Disemployment Effects of Unemployment Insurance: A Meta-Analysis *

Jonathan P. Cohen Peter Ganong

February 7, 2025

Abstract

We systematically review studies of how unemployment benefits affect unemployment duration. Statistically significant findings are eight times more likely to be published. Correcting for publication bias cuts the average elasticity by a third. Meta-analysis is a data-driven way to aggregate estimates across policy contexts and generalize sufficient statistics methods to compute the global optimal policy. Although existing consumption drop-based approaches typically imply an optimal replacement rate near zero, our corrected estimates imply an optimal replacement rate of 28% in the US. We are unable to reject the hypothesis that the "micro" elasticity is equal to the "macro" elasticity.

Keywords: Unemployment Benefits, Publication Bias, Meta-analysis, Baily-Chetty

JEL Codes: C13, E24, E64, J64, J65

*Cohen, MIT: jonpcohen@gmail.com. Ganong, UChicago and NBER: ganong@uchicago.edu (corresponding author). This research was conducted prior to the author's employment by Amazon. This paper is not sponsored or endorsed by, or associated with Amazon or any of its subsidiaries or affiliates. The views, opinions and positions included in this paper are the author's own and do not reflect the views, opinions and positions of Amazon. We thank Miguel Acosta, David Autor, Isaiah Andrews, Raj Chetty, Manasi Deshpande, Amy Finkelstein, Tomáš Havránek, Jonathan Heathcoate, Nathan Hendren, Simon Jaeger, Jonas Jessen, Max Kasy, Henrik Kleven, Bruce Meyer, Emi Nakamura, Pascal Noel, Matt Notowidigdo, and Charlie Rafkin for helpful conversations. We thank the staff of the Congressional Budget Office for their interest in the employment effects of unemployment insurance as well as seminar participants at the OECD, the Chicago Fed, and the NBER Public Economics group for useful suggestions. We thank Josephine Dodge, Avik Garg, John Spence, Nicolas Wuthenow, Fatima Yusuf, and Madeline Zuckerman for excellent research assistance.

1 Introduction

How much do more generous unemployment benefits prolong unemployment? Because the effect is not a universal constant, aggregating studies has two key advantages over relying on a single study. First, for understanding treatment effect heterogeneity, it may be difficult or impossible for a single study to contain the necessary variation. For example, knowing how responses vary with the level of benefits requires identifying variation at different baseline benefit levels. Second, for establishing external validity to a specific context, it is usually impossible to find a prior published paper in a closely matching context.

In this paper, we re-examine the published literature on how unemployment benefits impact unemployment duration, applying meta-analysis methods to 57 prior studies. We have four main findings.

First, publication bias toward statistically significant findings is pervasive. If researchers and journals apply more scrutiny to statistically insignificant findings, then estimates of small elasticities with large standard errors will be censored from the published record. We document evidence of this: estimates in the middle and top tercile of standard errors are about twice as large as estimates in the bottom tercile of standard errors. The published record overstates the mean of the latent elasticity distribution. Using a correction from Andrews and Kasy (2019), we find that statistically significant positive findings are 8 to 13 times more likely to be published. In addition, because many factors—not just the standard error—may influence a study's estimated elasticity, we conduct a meta-analysis using Bayesian model averaging (Irsova et al., 2023). After correcting for publication bias, we find that the elasticity with respect to replacement rates is 50% smaller using the first method and 37% smaller using the second method. The corrections to the elasticity with respect to potential benefit duration are even more drastic.

Second, elasticities meaningfully vary with the policy context. For the average worker in the US, the predicted publication-biased-corrected elasticity to a replacement rate change is 0.34 and to a change in potential benefit duration is 0.23. Elasticities are higher when benefits are more generous. In a policy context like Florida with limited benefits, the predicted elasticity to a replacement rate change is 0.29 and to a change in potential benefit duration is 0.14. In contrast, in a policy context like France with generous benefits, the predicted elasticities are 0.52 and 0.80 respectively.

Third, meta-analysis can generalize sufficient statistics methods to compute the global optimal policy. We re-evaluate the optimal replacement rate formula of Baily (1978) and Chetty (2006). Meta-analysis increases the credibility of the formula's elasticity input in three ways: correcting for publication bias, predicting a context-specific elasticity, and providing a data-driven approach to calculate the global optimum even if the reform would be large. Kleven (2021) discusses how sufficient statistic formulas are typically used only to assess *small* policy changes, but we show how heterogeneity in the sufficient statistic estimated using across-study moments can be a useful reference for *large* policy changes. Using this formula and auxiliary parameter estimates in the literature, we estimate an optimal replacement rate of 28% in the present-day US context. This is below the median replacement in all US states. But had we not corrected for publication bias, we would have reached the same conclusion as several other prior papers finding it optimal to have no benefits at all.[1]

Fourth, we find no evidence of a difference between the "micro" and "macro" duration elasticity. An active theoretical and empirical literature seeks to understand and compare the causal effect of increasing benefits for a single worker (the "micro" effect) to the causal effect of increasing benefits for all workers (the "macro" effect).[2] In principle, these two effects could differ due to a number of different general equilibrium channels. In our meta-analysis, we identify five studies where treatment occurs at the market level and the research design therefore captures a macro elasticity. We find no systematic difference in the elasticities from these studies compared to those from the rest of our review. However, our estimates are consistent with a macro elasticity which is modestly above or modestly below the micro elasticity.

Our analysis builds on prior excellent literature reviews (Krueger and Meyer, 2002; Meyer, 2002; Schmieder and von Wachter, 2016; Lopes, 2022) by applying recent methods from the meta-analysis literature. Section 2 contains the study collection procedure; we follow

[1]For example, a seminal paper by Gruber (1997) finds the optimal replacement rate is 2% for a relative risk aversion coefficient of 2 and a constant duration elasticity. Focusing on the social insurance motive for benefits, similar conclusions appear in Setty and Yedid-Levi (2021) and Krusell et al. (2010).

[2]See, e.g., Hagedorn et al. (2016), Michaillat (2012), Landais et al. (2018), and Jessen et al. (2023).

the best practices of meta-analysis described in Havránek et al. (2020). Section 3 presents evidence of publication bias. Section 4 documents heterogeneity in elasticities based on study characteristics. Section 5 draws out implications for the optimal level of benefits. Section 6 discusses the micro vs. macro elasticity. Section 7 concludes.

2 Data

Unemployment insurance (UI) pays claimants for each week they remain unemployed. One policy parameter is the share of prior weekly earnings replaced by payments: the replacement rate (RR). Another policy parameter is the maximum number of weeks claimants can receive benefits: the potential benefit duration (PBD). The number of weeks claimants receive UI payments is referred to as covered duration, while the number of weeks claimants remain not employed is referred to as total nonemployment duration.

We aim to survey all microeconomic papers that estimate the causal effect of UI generosity (either PBD or RR) on the duration of unemployment spells (either covered or total nonemployment duration). We collect journal publications from Google Scholar following the guidelines compiled by the Meta-Analysis in Economics Research Network (Havránek et al., 2020). We limit to papers published by the time the search date of August 15, 2022. We limit to the first 1,000 Google Scholar results.[3] Of these 1,000 search results, 407 are journal publications, and 57 papers use microeconomic methods while reporting enough information to calculate a UI duration elasticity and standard error. Appendix A details how we construct a comparable duration elasticity across study methods.

The 57 sample studies are high-quality based on quantifiable metrics. 72% of the studies use a quasi-experimental research design. Figure B-1 shows it is much more common for recently published papers to use quasi-experimental identification strategies. The sample's publications are in influential economics journals. The 25$^{\text{th}}$ percentile and median impact factors correspond to Oxford Economic Papers and the Journal of Public Economics, re-

[3]Specifically, we use the software Publish or Perish with the following query: `duration "standard error" OR "standard errors" OR PBD OR "benefit duration" OR WBA OR "weekly benefit amount" OR "replacement rate" "unemployment insurance"`. In words, this requires the paper's text to contain the word "duration", the exact phrase "unemployment insurance", and at least one of the other phrases.

spectively. Half of the estimates are from such field journals, one-fifth of the estimates are from one of the "Top-5" general interest journals, 7% are from econometric methods journals, and the rest are from other general interest journals. Additional information on study characteristics and journal classifications is reported in Appendix A.

We attempt to restrict attention to the author's preferred specification. We rely on the paper's discussion to identify what constitutes the main estimate. In the absence of such a discussion, we choose the estimate in the earliest table with the maximal set of controls. If the paper studies both RR and PBD variation or studies both covered duration and nonemployment, then we collect more than one estimate. Some papers do not include an overall elasticity, instead displaying only group-specific elasticities.[4] Because this disaggregation choice is potentially germane to publication bias, we collect each group-specific main estimate. We identify 91 estimates and standard errors. Table D-1 and Table D-2 contain the full list of included estimates and their sources for PBD elasticities and RR elasticities, respectively. We use a slightly more disaggregated dataset when testing for publication bias than when analyzing systematic determinants of heterogeneity; the exact level of aggregation we use is described in Appendix A. The unadjusted mean PBD elasticity is 0.46, while the unadjusted mean RR elasticity is 0.43.

In addition to the main duration elasticities and their standard errors, we collect economic and methodological characteristics associated with each study. Section 4.1 describes these characteristics. Here, we highlight some key patterns. In terms of economic characteristics, there is substantial variation across studies in the pre-reform level of the average replacement rate. It varies from 27% to 90% with an interquartile range of 15%. In terms of methodological characteristics, Table C-1 shows that quasi-experimental methods are more common for PBD elasticities (90%) compared to RR elasticities (45%).[5]

[4]For example, all statistical tests in Bennmarker et al. (2007) are gender-specific. The paper's elasticities for men are consistently positive and statistically significant while for women are consistently negative, of the same absolute magnitude, and marginally statistically significant.

[5]RR formulas typically depend on prior earnings with a minimum and maximum amount. Prior to the proliferation of regression kink designs, cross-sectional identification strategies for the RR elasticity would include parametric controls for the benefit formula's running variable.

3 Publication Bias Evidence and Corrections

3.1 Graphical Evidence of Publication Bias

The key assumption underlying our preferred test is that a study's elasticity estimate and its standard error should be orthogonal. Power calculations are one reason why elasticity magnitudes might be related to standard errors. However, we argue that this is implausible in our meta-analysis given the studies' exclusively observational variation and pre-existing datasets. Absent publication bias, both small and large estimated elasticities should come with small and large estimated standard errors. Alternatively, publication bias generates a correlation between published elasticities and their standard errors by truncating the distribution of latent elasticities.

Figure 1 illustrates evidence of publication bias with increasing degrees of parametric structure: as a scatter plot, as a non-parametric density by standard error tercile, and finally as the mean elasticity by standard error tercile. In the absence of publication bias, the mean elasticities in the bottom panel should be vertically aligned.[6] Evidence of publication bias is the positive correlation between published estimates and their standard errors. The vertical dashed line on the figure shows the mean unconditional elasticity. However, the figure shows that the mean elasticity is 0.24 for estimates in the bottom tercile, 0.49 in the middle tercile, and 0.58 for estimates in the top tercile. The estimates for each group are sufficiently precise that for two out of three terciles we can reject hypothesis that the group-specific means are equal to the overall mean.[7] This suggests that researchers and journals are indeed censoring small elasticity estimates with large standard errors.

In addition to the positive correlation, the figure also shows two other types of evidence of publication bias. First, 95% of elasticities are positive.[8] Many fall just above 0, but few fall just below 0.[9] Second, Figure 1a shows that there are many studies just on the statisti-

[6]The location of the line is not important; what matters is that the estimates grouped by standard error tercile should have estimates—allowing for sampling variability—that are on this line.

[7]If we drop the two noisy estimates with elasticity < -1, then we reject the hypothesis for all three groups.

[8]This 95% estimate includes some negative estimates which are omitted from the figure for legibility. In the figure itself, 96% of the elasticities are positive.

[9]A strong prior that the elasticity should be strictly positive is a potential mechanism behind publication bias. Indeed, Lancaster and Nickell (1980) write: "*The effect on unemployment durations of the relative level of unemployment benefit is consistent both with theoretical reasoning and a number of previous studies.*".

cally significant side (below the diagonal gray line) but few studies just on the statistically insignificant side (above the line). Figure B-2 emphasizes this by plotting the t-statistic, or the ratio between the elasticity (x-coordinate in Figure 1a) and standard error (y-coordinate in Figure 1a). Excess density is apparent to the right of 1.96.

Figure 1: Descriptive Evidence of Publication Bias

(a) Individual Estimates of Elasticities and Standard Errors

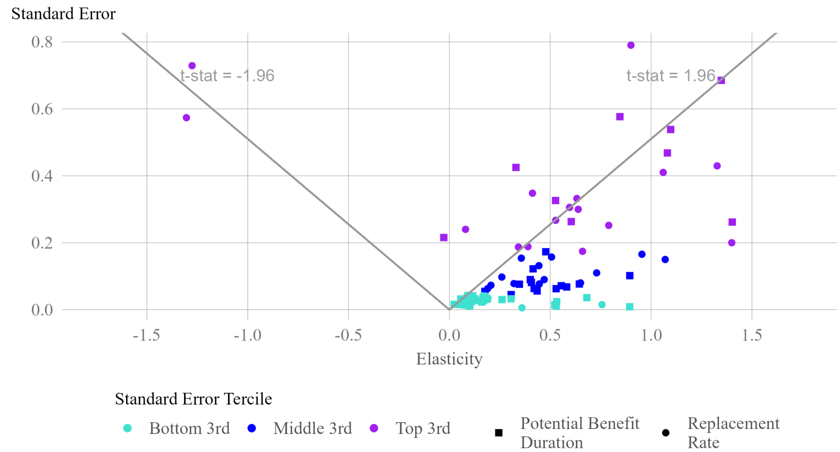

(b) Non-parametric Distribution of Elasticity Estimates

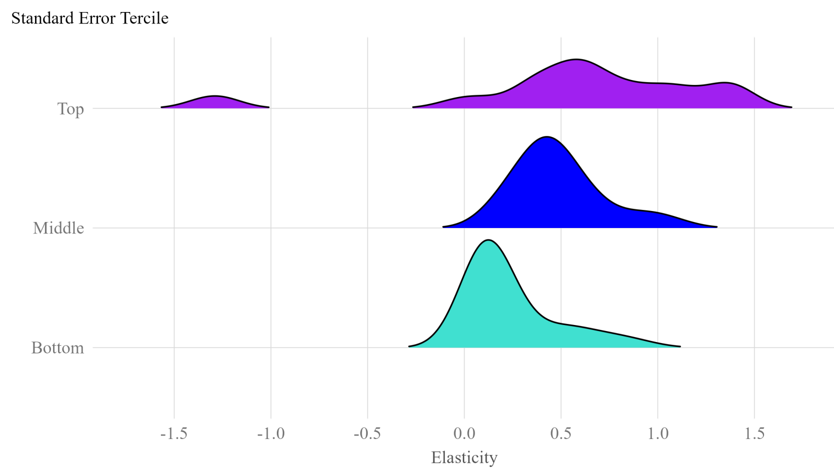

(c) Mean Elasticity Estimates

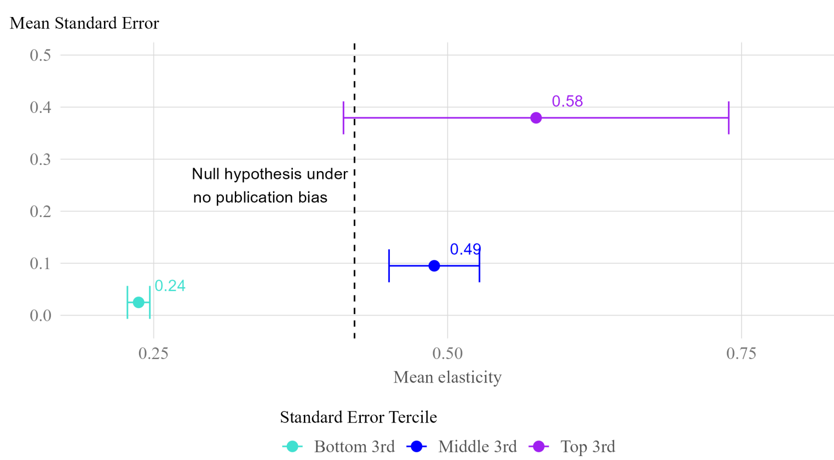

Notes: This figure describes the joint distribution of estimates of the elasticity and the standard error of unemployment duration with respect to unemployment benefit generosity. For visual clarity, eight estimates with elasticities < -1.5 or > 1.5 are excluded here, but are described in Appendix A. In panel (c), a 95 percent confidence interval is constructed for each tercile using the delta method.

3.2 Structural Bias Correction: Andrews and Kasy (2019)

We estimate the latent distribution of elasticities absent publication bias following Andrews and Kasy (2019). The approach jointly models the latent distribution of elasticities and the publication probability for different realizations. Latent estimates are defined as realizations of the estimated elasticity prior to the publication process. As discussed in Section 3.1, the key identification assumption is that the latent distribution of elasticities is independent of the latent distribution of standard errors.

Following Definition 1 in Section I of Andrews and Kasy (2019), we define the distribution of latent elasticities and their standard errors to be $(\Theta^*, \Sigma^*) \sim \mu_{\Theta,\Sigma}$. This distribution describes heterogeneity across studies and the noisiness of estimates. For a given study with distribution parameters (Θ, Σ), a noisy realization of the latent elasticity X^* is drawn: $X^* \mid \Theta^*, \Sigma^* \sim N(\Theta^*, \Sigma^{*2})$. Finally, a publication decision $D \mid X^*, \Theta^*, \Sigma^* \sim Ber(p(t^*))$ is drawn where $t^* = X^*/\Sigma^*$. We observe X, Θ, Σ if $D = 1$.

We make two functional form assumptions. First, in our baseline analysis, we assume that the latent distribution is a t-distribution

$$\Theta^* \sim \bar{\theta} + t(\nu) \cdot \tau$$

where $\bar{\theta}$ is the central tendency, ν is the degrees of freedom parameter determining tail fatness, and τ is the dispersion parameter determining spread. This is the same functional form used in the Andrews and Kasy (2019) minimum wage application. Second, motivated by the excess mass of t-statistics just above 1.96, we assume the each realization's publication probabilities as a function of its t-statistic:

$$p(t) = \begin{cases} \beta_p & \text{if } t < 1.96 \\ 1 & \text{if } t \geq 1.96 \end{cases} \tag{1}$$

Normalization of the publication probability for t-statistics above 1.96 is without loss of generality, as that pins down the total number of published estimates. Table 1 shows publication probability β_p for statistically insignificant elasticities is 8-12% of the publication

Table 1: Publication Bias Correction

	Average Published Elasticity	Publication Prob if $t < 1.96$	Latent Dist. Parameters $\Theta^* \sim \bar{\theta} + t(\nu) \cdot \tau$		
	$\hat{\epsilon}$	β_p	$\bar{\theta}$	τ	ν
Replacement Rate	0.43	0.12	0.21	0.26	3.80
	(0.12)	(0.08)	(0.14)	(0.09)	(1.44)
Potential Benefit Duration	0.46	0.08	0.09	0.25	2.92
	(0.05)	(0.05)	(0.07)	(0.06)	(1.54)

Notes: The table reports statistics on the published and latent distribution of elasticities for potential benefit duration and replacement rate. There are 49 potential benefit duration estimates and 42 replacement rate estimates. $\bar{\epsilon}$ is the sample mean of published elasticities. The remaining columns report estimated parameters from the Andrews and Kasy (2019) publication bias correction. β_p is the publication probabilities for t-statistics below 1.96. The latent distribution of elasticities is assumed to be distributed with $\Theta^* \sim \bar{\theta} + t(\nu) \cdot \tau$ so the mean of the latent elasticity distribution is $\bar{\theta}$ and the dispersion parameter (similar to the standard deviation) is τ. Standard errors are reported below parameter estimates. In some cases, multiple papers study the same benefit variation in the same region. We refer to these as "contexts" and cluster standard errors by context. Clustered standard errors for the sample mean come from a regression of the elasticity on a constant.

probability for statistically significant positive elasticities. Publication probability estimates are statistically precise; even at the upper end of the largest 95% confidence interval, the publication probability for insignificant estimates is only 28% of that for significant estimates. This extent of publication bias aligns with other literatures examined in Andrews and Kasy (2019).

Publication bias causes the average of published elasticities to overestimate the true average. The mean of the latent elasticity distribution ($\bar{\theta}$) is 0.09 for PBD (naive mean: 0.46) and 0.21 for RR (naive mean: 0.43). The corrected estimates for PBD are significantly different from the average naive elasticities at a 5% significance level, in the sense that the 95% confidence interval for $\bar{\theta}$ excludes the average published elasticity. This is not the case for RR.

Table 1 also shows that there is not a single latent elasticity value for either RR or PBD that can rationalize the distribution of estimates. Instead, there is substantial dispersion in the latent elasticity distribution. The dispersion parameter (similar to the standard deviation) of the latent t-distribution (τ) is 0.25 for PBD and 0.26 for RR. In other words, 90% of latent PBD estimates fall between -0.32 and 0.51 while 90% of the latent RR estimates fall between -0.22 and 0.65.

Table 2: Robustness Analysis of Publication Bias Correction

Margin	Difference from baseline	$\bar{\epsilon}$	β_p	$\bar{\theta}$	τ	$\bar{\theta} \pm 1.64\tau$
RR		0.43	0.12	0.21	0.26	[-0.22, 0.65]
RR	Normal Distribution	0.43	0.08	0.07	0.50	[-0.74, 0.88]
RR	Symmetric p(t)	0.43	0.19	0.33	0.23	[-0.05, 0.71]
RR	Extra p(t) cutoff	0.43	0.13	0.19	0.28	[-0.27, 0.65]
RR	Drop Hunt (1995)	0.52	0.11	0.21	0.26	[-0.21, 0.63]
RR	Non-Parametric GMM	0.43	0.22	0.01	1.01	[-1.65, 1.67]
RR	Non-Parametric GMM; Symmetric p(t)	0.43	0.20	0.04	1.00	[-1.60, 1.69]
RR	Non-Parametric GMM; Drop Hunt (1995)	0.52	0.21	0.21	0.81	[-1.12, 1.54]
RR		0.43	0.12	0.21	0.26	[-0.22, 0.65]
RR	Normal Distribution	0.43	0.08	0.07	0.50	[-0.74, 0.88]
RR	Symmetric p(t)	0.43	0.19	0.33	0.23	[-0.05, 0.71]
RR	Extra p(t) cutoff	0.43	0.13	0.19	0.28	[-0.27, 0.65]
RR	Drop Hunt (1995)	0.52	0.11	0.21	0.26	[-0.21, 0.63]
RR	Non-Parametric GMM	0.43	0.22	0.01	1.01	[-1.65, 1.67]
RR	Non-Parametric GMM; Symmetric p(t)	0.43	0.20	0.04	1.00	[-1.60, 1.69]
RR	Non-Parametric GMM; Drop Hunt (1995)	0.52	0.21	0.21	0.81	[-1.12, 1.54]

Notes: RR is replacement rate and PBD is potential benefit duration. $\bar{\epsilon}$ is the sample mean of realized elasticities. The remaining columns report estimated parameters from the Andrews and Kasy (2019) publication bias correction. β_p is the publication probabilities for t-statistics below 1.96. In our baseline specification, the latent distribution of elasticities is assumed to be distributed with $\Theta^* \sim \bar{\theta} + t(\nu) \cdot \tau$ so the mean of the latent elasticity distribution is $\bar{\theta}$ and the dispersion parameter (similar to the standard deviation) is τ. The second row assumes the latent distribution is normal, $N(\bar{\theta}, \tau^2)$. The third row returns to the t distribution but assumes that p(t) has three regions: t $<$ -1.96, $|t| < 1.96$, and t $>$ 1.96. In this row, β_p captures the estimate for $|t| < 1.96$ which is where the bulk of estimates occur. The fourth row also uses a t distribution and assumes p(t) has regions: t $<$ 0, $0 < t < 1.96$, and t $>$ 1.96. β_p captures the estimate for $0 < t < 1.96$. The fifth row drops Hunt (1995) which is a very negative estimate. The next three rows repeat (1), (3), and (5) but use a non-parametric specification of AK. The remaining rows repeat the analysis above, except looking at the PBD margin.

Estimated censoring (β_p) is quantitatively robust across specifications, but the mean latent elasticity ($\bar{\theta}$) is sensitive to specification. We explore robustness of the estimates to different functions for $p(t)$, distributions of latent estimates, and included studies in Table 2. Across all specifications, β_p remains far below one. However, $\bar{\theta}$ varies meaningfully. Across RR specifications, estimates vary from 0.01 to 0.33. Across PBD specifications, estimates vary from 0.03 to 0.33. In addition, specifications estimating lower $\bar{\theta}$ tend to estimate higher

τ. The GMM procedure is fragile: the inclusion or exclusion of Hunt (1995) produces widely differing estimates of $\bar{\theta}$.[10]

Overall, the corrections for publication bias are always substantial in economic terms, but both statistical and model uncertainty remain about the appropriate magnitude of the correction. In addition, across all specifications in Table 1 and Table 2, we find a high degree of dispersion in the latent elasticity distribution. This motivates characterizing systematic heterogeneity in elasticities in Section 4.

4 Predictors of Elasticity Heterogeneity

4.1 Motivation for Included Predictors

We document study-level characteristics that statistically predict heterogeneity in elasticities. We suggest caution in ascribing a causal interpretation to the predictors, as they are not randomly assigned across studies.

Collected study characteristics fall into two categories: economic characteristics and methodological characteristics. Economic characteristics should be interpreted as the key dimensions of heterogeneity. These are factors that policymakers can consider when setting UI benefit parameters in their own economic context. In contrast, methodological characteristics should be interpreted as auxiliary controls. We include them to account for estimation choices that could affect the estimated elasticity. Table C-3 contains the full set of predictors.

Economic characteristics: Meta-analysis's comparative advantage when applied to UI benefits is documenting heterogeneity by baseline benefit parameters (i.e., PBD or RR). Estimating an elasticity requires variation in the benefit parameter itself, but this typically comes from a single policy discontinuity or policy reform that in turn has a single baseline PBD or RR.

We define baseline PBD or RR as the control group's value in quasi-experimental designs when reported or the overall sample average otherwise. We interact the baseline PBD (in weeks) and baseline RR (fraction) with the policy margin.

[10] Hunt (1995) estimates an elasticity of -3.32 and standard error of 2.25.

There are several reasons to include these characteristcs. First, a PBD elasticity that increases with the baseline PBD is a prediction of the Shavell and Weiss (1979) model of job search under UI that has not been directly tested empirically. The closest evidence comes from two papers with multiple discontinuities in PBD at different ages or years of tenure (van Ours and Vodopivec, 2008; Schmieder et al., 2012). Second, we are not aware of existing evidence that the RR elasticity varies with the baseline RR, but we show in Section 5 that this is policy-relevant. Finally, the RR elasticity may vary with baseline PBD—and vice versa—because a given increase in benefit PBD increases total benefit entitlement more with a higher RR.

We also include economic dimensions studied by the literature. One is the relative unemployment rate as a proxy for business cycles. literature has found positive (Bell et al., 2024), negative (Kroft and Notowidigdo, 2016; Landais, 2015), and insignificant (Schmieder et al., 2012) correlations between the elasticity and unemployment rate over time within a single context. Another is whether the benefit variation affects the entire labor market. We discuss the interpretation of this covariate in more detail in Section 6. Finally, we include several other economic characteristics discussed in Table C-3 that prove to be unimportant for predicting heterogeneity.

These are exclusively study-level economic characteristics. This means that we rely on only across-study variation rather than any within-study heterogeneity analyses. This analysis therefore uses one estimate for each study and policy margin. When a study includes multiple main estimates for a policy margin, we take a precision-weighted average. See Appendix A for details.

Methodological characteristics: Most of the methodological characteristics are proxies for study quality: data source, research design, and publication bias susceptibility. First, administrative data is less prone to measurement error. Second, quasi-experimental research designs—particularly regression discontinuity designs—are less prone to selection bias due to policy endogeneity. Third, noisier estimates are more prone to publication bias. Following the meta-regression literature, we include the standard error as a control (Stanley, 2008).

Other methodological characteristics are less clearly related to study quality but are included to increase comparability across estimates. For example, it is unclear whether elasticities derived from hazard models are systematically different from those derived from duration regression. Additionally, elasticities with the outcome as either total nonemployment duration or covered unemployment are related but conceptually distinct estimands.

4.2 Documenting Heterogeneity Using Bayesian Model Averaging

Following recent developments in the meta-analysis literature, we use Bayesian model averaging (BMA) to highlight predictors of the elasticity (Havránek et al., 2024; Gechert et al., 2022; Zigraiova et al., 2021; Bajzik et al., 2020). The frequentist analogue to BMA is as follows: run separate regressions predicting the elasticity ϵ using all possible subsets of the study characteristics X

$$\epsilon = \alpha + \beta X \tag{2}$$

and average coefficients across regressions. BMA instead imposes priors over the set of included covariates and their coefficients, iteratively updates those priors based on the likelihood, and outputs Bayesian analogs to the regression coefficient (posterior mean) and p-value (posterior inclusion probability).[11] We use priors common in the BMA literature and estimate the model using the `bms` package in R (Eicher et al., 2011; Zeugner and Feldkircher, 2015).[12] In our initial analysis we found that the results were incredibly sensitive to one study (Hunt, 1995) which estimates a RR elasticity of -3.32 and standard error of 2.25; we therefore omit that study in what follows.

[11] BMA's priors can be thought of as regularization like in a data-driven machine learning approach. However, BMA's linear model coefficients are easily interpretable. These properties make it generally desirable for suggesting determinants with many possible covariates but small N, such as the macroeconomics literature on determinants of growth (Steel, 2020).

[12] In particular, we use Zellner's g-prior for the regression coefficients and a uniform model prior.

Table 3: Study Characteristics Correlated with Elasticities

	Posterior Inclusion Probability	Posterior Mean
(Intercept)	1.000	0.191
Baseline benefits		
Baseline RR (fraction) x RR estimate	0.639	0.390
Baseline RR (fraction) x PBD estimate	0.161	-0.029
Baseline PBD (weeks) x RR estimate	0.292	0.001
Baseline PBD (weeks) x PBD estimate	0.996	0.007
Policy variation		
PBD estimate (vs. RR estimate)	0.299	-0.090
Macro treatment	0.173	0.024
Study context		
Sample year (2023 = 0)	0.115	0.000
Relative unemployment (pp)	0.113	0.001
Labor tax wedge (pp)	0.143	-0.001
United States dummy	0.249	-0.041
Data and estimation		
Administrative data	0.127	-0.008
Nonemployment as outcome	0.659	-0.123
Hazard model	0.108	-0.002
Difference-in-Difference or Regression Kink Design	0.167	0.013
Regression Discontinuity Design	0.112	0.004
Standard errors		
Standard Error	1.000	1.334
Journal		
Impact Factor (z-score)	0.380	0.026

Notes: This table contains model output from Bayesian model averaging to predict elasticity estimates with study characteristics. The posterior mean is the Bayesian analog to the estimated regression coefficient. Baseline PBD or RR is the control group's value in quasi-experimental designs when reported or the overall sample average otherwise. "RR estimate" or "PBD estimate" refers to elasticity's policy variation. "Aggregate variation" is an indicator for the elasticity identified using a policy change affecting a large share of UI claimants: most claimants in the entire market or all claimants in segmented labor markets.cy reform or average sample year. It is relative to 2023, where larger positive values are further in the past. Relative unemployment is from the World Bank's World Development Indicators database, subtracting the average across all available years from the sample year's value. The labor tax wedge is defined as the ratio between taxes paid by an average worker and the corresponding total labor cost for the employer. It uses the latest available value from the OECD. The notes to Table C-3 describe variables in more detail.

We find strong evidence of systematic heterogeneity in the estimated elasticity. The model's posterior distribution for the variance of the elasticity based on observable study characteristics is 53%, meaning that about half of the variation in elasticities reflects systematic heterogeneity. Table 3 examines the relationship of each specific covariate with the predicted elasticity. The posterior inclusion probability in Column 1 captures a Bayesian notion of statistical significance. A covariate of random noise generally has a posterior inclusion probability of between 10 and 15%, so five of the covariates appear to have meaningful predictive power (defined as posterior inclusion probability > 30%). The posterior mean in Column 2 is the unconditional linear model coefficient (including zeros when model excludes the covariate).

We find strong evidence that the replacement rate elasticity is larger when the replacement rate is higher. The posterior mean of 0.390 implies that moving from the lower-end of US replacement rates to the upper-end of European replacement rates is associated with a $(0.86 - 0.33) \cdot 0.337 = 0.21$ increase in the elasticity.[13] Section 5 highlights this finding's policy implications. Similarly, when baseline PBD is 50 weeks longer, we find that the elasticity is 0.05 larger for changes in the replacement rate and 0.35 larger for changes in number of potential weeks of benefits.[14]

The two key methodological characteristics are the standard error and the unemployment duration definition. First, a study with a standard error of 1 yields an elasticity estimate 1.3 units higher than a study with a standard error of 0. This is consistent with censoring of insignificant estimates. Second, measuring unemployment as total nonemployment delivers a lower elasticity. This pattern holds even within studies, which implicitly controls for all contextual characteristics. This finding aligns with recent within-study evidence from Bell et al. (2024).

Perhaps the most valuable output from the meta-analysis is aggregating studies across different contexts to estimate the elasticity for a much-studied policy context: state-financed

[13]This pattern is also present in the the raw data, as shown in Figure B-3.

[14]The finding that the elasticity is larger when the baseline number of PBD weeks is larger is in part dependent on a single outlier context which studies a reform in Norway where the control group had PBD of 186 weeks (as compared to 100 weeks or less in all other studies) and the estimated elasticity is very high. Røed and Westlie (2012) estimate an elasticity of 1.71 w.r.t. PBD and Røed and Zhang (2003) estimate an elasticity for men of 0.95 w.r.t. RR. If we re-run BMA dropping these studies we get a posterior mean of 0.001. This means that when baseline PBD is 50 weeks longer, the elasticity is 0.05 larger.

UI benefits in the US. State-financed UI benefits are those that are available outside of recessions, have an average replacement rate of 43.5% (U.S. Department of Labor, 2024), and in most states last 26 weeks. The meta-analysis predicts that the covered duration elasticity to the replacement rate in this context is 0.34.[15] Applying the same method to calculate the covered duration elasticity with respect to changes in PBD delivers an elasticity of 0.23.

The predicted elasticity varies dramatically with the policy context. Across US states and OECD countries for 2023, the least generous policy context is Florida, where replacement rates average 33% and benefits last at most 12 weeks. In this context, our model predicts an RR elasticity of 0.29 and a PBD elasticity of 0.14. The most generous policy context is France, where replacement rates average 68% and benefits last up to 104 weeks (two years); there our model predicts an RR elasticity of 0.52 and a PBD elasticity of 0.80.

5 Application: Sufficient Statistics for Optimal UI Benefits

We show how meta-analysis can solve two challenges in calibrating sufficient statistics formulas. These formulas are increasingly popular in public economics (Kleven, 2018). Our application is the Baily (1978) - Chetty (2006) formula for the optimal replacement rate.

The first challenge is that researchers may not have a study with all sufficient statistic components in their context of interest. For example, the most precise and highest-quality elasticity estimates often come from European contexts but Gruber's (1997) consumption-smoothing estimates are for the US. More generally, researchers may be concerned about the external validity of any given study's estimate to their context.

The second challenge is that empirically calculating optimal policy often requires structural assumptions that sufficient statistics approaches aim to avoid. For example, while the Chetty (2006) derivation does not require a constant duration elasticity, the optimal replace-

[15] Specifically, allowing here for small rounding errors, the variables used are: 0.191 (intercept) + 0.024 (Aggregate variation) - 0.041 (United States dummy) - 0.001 * 30.5 (Labor tax wedge) - 0.008 (Administrative data) + 0.004 (RDD) + 0.001*26 (Baseline PBD) + 0.390*0.435 (Baseline RR) $\approx$ 0.34.

ment rate formula states it as a constant. However, in the context of studying *large* policy changes, Kleven (2021) shows that a constant labor supply elasticity embodies an isoelastic functional form assumption on structural labor supply cost preferences. This assumption is stronger when the optimal policy is further from the status quo. Researchers may instead prefer supplemental reduced-form moments over structural assumptions on preferences.

Meta-analysis provides a solution to both challenges. It does so by estimating systematic heterogeneity in the elasticity based on observable characteristics, as shown in Section 4.

To address the first challenge, calculate an elasticity in our specific context of interest. We predict a covered unemployment duration elasticity in the present-day US while drawing on all prior high-quality estimates from other contexts.

To address the second challenge, we allow the elasticity to vary with the replacement rate. Since we are treating the replacement rate as a policy to optimize, we must ascribe a causal interpretation to the estimated relationship between the elasticity and replacement rate. The ideal meta-analysis dataset would have many studies with randomly assigned baseline replacement rates—separate from random policy variation to estimate the elasticity—in order to flexibly estimate the relationship. In our application, we rely on auxiliary study characteristics to justify a conditional independence assumption for the baseline replacement rate. We additionally assume the relationship is linear and does not vary with the reform size. These assumptions allow us to use empirical moments instead of structural assumptions about preferences. In this respect, our analysis builds on other recent work that substitutes reduced-form moments for structural assumptions (Chiang and Żoch, 2022; Auclert et al., 2018).

Meta-analysis can also account for the fact that published findings used for sufficient statistic formulas may be subject to publication bias. We do so by predicting the elasticity when the standard error is zero (Card and Krueger, 1995).[16]

The Baily-Chetty optimal replacement rate formula equates the consumption-smoothing

[16] This is a linear correction for publication bias. However, Appendix D of Andrews and Kasy (2019) points out that the selection process is not necessarily linear in general.

gain and fiscal externality loss from an additional dollar of UI benefits:

$$\gamma \frac{\Delta c(r)}{c} = \frac{\epsilon_{1-e,r}(r)}{e}. \tag{3}$$

The right-hand side of the formula is the replacement rate elasticity $\epsilon_{1-e,r}(r)$. To address the first challenge, we predict it in the US policy context (US in 2023, non-recessionary period, covered benefit durations, 26-week PBD, purged of publication bias) using the Bayesian model averaging framework from Section 4. We assume an employment rate e of 0.95. To address the second challenge, we allow the elasticity to vary with the replacement rate. The posterior mean estimate in Table 3 means that a 10 percentage point increase in the replacement rate corresponds to a 0.039 increase in the elasticity.

To estimate the left-hand side, we assume a coefficient of relative risk aversion of $\gamma = 2$ and use estimates of the consumption drop $\frac{\Delta c(r)}{c}$ from Gruber (1997). These estimates are useful for calculating the optimal replacement rate because they allow consumption smoothing gains to vary with the replacement rate.[17] Gruber estimates $\frac{\Delta c(r)}{c} = 0.222 - 0.265r$. That is, if the replacement rate is 0 then consumption will fall by 22.2% and a replacement rate of $(0.222/0.265) = 84\%$ is sufficient to prevent any consumption drop. We implicitly solve for the optimal replacement rate:

$$r = \frac{1}{0.265}(0.222 - \frac{\epsilon_{1-e,r}}{e\gamma}) \tag{4}$$

Figure 2 jointly plots the left-hand side and right-hand side of Equation 3. The downward-sloping black line shows the left-hand side: the Gruber (1997) welfare gains from consumption-smoothing. The other lines show the right-hand side elasticity (scaled by the assumed employment rate $e = 0.95$). Our preferred estimate for the elasticity is the upward-sloping pink line. It is in the US context, corrects for publication bias, and allows for heterogeneity with the replacement rate. Given this, the optimal replacement rate given by Equation 4 is its intersection with the black line at 28%.

Figure 2 demonstrates the quantitative importance of correcting for publication bias and

[17] There are several methods for estimating workers' willingness to pay for UI benefits (Landais and Spinnewijn 2021, Chetty 2008, Hendren 2017). However, they estimate welfare gains from local changes in UI benefits and therefore cannot be used to estimate an optimal replacement rate.

Figure 2: Optimal replacement rates for UI benefits

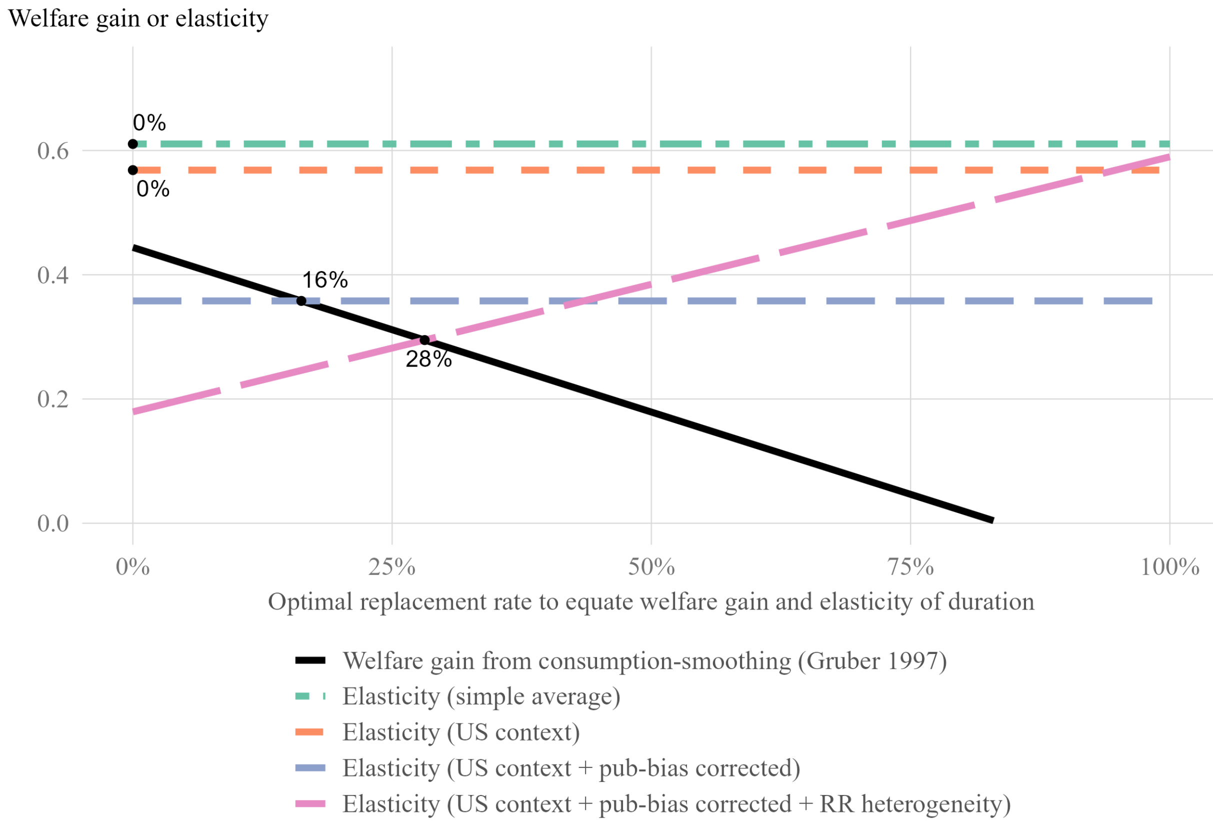

Notes: This figure shows the optimal replacement rates implied by the Baily-Chetty formula. Consumption-smoothing estimates assume constant relative risk aversion with parameter $\gamma = 2$ and a consumption drop during unemployment $\frac{\Delta c(r)}{c}$ from Gruber (1997), while elasticity estimates $\epsilon_{1-e,r}(r)$ come from the Bayesian model averaging (BMA) procedure from Section 4. The downward-sloping solid black line shows the Gruber (1997) welfare gains from consumption-smoothing. The other four lines show elasticity estimates. Our preferred estimate is the upward-sloping dashed pink line, which predicts the elasticity with parameter values from footnote 15. The line slopes upward because we find that the elasticity is higher when the replacement rate is higher. Its intersection with the black line is the optimal replacement rate given by Equation 4. The other three dashed lines illustrate the quantitative impacts of our three methodological innovations relative to the prior literature on optimal UI benefits. The green unevenly-dashed horizontal line shows the simple average elasticity in the BMA sample. The orange narrowly-dashed horizontal line shows the BMA sample's prediction for the US context, without correcting for publication bias nor allowing for elasticity heterogeneity with the replacement rate. The blue widely-dashed horizontal line uses the same parameter values from footnote 15 as the upward-sloping pink line, except it generates a single prediction using the BMA sample average replacement rate. This corrects for publication bias by making a context-specific prediction with SE=0 but does not allow for heterogeneity with the replacement rate.

allowing for heterogeneity with the replacement rate. If we no longer allow the elasticity to be lower at low replacement rates and predict the elasticity using the sample average replacement rate, then the optimal replacement rate falls to 16%. Both the simple average of elasticities and the context-specific predicted elasticity ignoring publication bias and heterogeneity by the baseline replacement rate lie above the consumption-smoothing gains at all replacement rates, implying that no UI is optimal.[18]

Figure B-4 shows that the alternative publication bias correction from Section 3 also delivers a similar optimal replacement rate. While the Andrews and Kasy (2019) procedure does not generate a context-specific elasticity that varies with the replacement rate, its correction for publication bias estimates an optimal replacement rate of 41%.

6 Application: Micro vs Macro Elasticity

Interest in the general equilibrium effects of UI surged in the wake of the Great Recession, when PBD was extended nearly quadrupled for all US UI claimants. Researchers typically summarize general equilibrium effects by comparing the causal effect of increasing benefits for a single worker (the "micro" effect) to the causal effect of increasing benefits for all workers (the "macro" effect).

Different models of the labor market make different sign predictions about the general equilibrium effects of benefits. In a Diamond-Mortensen-Pissarides model where UI reduces search effort through higher reservation wages, employers post fewer vacancies (Hagedorn et al., 2016). Such a vacancy posting response means the macro elasticity is larger. In a model with labor market congestion (Michaillat 2012 and Landais et al. 2018), UI also reduces search effort for some unemployed jobseekers. But this is offset by a higher offer arrival rate for other jobseekers, which Landais et al. call a "rat race" effect. In such a model, the macro elasticity is smaller. In addition, in models with a product market where monetary policy is constrained, more generous UI increases aggregate demand (Kekre, 2022). This makes the macro elasticity smaller.

[18] Using the median elasticity of covered duration with respect to replacement rate from Schmieder and von Wachter (2016) of 0.59 delivers the same conclusion. Instead using the median among US studies from Schmieder and von Wachter (2016) of 0.38 implies an optimal replacement rate of 8%.

Empirically measuring the difference between the micro elasticity and the macro elasticity in the same policy context is challenging. Estimating both in the same context requires "double randomization": (i) some *labor markets* are treated with higher benefits and compared to other control labor markets and (ii) some *workers* are treated with higher benefits compared to control workers in the same labor market (Lalive et al., 2015). The closest observational analog is Jessen et al. (2023) studying a labor market policy in Poland.

Overall, there is substantial uncertainty whether the micro elasticity is smaller than, larger than, or the same as the macro elasticity. Figure B-5 summarizes studies known to us that report both a micro and macro elasticity. Sufficient variation to also construct a confidence interval for the difference between the two elasticities is even rarer.

We use meta-analysis as an alternative methodology for comparing the micro and the macro elasticity. Most prior macro estimates rely on measuring the response of state- or region-level unemployment to a state- or region-level policy change (Chodorow-Reich et al. 2018, Acosta et al. 2023, Boone et al. 2021). Our contribution is to use recent theoretical advances in understanding the distinction between micro and macro elasticities to re-interpret some of the older elasticity estimates which rely on microeconomic data. In our meta-analysis of 57 studies, we identify five studies whose design captures a macro elasticity rather than a micro elasticity. These are studies where the variation in benefits is at the market level rather than the individual level and the control group is an *untreated market*. Two studies (Lalive et al. 2015; Topel 1983) compare regions with more and less generous benefits. Three studies (Card and Levine 2000; van Ours and Vodopivec 2006; Arranz et al. 2008) compare data on claimants from years where state- or country-wide policy was more or less generous. Section A.2 provides additional detail on these studies.

One strength of our approach is that it systematically draws on evidence from several different policy contexts rather than relying on a single case study. Meta-analysis methods provide a principled framework for comparing macro estimates from one context to micro estimates from another context. A related strength is the ability to construct a confidence interval. Confidence intervals are impossible to construct in comparisons that rely on a single case study and can be quite imprecise when constructed based on estimates from a small number of regions. A limitation of the microeconomic data approach is that it focus on *already* unemployed workers. It is therefore silent on how UI impacts flows into unemployment.[19]

We find no systematic difference between micro and macro elasticities. More precisely, the "macro treatment" row in Table 3 shows that the studies with market-level treatments have an elasticity 0.024 units higher than studies with individual-level treatments. The posterior inclusion probability is 0.173, meaning the covariate is no different from random noise.

Our estimates are consistent with a macro elasticity which is modestly above or modestly below the micro elasticity. Figure B-5 shows a 95% confidence interval for the US macro PBD elasticity that ranges from 0.08 to 0.38. This interval rules out the possibility that UI benefits have no effect in general equilibrium and also that their general equilibrium effect is much larger than their micro effect.

7 Conclusion

We confirm prior reviews' finding that UI benefit expansions increase unemployment duration, but we find evidence that publication bias exaggerates the magnitude. We use Bayesian modeling averaging to document how elasticities vary across studies and show the usefulness of these across-study moments in two applications. First, we calibrate the optimal replacement rate using a sufficient statistics formula. We extrapolate how the elasticity would change under a large policy reform using the conditional relationship across studies between the replacement rate elasticity and the baseline replacement rate. This expands the remit of sufficient statistics from assessing local optimality to global optimality. Second, we test whether micro elasticities differ from macro elasticities. We account for a variety of other observable study characteristics. This provides a principled way to make an apples-to-apples comparison across different studies.

[19] This is a burgeoning literature, reviewed in Le Barbanchon et al. (2024). See Winter-Ebmer (2003) for an early example and Jessen et al. (2023) for a recent example.

References

Acosta, Miguel, Andreas I. Mueller, Emi Nakamura, and Jón Steinsson. 2023. "Macroeconomic Effects of UI Extensions at Short and Long Durations." Working Paper 31784, National Bureau of Economic Research. 10.3386/w31784.

Andrews, Isaiah, and Maximilian Kasy. 2019. "Identification of and Correction for Publication Bias." *American Economic Review* 109 (8): 2766–2794. 10.1257/aer.20180310.

Arranz, Jose, Juan Muro, and Juan Es. 2008. "Do unemployment benefit legislative changes affect job finding?".

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2018. "The Intertemporal Keynesian Cross." Working Paper 25020, National Bureau of Economic Research. 10.3386/w25020.

Baily, Martin Neil. 1978. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics* 10 379–402.

Bajzik, Josef, Tomáš Havránek, Zuzana Irsova, and Jiri Schwarz. 2020. "Estimating the Armington Elasticity: The Importance of Study Design and Publication Bias." *Journal of International Economics* 127 103383. 10.1016/j.jinteco.2020.103383.

Bell, Alex, T.J. Hedin, Geoffrey Schnorr, and Till von Wachter. 2024. "Unemployment Insurance (UI) Benefit Generosity and Labor Supply from 2002 to 2020: Evidence from California UI Records." *Journal of Labor Economics* 42 (S1): S379–S416. 10.1086/728808.

Bennmarker, Helge, Kenneth Carling, and Bertil Holmlund. 2007. "Do Benefit Hikes Damage Job Finding? Evidence from Swedish Unemployment Insurance Reforms." *LABOUR* 21 (1): 85–120. 10.1111/j.1467-9914.2006.00363.x.

Boone, Christopher, Arindrajit Dube, Lucas Goodman, and Ethan Kaplan. 2021. "Unemployment Insurance Generosity and Aggregate Employment." *American Economic Journal: Economic Policy* 13 (2): 58–99. 10.1257/pol.20160613.

Card, David, and Alan B. Krueger. 1995. "Time-Series Minimum-Wage Studies: A Meta-Analysis." *The American Economic Review* 85 (2): 238–243.

Card, David, and Phillip B Levine. 2000. "Extended benefits and the duration of UI spells: evidence from the New Jersey extended benefit program." *Journal of Public Economics* 78 (1): 107–138. 10.1016/S0047-2727(99)00113-9.

Chetty, Raj. 2006. "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics* 1879–1901.

Chetty, Raj. 2008. "Moral Hazard versus Liquidity and Optimal Unemployment Insurance." *Journal of Political Economy* 116 173–234.

Chiang, Yu-Ting, and Piotr Żoch. 2022. "Financial Intermediation and Aggregate Demand: A Sufficient Statistics Approach." Working Papers 2022-038, Federal Reserve Bank of St. Louis. 10.20955/wp.2022.038.

Chodorow-Reich, Gabriel, John Coglianese, and Loukas Karabarbounis. 2018. "The Macro Effects of Unemployment Benefit Extensions: a Measurement Error Approach." *The Quarterly Journal of Economics* 134 (1): 227–279. 10.1093/qje/qjy018.

Eicher, Theo S., Chris Papageorgiou, and Adrian E. Raftery. 2011. "Default Priors and Predictive Performance in Bayesian Model Averaging, with Application to Growth Determinants." *Journal of Applied Econometrics* 26 (1): 30–55. 10.1002/jae.1112.

Gechert, Sebastian, Tomáš Havránek, Zuzana Irsova, and Dominika Kolcunova. 2022. "Measuring Capital-Labor Substitution: The Importance of Method Choices and Publication Bias." *Review of Economic Dynamics* 45 55–82. 10.1016/j.red.2021.05.003.

Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *The American Economic Review* 87 (1): 192–205.

Hagedorn, Marcus, Fatih Karahan, Iourii Manovskii, and Kurt Mitman. 2016. "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects." Working Paper 19499, National Bureau of Economic Research. 10.3386/w19499.

Havránek, Tomáš, Zuzana Irsova, Lubica Laslopova, and Olesia Zeynalova. 2024. "Publication and Attenuation Biases in Measuring Skill Substitution." *The Review of Economics and Statistics* 1–37. 10.1162/rest_a_01227.

Havránek, Tomáš, T.D. Stanley, Hristos Doucouliagos et al. 2020. "Reporting Guidelines for Meta-Analysis in Economics." *Journal of Economic Surveys* 34 (3): 469–475. 10.1111/joes.12363.

Hendren, Nathaniel. 2017. "Knowledge of Future Job Loss and Implications for Unemployment Insurance." *American Economic Review* 107 1778–1823.

Hunt, Jennifer. 1995. "The Effect of Unemployment Compensation on Unemployment Duration in Germany." *Journal of Labor Economics* 13 (1): 88–120. 10.1086/298369, Publisher: The University of Chicago Press.

Irsova, Zuzana, Hristos Doucouliagos, Tomáš Havránek, and T.D. Stanley. 2023. "Meta-Analysis of Social Science Research: A Practitioner's Guide." *Journal of Economic Surveys.*

Jessen, Jonas, Robin Jessen, Ewa Galecka-Burdziak, Marek Góra, and Jochen Kluve. 2023. "The Micro and Macro Effects of Changes in the Potential Benefit Duration." IZA Discussion Papers 15978, Institute of Labor Economics (IZA).

Kekre, Rohan. 2022. "Unemployment Insurance in Macroeconomic Stabilization." *The Review of Economic Studies* 90 (5): 2439–2480. 10.1093/restud/rdac080.

Kleven, Henrik J. 2018. "Language Trends in Public Economics." https://www.henrikkleven.com/uploads/3/7/3/1/37310663/languagetrends_slides_kleven.pdf.

Kleven, Henrik J. 2021. "Sufficient Statistics Revisited." *Annual Review of Economics* 13 (Volume 13, 2021): 515–538. /10.1146/annurev-economics-060220-023547.

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. 10.1093/restud/rdw009.

Krueger, Alan B., and Bruce D. Meyer. 2002. "Chapter 33 Labor Supply Effects of Social Insurance." In *Handbook of Public Economics*, Volume 4. 2327–2392, Elsevier, . 10.1016/S1573-4420(02)80012-X.

Krusell, Per, Toshihiko Mukoyama, and Ayşegül Şahin. 2010. "Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations." *The Review of Economic Studies* 77 1477–1507.

Lalive, Rafael, Camille Landais, and Josef Zweimüller. 2015. "Market Externalities of Large Unemployment Insurance Extension Programs." *American Economic Review* 105 (12): 3564–3596. 10.1257/aer.20131273.

Lancaster, Tony, and Stephen Nickell. 1980. "The Analysis of Re-Employment Probabilities for the Unemployed." *Journal of the Royal Statistical Society Series A: Statistics in Society* 143 (2): 141–152.

Landais, Camille. 2015. "Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design." *American Economic Journal: Economic Policy* 7 (4): 243–278. 10.1257/pol.20130248.

Landais, Camille, Pascal Michaillat, and Emmanuel Saez. 2018. "A Macroeconomic Approach to Optimal Unemployment Insurance: Theory." *American Economic Journal: Economic Policy* 10 (2): 152–81.

Landais, Camille, and Johannes Spinnewijn. 2021. "The Value of Unemployment Insurance." *The Review of Economic Studies* 88 3041–3085.

Le Barbanchon, Thomas, Johannes F. Schmieder, and Andrea Weber. 2024. "Job Search, Unemployment Insurance, and Active Labor Market Policies." Working Paper 32720, National Bureau of Economic Research. 10.3386/w32720.

Lee, David S., Pauline Leung, Christopher J. O'Leary, Zhuan Pei, and Simon Quach. 2021. "Are Sufficient Statistics Necessary? Nonparametric Measurement of Deadweight Loss from Unemployment Insurance." *Journal of Labor Economics* 39 (S2): S455–S506. 10.1086/711594.

Lopes, Marta C. 2022. "A Review on the Elasticity of Unemployment Duration to the Potential Duration of Unemployment Benefits." *Journal of Economic Surveys* 36 (4): 1212–1224. 10.1111/joes.12479.

Meyer, Bruce D. 2002. "Unemployment and Workers' Compensation Programmes: Rationale, Design, Labour Supply and Income Support." *Fiscal Studies* 23 (1): 1–49, Publisher: Wiley.

Michaillat, Pascal. 2012. "Do Matching Frictions Explain Unemployment? Not in Bad Times." *American Economic Review* 102 (4): 1721–1750.

van Ours, Jan C., and Milan Vodopivec. 2008. "Unemployment Insurance and the Distribution of Unemployment Spells?" *Journal of Public Economics* 92 (3): 684–695. 10.1016/j.jpubeco.2007.05.006.

van Ours, Jan C., and Milan Vodopivec. 2006. "How Shortening the Potential Duration of Unemployment Benefits Affects the Duration of Unemployment: Evidence from a Natural Experiment." *Journal of Labor Economics* 24 (2): 351–378. 10.1086/499976.

Røed, Knut, and Lars Westlie. 2012. "Unemployment Insurance in Welfare States: The Impacts of Soft Duration Constraints." *Journal of the European Economic Association* 10 (3): 518–554. 10.1111/j.1542-4774.2011.01064.x.

Røed, Knut, and Tao Zhang. 2003. "Does Unemployment Compensation Affect Unemployment Duration?" *The Economic Journal* 113 190–206.

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. 10.1146/annurev-economics-080614-115758.

Schmieder, Johannes F., Till von Wachter, and Stefan 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. 10.1093/qje/qjs010.

Setty, Ofer, and Yaniv Yedid-Levi. 2021. "On the Provision of Unemployment Insurance when Workers are Ex-Ante Heterogeneous." *Journal of the European Economic Association* 19 664–706.

Shavell, Steven, and Laurence Weiss. 1979. "The Optimal Payment of Unemployment Insurance Benefits over Time." *Journal of political Economy* 87 (6): 1347–1362.

Stanley, Tom D. 2008. "Meta-Regression Methods for Detecting and Estimating Empirical Effects in the Presence of Publication Selection." *Oxford Bulletin of Economics and statistics* 70 (1): 103–127.

Steel, Mark FJ. 2020. "Model Averaging and Its Use in Economics." *Journal of Economic Literature* 58 (3): 644–719.

Topel, Robert H. 1983. "On Layoffs and Unemployment Insurance." *The American Economic Review* 73 (4): 541–559, `http://www.jstor.org/stable/1816558`.

U.S. Department of Labor. 2024. "UI Replacement Rates Report." `https://oui.doleta.gov/unemploy/ui_replacement_rates.asp`.

Winter-Ebmer, R. 2003. "Benefit Duration and Unemployment Entry: A Quasi- Experiment in Austria." *European Economic Review* 47 259–273.

Zeugner, Stefan, and Martin Feldkircher. 2015. "Bayesian Model Averaging Employing Fixed and Flexible Priors: The BMS Package for R." *Journal of Statistical Software* 68 1–37. 10.18637/jss.v068.i04.

Zigraiova, Diana, Tomáš Havránek, Zuzana Irsova, and Jiri Novak. 2021. "Ow Puzzling Is the Forward Premium Puzzle? A Meta-Analysis." *European Economic Review* 134 103714. 10.1016/j.euroecorev.2021.103714.