandemic transitions from Krueger, Mitman, and Perri (2016). We set the expected recall rate $recall_t$ to be constant at its average value of 0.08 but use the actual evolution over the pandemic where relevant. We set $b = 0.21, h = 0.7, m = 0.35, \alpha = 2.35$ to match household income series for the waiting and receiving UI groups over the pandemic. Figure F-1 illustrates the environment by showing income for the unemployed relative to employed in the model and data for a newly unemployed worker with and without a benefit delay in our calibrated model. We solve the model for five different w_i groups and choose the variation to match household income by five quintiles of the replacement rate in JPMCI data.

Figure F-1: Income: Model vs. Data



Notes: This figure shows income of unemployed (relative to employed) households in the model and data for unemployed workers receiving benefits immediately in April 2020 as well as those who face a delay in benefit receipt until June 2020.

We assume that $\psi(search_t) = k_0 \frac{(search_t)^{1+\phi}}{1+\phi} + k_1$ and pick the parameters of this search cost function in one of two ways. In our “pre-pandemic” calibration, we calibrate search costs to generate

a pre-pandemic job-finding rate of 0.28 and an elasticity of average unemployment duration to a small 6 month change in benefits of 0.5. This is the median estimate from the Schmieder and von Wachter (2016) meta-analysis. In our “best-fit” calibration, we instead calibrate search costs to target the time-series of job-finding over the course of the pandemic. Since we have two moments and three parameters, the pre-pandemic calibration is not identified without additional restrictions and so we impose $k_1 = 0$. If we instead impose the value of k_1 we estimate for the pandemic, qualitative conclusions are unchanged. The search cost parameters are most interpretable in terms of implications for job search elasticities. The calibrated pre-pandemic search cost parameters imply a job search hazard elasticity to a small benefit change for 6 months of -0.66 and the best-fit search cost parameters imply a hazard elasticity of -0.24.⁵⁸

We calibrate the discount factor in one of two ways. In the pre-pandemic calibration, we pick β so the model matches pre-pandemic evidence on the response of spending to stimulus checks summarized in Kaplan and Violante (Forthcoming). Specifically, we pick set $\beta = .99$ monthly to generate a 3-month MPC of 0.25 in response to a 500 stimulus check sent to all households. In our alternative best-fit calibration, we instead pick $\beta = 0.978$ to target the MPC out of UI payments in our waiting design.

We solve the model using the endogenous grid method with linear interpolation for policy functions off grid points. We use 100 grid points for assets distributed exponentially from 0 to 2000 times median household income. The model must be solved for several different benefit profiles with length up to 13 months (pandemic era benefits last for 12 months and then expire as an absorbing state; regular benefits last for 6 months) as well as the different delay statuses. We solve the model separately for each of the five wage groups. We solve for the value function for employment and regular unemployment benefits iterating to stationary policy functions and then solve for the pandemic-era temporary policies using backward induction from these stationary value functions. We similarly backward induct one period from the stationary value functions to solve for the solutions with the discount factor shock. To solve for optimal search, we use the first order condition given next period’s value functions of employment and unemployment. Given optimal policies, we then simulate the model for 1,000 households of each of each wage group (5,000 households total) and compute average statistics.

F.2 Model Robustness and Additional Model Results

This section discusses model extensions and robustness checks which are mentioned in the main text.


F.2.1 Time-Aggregation Issues

Our baseline model is monthly and assumes that benefits start immediately upon job loss: households are employed until March 2020 and then become unemployed and start receiving regular benefits plus \$600 supplements immediately in April. In the data, there are some high frequency changes around job loss not captured by this simple model. Our data sample selects households who receive their last paychecks at the end of March and receive benefits starting in April. Furthermore, most states did not begin paying supplements until the second half of April even though the amount of supplements covered benefit weeks for the entire month of April. What this means in practice is that in the data, households who become unemployed at the end of March see a small decline in income at the end of March and

⁵⁸Note that the hazard elasticity of 0.65 in the former case implies the targeted duration elasticity of 0.5.

start of April before benefits start, and then a jump up in income in the second half of April that makes up for the decline in the first half of the month. When aggregated to calendar months, this manifests as a small decline in income in March and a jump up in income in April as households start receiving supplements, but at the weekly frequency, this jump up in April is concentrated in the second half of the month. This means that although April income is above that when employed, this increase occurs primarily in the second half of the month and so does not allow for a full month of spending opportunities.

Figure F-2: Spending in Model that Accounts for Within-Month Time-Aggregation



Notes: This figure shows the behavior of spending in the model more closely matches high-frequency data patterns after accounting for time-aggregation issues.

To address both high frequency timing issues around the start of unemployment and the start of supplements, we make two changes while still maintaining the monthly model period for tractability.

First, we assume that both in regular times as well as during the pandemic, households have a decline in income in the month that they lose their job but do not start unemployment insurance until the following month. We choose this drop in income in the month of job loss to match the observed decline in income in the data. In practice, this drop in income is small since we define the month of job loss as the month when the last paycheck is received. However, this small decline in income in the month of job loss is sufficient to generate declines in spending in the month of job loss. Given this adjusted income process, we assume that households in this extended model lose their jobs in March rather than in April. This model change is sufficient to generate the spending declines in March 2020 observed in the data, as shown in Figure F-2.

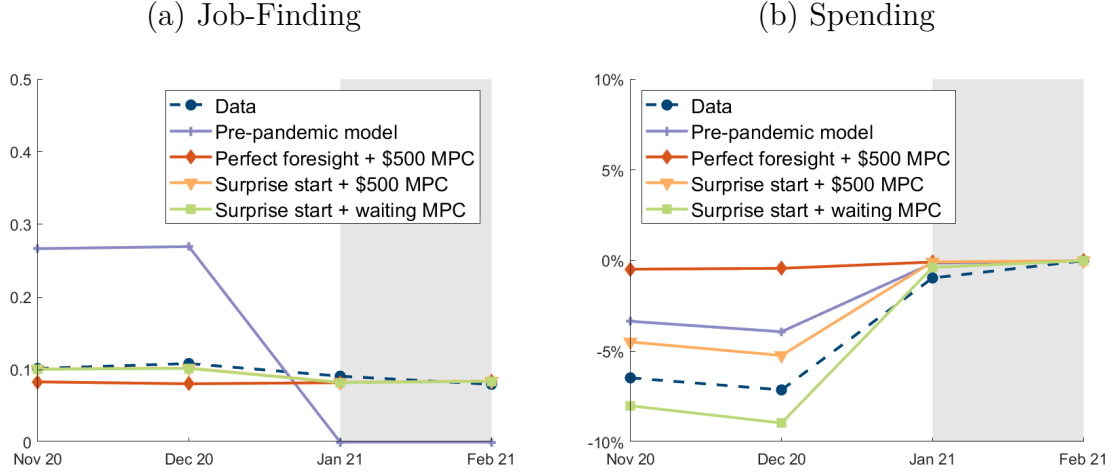
Second, we assume the income increase arising from supplements occurs in the second half of April. Specifically, we assume that in the first half of the month, households spend as if they are receiving the regular benefits profile but not supplements. Concretely, we compute spending in April in time-aggregation model as an equal weighted average of spending under the regular benefits profile and spending under the benefits profile with supplements. Figure F-2 shows that the model with this extension is a good fit to spending in April. Since it is more parsimonious, we use the simpler model in the main text, but this figure shows that that all of our conclusions about our “best fit” model relative to alternative models also hold in this extended model so our results are unchanged if we

complicate the model to match these high frequency patterns around the first two weeks of job loss.

F.2.2 Best Fit Model-Additional Results

The main text focuses on the model-fit for the \$600 supplements, but Figure F-3 shows that our conclusions also apply for the model fit to the \$300 supplements.

Figure F-3: Job-Finding and Spending Responses to \$300 Supplement: Models vs. Data



Notes: This figure repeats Figure 6 but showing models vs. data in response to the \$300 supplement.

We calibrate search costs in our best fit model to target time-series variation in the job-finding rate. Since our model features heterogeneity in wages it also has implications for job-finding over the wage distribution which can be compared to our difference-in-difference research design. Figure F-4 shows that the model calibrated to time-series evidence produces difference in difference results which align closely with those from the data. This means that the choice of targeting time-series vs. cross-sectional variation makes little difference for our model conclusions. This figure also provides additional support for the linearity assumption imposed by our empirical difference-in-difference research design.

Table F-1 also shows how employment distortions vary across group and that the employment effects of supplements were larger for lower wage workers.


Table F-1: Distortions by Group

	Bottom wage quintile workers	Top wage quintile workers
Average change in employment during \$600 supplement period	-0.93%	-0.39%
Average change in employment during \$300 supplement period	-0.51%	-0.27%

Notes: This table computes employment losses caused by supplements over the wage distribution in the best-fit model.


Figure F-5 shows that liquidity effects on job-finding after supplements expire are small demonstrating that dynamic liquidity effects imply little bias for our reduced-form specifications.

Figure F-4: Model vs. Data Dose-Response Difference in Difference



Notes: This figure compares the dose-response DiD for the \$600 expiration in our best fit model to the empirical DiD.

Figure F-5: Job-Finding Rates with and without \$600 supplements



Notes: This figure shows the job-finding rate in the best-fit model with a \$600 supplement compared to a counterfactual job-finding rate had there been no supplement. This shows that liquidity accumulation caused by the supplement slightly lowers the job finding rate even after the supplement expires.

In our main model specification, we target mean income, and the mean MPC from the waiting design and then compare results to mean spending series. We use means since the existing literature studying MPCs focuses primarily on regressions which identify the mean MPC, so it is easier to compare to this literature. Nevertheless, Table A-2 shows that the median MPC in the data is even larger. Figure F-6 shows that results are very similar for a version of the model which targets this median MPC, inputs median income series and compares to median spending series.

Figure F-6: Models vs. Data: Targeting Medians instead of Means



Notes: This repeats our main spending figures but targeting median instead of mean income and waiting MPCs.


F.2.3 Alternative Assumptions About Waiting Expectations

Our baseline model assumes that households who are waiting for benefits correctly anticipate that they will eventually receive them. In this section we show results under an alternative specification where waiting households think they will never receive any benefits.⁵⁹ Importantly, we recalibrate this model with alternative expectations so that it continues to hit the waiting MPC from the data.

Figure F-7 shows job-finding and spending in red for the waiting group in this model compared to our baseline model in purple. If households expect to never get benefits, the job-finding rate starts high and then drops dramatically when benefits start in June and then jumps back up to an intermediate value when supplements expire in August. Since we condition on receiving benefits in June in the data when defining this group and we measure unemployment exit based on ending unemployment benefits, we do not have an empirical job-finding rate to compare to for this group. Nevertheless,


⁵⁹We have also computed results in versions where households are uncertain about eventual benefit receipt, and the results of that model are in between these two specifications.

Figure F-7: Effects of Expectations During Benefit Delays



Notes: This figure compares the job-finding behavior and spending of households waiting for benefits in our baseline model (where waiting households expect to receive benefits) to that in an alternative model in which households expect to never get benefits. The baseline model is in purple and the alternative model is in red.

Figure F-8: Spending Responses in Calibrations with Different Expectations



Notes: This figure shows that the model calibrated to hit the waiting MPC when waiting households expect to never hit benefits generates a spending response which is below that in our baseline specification as well as the data.

this job-finding pattern, when combined with the levels of recall would imply that 40% of the group receiving a first benefit payment June 1 has already gone back to work prior to receiving this payment.

The right panel of this figure shows that when households do not expect benefits to ever come, they cut spending by more than in the model where households anticipate eventual benefit receipt. That is, households engage in less consumption smoothing and then have a larger jump in spending at benefit receipt. When households correctly anticipate that they will eventually receive benefits, the waiting MPC reflects only liquidity effects from the timing of benefits. When households instead


think they will never receive benefits and then are surprised when they receive them, the waiting MPC will capture the full effects of benefits including both liquidity and income effects. This means that the waiting MPC under our baseline assumption is relatively conservative in terms of reflecting spending responses to benefits. The waiting design MPC reflecting just liquidity effects is 0.43 while the full MPC out of benefits is 0.47. In contrast, when households are surprised by benefit receipt, the waiting design and full MPC are identical at 0.43.

Since spending of the waiting group falls more when households expect to never get benefits, this increases in the waiting MPC. Thus, in order to maintain the same waiting MPC of 0.43, this model requires a higher amount of patience. This in turn reduces the spending response to supplements for unemployed households who do *not* face delays. As shown in Figure F-8 this specification then falls a little bit short of the spending responses in the time-series data. This is one reason we do not use this specification for expectations in our baseline model. In addition, most households filing for benefits in this time period were indeed eligible, and so most households facing delay would receive benefits eventually.

F.3 Policy Counterfactuals – Additional Results

Figure F-9 repeats that the result that MPCs decline with supplement size does not depend importantly on the particular horizon over which MPCs are computed.

Figure F-9: Spending Effects at Different Horizons




Notes: This figure repeats the left panel from Figure 11, but showing MPCs at various horizons rather than just the MPC in the first month of the supplement.

Sending unconditional stimulus checks is an order of magnitude more expensive than providing temporary UI supplements of the same size, since unconditional stimulus checks flow to many more people. Thus in addition to comparing severance and stimulus checks of equal size like in the main text, it is interesting to compare the relative spending effects of severance to stimulus checks with the same aggregate cost rather than the same size. If the unemployment rate is $X\%$, then this means comparing the spending implications of a UI supplement of size Y sent to $X\%$ of the population to an unconditional stimulus check of size XY sent to the entire population. The solid line in Panel (b)

of Figure F-10 compares the quarterly MPC out of various sized UI transfers to the quarterly MPC out of equal cost stimulus checks, assuming an unemployment rate of 10%. The dashed line shows the relative marginal effectiveness of the last dollars of spending. This figure shows that up to a one month UI supplement of \$2,400 (\$600 weekly), there is more bang for the buck from expanding UI supplements than from sending stimulus checks. However, as supplements become larger, the relative marginal spending effect falls below 1 and it becomes more effective to spend the marginal dollar on stimulus checks. This is because by this point, UI supplements are already large enough that they have substantially relaxed liquidity constraints, while much smaller equal cost stimulus checks have not.

Figure F-10: UI Severance vs. Untargeted Stimulus – Equal Cost



Notes: This figure shows quarterly spending responses to UI severance as well as to one-time stimulus checks of various sizes in a non-pandemic environment with normal liquidity levels. We compute these responses in the best fit model with discount factor heterogeneity. Note that we calibrate the degree of heterogeneity in this model so that it still produces a quarterly MPC out of \$500 stimulus checks of 0.25. The right panel reports the ratio of the MPC out of UI supplements to the MPC out of a smaller stimulus check of equal cost. Solid lines compute MPCs out of the entire transfer while dashed lines compute MPCs out of the last \$50 of transfers.

References

- Conley, T.G. 1999. “GMM estimation with cross sectional dependence.” *Journal of Econometrics* 92 (1):1–45.
- Ganong, Peter, Damon Jones, Pascal Noel, Fiona Greig, Diana Farrell, and Chris Wheat. 2020. “Wealth, Race, and Consumption Smoothing of Typical Income Shocks.” WP 27552, NBER.
- Kaplan, Greg and Giovanni L. Violante. Forthcoming. “The Marginal Propensity to Consume in Heterogeneous Agent Models.” *Annual Reviews of Economics* .
- Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn. 2018. “The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden.” *American Economic Review* 108 (4-5):985–1033.
- Krueger, D., K. Mitman, and F. Perri. 2016. “Macroeconomics and Household Heterogeneity.” In *Handbook of Macroeconomics*, vol. 2. Elsevier, 843–921.
- Schmieder, J. F., T. von Wachter, and S. Bender. 2012. “The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years.” *The Quarterly Journal of Economics* 127 (2):701–752.
- Schmieder, Johannes F. and Till von Wachter. 2016. “The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation.” *Annual Review of Economics* 8 (1):547–581.
- Whittaker, Julie M. and Katelin P. Isaacs. 2022. “Taxing Unemployment Insurance (UI) Benefits: Federal- and State-Level Tax Treatment During the COVID-19 Pandemic.” CRS report, Congressional Research Service.