The Targeting–Timing Tradeoff of Fiscal Transfers*

Michael Boutros†

September 2026

Preliminary and Incomplete

Abstract

This paper studies the trade-off between the targeting and the timeliness of fiscal transfers during a crisis. I measure the terms of this trade-off from the pandemic programs: the universal check paid every household in under a month, and unemployment insurance took three and a half months to pay only households that lost employment income. I then evaluate the trade-off in a two-sector HANK model with a lockdown calibrated to the industries that were shut down. Most of the welfare loss falls on households in the sector that was never restricted, and the return to targeting is positive in the crisis and negative in normal times. The universal check reduces the welfare loss of both sectors by a small amount, whereas full targeting removes 78% of the shut sector's loss, raises the open sector's loss by 43%, and reduces the population's loss by only 8%. Since the demand linkage runs from the shut sector to the open sector, whom a program pays matters less than in a tagging model. Full targeting yields the smallest welfare loss, and its delay forfeits only about a fifth of the return to targeting.

*This paper subsumes "Evaluating the Impact of Economic Impact Payments." I thank Ayaka Behro for excellent research assistance.

†University of Toronto. Email: michael.boutros@utoronto.ca.

1 Introduction

A central problem in the design of fiscal transfers during a crisis is deciding which households should receive them. This problem is particularly difficult when the incidence of a crisis extends beyond the households that are directly subject to policy. During the pandemic, for example, public health restrictions applied to a limited set of industries, but the resulting losses of employment income extended well beyond them to households in industries that were never restricted, and a transfer program must determine whether these households are eligible for relief.

Programs make this determination through a screen, which conditions payment on an observable characteristic of the household in the spirit of the tagging literature (Akerlof, 1978). Every screen can be characterized along two dimensions: its precision, or how well it discriminates between affected and unaffected households, and its delay, or how long it takes to verify eligibility and disburse payment. A universal check based on the previous year's tax return can be disbursed to every household within weeks, but it provides no more relief to affected households than to anyone else. Unemployment insurance, on the other hand, conditions payment on a verified loss of employment income, and so cannot pay until that loss has occurred and been claimed. More precise screens therefore take longer to run, and the terms of this trade-off between precision and delay are the price of the screen.

The pandemic made this trade-off explicit. At the onset of the pandemic in March 2020, Congress enacted the CARES Act, which included both a universal check and an expansion of unemployment insurance, and the record of the debate shows that the speed with which the check could be disbursed was the reason for including it. As Schild and Garner (2021) discuss, this included choosing to structure the payments as tax rebates so as to utilize the Internal Revenue Service's existing infrastructure for making transfers to households, and the Senate majority leader described the purpose of the payments as providing immediate relief to households facing cash-flow problems as they stayed home.[1] Even on this infrastructure, which was the fastest that the government had, the dollars of the first round of payments reached households 0.91 months after enactment on average, and so paying households at the onset of the crisis was never an option.

The expansion of unemployment insurance, on the other hand, paid only households that had lost employment income, and could not pay them until that loss had occurred

[1] Congressional Record Vol. 166, No. 53. Senator Booker added that "we need to be doing it quickly, getting payments to people as soon as possible," and that relief should reach "everyone, including those who, through no fault of their own, now find themselves on that financial brink" (Congressional Record Vol. 166, No. 54).

and been claimed, which took 3.45 months. In fact, the same bill enacted a program with no targeting and a program with full targeting. As such, the time that full targeting costs is measured from two programs that actually ran. However, Congress did not settle what the targeting was worth, what the delay cost, or whether either program was the better use of the marginal dollar. Answering these questions is the main exercise in this paper.

This paper makes two contributions. The first is to measure the screening frontier. Table 1 reports the precision and the delay of every transfer program that was enacted during the pandemic. The delay of a program is measured from the monthly flows of its disbursements, which the Bureau of Economic Analysis publishes, while its precision is measured on the loss of employment income that the Household Pulse Survey records. Note that most of the households that lost employment income worked in an industry that was never restricted. As such, the losses that the programs screened on extended well beyond the industries that were shut down. It is this spread of the loss that the model in this paper is built to price.

Second, I price the frontier in a quantitative model. The value of precision and the cost of delay are both equilibrium objects. As such, the frontier describes the combinations of precision and delay that were available to the government, but it does not say which combination the government should have chosen. The classification error that a screen makes has been priced in static models of tagging and costly verification, in the spirit of Akerlof (1978), Townsend (1979) and Kleven and Kopczuk (2011). In these models, however, a dollar transferred to the "wrong" household is a pure redistribution towards that household, and there is no cost of delaying the transfer. Neither of these features is appropriate in a crisis in which one sector is shut down. A dollar transferred to a household in the shut sector is spent on goods produced in the sector that remains open, and a dollar that is transferred late arrives after the household has already borrowed against it.

In this paper, I build a two-sector heterogeneous-agent New Keynesian (HANK) model in which the shut sector is rationed directly. The model is calibrated to the industries that were shut down during the pandemic, and welfare is computed for every combination of precision and delay. The key reason for modeling two sectors is the linkage between them, rather than to have something for the government to screen between. In fact, a screening error only matters when the misdirected dollar has an impact on the households that the program was intended to reach in the first place, and in a two-sector economy this occurs through general equilibrium.

The model delivers four key results. The universal check reduces the loss of both sectors by a small amount, 4% of the population's loss, whereas full targeting removes 78%

of the shut sector's loss, raises the open sector's loss by 43%, and reduces the population's loss by 8%. Targeting in a lockdown is therefore mostly a transfer between the two sectors, and the four results account for the small net that remains. First, most of the welfare cost of the lockdown is borne by households in the sector that was never shut down, who were not restricted by the lockdown but are affected by it nonetheless. In fact, this sector bears 57.5% of the aggregate welfare loss, and all of it comes through the demand linkage from the shut sector. Since no targeted program reaches these "unaffected" households, this is the justification for a universal check. Second, the return to targeting depends on the state of the economy. In the lockdown, concentrating the check on households in the shut sector yields a smaller welfare loss for the population as a whole, but in normal times the same policy yields a larger welfare loss, since concentrating a program of fixed size on a third of households runs into diminishing returns. The sign of the return to targeting thus depends on whether or not there is a crisis, which is not something that a static model of targeting is able to generate.

Third, the linkage between the two sectors runs in only one direction during the lockdown. Since households in the shut sector cannot spend a transfer on their own good, which is rationed, any payment directed towards them is spent on the good produced by the open sector, where it arrives as demand. In fact, 44.2% of every dollar that is taken away from the open sector returns to the open sector through this channel. The same is not true in the other direction, however, since a payment directed towards households in the open sector has no way of reaching the shut sector for as long as that sector is unable to sell its good. The question of which households receive the transfer therefore matters less than it would in a tagging model, and in the cases where it does matter, the direction depends on which sector is still able to sell. Overall, the population as a whole narrowly prefers that the transfer be directed towards the shut sector, while households in the open sector, who bear most of the damage, prefer the check.

Fourth, the delay of the screen is not the key cost of targeting. With both programs paid at once, full targeting reduces the population's loss by 0.47 pp more than the universal check, compared to 0.38 pp when each program is paid at the date that its screen allows, so that the months the screen took forfeit 19.3% of the return to targeting. Full targeting remains the design that reduces the loss most at every delay that a real screen took.[2]

Related Literature This paper contributes to the literature on fiscal transfers to house-

[2]The model does not include the hand-to-mouth households for whom a month of waiting is most costly, so this is a statement about the calibrated population rather than about the most constrained households, and Section 5.3.2 returns to this. Note that the case for the universal check in the model is its reach, which is the first result above, and not its speed, which was the reason that Congress gave for it.

holds in heterogeneous-agent economies. Oh and Reis (2012) document that a large share of fiscal support during the Great Recession was delivered through direct transfers to households, and build a model which demonstrates the positive welfare impacts of targeted transfers from high- to low-wealth households. Bilbiie, Monacelli and Perotti (2013) and Mehrotra (2018) study the impact of fiscal transfers vs. fiscal spending in models that also break Ricardian equivalence, while di Giovanni, Kalemli-Özcan, Silva and Yildirim (2023) and Hale, Leer and Nechio (2023) focus on the inflationary consequences of stimulative transfers.

Most closely related to this paper are Bayer, Born, Luetticke and Müller (2023), who evaluate the targeting of pandemic transfers in a HANK model with a lockdown, and Carroll, Crawley, Slacalek and White (2026), who model the consumption response to the CARES Act, and both of these papers establish that targeting raises welfare. Relative to all of these papers, I focus on the trade-off between the targeting of fiscal transfers and the time it takes for the targeting to occur. The screen in this paper also connects closely to the literature on tagging and costly verification, in which Akerlof (1978) studies transfers that are conditioned on an observable, Townsend (1979) studies the cost of verifying a state, and Kleven and Kopczuk (2011) study the trade-off between the complexity of a program and its take-up. In these papers, the cost of the screen is a resource cost. In this setting, however, the cost of the screen is the calendar time it takes to apply, which I measure from the transfer programs that were actually enacted. The welfare cost of this delay is then determined in general equilibrium. The lockdown in the model closely follows Guerrieri, Lorenzoni, Straub and Werning (2022) in treating the episode as a constraint on activity rather than as a change in preferences, and rationing at a fixed price follows Barro and Grossman (1971) and Malinvaud (1977).

On the empirical side, the paper builds on estimates of the marginal propensity to consume out of transfers. Revealed- and reported-preference estimates are largely in agreement (Parker and Souleles, 2019), and a large literature estimates the consumption response to the transfers of the Great Recession (Shapiro and Slemrod, 2009; Sahm, Shapiro and Slemrod, 2010; Parker, Souleles, Johnson and McClelland, 2013; Broda and Parker, 2014), the pandemic (Parker, Schild, Erhard and Johnson, Forthcoming; Dunn, Hood and Driessen, 2020; Chen, Qian and Wen, 2020; Andersen, Hansen, Johannesen and Sheridan, 2020; Baker, Farrokhnia, Meyer, Pagel and Yannelis, 2020; Karger and Rajan, 2020; Coibion, Gorodnichenko and Weber, 2020), and earlier programs (Johnson, Parker and Souleles, 2006, 2009). Orchard, Ramey and Wieland (2025) study the 2008 stimulus in a New Keynesian model and show that general equilibrium reconciles a large household-level consumption response with a muted aggregate one. This paper also draws on the

Household Pulse Survey, which asked households how they used their payments. The model takes the aggregate propensity to consume from the calibration. As such, I use the survey to measure which households the screen reached and what a household does with a payment that arrives late. The model builds on Kaplan and Violante (2014) and Kaplan, Moll and Violante (2018) and is solved with the toolkit of Auclert, Bardóczy, Rognlie and Straub (2021); its two-sector production structure follows Cantelmo and Melina (2023), and its treatment of transfers arriving during rather than after a shock connects to the intertemporal marginal propensities of Auclert, Rognlie and Straub (2024).

Outline Section 2 measures the screening frontier. Section 3 presents the model and Section 4 the calibration of the general economy. Section 5 calibrates the lockdown and prices the frontier in it. Section 6 concludes.

2 The Screening Frontier

This section measures the precision and the delay of the transfer programs of the pandemic, which differed in their screens, the records or claims on which each program decided whom to pay. The first round of Economic Impact Payments was paid on the 2018 or 2019 tax return, and the unemployment insurance programs of the CARES Act were paid on a claim that the household filed after losing employment income and that the state verified. Table 1 reports both properties for every program.

2.1 Screens

The focus of this paper is on targeting fiscal transfers to households directly affected by a negative shock. To do this, the fiscal authority must screen households, which is an administrative process that can take a long period of time. In the COVID-19 pandemic, I focus on six "shut down" sectors representing almost 30% of households. While these sectors were essentially mandated into shutting down by the federal government, it was nearly impossible for the government to determine which households worked in these sectors since tax returns do not record the industry a household works in. Instead, the pandemic's targeted household programs took the form of expanded unemployment insurance, which reached households in the affected sector through a lengthy administrative process often requiring verification of a lost job.

Screening increases precision, which is the share of fiscal transfers that reaches affected households, measured over and above their share of the population. Let s denote the share of a program's dollars that reaches households that lost employment income, and

let χ_L denote the share of these households in the population as a whole. Precision is then defined as $\theta = (s - \chi_L)/(1 - \chi_L)$, so that a program which pays every household alike has $\theta = 0$, while a program which pays only the affected households has $\theta = 1$. Table 1 reports the precision of every program.

The CARES Act was enacted in March 2020, and its first round of Economic Impact Payments was paid on the household's 2018 or 2019 tax return.[3] Individual filers received $1,200, compared to $2,400 for joint filers, plus $500 per qualifying child, and the payment was reduced by 5% of adjusted gross income above $75,000 for individual filers and $150,000 for joint filers (so that individual filers without children received nothing above $99,000). Non-filers receiving Social Security were paid automatically. The program paid within a month of enactment, and its precision is 0.038. The same Act expanded unemployment insurance through the $600 weekly supplement (FPUC), Pandemic Unemployment Assistance (PUA) for workers outside regular unemployment insurance and the extended-benefit tier (PEUC). Each of these programs was paid on a claim that the household filed after losing employment income and that the state verified, and none of them could pay until the loss had occurred and been claimed.

2.2 Measurement

Precision and delay are measured from different sources, and neither is a survey estimate of household behavior. Delay is measured from the monthly flows of each program's disbursements, which the Bureau of Economic Analysis publishes, with the flow within a month placed at mid-month, and it is expressed in months from the enactment of the statute that created the program. For every program, I report two statistics of the dollar-weighted disbursement path, the dollar-weighted mean delay and the month by which 90% of the program's dollars had been paid. Section 2.4 explains which of the two enters the frontier.

Precision requires the share of a program's dollars that reached each group of households. For the Economic Impact Payments, the Internal Revenue Service publishes the dollars paid within each bracket of adjusted gross income, and I cross these brackets with the probability that a household in a given bracket lost employment income, which I obtain from the Household Pulse Survey, under the assumption of a uniform density within each bracket. For the unemployment insurance programs, a household is eligible only if it has lost employment income, and precision is therefore one by construction. The

[3]This section subsumes the empirical part of an earlier working paper, "Evaluating the Impact of Economic Impact Payments: Timely vs. Targeted Fiscal Stimulus," whose data are reused here and whose usage results are reproduced and extended in Section 2.5.

population share χ_L that enters θ is 0.454. This is the share of households that had lost employment income in the Household Pulse Survey during the weeks in which the first round of payments was arriving.

I use data from the Household Pulse Survey for both the loss criterion and household behavior. The Household Pulse Survey was designed by the Census Bureau to quickly measure the effects of the pandemic on households. It is fielded by text message and email to a sample drawn from the Master Address File, and the survey weights allow for national estimates despite response rates below 5%. I use 33 weekly waves of the survey, spanning the first week in June 2020 in which households were asked about the payments through the summer of 2021, when households were last asked about the third round, and household weights are applied throughout. The analysis uses three questions from the survey. First, households were asked whether anyone in the household had experienced a loss of employment income since March 13, 2020. This question defines the loss criterion. Second, households that had received or expected to receive a payment were asked whether they did or would use it mostly to pay for expenses, mostly to pay off debt, or mostly to add to savings. This is the usage question studied in Section 2.5. Third, households that selected one of these three responses were then asked to select every category on which the payment was spent, including food and clothing, rent, utilities and vehicle payments, debt, donations and savings. This is the composition question studied in Section 5.2. The wording of the usage question differs across rounds, and Section 2.5 discusses what these differences permit, while Appendix C reproduces the questions in full.

2.3 The Pandemic Programs

Table 1 reports the precision and the delay of the three rounds of Economic Impact Payments, which were paid on the tax record, and of the unemployment insurance programs of the CARES Act, which were paid on a verified claim. The second and third rounds of payments were disbursed in 0.7 and 0.5 months, respectively, compared to 0.91 months for the first round, since they were paid on the rail that had been built for the first round, and for this reason the frontier is fit only on the programs that had to build their own screens.

2.4 The Screening Frontier

I define the screening frontier as the lower boundary of the set of targeting and delay pairs, (θ, d), that a government could have chosen. Note that the frontier is a constraint

Table 1: The pandemic programs in screening coordinates

		Precision	Delay, Months	
Program	\$bn	θ	Mean	90th
Economic Impact Payments, on the Tax Record				
Round 1 (CARES Act)	275	0.04	0.91	1.24
Round 2	146	0.05	0.68	0.60
Round 3	423	0.04	0.53	0.90
Unemployment Insurance, on a Verified Claim				
CARES Package	**375**	**1.00**	**3.45**	**6.10**
\$600 Weekly Supplement (FPUC)	253	1.00	2.69	3.53
Pandemic Unemployment Assistance (PUA)	98	1.00	4.67	7.53
Extended Benefits (PEUC)	24	1.00	6.58	8.22

Notes: θ is the concentration of a program's dollars on the affected, $(s-\chi_L)/(1-\chi_L)$, where s is the share of dollars reaching them and $\chi_L = 0.454$ is their share of all households. The affected are the households that lost employment income (Household Pulse Survey, weeks 7–12, household-weighted), and for unemployment insurance θ is one by eligibility. Delay is measured on the dollar-weighted disbursement path from enactment (BEA monthly flows, mid-month convention). The mean is the statistic the frontier prices, and the 90th percentile of dollars is the month by which a program finished paying. The second and third rounds of payments were enacted in December 2020 and March 2021. The CARES package pools its three tiers, which were enacted together and paid the same claimant, and its 90th percentile (6.10) exceeds the dollar-weighted average of the 90th percentiles of the tiers (4.87).

on policy and not a description of any one program, and its shape follows from the way in which programs combine.[4] With two enacted designs, the frontier is the segment between them, and it is therefore linear.

Both endpoints of the frontier are programs that were enacted in the CARES Act, whose disbursement paths are measured in Section 2.2. Neither end of the frontier is therefore extrapolated. The endpoint at $\theta = 0$ is the first round of Economic Impact Payments, which was paid through the existing tax-filing rail with a precision of 0.038 and a dollar-weighted mean delay of 0.91 months from enactment. On the other hand, the endpoint at $\theta = 1$ is the unemployment insurance package of the CARES Act, which pays only households that lost a job and has a dollar-weighted mean delay of 3.45 months.

For each program, Table 1 reports two summary statistics of the disbursement path, but only the dollar-weighted mean enters the frontier. The 90th percentile of dollars is the month by which a program had finished paying, and it is the relevant statistic for how long the last recipients waited. It has the disadvantage, however, that it depends on how a program is divided up. For example, the 90th percentile of the pooled unemployment insurance package is 6.10 months, whereas the dollar-weighted average of the 90th percentiles of its three tiers is 4.9 months, since the quantile of a mixture is not the mixture of the quantiles. The mean does not have this problem, as it is the same however the tiers are grouped. Further, I show in Section 5 that welfare in the model is close to linear in the delay, which means that the mean is also the statistic that collapses a disbursement path to a single model date without material error. Delays are measured from the enactment of the statute rather than from the onset of the shock. The later rounds of both programs were carried by the same statutes, in December 2020 and March 2021, so measuring delay relative to the shock would record when Congress acted, which is common to both programs, and not how long each screen took to run.

The frontier is therefore

$$D(\theta) = 0.81 + 2.65\,\theta, \tag{1}$$

with $D(0) = 0.81$ and $D(1) = 3.45$. The intercept of the frontier is the time it takes the fastest existing rail to pay everyone, and the slope is the additional time it takes a verified-loss screen to pay only the affected. The model takes equation (1) as its constraint in the analysis that follows. I also report, alongside it, the variant in which the 90th percentile is used in place of the mean at both ends, which raises $D(1)$ to 6.10 months.

[4] Since a portfolio of two programs has a precision and a delay equal to the dollar-weighted averages of its two programs, the boundary of the feasible set is the lower convex hull of the enacted designs. No functional form is assumed for it.

2.5 What the Frontier Cannot Say

Table 2 reports how recipients used the payment in each of the three rounds of Economic Impact Payments, which were all paid on the same screen at different dates, and which the frontier records only by what each round achieved and how long it took. Using the usage question of the Household Pulse Survey, I estimate the gap in reported use between recipients that lost employment income and recipients that did not in each round, for 601,182 recipients. In the first round, which was paid within a month of enactment, households that had lost income were 19.6 pp more likely to report using the payment mostly for expenses, compared to 4.8 pp less likely in the second round, paid nine months into the pandemic, and 8.1 pp less likely in the third round. The gap in paying down debt moved in the opposite direction, from -8.7 pp in the first round to $+15.4$ pp and $+18.8$ pp in the second and third rounds, respectively. On the other hand, households that had lost income were about 11 pp less likely to report saving the payment in every round, and the round interactions are zero to within 0.5 pp.

Table 2: What a late check does: the work-loss gap in reported use, by round

	Gap, pp			Change from Round 1	
Mostly Used the Payment to	Round 1	Round 2	Round 3	Round 2	Round 3
Pay for Expenses	+19.6	-4.8	-8.1	-24.3 (0.5)	-27.6 (0.9)
Pay Down Debt	-8.7	+15.4	+18.8	+24.1 (0.5)	+27.5 (1.1)
Add to Savings	-10.8	-10.6	-10.7	+0.2 (0.4)	+0.1 (0.8)

Notes: Household Pulse Survey, 601,182 recipients of an Economic Impact Payment in weeks 7–12 (round 1), 22–27 (round 2) and 28–33 (round 3), with household weights. Each entry is the difference in the share reporting that use between households that lost employment income and households that did not, from a weighted least-squares regression with week fixed effects and controls for income bracket, education, sex, Hispanic origin, race, marital status, household size, children, age and age squared. The last two columns are the round interactions, with standard errors from the Census successive-difference replicate weights in parentheses. Rounds 2 and 3 ask about payments received in the last seven days, whereas round 1 asks whether a payment was ever received. Levels are therefore not comparable across rounds, and the table reports the gap within each round.

The flat saving gap is inconsistent with the interpretation that by 2021 the affected were no longer in need. Since the first round asked about a payment ever received while the later rounds asked about the last seven days, levels of the usage question are not comparable across rounds. The gap between the two groups is a within-round difference, however, and is not exposed to this change in wording. What changed across rounds is the use to which the money was put, with the reversal being one-for-one between spending and debt repayment. Thus, a payment that arrives a year after the loss still reaches the affected, who do not save it and instead use it to repair the balance sheet that the wait

itself damaged.

The frontier therefore describes the combinations of precision and delay that were available to the government, whereas the value of each combination depends on what a dollar does in general equilibrium once it arrives. In what follows, I build a model in which a dollar paid to the "wrong" household has such effects, and I use it to price the frontier.

3 Model

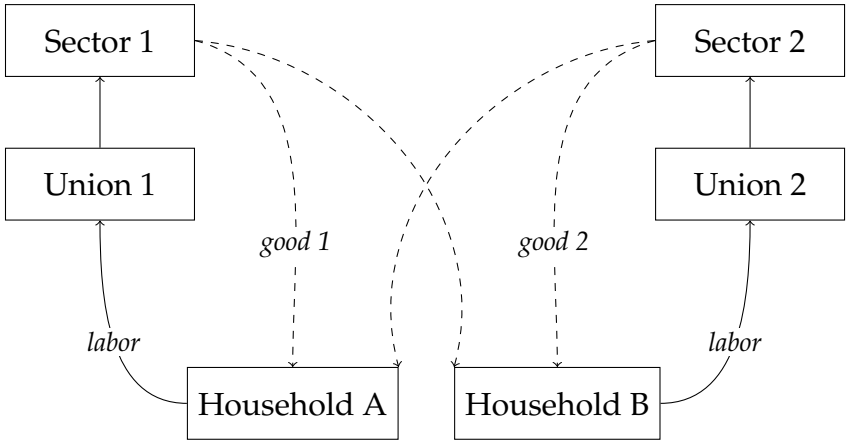
I construct a medium-scale HANK model with two sectors and two types of households. The model extends the two-sector structure of Cantelmo and Melina (2023) to feature two types of households, each of which supplies labor to one sector. Households have access to two different assets, and the rich production and household sectors are embedded into a canonical HANK model, e.g., the one studied in Auclert et al. (2021). The model is monthly, matching the frequency at which crisis income losses and policy responses unfold.

The structure of the economy is depicted in Figure 1. Households purchase both sectoral goods directly through type-specific consumption baskets, so there is no final goods firm. Good 1 serves as the numeraire, so all prices, wages, returns, and asset positions are quoted in units of good 1, and $q_t \equiv P_{2,t}/P_{1,t}$ denotes the relative price of good 2. Each sector contains one sectoral aggregator firm and a continuum of monopolistically competitive producers. The producers face aggregate and sectoral productivity shocks, and use sectoral capital and labor. Given demand for labor by each sectoral firm, labor is inelastically supplied by a sectoral labor union. Each union sets wages on behalf of all of its members.

In the baseline, type A households work in sector 1 and type B households work in sector 2. Total labor demand is dictated by the sectoral production firms to the unions, and the economy operates in one of two regimes. In the normal regime, wages are set by the unions on behalf of their members and employment is determined by labor demand. In extreme episodes, activity in the affected sector is rationed and different pricing and wage-setting conditions apply. Section 3.8 describes the negative shocks that the model admits and the regime under which each operates.

Omitted from the figure but described below are the fiscal and monetary policy rules, the capital and investment decisions, and the financial intermediaries. In what follows, index j refers to the firm (1 or 2), index h to the household type (A or B), and index k to individual firms and/or households.

Figure 1: Partial Illustration of Model Structure



3.1 Households

A continuum $k \in [0, 1]$ of infinitely lived households consuming both sectoral goods and supplying labor. Households inelastically supply labor to their sector's union. There are two types of households, indexed by $h \in \{A, B\}$, with a fraction $\chi_A \in [0, 1]$ of type A and $\chi_B = 1 - \chi_A$ of type B.[5] Households face idiosyncratic productivity risk e, which follows a stationary Markov process with $E[e] = 1$.

Each household of type h solves:

$$V_{h,t}(e, b_-, a_-) = \max_{c,b,a} \; u(c) - \varphi_h \frac{n_{h,t}^{1+1/\nu}}{1+1/\nu} + \beta_h E_t[V_{h,t+1}(e', b, a)] \tag{2}$$
$$\text{s.t. } P_{h,t}\, c + b + a = z_{h,t}(e) + T_{h,t} + (1 + r_t^a)a_- + (1 + r_t^b)b_- - \Psi(a, a_-).$$

The utility function is $u(c) = c^{1-\gamma}/(1-\gamma)$, where γ is the inverse elasticity of intertemporal substitution, c is a basket of the two sectoral goods with type-specific price index $P_{h,t}$, defined below, ν is the Frisch elasticity, and φ_h is a type-specific weight on the disutility of work. Note that labor supply is determined by the union and does not enter the household's optimization problem directly. The household nonetheless experiences the disutility of the hours $n_{h,t}$ selected on its behalf, where $n_{h,t} \equiv n_{1,h,t} + n_{2,h,t}$ is total hours worked across sectors. Labor income, $z_{h,t}(e)$, is given by:

$$z_{h,t}(e) = (1 - \tau_t)(w_{1,t}\, n_{1,h,t} + w_{2,t}\, n_{2,h,t})e, \tag{3}$$

[5] In the data, types are defined by industry attachment; Section 5.1 details the classification. The notation below allows households to work in both sectors, and I examine robustness to the classification cutoff and to temporary reallocation of workers across sectors during rationing episodes, a well-documented response to hours constraints on a primary job (Paxson and Sicherman, 1996; Barrero, Bloom and Davis, 2020).

where $T_{h,t}$ is a tax-free transfer to each household. Hours are assigned by union membership. Each household of type h holds memberships in the two sectors' unions in fixed shares $m_{j,h} \in [0,1]$, with $m_{1,h} + m_{2,h} = 1$, and each sector's unions allocate the sector's total hours uniformly across their members, so that $n_{j,h,t} = m_{j,h}\,\tilde{n}_{j,t}$, where $\tilde{n}_{j,t} = N_{j,t}/(\chi_A m_{j,A} + \chi_B m_{j,B})$ is hours per unit of membership in sector j. In the baseline, $m_{1,A} = 1$ and $m_{1,B} = 0$: type A works only in sector 1 and type B only in sector 2, so $n_{1,A,t} = n_{A,t}$ and $n_{2,B,t} = n_{B,t}$. Liquid positions may be negative up to a type-specific limit, $b \geq \underline{b}_h$, set to 2.22 months of type h's mean net income following the unsecured-credit-limit evidence in Kaplan and Violante (2014); negative balances pay the liquid rate plus a borrowing wedge, calibrated to the spread between unsecured borrowing and deposit rates.

$$\Psi(a, a_-) = \frac{\chi_1}{\chi_2}\left|\frac{a - (1 + r_t^a)a_-}{(1 + r_t^a)a_- + \chi_0}\right|^{\chi_2}\left[(1 + r_t^a)a_- + \chi_0\right]. \tag{4}$$

Consumption is a constant elasticity of substitution basket of the two sectoral goods with type-specific weights:

$$c_{h,t} = \left[\alpha_h^{\frac{1}{\eta}}\left(c_{1,h,t}\right)^{\frac{\eta-1}{\eta}} + (1 - \alpha_h)^{\frac{1}{\eta}}\left(c_{2,h,t}\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}, \tag{5}$$

where α_h is the steady-state expenditure share of type h on good 1 and η is the elasticity of substitution between the two goods. The corresponding price index, in numeraire units, is:

$$P_{h,t} = \left[\alpha_h + (1 - \alpha_h)\, q_t^{1-\eta}\right]^{\frac{1}{1-\eta}}, \tag{6}$$

and intratemporal optimization yields the sectoral demands $c_{1,h,t} = \alpha_h P_{h,t}^{\eta}\, c_{h,t}$ and $c_{2,h,t} = (1 - \alpha_h)\left(q_t/P_{h,t}\right)^{-\eta} c_{h,t}$, derived in Appendix A.1. Denoting by $\pi_{j,t}$ the inflation rate of the sector-j price, the relative price evolves according to:

$$q_t = q_{t-1}\,\frac{1 + \pi_{2,t}}{1 + \pi_{1,t}}. \tag{7}$$

Since the two types have different expenditure weights, a change in the relative price moves the cost of living differently across types, and this incidence channel is active whenever a negative shock moves q_t.

3.2 Firms

Each sector $j \in \{1, 2\}$ consists of a sectoral aggregator and a continuum of intermediate goods firms. The sectoral aggregator combines the intermediate goods into the sectoral

good, $Y_{j,t}$, with a constant elasticity of substitution of $\frac{\mu}{\mu-1}$. Each intermediate goods firm k in sector j has a Cobb-Douglas production function over capital and labor, subject to an aggregate and sectoral shock:

$$y_{k,j,t} = F(k_{j,k,t-1}, n_{k,j,t}) = Z_t Z_{j,t} k_{j,k,t-1}^{\alpha_j} n_{k,j,t}^{1-\alpha_j}. \tag{8}$$

The parameter α_j is the capital share in sector j. Each firm sets its price subject to quadratic adjustment costs and chooses capital subject to quadratic adjustment costs. Aggregating across the sector produces a sectoral Phillips curve and sectoral Tobin's Q. Appendix A.2.1 details the derivation of these objects. Since wages are quoted in the numeraire, sector-2 marginal cost is deflated by the relative price, and sector-2 revenues and investment are valued at q_t, so that dividends are $d_{1,t} = Y_{1,t} - w_{1,t}N_{1,t} - I_{1,t}$ and $d_{2,t} = q_t\left(Y_{2,t} - I_{2,t}\right) - w_{2,t}N_{2,t}$, where $I_{j,t}$ denotes investment inclusive of its adjustment costs. The price adjustment penalty disciplines price setting but is not a resource cost.[6]

3.3 Labor Unions

In each sector, a labor packer purchases employment services from a continuum of labor unions and packs them into a composite labor good, $N_{j,t}$, with a constant elasticity of substitution, $\frac{\mu_w}{\mu_w-1}$. The composite labor good is then leased to the continuum of intermediate goods in that sector for wage $W_{j,t}$.

Households provide differentiated labor input to each sector's unions, which gives them some pricing power in setting their own wage. Each differentiated labor input's wage is set by a corresponding union, and each sector's unions represent their member households, with membership shares $m_{j,h}$ as described above. Therefore, unions set wages to maximize the utility of their members, taking as given their consumption and savings decisions. Unions face quadratic utility costs of adjusting wages, in the spirit of Auclert et al. (2024). This generates a sectoral wage Phillips curve in each sector, derived in Appendix A.3. As such, sectoral wage inflation is given by:

$$1 + \pi_{j,t}^w = \frac{w_{j,t}}{w_{j,t-1}}(1 + \pi_{1,t}). \tag{9}$$

In the baseline assignment, sector j's members are the households of type $h(j)$, with

[6]Price and wage adjustment penalties enter the pricing conditions but are not subtracted from dividends or from the resource constraints. Wage adjustment penalties are utility costs to the union, following Auclert et al. (2024), and I adopt the same convention for price setting.

$h(1) = A$ and $h(2) = B$, and the wage Phillips curve in sector j is:

$$\log(1 + \pi^w_{j,t}) = \kappa_w \, n_{h,t} \left(\varphi_h \, n^{1/\nu}_{h,t} - \frac{1 - \tau_t}{\mu_w} \, U_{h,t} \, w_{j,t} \right) + \beta_{u,j} \log(1 + \pi^w_{j,t+1}), \tag{10}$$

where $U_{h,t} \equiv \int e \, c_{h,t}(e, b, a)^{-\gamma} \, dD_{h,t}$ is the average marginal utility of income across the distribution $D_{h,t}$ of type-h households over idiosyncratic states, and $\beta_{u,j}$ is the membership-weighted discount factor of sector j's unions, $\beta_{u,j} \equiv \vartheta_{A,j} \, \beta_A + \vartheta_{B,j} \, \beta_B$, where $\vartheta_{h,j} = \chi_h m_{j,h} / (\chi_A m_{j,A} + \chi_B m_{j,B})$ is the share of sector j's membership from type h. In the baseline, $\beta_{u,1} = \beta_A$ and $\beta_{u,2} = \beta_B$. Wage inflation rises when the marginal disutility of work falls short of the marked-down utility value of the marginal wage payment, so that a demand expansion which lowers members' marginal utility of income generates wage pressure. The weights φ_h are set so that equation (10) holds with zero wage inflation in steady state. For interior membership shares, a sector's unions draw members from both types; Appendix A.3 states the wage Phillips curve for that general case. Section 3.8 describes how wage setting changes during episodes in which hours in a sector are rationed.

3.4 Financial Intermediary

A representative financial intermediary takes deposits from households. Illiquid assets and liquid assets, transformed at proportional cost ω, are invested in government bonds, $B_{G,t}$, and firm equity in each sector, $p_{j,t}$ (see Appendix A.7 for details). The intermediary also provides unsecured credit: negative liquid balances pay the liquid rate plus a wedge, and the wedge payments are the real cost of providing credit, incurred by the intermediary in the same way as the transformation cost ω.

By assuming that capital gains accrue to the illiquid asset held by the household, ex-post returns are equal to:

$$1 + r_t = \frac{1 + i_{t-1}}{1 + \pi_{1,t}} = 1 + r^b_t + \omega, \tag{11}$$

and:

$$1 + r^a_t = \Theta_{1,t} \left(\frac{d_{1,t} + p_{1,t}}{p_{1,t-1}} \right) + \Theta_{2,t} \left(\frac{d_{2,t} + p_{2,t}}{p_{2,t-1}} \right) + (1 - \Theta_{1,t} - \Theta_{2,t})(1 + r_t), \tag{12}$$

where $\Theta_{j,t}$ is the share of firm j equity in the illiquid portfolio, and the residual share of the portfolio is held in government bonds.

3.5 Fiscal Policy

The fiscal authority uses a proportional labor income tax, τ_t, and one-period nominal debt, $B_{G,t}$, to finance its spending, G_t.

$$B_{G,t} + \tau_t \sum_{h\in\{A,B\}} \sum_{j\in\{1,2\}} w_{j,t} N_{j,h,t} = (1 + r_t) B_{G,t-1} + G_t + \chi_A T_{A,t} + \chi_B T_{B,t}. \tag{13}$$

Government spending is allocated across the two goods with a share α_G, so that $G_{1,t} = \alpha_G\, G_t$ and $G_{2,t} = (1 - \alpha_G)\, G_t/q_t$; α_G is determined residually by sectoral goods-market clearing at the steady state (Appendix B.1), since the sectoral size and the household expenditure weights are separate measurements. Debt follows the rule:

$$B_{G,t} - B_G^{ss} = \rho_B \left(B_{G,t-1} - B_G^{ss}\right) + \phi_T\, \mathcal{D}_t, \qquad \mathcal{D}_t \equiv \chi_A T_{A,t} + \chi_B T_{B,t} + \tau^{ss} \left(W^{ss} - W_t\right), \tag{14}$$

where $W_t \equiv w_{1,t} N_{1,t} + w_{2,t} N_{2,t}$ is the aggregate wage bill. The parameter ϕ_T is the share of new primary deficits, whether from stimulus transfers or from a recession's erosion of the tax base, that is debt-financed at impact, and ρ_B governs the pace at which the resulting debt is repaid. Given the debt path, the labor tax τ_t adjusts to balance the flow budget constraint. As a result, both the negative shock and the stimulus are deficit-financed on impact and repaid gradually through the labor tax.

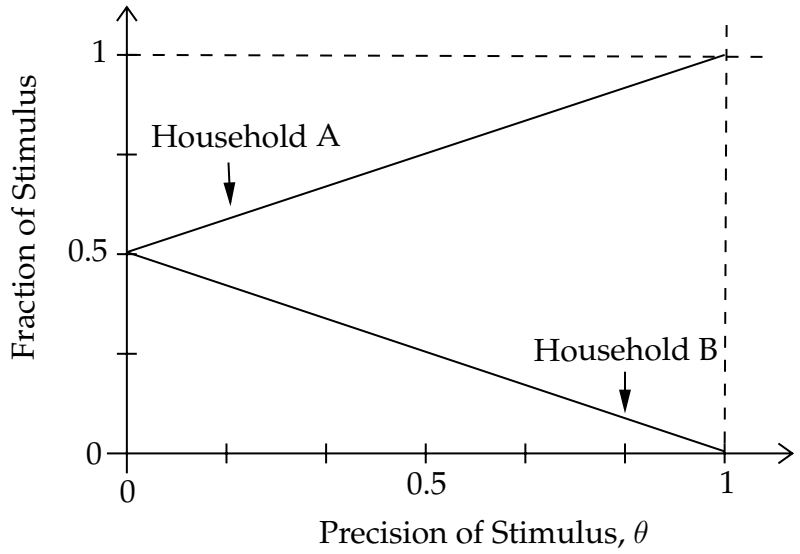
3.5.1 The Screening Problem

The stimulus programs studied in this paper are lump-sum transfers that are zero in steady state and are announced when the shock arrives. A program has three elements: its size, S, the precision with which it is directed toward the affected sector, θ, and the month in which it is paid, d. To parameterize precision, let $\theta \in [0, 1]$ and define the transfer to each type as

$$\frac{S_A}{S} = \theta + (1 - \theta)\chi_A$$
$$\frac{S_B}{S} = (1 - \theta)\chi_B.$$

Figure 2 illustrates how stimulus is allocated between the two types as precision, θ, varies. This specification is chosen to deliver intuitive values of total transfers for the polar cases of targeting. When $\theta = 0$, stimulus is not targeted toward either type but is distributed evenly across the population, $S_A = \chi_A S$ and $S_B = \chi_B S$. At the other extreme, when $\theta = 1$, stimulus is directed entirely toward the affected sector and type B receives no transfer,

Figure 2: Relation Between Stimulus Precision and Distribution to Each Household



Notes: Illustration of the relationship between the precision of stimulus, θ, and the fraction of stimulus allocated to each household.

$S_A = S$ and $S_B = 0$. In prose, θ is the share of the program that reaches the affected sector, typically the households that lost their jobs, rather than the sector in which households did not.

Fiscal transfers are disbursed to each type as one-time payments. The total stimulus to each type, S_A and S_B, is paid as a single lump-sum transfer in month d, $T_{A,t} = S_A \mathbb{1}[t = d]$ and $T_{B,t} = S_B \mathbb{1}[t = d]$, and the size and timing of the program are announced in the same month that the shock is realized. In order to compare programs with different delays, the present value of the program is held fixed, so that a program disbursed in month d has face value $S(1 + r)^d$. Every (θ, d) pair therefore costs the government the same amount in steady-state present-value terms, and programs differ only in who is paid and when they are paid. The transfer itself is financed according to the debt rule in equation (14), so that it is deficit-financed at impact and repaid through the labor income tax.

A program with precision θ cannot be paid before its screen has run, and so the feasible programs satisfy $d \geq D(\theta)$, where $D(\theta)$ is the screening frontier in equation (1) that is measured for the pandemic programs in Section 2.

The precision parameter θ in the model is mapped to the precision measured in Section 2 through the endpoints of each. In the data, precision is measured on the loss criterion, that is, the concentration of program dollars on households that lost employment income, since this is the criterion on which the enacted programs screened, whereas the transfer rule in the model concentrates dollars on the affected sector. The two definitions coincide

at both endpoints: at $\theta = 0$ both are the universal check, and at $\theta = 1$ the data object is unemployment insurance, which pays only households that lost a job, while the model object is a program that pays only the affected sector. The mapping between the two rests on the magnitudes involved. In the lockdown episode of Section 5, gross labor income in the affected sector falls by 20.8% at impact and by 6.7% on average over the window, compared to 3.2% and 1.1% in the other sector, which is a factor of 6.4 at impact. Similarly, 28.3% of the labor force drew unemployment insurance in 2020, which compares to an affected-sector share of 0.294 in the model. In this sense, the affected sector in the model corresponds to the households that lost their jobs in the data.

The affected sector is therefore the model's representation of the households that lost their jobs. Note that the open sector's loss is priced by this representation, and not set aside by it, since whatever a screen on the affected sector leaves uncompensated is exactly the cost of screening that Section 5 measures.

3.6 Monetary Policy

The monetary authority uses a Taylor rule to set the nominal rate on bonds:

$$i_t = \rho_i \, i_{t-1} + (1 - \rho_i) \left(r^* + \phi_\pi \, \pi_t^* \right), \qquad \pi_t^* = \varpi_t \, \pi_{1,t} + (1 - \varpi_t) \, \pi_{2,t}. \tag{15}$$

The target π_t^* is a stickiness-weighted index of sectoral inflation, with the weight on each sector proportional to its size in the population consumption basket and to the degree of its price stickiness:

$$\varpi_t = \frac{\bar{\alpha}/\kappa_{1,t}}{\bar{\alpha}/\kappa_{1,t} + (1 - \bar{\alpha})/\kappa_{2,t}}, \tag{16}$$

where $\kappa_{j,t}$ is the slope of the sector-j price Phillips curve. This weighting follows the welfare-based loss functions of Aoki (2001) and Woodford (2003), in which the weight on a sector's inflation is proportional to the distortions its price dispersion creates, and Benigno (2004) generalizes the principle to arbitrary differences in stickiness. Since the two sectors share a common slope, the weight reduces to $\varpi_t = \bar{\alpha}$: the target is headline inflation $\bar{\pi}_t$, the growth rate of the population price index $\bar{P}_t$, and the rule is a standard Taylor rule. Note that these results establish the weighting principle in representative-agent economies; I adopt the rule as a policy specification rather than derive it as optimal in the present model.

Partial adjustment of the policy rate, with smoothing coefficient ρ_i, is a standard feature of estimated policy rules (Clarida, Galí and Gertler, 2000; Coibion and Gorodnichenko, 2012), and the long-run response to inflation is ϕ_π for any value of ρ_i. To discipline the cal-

ibrated value, I estimate the partial-adjustment rule above on monthly data, in which the policy rate is regressed on its own lag, twelve-month core PCE inflation, and the unemployment gap. The estimates of ρ_i range between 0.96 and 0.99 across every pre-pandemic sample, and the baseline sets $\rho_i = 0.95$, at the lower edge of this range. Two caveats apply to this mapping, and both cut in the same direction. The estimated rules include an unemployment-gap term, whereas the rule in the model does not, and Rudebusch (2002) argues that estimated inertia is inflated by serially correlated omitted variables, so that the estimates are an upper bound on genuine partial adjustment. As such, the baseline takes the lower edge of the estimated range.

Note that with $\rho_i = 0$ the rule is purely contemporaneous and the policy block contributes no predetermined state, and in a two-sector economy with sticky prices *and* sticky wages in both sectors, the Taylor principle alone does not then deliver a locally unique equilibrium. Appendix A.8 documents the indeterminacy along with a horizon-invariance test, which the calibrated rule passes and $\rho_i = 0$ fails.

3.7 Market Clearing

The joint decision of the labor unions to decide total labor supplied to each firm and, during rationing episodes, the rationing constraint ensure that the sectoral labor markets clear in each period:

$$N_{1,t} = \chi_A \, n_{1,A,t} + \chi_B \, n_{1,B,t}, \tag{17}$$
$$N_{2,t} = \chi_A \, n_{2,A,t} + \chi_B \, n_{2,B,t},$$

which in the baseline reduces to $N_{1,t} = \chi_A n_{A,t}$ and $N_{2,t} = \chi_B n_{B,t}$. Denote by $C_{j,t}$ the total household demand for good j, aggregated across both types with population masses, and by Ξ_t the sum of household portfolio adjustment costs and intermediation costs (the transformation cost ω on liquid balances and the borrowing wedge on negative balances, both evaluated on beginning-of-period positions $b_{h,t-1}$, since that is the stock on which households pay them), which are incurred in units of the population basket with expenditure shares $s_{1,t}$ and $s_{2,t}$. The two goods markets clear:

$$Y_{1,t} = C_{1,t} + I_{1,t} + G_{1,t} + s_{1,t} \, \Xi_t, \tag{18}$$
$$Y_{2,t} = C_{2,t} + I_{2,t} + G_{2,t} + \frac{s_{2,t}}{q_t} \, \Xi_t. \tag{19}$$

Finally, the asset market clears when total household wealth equals the value of equity and government debt:

$$A_t + B_t = p_{1,t} + p_{2,t} + B_{G,t}, \tag{20}$$

where A_t and B_t are aggregate illiquid and liquid holdings.

3.8 Rationing Episodes

The negative shock that the model admits is an extreme episode in which price-mediated adjustment in the affected sector is suspended and quantities are rationed directly. The quantitative analysis applies this regime to the pandemic lockdown, which is calibrated in Section 5.1, and Appendix A.6 relates the regime to fixed-price equilibria and to the closures used in the pandemic literature.

3.8.1 The Constraint

Activity is rationed by an outside constraint, and the model carries an explicit flag for the periods in which the constraint is active. The episode is a window $\mathcal{T} = \{0, 1, \ldots, t^*\}$ during which hours in sector 1 are constrained to an exogenous path:

$$N_{1,t} = \mathcal{N}_t, \qquad t \in \mathcal{T}, \tag{21}$$

where $\mathcal{N}_t$ is the imposed hours path of the affected industries. The constraint binds on sector totals: hours per unit of membership, $\tilde{n}_{1,t}$, fall equiproportionally for all sector-1 members, leaving members' hours in the unaffected sector unconstrained, and in the baseline the episode falls entirely on type A. Sectoral productivity is unchanged throughout the episode: the constraint restricts labor input, and output falls through the production function in equation (8). The constraint disturbs two markets, the market for the affected sector's output and the sector's labor market, and the episode is completed by specifying how each operates during the window. In a lockdown neither side of the affected market operates: firms cannot produce and households cannot buy.

3.8.2 The Rationed Market

The sector-1 price is predetermined during the window,

$$\pi_{1,t} = 0, \qquad t \in \mathcal{T}, \tag{22}$$

and so the goods market clears by rationing quantities, and households cannot purchase their notional demand at the posted price.[7] As in Neary and Roberts (1980), the ration is allocated through a single market-wide scarcity, in which all households choose freely at a common *virtual* price $\tilde{p}_{1,t} \geq p_{1,t}$ but pay the posted price, and Appendix A.4 gives the construction. The virtual price is the new equilibrium object, and it adjusts so that rationed demand absorbs the rationed supply,

$$\sum_h \chi_h \, c_{1,h}(\tilde{p}_{1,t}) + G_{1,t} + I_{1,t} = Y_{1,t}, \qquad t \in \mathcal{T}, \tag{23}$$

and it is the shadow price of the constrained good. In a full shutdown Guerrieri et al. (2022) describe this shadow price as "prohibitively high", whereas it is finite in this model since the restriction is partial.

The episode generates forced saving as a result, which is the mechanism by which a supply restriction transmits as a demand shortage in Guerrieri et al. (2022). Since the two goods are complements, households cannot redirect the expenditure that the ration frees toward close substitutes, and part of it is postponed.[8]

I verify ex post that the ration binds. The virtual price strictly exceeds the posted price at every date of the window (its minimum over the window is 1.02 in the baseline calibration), and a solution in which it does not is rejected, since a virtual price below the posted price would describe excess supply.

3.8.3 Wages and Policy in the Window

During the window, hours in sector 1 are fixed by the rationing constraint, so no adjustment of the sector-1 wage can recruit or withdraw employment, and since the sector-1 price is also predetermined, there is no firm optimality condition which determines the wage. As such, the wage has no allocative role in the window, and wage-setting theory delivers no unique prediction for it. I follow Hall (2005) in treating the wage as an equilibrium selection, in which any wage in the bargaining set is consistent with equilibrium and a sticky wage norm selects one of them. As Barro (1977) argues, the wage in an ongoing employment relationship is an installment payment within that relationship and is not a price that clears a spot market, and this logic applies with particular force when employ-

[7] During the pandemic, the prices of locked-down services did not spike. Spending on these services collapsed and household saving surged, which is the pattern of quantity rationing at a fixed price that the regime in the model follows.

[8] Under strong complementarity, a market-clearing counterfactual in which the good remains freely traded admits no equilibrium relative price for hours restrictions larger than roughly half the observed pandemic decline, and so the demand side has to be rationed.

ment is fixed by an outside constraint. In the window, sector-1 wage inflation therefore follows an error-correction norm anchored to the steady-state real wage (Blanchard and Katz, 1999; Blanchard and Galí, 2007):

$$\log(1+\pi^{w}_{1,t}) = \kappa_w \left(\frac{w_{1,t}}{w^{ss}_{1}} - 1 \right) + \beta_{u,1} \log(1+\pi^{w}_{1,t+1}), \qquad t \in \mathcal{T}. \tag{24}$$

The sector-2 union, whose labor market is never rationed, sets wages according to equation (10) at every date, and sector 1 returns to equation (10) when the window closes; all normal-regime conditions then resume, and households and firms have perfect foresight over the entire transition path.

The rationing episode also affects the monetary rule, since the pegged sector-1 price generates no sector-1 inflation during the window. The rule in equation (15) reduces to $i_t = \rho_i \, i_{t-1} + (1-\rho_i)\,(r^* + \phi_\pi \, \pi_{2,t})$ inside the window, and the sector-2 and headline rules coincide mechanically. Note that the scarcity which the market cannot express through the posted price appears instead in the virtual price, which is not a transacted price and to which no policy rule should respond (Aoki, 2001; Guerrieri, Lorenzoni, Straub and Werning, 2021).

3.9 Welfare

Welfare is measured from the value function at the date the shock arrives. For each type, I partition households into quartiles of steady-state cash on hand and hold the partition fixed at its steady-state definition, so that the expected value of a cohort at date 0 is the expected lifetime utility of the households who begin the transition in that cohort. For cohort k, the consumption-equivalent welfare change λ_k solves:

$$(1+\lambda_k)^{1-\gamma} \, E\left[V^{c,ss} \,|\, k\right] = E\left[V^{c}_{0} \,|\, k\right], \tag{25}$$

where V^c denotes the consumption component of the value function, i.e., expected lifetime utility in equation (2) with the labor disutility terms removed. The baseline welfare measure is consumption-only. In normal times, hours are demand-determined under union wage setting, and during extreme episodes they are rationed, so in both regimes movements in hours occur off of households' labor supply curves. For this reason, reduced hours are not credited as leisure. This follows the convention in Guerrieri et al. (2022), and it is consistent with evidence that the value of forgone work time falls in downturns (Chodorow-Reich and Karabarbounis, 2016). As a robustness case, I report a full welfare measure that adds to the consumption value the leisure value of the rationed

hours shortfall (sector-1 members, window dates only) computed with the labor disutility in equation (2) and scaled by $\xi \in [0,1]$; I report $\xi \in \{0,1\}$, the consumption-only baseline and an upper bound.[9] As a second robustness case, I relax the assumption on how the rationed quantity is distributed across buyers. The baseline implicitly assumes that all buyers receive an equal ration, while Holden (2025) shows that sticky-price firms ration by exclusion, in which orders are filled in full or not at all, so that variety is destroyed and the welfare cost is deeper at identical aggregate quantities and prices. In Appendix A.5, I derive the resulting wedge and show that it alters no equilibrium object, as with ξ. The wedge is reported as a robustness column alongside the ξ columns. Since both are welfare-only overlays on the same equilibrium, and since the headline measure is the consumption-only, equal-rations case ($\xi = 0$, $\Delta = 1$), it would be inconsistent to promote one overlay to the headline and not the other. Aggregate welfare is the population-weighted average of cohort welfare changes. I report every welfare change as a percentage of annual consumption, multiplying λ by $\bar{H} = \frac{1}{12}\sum_h \chi_h/(1-\beta_h) = 14.06$, the population-weighted discounted number of years over which a permanent change in consumption applies. Under logarithmic utility the result is the present value of the consumption-equivalent change, discounted at the households' rate of time preference, as a share of one year's normal consumption.Stimulus designs are compared on this measure in the crisis economy. Section 5 also reports the welfare of the same transfer in the economy without the shock, which decomposes the return to a design into the part that is specific to the crisis and the part that would obtain in normal times. This comparison is read across precision at a fixed payment date, and it is never itself the objective.

3.10 Solution

The model is coded and solved using the toolkit developed in Auclert et al. (2021). It is solved at a monthly frequency over a horizon of 600 months, by the end of which every variable is back at its steady-state value. Each transition is a nonlinear perfect-foresight path, and I solve for the path using a Newton method in which the steady-state Jacobian serves as the preconditioner. Since the steady-state Jacobian contains no information about the rationing constraint, the rationing window is spliced into the preconditioner. Each of the 2,525 programs on the grid of precision and delay is solved as a separate nonlinear transition, with welfare computed from the value function at the onset of the shock. Since the effect of a transfer in a first-order solution does not depend on whether

[9]Since preferences are separable and hours are never a household choice, the value of ξ affects no equilibrium object, and the welfare-only measure is exact. Under non-separable preferences (e.g., GHH), the leisure value would instead enter labor supply.

it is paid during a crisis or in normal times, a nonlinear solution is required to generate the crisis-specific return to targeting in Section 5.3.1.

4 Calibration

The model is calibrated at a monthly frequency to the U.S. economy before the pandemic. Table 3 lists the parameters that would keep their values in any application of the two-sector economy, with the value of each and the target or convention that sets it, whereas the parameters that depend on which industries were shut, including the two discount factors, are calibrated in Section 5.1 along with the shock itself. Only one parameter of the general economy is targeted, the scale of the portfolio adjustment cost, which is set to the quarterly marginal propensity to consume out of a payment the size of the first Economic Impact Payment, and this moment is 0.23 in the model compared to a range of 0.15 to 0.25 in the data. The remaining parameters are either external, including the two parameters of the debt rule,[10] or determined by internal consistency, and Appendix B.1 documents each of the latter.

Nothing in the calibration targets the population-level moments in the last panel of the table. The annual marginal propensity to consume is 0.42, which lies within the range that the literature reports, and the labor share is 0.66, which lies at the top of its range. The latter reflects the fact that the price markup is implied by the value of equity, and is not set by convention.

The calibration does not match the wealth distribution. Mean wealth in the data is about seven times median wealth, and the asset market forces the model to match mean wealth. Since the model has no right tail, the median household in the model is far richer than the median household in the data, and the shares of households that hold little or no liquid wealth are too low as a result. This miss matters for the results in one direction only. Households with no liquid wealth are the households for whom a transfer during a crisis is worth the most, and it follows that the value that the model assigns to any transfer in Section 5 is a lower bound.

[10] The elasticity of substitution between the two goods is the estimated value of Cantelmo and Melina (2023), Section 3.6 discusses the smoothing coefficient of the policy rule, and Table 3 gives the source of each of the other external parameters. A recession financed with a larger tax response at impact would repay faster than the debt rule assumes.

Table 3: Calibration of the two-sector two-asset HANK economy

Parameter		Value	Target
Households			
γ	Inverse IES (log utility)	1.0	log utility
φ_A, φ_B	Disutility of labor, by type	*implied*	union FOC at $\pi^w = 0$
ν	Frisch elasticity	1.0	convention
χ_0	Portfolio adj. cost pivot	0.25	irrelevant; χ_1 absorbs it
χ_1	Portfolio adj. cost scale	60.0	quarterly MPC, EIP-sized: data 0.15–0.25, model 0.233
χ_2	Portfolio adj. cost curvature	2.0	imposed by the solver
$\underline{b}_h$	Borrowing limit, by type	$-2.22\,E[y_h]$	SCF credit limit / income
κ_b	Borrowing wedge	0.10	unsecured borrowing cost
ρ_e	Autocorrelation of earnings	0.966	standard quarterly value
σ_e	Cross-sectional std of log earnings	0.92	PSID $p90/p10$: data 11.2, model 11.18
η	Elasticity of substitution across goods	0.349	estimated demand system
Labor Unions			
μ_w	Steady-state wage markup	1.10	convention
κ_w	Slope of wage Phillips curve	0.10	convention, quarterly
Firms			
α_j	Capital share, sector j	0.276	firms' FOC
Z_j	Sectoral productivity	*implied*	firms' FOC
μ_p	Steady-state price markup	1.096	$p = \text{wealth} - B_G$
δ	Depreciation rate	0.0580	BEA depreciation / net stock
κ_p	Slope of price Phillips curve	0.10	convention, quarterly
ϵ_I	Capital adjustment cost	4.0	convention
K	Capital stock	$2.304\,Y^{\text{ann}}$	BEA fixed assets / GDP
Financial Intermediary			
r	Real return on illiquid assets	0.0509	illiquid return, 5–6%
ω	Liquidity premium	0.0550	real liquid return: data $\approx -0.5\%$, model -0.0052
Θ_j	Equity shares in the illiquid portfolio	*implied*	asset-market clearing
Policy			
τ	Labor income tax	*implied*	government flow budget
G	Government spending	$0.176\,Y$	NIPA government $C + I$ / GDP
B_G	Government debt	$0.784\,Y^{\text{ann}}$	federal debt held by the public
ϕ_π	Taylor rule coefficient	1.5	convention
ρ_i	Policy-rate smoothing (monthly)	0.95	estimated monthly rules give 0.96–0.99; set conservatively
ϖ	Taylor weight on sector-1 inflation	$\bar{\alpha}$	stickiness weights
ρ_B	Debt repayment persistence	0.98	fitted: 2020–23 debt consolidation
ϕ_T	Deficit-financed share at impact	1.0	measured: 2020 deficit financing
Untargeted Outcomes			
	Annual MPC, EIP-Sized	0.416	data 0.4–0.5
	Labor Share	0.661	data 0.62–0.66
	Non-Labor Share (Capital + Profit)	0.339	data 0.35–0.38

Notes: Monthly model, $Y = 1$ per period. The parameters that descend from the partition of industries into the shut sector and the rest, including the two discount factors, are in Table 4 with the moments they are set to. One parameter is set to a target, χ_1 to the quarterly MPC, and the targeted rows report the data and model values side by side; everything else is measured, external, or implied by the requirement that the steady state be an equilibrium. The debt rule is fitted to the financing of the 2020 deficit and the 2020–23 consolidation. At the steady state, household wealth equals $K + B_G$ to within 10^{-7} and $I/Y = \delta K/Y = 0.137$.

5 The Lockdown

5.1 Calibrating the Lockdown

Table 4 reports the calibration of the lockdown, every entry of which follows from the partition of industries into the sectors that were shut and the rest. Sector 1 is the six industries in which more than 40% of employment was in occupations that the pandemic restricted, namely construction, real estate and rental, administrative and support services, arts and entertainment, accommodation and food services, and other services, as classified in the Panel Study of Income Dynamics. Type A is the households attached to the shut industries, who account for 0.294 of the population and 0.265 of output, and this income share is the size of the affected sector. The two expenditure weights on good 1 are the types' expenditure shares on the output of the same industries, which are measured from the PSID expenditure crosswalk, and the government's basket share on good 1 is then determined by sectoral goods-market clearing. The two discount factors are set so that each type's total wealth matches its ratio to annual income in the data, which is 3.43 for type A and 5.35 for type B, respectively.

Table 4: The lockdown economy: the partition and the discount factors

Parameter	Value	Description		
The Partition				
χ_A	0.2942	population share of the shut sectors, PSID 2019		
Y_1	0.2647	income share of the shut sectors, PSID 2019		
α_A, α_B	0.2808, 0.2724	expenditure shares on good 1, PSID crosswalk		
Household Moments				
		Moment	Data	Model
β_A	0.99379	Wealth / Annual Income, Type A	3.430	3.430
β_B	0.99418	Wealth / Annual Income, Type B	5.350	5.350

Notes: Sector 1 is the six industries with more than 40% pandemic attachment in the PSID 2019 classification, and type A the households attached to them. The first panel lists the measurements the partition fixes directly. The second lists the two parameters set to household moments: each discount factor is solved so that its type's total wealth matches its ratio to annual income (PSID 2019, weighted, KMV definition, housing-inclusive). The government basket share on good 1 is determined by sectoral goods-market clearing given these measurements (Appendix B.1).

The lockdown enters the model as the rationing episode of Section 3.8, and the path that it imposes on sector-1 hours is the measured collapse of the shut industries. Figure 3 plots the real wage bill of the six shut industries as a deviation from February 2020. At the trough, employment in these industries fell by 29.8% compared to 20.8% for the real wage bill, and the gap between the two is due to composition, since low-wage workers lost jobs disproportionately and the average wage of those who remained rose. Transfers

substitute for lost income, and so the wage bill is the moment to which the constraint is sized, as imposing the employment path would destroy more income than the data show. The imposed path is a smooth pulse fitted to the de-trended series, pinned at the trough and decaying at the series' own recovery rate, and the constraint binds for the $t^* = 24$ months from April 2020 in which the pulse is more than 1% below zero.[11]

5.2 The Lockdown in the Model

Figure 4 plots the impulse responses to the lockdown when there is no transfer. Hours in sector 1 are imposed, and so the response of hours in sector 2 is the first of two untargeted checks on the model. Hours in the sector that was never shut fall by 3.2% at the onset of the episode and recover over the window, compared to a decline of 8.5% in the de-trended data, so the model produces the correct sign and about a third of the magnitude. On a mechanical level, this reflects the fact that the rationed households lose income and, since the two goods are complements, reduce their spending on both goods, which lowers demand for the open sector's good and with it hours and wages in the open sector. It is this linkage between the two sectors that the rest of the paper prices. The right panel of Figure 4 plots the incidence of the episode by type. At the onset, gross labor income falls by 20.8% for type A compared to 3.2% for type B, a factor of 6.4, and the average declines over the window are 6.7% and 1.1%, respectively. However, consumption and income move differently for the two types. Consumption of the rationed type falls by far less than its income, 3.9% at the trough, since this type borrows against the reopening, whereas consumption of the other type falls by about as much as its income, 3.0%. After the window, both consumption paths settle below steady state, since the labor tax rises in order to repay the debt that the episode ran up.

Prices do not appear in the figure since they play no role in the comparison. Inside the window, the price of the sector-1 good is pegged and the price of the open sector's good falls, so that no transfer paid within the window is constrained by inflation, which is the key difference between a lockdown and an ordinary recession. The kink at the close of the window corresponds to the reopening, at which point the peg is lifted and the sector-1 wage returns from the norm to the union's first-order condition.

The Household Pulse Survey provides the second untargeted check on the model, since it records what recipients spent the first payments on. Table 5 reports, for each

[11] Given the measured employment and wage-bill paths, the productivity path that the production function implies is approximately flat, and so holding technology fixed is consistent with the measured series. The rationing wedge, the ratio of the virtual price to the posted price, is 1.22 at the onset and falls to 1.02 by the close of the window.

round of Economic Impact Payments, the share of recipients who reported spending the payment on each category, separately for households that had lost employment income and households that had not. In the first round, households that had lost income were more likely to report spending the payment on rent, utilities, vehicle payments and food, by 19, 22, 14 and 17 pp, respectively, and less likely to report spending it on every discretionary category or saving it, and the sign of the gap is the same in every category and every round. This pattern is (indirectly) consistent with the model, as a household in a lockdown cannot spend a transfer on the rationed good, and so the spending of affected households goes to what can still be bought.

Table 5: What the payment was spent on, by work-loss status

	Round 1		Round 2		Round 3	
Category	Lost Income	Did Not	Lost Income	Did Not	Lost Income	Did Not
Rent	38.3	19.4	32.8	15.7	45.5	24.1
Utilities	62.3	40.4	50.8	31.5	58.8	39.0
Vehicle Payments	31.4	17.5	25.0	13.5	32.3	20.0
Mortgage	29.3	18.8	20.4	13.5	21.5	15.6
Food	77.2	59.9	59.1	46.6	66.4	53.8
Supplies	56.2	44.0	35.7	27.7	47.0	35.3
Clothing	18.5	17.7	10.3	9.5	18.4	17.4
Debt Paydown	23.3	23.9	32.0	28.9	33.7	32.4
Household Items	5.4	8.3	4.5	6.5	6.4	8.7
Recreational Goods	2.4	3.6	1.8	2.5	2.5	3.1
Donations	3.5	7.4	2.5	5.1	3.2	4.5
Savings	9.6	19.9	11.3	20.1	10.4	21.0

Notes: Household Pulse Survey, household weights, over the weeks of each round. Each entry is the share of recipients that selected the category, with multiple selections allowed, separately by whether the household lost employment income. The categories above the rule are necessities and recurring bills, and the categories below it are discretionary purchases, transfers and saving.

With no transfer, the episode costs the population 7.7% of annual consumption, as the first row of Table 7 in the next subsection reports. Type A, the rationed type, loses 11.1% compared to 6.3% for type B, which no restriction named and whose sector never shut, and since type B accounts for 0.706 of households, it bears 57.5% of the aggregate loss, all of which comes through the linkage. The data show the same spread. Table 6 crosses sector and loss of employment income for the pandemic, and 0.474 of households in the shut sectors lost employment income compared to 0.302 in the open sectors. Since the open sectors are larger, 60% of the households that lost income worked in one of them, compared to the 57.5% of the welfare loss that the open sector bears in the model.

Table 6: Who lost employment income, by sector: share of households

	Lost Income	Did Not	All
Shut Sector	0.139	0.155	0.294
Open Sector	0.213	0.493	0.706
All	0.353	0.647	1.000

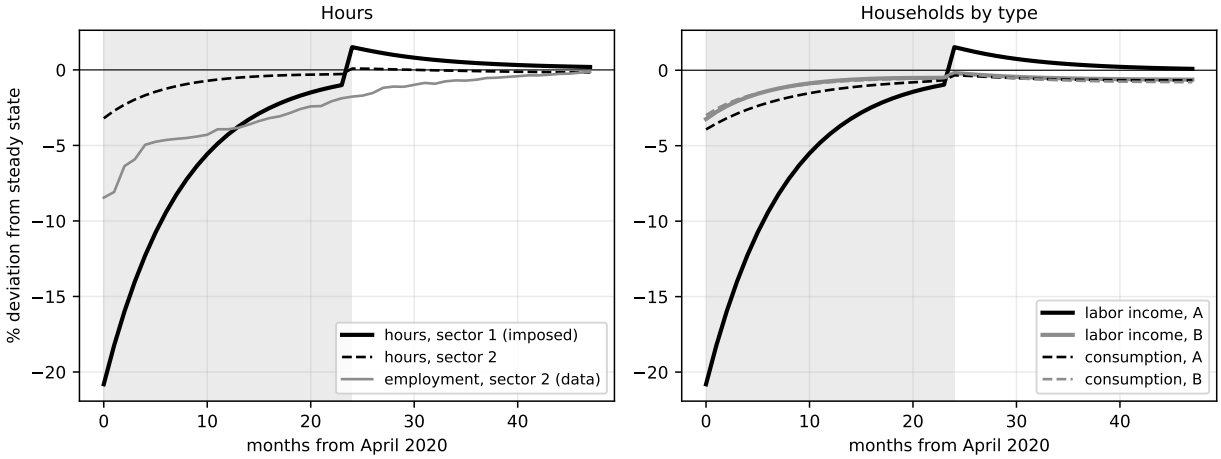
Notes: PSID 2021. The shut sector is the six-industry partition with more than 40% pandemic attachment (Section 5.1), and loss is any loss of employment income since March 2020. Each entry is a share of all households. Within the shut sector, a share 0.474 of households lost employment income, compared to 0.302 in the open sector.

Figure 3: The lockdown episode: the imposed path



Notes: The measured real wage bill of the six shut industries, as a deviation from February 2020 (BLS CES, CPI-deflated), and the fitted de-trended pulse that the model imposes on sector-1 hours. The shaded region is the window of $t^* = 24$ months in which the constraint binds.

Figure 4: Impulse responses to the lockdown with no stimulus



Notes: Nonlinear perfect-foresight paths with no transfer, as percentage deviations from steady state. Left: hours in the shut sector, which are imposed, and hours in the open sector, which are predicted, against the de-trended measured employment of the open industries. Right: gross labor income, the wage times hours, and consumption, by type. The shaded region is the window; the kink at its close is the reopening.

5.3 Economic Impact Payments and Unemployment Insurance

The two programs I study are the first round of Economic Impact Payments and the unemployment insurance package, both of which were enacted at once in the CARES Act and which correspond to the two endpoints of the frontier, and each is evaluated as the model represents it. Table 7 reports the loss by type with no transfer and under each program as paid.[12]

The upper panel of Table 7 reports the loss under each design and the lower panel the change relative to no transfer, and the contrast between the two designs is the paper in one table. The universal check changes little. It reduces the loss of the shut sector by 0.19 pp and the loss of the open sector by 0.31 pp, so that the population's loss falls from 7.7% of annual consumption to 7.4%, a reduction of 4%. The check reaches the households that bear most of the damage, and it does more per household for the open sector than for the shut sector, since the shut sector cannot spend the payment on its own rationed good, which is why the check is the design the open sector prefers. Full targeting, on the other hand, changes a great deal for each sector and little for the population. It removes 78% of the shut sector's loss, from 11.1% to 2.4%, and raises the open sector's loss by 43%, from 6.3% to 9.0%, since the open sector receives nothing while its customers in the shut sector are paid. The population's loss falls to 7.1%, a reduction of 8%, which is larger than under

[12] Appendix D.6 shows that the two welfare-only overlays of Section 3.9 shift every design by a common amount and change no comparison.

Table 7: The lockdown, the universal check, and the screen: welfare loss by type

	Type A (Shut)	Type B (Open)	Population
No Transfer	-11.14	-6.28	-7.71
Universal Check ($\theta = 0$, $d = 0.81$)	-10.96	-5.97	-7.44
Fully Targeted ($\theta = 1$, $d = 3.45$)	-2.44	-9.00	-7.06
Change Relative to No Transfer, pp			
Universal Check	+0.19	+0.31	+0.27
Fully Targeted	+8.70	-2.72	+0.65

Notes: Consumption-equivalent welfare change from the lockdown, as a percentage of annual consumption, by type and for the population, measured from the value function at the onset (Section 3.9), consumption-only. Type A is the shut sector, 0.294 of households, and type B is the open sector. The two designs are the same PV-equalized transfer (1.28% of annual output), with the universal check paid at the frontier's floor $D(0)$ and the fully targeted design paid at $D(1)$. Values at non-integer delays are interpolated in d. The lower panel is each design's row less the no-transfer row.

the check by 0.38 pp.

Targeting in a lockdown is therefore mostly a transfer between the two sectors: what the shut sector gains, the open sector largely loses, and the net for the population is small. Both programs reduce the loss, and the screen reduces it by more, but the case for the screen rests on that small net, and the case for the check rests on the incidence. The remainder of this section accounts for the net, which requires explaining why it is positive and why it is small.

The programs differ along two dimensions, precision and payment date, and so the model separates the net into a component due to each. I first hold the payment date fixed and study what is bought by moving the dollar from the open sector to the shut sector, which is the key mechanism of the paper (Section 5.3.1), and then study the cost of the months that the screen took (Section 5.3.2). With both designs paid at the onset, the return to precision is 0.47 pp, of which 0.09 pp is forfeited by the screen's delay, and the remaining 0.38 pp is the difference between the two enacted rows of Table 7.

5.3.1 From the Check to the Screen

I first study precision on its own by holding the payment date at the onset for both designs, so that the screen is free and only precision varies. Section 5.3.2 then restores the dates at which the programs were actually paid. As the notes to Table 7 indicate, paying the check at its actual date rather than at the onset changes its value by a tenth of what precision is worth. Table 8 reports the effect on the welfare of each group of moving the dollar from type B to type A, and Appendix Figure A.5 plots the same objects over the full

range of θ. The first column reports the effects in the lockdown. Concentrating the transfer entirely on the rationed type reduces the loss of that type by 8.74 pp and raises the loss of type B by 2.99 pp, and the population's loss falls by 0.47 pp. Full precision is the design that makes the loss smallest. The amount that it recovers is 2.6 times the amount recovered by the universal check, which is small relative to a crisis that costs 7.7% of a year's consumption but large relative to the size of the transfer.

The interpretation of the first column depends on the second, which repeats the same reallocation in the economy with no shock. In normal times, concentrating a check of this size on 0.294 of households raises the population's loss by 1.11 pp. On a mechanical level, this reflects the fact that the marginal utility of consumption is decreasing, so that the households who are left out of the program would have valued the transfer more than the households who receive a larger share of it. Net of this normal-times cost, the return to precision that is specific to the crisis is 1.58 pp in the third column, which is 3.4 times the return that the crisis surface itself displays.

The third column also identifies where the crisis-specific return sits, and it lies entirely with type B. The reduction in type A's loss from being targeted is the same in the crisis as in normal times, 8.74 against 9.07 pp, and is if anything slightly smaller in the crisis, since type A cannot spend the money on the rationed good. On the other hand, taking type B's check and giving it to type A costs type B 5.36 pp in normal times but only 2.99 pp in the crisis, so that the demand linkage returns 44.2% of every reallocated dollar to the household that gave it up. This figure is a lower bound on the spillback, since the marginal-utility effect runs in the opposite direction: type B is poorer in the crisis, and a lost dollar should therefore hurt it more.

Appendix Table A.11 reports the paths of the same two designs solved with their full responses. When the shut sector is paid, its consumption response is the same in the crisis as in normal times. Since the shut sector's own good is rationed inside the window, its extra purchases of that good come partly out of the open sector's ration, and the rest of its spending flows to the output of the open sector, which is why the open sector's wage rises in the crisis whereas it falls in normal times, and why its consumption loss is a third of what the same reallocation costs it in normal times. A dollar paid to the open sector, however, cannot reach the shut sector for as long as the shutdown lasts, since that sector cannot sell more, and the open sector's extra demand for the rationed good only bids the ration away from the shut sector's own households. Thus, the linkage between the two sectors runs in only one direction during the lockdown.

Overall, moving from the check to the screen reduces the loss, and full precision reduces it by the most. However, the return to precision is only about a third as large as a

tagging model would attribute to it, since the dollar that leaks to the open sector is not lost. Further, the optimum at full precision is unusually robust, and Appendix Figure A.5 illustrates why. I fit each column as a quadratic in θ and find that the crisis-specific return is nearly linear, with a curvature of 3.4% of its slope, whereas the crisis surface itself has a curvature of 21.7% of its slope. The crisis surface is the small residual of two large components with opposite signs, which is why it appears far more curved than either of them, and the normal-times concentration cost accounts for 56% of its curvature. There is thus little curvature in the return to precision for a frontier to work against.

Table 8: The return to targeting with free screening ($d = 0$): moving the dollar from B to A

	Crisis	Normal Times	Crisis-Specific
Type A (Shut)	+8.74	+9.07	-0.33
Type B (Open)	-2.99	-5.36	+2.37
Population	+0.47	-1.11	+1.58

Notes: Change in the consumption-equivalent welfare of each group, in percentage points of annual consumption, from raising θ from 0 to 1 with the transfer paid at $d = 0$. The transfer is the same PV-equalized transfer, sized to the first round of Economic Impact Payments at 1.28% of annual output, and it is concentrated entirely on type A. *In the crisis* is the lockdown surface, *in normal times* is the identical transfer in the economy with no shock, and the last column is their difference.

Appendix Figure A.5 shows that the loss of the open sector rises with precision at every point, both in the crisis and in normal times, so that a planner who counted only the open sector would choose the universal check at every delay. A planner who weights the two types by population, with a weight of 0.294 on the shut sector, chooses full precision, but only narrowly, since the weight at which the optimum on the frontier reaches full precision is 0.024 below the population share (Appendix Figure A.6).[13] A planner who weights households by the marginal utility of consumption lies further inside this region, at a weight of 0.356, and as such full precision reduces the loss most under both of these weightings.

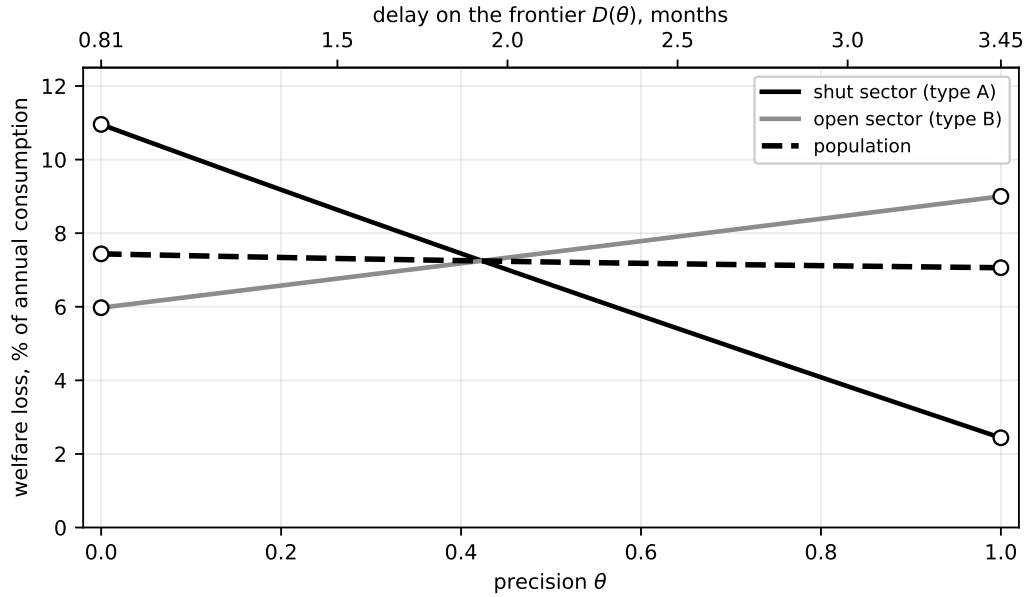
5.3.2 The Price of the Screen

Figure 5 plots the welfare loss along the frontier for the shut sector, open sector, and total population. The lower horizontal axis plots the precision of each program and the upper the delay. The vertical axis plots the welfare loss as a percentage of annual consumption.

[13] As the weight on the shut sector rises from zero, the optimum on the frontier leaves the universal check at 0.245 and reaches full precision at 0.270.

As precision rises, the shut sector's loss falls steeply, and the open sector's loss increases, but at a far slower rate. This figure highlights two features which are apparent in the full plane plotted in Appendix Figure A.3. First, the loss rises with delay at every precision, since no household in the model benefits from delay. Second, the loss decreases with precision at every delay, so that the unconstrained optimum is, naturally, the corner at which the transfer is fully targeted and has zero delay. The key tradeoff in the model is therefore that the precision of the transfer is increasing in the delay until it can be paid, and, along the frontier, the fully targeted design for any delay is the design that reduces the population's loss by the largest amount.

Figure 5: The welfare loss along the frontier



Notes: Welfare loss from the lockdown as a percentage of annual consumption for a transfer with precision θ paid at the date the frontier of Section 2 allows.

Table 9 prices the screen by decomposing the net of Table 7. With both designs paid at the onset, full targeting reduces the population's loss by 0.47 pp more than the universal check. Paying the fully targeted design when its screen allows, 3.45 months after the onset, forfeits 0.09 pp of that return, which is 19.3% of it, and the remaining 0.38 pp is the net between the two enacted rows. Appendix Table A.9 prices the frontier on the 90th percentile of dollars, which raises the share forfeited to 40.9%, and sets its endpoint to each of the three tiers of the package in turn, for which the share ranges from 13.3% to 46.9%. In every case, the precision that minimizes the loss along the frontier is one.

Appendix Table A.10 reports the cost of a month of delay for the fully targeted transfer, both in the crisis and in normal times. Over the first four months, which is the range of

Table 9: The price of the screen

	pp
Return to Targeting, Both Paid at the Onset	+0.47
Forfeited by Delay, $D(1) = 3.45$ Months	-0.09
Net	+0.38
Share Forfeited	19.3%

Notes: Population welfare on the lockdown surface, with W interpolated linearly in d at each θ. The first row is $W(1,0) - W(0,0)$; the net row is $W(1, D(1)) - W(0, D(0))$, the difference between the two enacted rows of Table 7; the forfeited row is their difference, and the share divides it by the first row. Appendix Table A.9 prices the frontier on the 90th percentile of dollars and on each tier of the package.

delays that the frontier spans, the population loses 0.033 pp per month of delay in the crisis compared to 0.052 pp in normal times. The recipient loses 0.080 pp per month in the crisis compared to 0.043 pp in normal times, nearly 1.9 times as much, since it is poorer and its marginal utility is higher, while type B, which is never paid, loses 0.013 pp per month of the recipient's delay in the crisis compared to 0.056 pp in normal times, a factor of 4.3 less. Since the ration is fixed inside the window, a transfer paid there reallocates the ration among buyers through the virtual price, which is forced saving at the level of the aggregate ration, and not of the recipient. As a result, the population's delay cost is lower in the crisis, since the 0.706 of households that are not paid have little left to lose from the wait.

On the surface of Appendix Figure A.3, the loss under the fully targeted design is smaller than the loss under the universal check paid at its own floor at every delay up to 11.8 months, compared to the 3.45 months that the screen took. On the other hand, holding the floor fixed, the optimum along the frontier leaves full precision once the endpoint exceeds 8.36 months, which is longer than the delay of every tier of the package.

At the delays that real screens took, delay was not the term that decided the design. Note that this is a statement about the calibrated population, and that it is narrower than a statement about the most constrained households in 2020, since the calibration has too few households with little or no liquid wealth, as Section 4 records, and these are the households for whom a month of waiting is most costly. The delay costs above therefore understate what waiting cost them.

5.4 Implications for Policy

The key result of the paper is that even though the lockdown restricted only one sector, most of the welfare loss is borne by households employed in the other sector, which

bears 57.5% of the aggregate loss through the demand linkage. This is the justification for the universal check, since despite never being restricted, these "unaffected" households are affected by the lockdown nonetheless, and a program that reaches every household reaches the majority of the damage. At the same time, the linkage runs in only one direction in a lockdown. A dollar paid to the shut sector is spent on the good produced by the open sector, and 44.2% of every dollar taken from the open sector returns to it as demand, whereas a dollar paid to the open sector stays there. Which households a program pays therefore matters less than a tagging model implies, and when it does matter, the direction is set by which sector can still sell, since paying the sector that cannot sell also pays the sector that can. Overall, the population as a whole narrowly prefers that the transfer be paid to the shut sector, while the open sector, which bears most of the damage, prefers the check.

6 Conclusion

I measure the trade off between the precision and delay of pandemic transfer programs of 2020 in a two-sector HANK model. In a lockdown, the sector that was never restricted bears most of the welfare loss, so the justification of untargeted transfers is their reach, not their speediness. While full screening is the best program, the screen's delay, which is the cost of precision, forfeits nearly 20% of the benefits to precision.

The analysis in this paper is undertaken in a lockdown, in which the shut sector is unable to sell for as long as the restriction lasts. In a supply-side recession, however, the hit sector is still able to sell, and so the linkage between the two sectors may run in both directions, in which case the direction of the design lesson would no longer be given by which sector can sell. Computing the same surface for such a recession would say whether the lesson of the lockdown reverses.[14] Of course, the results in this paper also obtain in a model that averages over the households that the lockdown devastated and those that it barely touched. A model with a discrete employment margin within each sector would separate these households, and a program in that model could screen on the loss itself instead of on the sector in which it fell.

[14]Since both prices adjust in such a recession, precision may also buy the same relief with less aggregate demand. In that case the surface would price a second reason to screen, which is that screening economizes on inflation as well as on dollars.

References

Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles, "The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data," *Journal of Political Economy*, Dec. 2007, *115* (6), 986–1019.

Akerlof, George A., "The Economics of "Tagging" as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning," *American Economic Review*, 1978, *68* (1), 8–19.

Andersen, Asger, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan, "Consumer Responses to the COVID-19 Crisis: Evidence from Bank Account Transaction Data," Technical Report, Working Paper. 2020.

Aoki, Kosuke, "Optimal Monetary Policy Responses to Relative-Price Changes," *Journal of Monetary Economics*, 2001, *48* (1), 55–80.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, Sept. 2021, *89* (5), 2375–2408.

_ , Matthew Rognlie, and Ludwig Straub, "The Intertemporal Keynesian Cross," *Journal of Political Economy*, 2024, *132* (12).

Baker, Scott R., R. A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis, "Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments," September 2020. NBER Working Paper No. 27097.

Baqaee, David and Emmanuel Farhi, "Supply and Demand in Disaggregated Keynesian Economies with an Application to the COVID-19 Crisis," *American Economic Review*, 2022, *112* (5), 1397–1436.

Barrero, Jose Maria, Nicholas Bloom, and Steven J. Davis, "COVID-19 Is Also a Reallocation Shock," *Brookings Papers on Economic Activity*, 2020. COVID-19 and Reallocation special edition.

Barro, Robert J., "Long-Term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics*, 1977, *3* (3), 305–316.

_ and Herschel I. Grossman, "A General Disequilibrium Model of Income and Employment," *American Economic Review*, 1971, *61* (1), 82–93.

Bayer, Christian, Benjamin Born, Ralph Luetticke, and Gernot J. Müller, "The Coronavirus Stimulus Package: How Large is the Transfer Multiplier?," *The Economic Journal*, 2023, *133* (652), 1318–1347.

Benigno, Pierpaolo, "Optimal Monetary Policy in a Currency Area," *Journal of International Economics*, 2004, *63* (2), 293–320.

Bilbiie, Florin O., Tommaso Monacelli, and Roberto Perotti, "Public Debt and Redistribution with Borrowing Constraints," *The Economic Journal*, Feb. 2013, *123* (566).

Blanchard, Olivier and Jordi Galí, "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, 2007, *39* (s1), 35–65.

_ **and Lawrence F. Katz**, "Wage Dynamics: Reconciling Theory and Evidence," *American Economic Review*, 1999, *89* (2), 69–74.

Broda, Christian and David E. Weinstein, "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 2006, *121* (2), 541–585.

_ **and Jonathan A. Parker**, "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption," *Journal of Monetary Economics*, Fall 2014, *68.*

Cantelmo, Alessandro and Giovanni Melina, "Sectoral labor mobility and optimal monetary policy," *Macroeconomic Dynamics*, Jan. 2023, *27* (1), 1–26.

Carroll, Christopher D., Edmund Crawley, Jiri Slacalek, and Matthew N. White, "Modeling the Consumption Response to the CARES Act," *Quantitative Economics*, 2026. TODO verify volume/pages before submission.

Chen, Haiqiang, Wenlan Qian, and Qiang Wen, "The Impact of the COVID-19 Pandemic on Consumption: Learning from High Frequency Transaction Data," Technical Report, Working Paper. 2020.

Chodorow-Reich, Gabriel and Loukas Karabarbounis, "The Cyclicality of the Opportunity Cost of Employment," *Journal of Political Economy*, 2016, *124* (6), 1563–1618.

Clarida, Richard, Jordi Galí, and Mark Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *The Quarterly Journal of Economics*, 2000, *115* (1), 147–180.

Coibion, Olivier and Yuriy Gorodnichenko, "Why Are Target Interest Rate Changes So Persistent?," *American Economic Journal: Macroeconomics*, 2012, *4* (4), 126–162.

_ **,** _ **, and Michael Weber**, "How Did U.S. Consumers Use Their Stimulus Payments?," August 2020. Working Paper.

di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A. Yildirim, "Quantifying the Inflationary Impact of Fiscal Stimulus Under Supply Constraints," Technical Report Jan. 2023. NBER Working Paper No. 30892.

Dunn, Abe C., Kyle K. Hood, and Alexander Driessen, "Measuring the Effects of the COVID-19 Pandemic on Consumer Spending Using Card Transaction Data," Technical Report, Working Paper. 2020.

Fields, JF, J Hunter-Childs, A Tersine, J Sisson, E Parker, V Velkoff, C Logan, and H Shin, "Design and Operation of the 2020 Household Pulse Survey," *U.S. Census Bureau*, Forthcoming.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, "Monetary Policy in Times of Structural Reallocation," in "Macroeconomic Policy in an Uneven Economy, Jackson Hole Economic Policy Symposium," Federal Reserve Bank of Kansas City, 2021.

_ , _ , _ , and _ , "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?," *American Economic Review*, May 2022, *112* (5), 1437–1474.

Hale, Galina, John C. Leer, and Fernanda Nechio, "Inflationary Effects of Fiscal Support to Households and Firms," Technical Report Jan. 2023. NBER Working Paper No. 30906.

Hall, Robert E., "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 2005, *95* (1), 50–65.

Holden, Tom D., "Rationing Under Sticky Prices," 2025. Working paper, Deutsche Bundesbank.

International Monetary Fund, "Fiscal Monitor: Policies for the Recovery," Washington, October 2020.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, Dec. 2006, *96* (5).

_ , _ , and _ , "The Response of Consumer Spending to Rebates During an Expansion: Evidence from the 2003 Child Tax Credit," April 2009. Working Paper.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, July 2014, *82* (4), 1199–1239.

_ , Ben Moll, and Gianluca Violante, "Monetary Policy According to HANK," *American Economic Review*, 2018, *3* (108), 697–743.

Karger, Ezra and Aastha Rajan, "Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments," May 2020. FRB of Chicago Working Paper No. WP 2020-15.

Kleven, Henrik Jacobsen and Wojciech Kopczuk, "Transfer Program Complexity and the Take-Up of Social Benefits," *American Economic Journal: Economic Policy*, 2011, *3* (1), 54–90.

Malinvaud, Edmond, *The Theory of Unemployment Reconsidered*, Oxford: Basil Blackwell, 1977.

Mehrotra, Neil R., "Fiscal Policy Stabilization: Purchases or Transfers?," *International Journal of Central Banking*, March 2018, *14* (2).

Neary, J. Peter and Kevin W. S. Roberts, "The Theory of Household Behaviour under Rationing," *European Economic Review*, 1980, *13*, 25–42.

Oh, Hyunseung and Ricardo Reis, "Targeted Transfers and the Fiscal Response to the Great Recession," *Journal of Monetary Economics*, Dec. 2012, *59* (Supplement), S50–S64.

Orchard, Jacob, Valerie A. Ramey, and Johannes F. Wieland, "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates," *The Quarterly Journal of Economics*, 2025, *140* (3).

Parker, Jonathan A., "On Measuring the Effects of Fiscal Policy in Recessions," *Journal of Economic Literature*, September 2011, *49* (3), 703–718.

_ , Nicholas S. Souleles, David S. Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, October 2013, *103* (6), 2530–2553.

Parker, Jonathan and Nicholas Souleles, "Reported Effects vs. Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates," *American Economic Review: Insights*, 2019, *1* (3), 273–290.

_ , Jake Schild, Laura Erhard, and David Johnson, "Economic Impact Payments and Household Spending During the Pandemic," *Brookings Papers on Economic Activity*, Forthcoming.

Paxson, Christina H. and Nachum Sicherman, "The Dynamics of Dual Job Holding and Job Mobility," *Journal of Labor Economics*, 1996, *14* (3), 357–393.

Rudebusch, Glenn D., "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia," *Journal of Monetary Economics*, 2002, *49* (6), 1161–1187.

Sahm, Claudia, Matthew Shapiro, and Joel Slemrod, "Direct Stimulus Payments to Individuals in the Covid Pandemic," 2020. Presented in Virtual Macro Seminar on 06/11/2020. See `https://www.youtube.com/watch?v=lnzHSSDbVF0`.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod, "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," *Tax Policy and the Economy*, 2010, *24* (1), 69–110.

Schild, Jake and Thesia Garner, "Consumer Response to Economic Impact Payments during the COVID-19 Pandemic and the Role of Subjective Assessments of Well-Being: A View from the U.S. Using a Rapid Response Survey," Technical Report, Working Paper. 2021.

Schmitt-Grohé, Stephanie and Martín Uribe, "Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment," *Journal of Political Economy*, 2016, *124* (5), 1466–1514.

Shapiro, Matthew D. and Joel Slemrod, "Consumer Response to Tax Rebates," *American Economic Review*, 2003, *93* (1), 381–396.

_ **and** _ , "Did the 2008 Tax Rebates Stimulate Spending?," *American Economic Review: Papers & Proceedings*, 2009, *99* (2), 374–379.

Townsend, Robert M., "Optimal Contracts and Competitive Markets with Costly State Verification," *Journal of Economic Theory*, 1979, *21* (2), 265–293.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ: Princeton University Press, 2003.

A Model

A.1 Households

Given total consumption expenditure $P_{h,t}c_{h,t}$, the household allocates spending across the two goods to minimize cost. In numeraire units, the first-order conditions of the expenditure minimization problem yield the sectoral demands:

$$c_{1,h,t} = \alpha_h \left(\frac{1}{P_{h,t}}\right)^{-\eta} c_{h,t} = \alpha_h P_{h,t}^{\eta}\, c_{h,t},$$
$$c_{2,h,t} = (1-\alpha_h)\left(\frac{q_t}{P_{h,t}}\right)^{-\eta} c_{h,t},$$

and substituting the demands into the basket delivers the type-specific price index in equation (6). Since dividing the second condition by the first shows that relative expenditure on the two goods depends only on α_h and $q_t^{1-\eta}$, it follows that the two types respond identically to a movement in the relative price, up to their expenditure weights. The population price index aggregates the two goods with the weight $\bar{\alpha} = \chi_A\alpha_A + \chi_B\alpha_B$, and the relative price accumulates sectoral inflation differentials according to equation (7).

A.2 Firms

A.2.1 Sectoral Goods Firms

Each firm sets its price subject to quadratic adjustment costs relative to the aggregate price index in the previous period:

$$\psi_{k,j,t} = \frac{\mu}{\mu-1}\frac{1}{2\kappa}\left[\log\left(\frac{P_{k,j,t}}{P_{t-1}}\right)\right]^2 Y_{j,t}.$$

The penalty disciplines price setting but is not a resource cost: it is excluded from dividends and from the sectoral resource constraints. This yields a Phillips curve in each sector:

$$\log(1+\pi_{j,t}) = \kappa\left(mc_{j,t} - \frac{1}{\mu}\right) + \frac{1}{1+r_{t+1}}\frac{Y_{j,t+1}}{Y_{j,t}}\log(1+\pi_{j,t+1}),$$

with:

$$mc_{1,t} = \frac{w_{1,t}}{F_N(K_{1,t-1}, N_{1,t})}, \quad mc_{2,t} = \frac{w_{2,t}/q_t}{F_N(K_{2,t-1}, N_{2,t})}, \quad 1+\pi_{2,t} = \frac{q_t}{q_{t-1}}(1+\pi_{1,t}).$$

Firms choose capital subject to quadratic adjustment costs:

$$\varsigma\left(\frac{k_{k,j,t}}{k_{k,j,t-1}}\right)k_{k,j,t-1}, \quad \varsigma(x) = x - (1-\delta) + \frac{1}{2\delta\epsilon_I}(x-1)^2,$$

with depreciation $\delta \in (0,1)$. Dividends equal output net of labor costs and investment:

$$d_{1,t} = Y_{1,t} - w_{1,t}N_{1,t} - I_{1,t}, \quad d_{2,t} = q_t\left(Y_{2,t} - I_{2,t}\right) - w_{2,t}N_{2,t},$$

with investment defined as:

$$I_{j,t} = K_{j,t} - (1-\delta)K_{j,t-1} + \varsigma\left(\frac{K_{j,t}}{K_{j,t-1}}\right)K_{j,t-1}.$$

Then, capital evolves according to:

$$(1+r_{t+1})Q_{j,t} = \alpha_j \frac{Y_{j,t+1}}{K_{j,t}} mc_{j,t+1} - \left[\frac{K_{j,t+1}}{K_{j,t}} - (1-\delta) + \frac{1}{2\delta\epsilon_I}\left(\frac{K_{j,t+1}-K_{j,t}}{K_{j,t}}\right)^2\right] + \frac{K_{j,t+1}}{K_{j,t}}Q_{j,t+1},$$

with Tobin's Q in each sector given by:

$$Q_{j,t} = 1 + \frac{1}{\delta\epsilon_I}\frac{K_{j,t}-K_{j,t-1}}{K_{j,t-1}}.$$

For sector 2, the valuation and investment conditions hold in numeraire units, so output and installation terms are valued at q_t; the model code states these conditions in full. During a rationing episode the sector-1 conditions carry the virtual price in the same place: investment goods for sector 1 *are* sector-1 output, so a firm buying them is rationed exactly as a household is, and the acquisition and installation terms are valued at $\tilde{p}_{1,t}$ rather than at the posted price. As in the household problem, dividends charge the posted price for the units actually obtained, since the virtual price governs the choice while the posted price settles the bill. Since government purchases are the one component of sector-1 demand that is not rationed, households and firms absorb the whole restriction, and the household ration is correspondingly more severe than the aggregate one.

A.3 Labor Unions

Unions impose a quadratic wage adjustment cost in utility terms:

$$\varrho(W_{k,j,t}, W_{k,j,t-1}) = \frac{\mu_w}{\mu_w - 1}\frac{1}{2\kappa_w}\left[\log\left(\frac{W_{k,j,t}}{W_{k,j,t-1}}\right)\right]^2.$$

In the symmetric equilibrium, the wage Phillips curve in each sector is given by:

$$\log(1+\pi^w_{j,t}) = \kappa_w \sum_{h\in\{A,B\}} \chi_h\, N_{j,h,t}\left(\varphi_h N_{h,t}^{1/\nu} - \frac{(1-\tau_t)w_{j,t}}{\mu_w}\int e_k\, u'(c_{h,k})\, dk\right) + \beta_{u,j}\log(1+\pi^w_{j,t+1}),$$

where hours are allocated across sectors by the membership shares $m_{j,h}$, so that $N_{j,h,t} = \chi_h\, m_{j,h}\, \tilde{n}_{j,t}$, and total hours $N_{h,t} = N_{1,h,t} + N_{2,h,t}$ enter the disutility term, so the marginal disutility of an hour is equalized across the sectors a household works in. In the baseline assignment, in which each sector's members are of a single type, only the type employed

in sector j appears in the sum, and writing hours per capita delivers equation (10) in the main text. During a rationing episode, however, there is no labor demand schedule for the union to optimize against, since equation (21) fixes sector-1 hours independently of the wage. Within the window, the wage is instead selected from the bargaining set by the norm in equation (24).

A.4 Rationed Demand and Virtual Prices

This appendix constructs the demand side of the lockdown. Proportional rationing of a good under homothetic CES preferences is equivalent to unconstrained choice at a common virtual price (Neary and Roberts, 1980): household h's sectoral demands during the window are the CES demands evaluated at the virtual price $\tilde{p}_{1,t}$ rather than the posted price $p_{1,t}$, with the virtual price common across households and chosen so that equation (23) holds.[15] Define the virtual price index and the good-1 quantity share

$$\tilde{P}_{h,t} = \left[\alpha_h\, \tilde{p}_{1,t}^{1-\eta} + (1-\alpha_h)\, p_{2,t}^{1-\eta}\right]^{\frac{1}{1-\eta}}, \qquad \phi_{h,t} = \alpha_h \left(\frac{\tilde{p}_{1,t}}{\tilde{P}_{h,t}}\right)^{-\eta}. \tag{26}$$

The household chooses its bundle at virtual prices but pays posted prices, so its outlay per unit of the consumption composite $C_{h,t}$ is

$$p_{1,t}\, c_{1,h,t} + p_{2,t}\, c_{2,h,t} = \left[\tilde{P}_{h,t} - (\tilde{p}_{1,t} - p_{1,t})\, \phi_{h,t}\right] C_{h,t} \equiv \hat{P}_{h,t}\, C_{h,t}, \tag{27}$$

and the composite the household attains is $C_{h,t} = X_{h,t}/\hat{P}_{h,t}$, where $X_{h,t}$ denotes its expenditure.

The intertemporal problem, however, does not keep its unrationed form; this is the one place where the ration does more than relabel a price. Two price indices appear in the household's problem. The effective outlay index $\hat{P}_{h,t}$ is the average cost of a unit of the composite, and is the correct index for the level of consumption (as in equation (27)). Since a household that cannot buy more of good 1 can only raise its composite by buying more of good 2, the cost of one further unit of the composite is instead given by the virtual index $\tilde{P}_{h,t}$:

$$\frac{\mathrm{d}X_{h,t}}{\mathrm{d}C_{h,t}} = \tilde{P}_{h,t}, \tag{28}$$

It follows that the marginal utility of expenditure is $u'(C_{h,t})/\tilde{P}_{h,t}$ and not $u'(C_{h,t})/\hat{P}_{h,t}$. The difference between the two indices,

$$\tilde{P}_{h,t} - \hat{P}_{h,t} = (\tilde{p}_{1,t} - p_{1,t})\, \phi_{h,t}, \tag{29}$$

is the shadow rent on the ration. When the ration does not bind, $\tilde{p}_{1,t} = p_{1,t}$ and both

[15] Note that with homothetic CES preferences a common virtual price cuts the purchases of every household by the same proportion only when $\alpha_A = \alpha_B$, since the cut depends on the expenditure weight α_h of each type. In the calibration, the two weights differ by less than a percentage point, 0.281 compared to 0.272, so the difference is quantitatively negligible here.

indices equal $P_{h,t}$, so the unrationed model is nested exactly.

Carrying $\hat{P}_{h,t}$ into the Euler equation as well would price a marginal unit at average cost. The error is second order in the depth of the ration ($\ln \hat{P}_{h,t}$ is of order Δ^2 whereas $\ln \tilde{P}_{h,t}$ is of order Δ), and so it is negligible for a shallow ration, but not for a lockdown. Under log utility, the single-index form reduces the marginal utility of expenditure to $1/X_{h,t}$, i.e., the ration drops out of the intertemporal margin altogether. In that case expenditure does not respond to the constraint, and since the budget identity routes the entire quantity freed from good 1 into good 2, the unaffected sector expands.

Welfare is evaluated with the household's true preferences at the constrained bundle. Utility is $u(C_{h,t})$, where $C_{h,t}$ is the composite the household actually attains, and the cost of the ration appears in the fact that $\hat{P}_{h,t} > P_{h,t}$, so that each unit of expenditure buys less of the composite than it would without the ration. The virtual price is a sufficient statistic for how binding the episode is, and I report the wedge $\tilde{p}_{1,t}/p_{1,t}$ with the quantitative results.

A.5 Exclusion Rationing and the Variety Wedge

The virtual-price construction above determines the total quantity of the rationed good that is sold, but it does not say anything about which buyers receive it. Implicitly, the construction assumes equal rations, i.e., that every buyer of every variety has their order scaled down by the same proportion, which is the least costly way of distributing the shortfall. Holden (2025) shows that when prices are sticky below marginal cost, firms do not ration proportionally, but instead ration by *exclusion*: some orders go unfilled and some varieties stock out, so that a buyer's order for a given variety is either filled in full or not at all. Under both distributions, the aggregate quantity, the posted prices, and measured expenditure are exactly the same, and the only difference is in welfare, since exclusion destroys variety.[16] In this appendix, I derive the welfare wedge between the two distributions, show that it does not alter any equilibrium object, and describe how it is measured and reported.

Sector 1 output is a CES composite of a continuum of varieties with elasticity of substitution $\varepsilon_1 > 1$, which in the model of Section 3.2 is the aggregator elasticity $\mu/(\mu-1)$; its measurement is discussed below. Let $m_t \in (0,1]$ denote the fraction of orders served at date t, where exclusion is independent across varieties, and fix the aggregate quantity Q_t of the sector-1 composite. With equal rations, every variety is scaled down uniformly, and so the household's effective composite is simply Q_t. Under exclusion, the household is only served on a measure m_t of varieties. It reallocates its sector-1 expenditure to the varieties that it can obtain and therefore receives Q_t/m_t of each of them, and so its effective composite is

$$\left[\int_0^{m_t} \left(\frac{Q_t}{m_t}\right)^{\frac{\varepsilon_1-1}{\varepsilon_1}} di\right]^{\frac{\varepsilon_1}{\varepsilon_1-1}} = m_t^{\frac{\varepsilon_1}{\varepsilon_1-1}} \cdot \frac{Q_t}{m_t} = m_t^{\frac{1}{\varepsilon_1-1}} Q_t. \tag{30}$$

[16] Holden (2025) derives a further implication that the model inherits: under exclusion, measured real output, which counts transacted quantities at posted prices, exceeds welfare-relevant output, and statistical-agency imputation of stocked-out items biases measured prices toward the posted price. The divergence between the model's measured and welfare-relevant sector-1 output during the window is the model counterpart of that observation.

The factor $m_t^{1/(\varepsilon_1-1)} < 1$ is the love-of-variety cost of concentrating the same aggregate quantity on fewer varieties. The wedge already accounts for ex-post substitution, since the excluded expenditure is not lost but reallocated to the varieties the household can obtain. It follows that the exponent is governed by the substitutability of varieties ε_1, so that a high ε_1 means easy substitution and a small loss, and that risk aversion does not enter.

The wedge enters welfare as follows. The equilibrium bundle $(c_{1,h,t}, c_{2,h,t})$ is unchanged, since the wedge is a measurement overlay on the same equilibrium, so the effective composite is the CES aggregate of $(m_t^{1/(\varepsilon_1-1)} c_{1,h,t},\ c_{2,h,t})$. Preferences are homothetic and the intratemporal split is the same for every household of a given type, so this amounts to scaling the composite of every type-h household by the same factor,

$$\Delta_{h,t} = \left[s_{1,h,t}\, m_t^{\frac{1}{\varepsilon_1-1}\cdot\frac{\eta-1}{\eta}} + (1 - s_{1,h,t}) \right]^{\frac{\eta}{\eta-1}}, \qquad s_{1,h,t} = \alpha_h \left(\frac{\tilde{p}_{1,t}}{\tilde{P}_{h,t}} \right)^{1-\eta}, \tag{31}$$

where $s_{1,h,t}$ is the weight on good 1 in the aggregator, evaluated at the virtual prices that support the equilibrium bundle. Under log utility, $\log(\Delta_{h,t} C_{h,t}) = \log C_{h,t} + \log \Delta_{h,t}$, i.e., the wedge is an additive term in flow utility that is deterministic and common within a type. Its present value, $\sum_{t \in \mathcal{T}} \beta_h^t \log \Delta_{h,t} \leq 0$, is therefore simply added to the consumption value V^c inside the transform in equation (25), exactly as the leisure term is.

I take the serving fraction from the rationing constraint itself, and set $m_t = \mathcal{N}_t / N_1^{ss}$ on window dates and $m_t = 1$ elsewhere. Note that the measured shortfall that rations hours in equation (21) is the same one that rations orders. For the variety elasticity, I report a grid of values, $\varepsilon_1 \in \{4, 7, 10\}$.[17] The wedge is reported alongside the ξ columns of the welfare tables, with the equal-rations case $\Delta = 1$ as the nested benchmark.

A.6 Relation to Fixed-Price and Lockdown Closures

Rationing at a fixed price has a long lineage in the fixed-price equilibria of Barro and Grossman (1971) and Malinvaud (1977), in which the price is predetermined and quantities take the short side of the market, and which the lockdown implements directly. The regime structure is similar to that in models of occasionally binding downward nominal wage rigidity, in which the economy moves between a regime where a wage rule binds and employment is rationed and a regime where wages are set freely (Schmitt-Grohé and Uribe, 2016; Baqaee and Farhi, 2022). In these papers, the switch between the two regimes is triggered endogenously by a wage floor. In this setting, however, the flag marks an observed episode. The window thus enters the model as dated, and is not triggered by a threshold condition.

As in the pandemic literature, the closure models the episode as a constraint on activity, and not as a change in preferences. Guerrieri et al. (2022) model the lockdown as the

[17] The model-consistent value is the aggregator elasticity $\mu/(\mu - 1)$ implied by the calibrated markup. Further, the goods-level estimates in Broda and Weinstein (2006) put median elasticities near the low end of this grid, with substantial mass above it. Since a higher ε_1 shrinks the wedge, the upper entries of the grid are conservative.

exclusion of workers from the affected sector, with technology held fixed and the sector's output correspondingly not consumed, and Guerrieri et al. (2021) make the corresponding argument that monetary policy should look through the relative-price movements that such episodes produce. The episode can also be modeled as a shock to the disutility of work, which preserves the wage schedule throughout (Bayer et al., 2023). Relative to all of these papers, I impose the rationing constraint directly. Since a closure based on preferences would treat nonemployment in the crisis as voluntary leisure, it would understate the welfare cost of the shock.

A.7 Financial Intermediary

In expectation, the return on nominal government bonds, the return on equity for each firm, and the economy-wide return, $E_t[1 + r_{t+1}]$, are all equal due to the absence of arbitrage:

$$\frac{1 + i_t}{1 + E_t[\pi_{1,t+1}]} = \frac{E_t[d_{1,t+1} + p_{1,t+1}]}{p_{1,t}} = \frac{E_t[d_{2,t+1} + p_{2,t+1}]}{p_{2,t}} = E_t[1 + r_{t+1}].$$

The financial intermediary passes these returns on to households net of intermediation costs:

$$E_t[1 + r^a_{t+1}] = E_t[1 + r^b_{t+1}] + \omega = E_t[1 + r_{t+1}].$$

The transformation cost and the borrowing wedge are earned on the positions that households hold entering the period, and so the corresponding resource cost in equations (18) and (19) is charged on $b_{h,t-1}$ and not on the current stock. With this beginning-of-period convention, the aggregate identity closes at every date along the transition, and not only in the steady state.

A.8 Determinacy and Horizon Invariance

The model is solved in the sequence space over a finite horizon T, and every variable is set to its steady state value beyond T. The terminal condition is an economic assertion as well as a numerical device, since it is innocuous only if the solution it selects is the one that an infinite-horizon agent would choose. In this appendix, I document a case in which this is not true, and the restriction on the policy rule that removes it.

With no smoothing in equation (15), the policy block contributes no predetermined state, and the price–wage–monetary subsystem has a complex pair of roots with modulus just inside the unit circle, approximately $0.997\, e^{\pm 0.080 i}$.[18] The requirement that the solution remain bounded does not select a root inside the unit circle, and there is no initial condition attached to it, so that its amplitude and phase are free and the equilibrium is locally indeterminate. Note that finite-horizon solution methods do not report this as a failure, and instead return a solution which depends on where the terminal condition is placed. On a mechanical level, this reflects the fact that the terminal condition supplies exactly the two restrictions that the model is missing.

[18] This pair implies a cycle of about seventy-eight months, with a decay half-life measured in decades.

The symptom of the indeterminacy is that the solution at date zero depends on the truncation horizon. Solving the sectoral productivity shock at $T = 300$, 600, 900 and 1200 months gives impact responses of affected-sector output of -7.9%, -11.0%, -18.8% and -39.8%, respectively, i.e., the answer diverges as the horizon grows. This is despite the fact that the shock itself is over within a year and is numerically identical across the four solves. Neither of the two natural explanations applies here. The first is that the system is poorly conditioned, but at $T = 900$ the linear system solves to a backward error of 10^{-12}, and re-solving it after full equilibration changes the answer in the twelfth digit. The second is the slow real transition of the capital stock, but the relative price, which is what governs the sectoral allocation, is the single most horizon-stable object in the model.

Interest-rate smoothing supplies the missing state. At $\rho_i = 0.8$ the complex pair moves outside the unit circle, boundedness selects the solution again, and the horizon dependence disappears. I verify this with a test, reported for every specification, that solves the model at $T = 600$, 900 and 1200 and compares the paths and welfare over a common window that excludes each solve's own terminal boundary layer. The calibrated rule passes this test, with no change of sign in the welfare measure, and $\rho_i = 0$ fails it.[19] What separates the two cases is the profile of the disagreement as the comparison window lengthens, since a terminal boundary layer grows toward the end of the window and vanishes at its front, whereas genuine horizon dependence is already present at date zero, and so its profile is flat.[20]

B Calibration

The parameters of the model are set in three distinct ways, and the distinction matters for interpreting the calibration. A first group is measured directly in the data and imposed. A second group is chosen so that the steady state reproduces a set of targeted moments. A third group is not free at all: once the first two groups are set, these parameters are determined by the requirement that the steady state be an equilibrium. This appendix documents each group in turn, along with additional results on the benchmark calibration.

B.1 Internal Consistency

The third group of parameters is worth documenting separately, since its members are often reported alongside calibrated parameters without the distinction being drawn, and since the values they take are a useful diagnostic of whether the rest of the calibration hangs together.

[19] Under the calibrated rule the paths agree to within 0.4% of each variable's own peak response and the consumption-equivalent welfare measure agrees to 0.002 percentage points, whereas with $\rho_i = 0$ the same test gives disagreements of 180% and 0.9 percentage points, respectively.

[20] Since the indeterminate direction is a nominal one, welfare *levels* are the quantity most exposed to it. Differences netted across policies at a common horizon are far better protected, since a common error cancels.

Table A.1 lists the parameters determined by internal consistency and the condition that determines each. Three cases are worth describing in more detail.

Sectoral goods-market clearing and the government basket. The government expenditure weight α_G does not target any moment. It is determined by the requirement that the market for the affected good clear in steady state, given the size of the affected sector, Y_1, which is measured as the affected type's share of labor income, and the household expenditure weights α_A and α_B, which are measured from the expenditure crosswalk. Since the relative price is normalized, the composition of government purchases is the natural residual, and α_G is determined by

$$Y_1 = C_1 + I_1 + \alpha_G G + s_1 \Xi, \tag{32}$$

rather than assumed to equal the population expenditure share $\bar{\alpha}$.

Labor disutility weights. The weights φ_A and φ_B do not target any moment. They are determined by the requirement that the union first-order condition in equation (10) hold with zero wage inflation in steady state, given the Frisch elasticity, the wage markup, and the steady-state wage, hours, tax rate, and distribution of marginal utility. Since the baseline welfare measure is consumption-only, these weights affect the reported welfare numbers only through the robustness exercise in Section 3.9.

Production parameters. Given the sectoral outputs, capital stocks, hours, and the return on capital, the sectoral capital shares, productivity levels, wages, marginal costs, and the price markup are all implied by the firms' first-order conditions. None of them is free once the targets above are set.

Table A.1: Parameters Determined by Internal Consistency

Parameter	Determined by	Value
α_j, Z_j, w_j, mc_j, μ_p	Firms' first-order conditions given Y_j, K_j, N_j, r	—
τ	Government flow budget, equation (13)	—
α_G	Sectoral goods-market clearing, equation (32)	—
φ_A, φ_B	Union first-order condition at zero wage inflation	—
Θ_j	Asset-market clearing, equation (20)	—

Notes: Parameters that are not measured and not targeted, but implied by the requirement that the steady state be an equilibrium. Values are reported in the calibration table in Section 4.

Table A.2 is the full moment ledger of the steady state: the identities the equilibrium forces, the moments the calibration targets, and the moments it reports as outcomes, including the type-level wealth moments on which the limitation discussed in Section 4 is visible.

Table A.2: Moment ledger: what the equilibrium forces, what is targeted, and what is reported

Moment	Data	Model	Role
Forced by Equilibrium			
Aggregate Household Wealth	$\geq K + B_G$	matched	asset-market clearing
Asset-Market Residual	0	$< 10^{-7}$	equilibrium
I/Y	$\delta K/Y$	0.137	stationarity identity
Targeted			
Quarterly MPC, EIP-Sized	0.15–0.25	0.233	targets χ_1
Interdecile Range of Log Income	11.2	11.18	targets σ_e, n_e
Real Return on Liquid Assets	$\approx -0.5\%$	-0.0052	sets ω externally
Reported as Outcomes			
Annual MPC, EIP-Sized	0.4–0.5	0.416	outcome
Borrower Share, A / B	15.4% / 14.8%	19.1% / 13.4%	**outcome**; κ_b is external
Liquid Share of Wealth, A / B	0.0814, 0.0858	0.0384, 0.0185	**unattainable**: peaks at ≈ 0.05
Poor Hand-to-Mouth Share, A / B	15.4% / 11.2%	0.5% / 0.1%	known miss; no low-wealth mass
Wealthy Hand-to-Mouth Share, A / B	18.8% / 14.7%	50.5% / 61.7%	known miss; same cause
Median Wealth / Annual Income, A / B	0.952 / 1.626	far higher	known miss; no rich tail
Labor Share	0.62–0.66	0.661	implied
Price Markup	1.10–1.25	1.096	implied by $p = \text{wealth} - B_G$
Non-Labor Share (Capital + Profit)	0.35–0.38	0.339	the object the data measures

Notes: Data are from the 2019 PSID, weighted, on the KMV liquid definition, housing-inclusive. Model values are evaluated at the baseline calibration. Rows marked as known misses are the wealth-distribution limitation discussed in Section 4.

C Economic Impact Payments in the Household Pulse Survey

This appendix presents the survey evidence on how households used the Economic Impact Payments. Section 2 draws on this evidence for two objects, the loss criterion and the usage gradient across rounds. The remaining evidence is included since the Economic Impact Payments are the program that the paper takes as its anchor. Section C.1 describes the survey and its questions, Sections C.2 and C.3 document who received the first payment, what they did with it, and how households that had lost income differed from households that had not, and Sections C.4 and C.5 construct the share spent and report additional regressions.

C.1 The Survey and the Questions

The Household Pulse Survey is drawn from the Census Bureau's Master Address File, with sampled addresses contacted by text message and email and, in the first phase of the survey, invited to respond for up to three consecutive weekly waves. Response rates are below 5%, and the Bureau constructs its weights so as to deliver national estimates at those rates (Fields, Hunter-Childs, Tersine, Sisson, Parker, Velkoff, Logan and Shin, Forthcoming). Questions about the payments were added in the seventh week of the survey. For the first-round evidence in this appendix, I use weeks 7 to 12, June 11 to July 21, 2020. The regressions pool these six weeks, about 396,000 recipient households, while the shares quoted for a single week are from week 12, in which 86,792 households responded, representing 249 million adults under the person weights published for the first phase. For the cross-round evidence of Section 2.2, I use all 33 weeks under household weights. Further, more than 99% of respondents answered the employment and payment questions. Since the wording of the questions determines what is measured, the four questions are reproduced below, along with the labels that this paper attaches to them.

> **Question 1**: Have you, or has anyone in your household experienced a loss of employment income since March 13, 2020?
>
> **Question 2**: Do you expect that you or anyone in your household will experience a loss of employment income in the next 4 weeks because of the coronavirus pandemic?

Both questions refer to anyone in the household rather than only the respondent, and both are answered yes or no. I label households that answer no to both questions as *employed*, households that answer no to the first and yes to the second as *anticipating*, and households that answer yes to the first as having *lost income*. Note that a household whose members lost employment income need not contain anyone who is unemployed in the usual sense. In week 12, 51% of households had lost employment income, 45% had neither lost it nor expected to, and 3.5% expected to lose it. The anticipating share was largest in the first week of the survey, at 6%, when the lost-income share was at its smallest, at 47%, and in the weeks that followed households moved from the anticipating

group to the lost-income group. The same movement is visible within the households that answered more than one wave.

> **Question 3**: If you, or anyone in your household, already received, or plan to receive a "stimulus payment," that is the coronavirus related Economic Impact Payment from the Federal Government, did or will you use it:
>
> 1. Mostly to pay for expenses (food, clothing, shelter, etc.)
> 2. Mostly to pay off debt (car loans, student loans, credit cards)
> 3. Mostly to add to savings
> 4. Not applicable, I did not and do not expect to receive the stimulus payment
>
> **Question 4**: What did, or will, you and your household spend the "stimulus payment" on? Categories: (1) Food (groceries, eating out, take out), (2) Clothing (clothing, accessories, shoes), (3) Household supplies and personal care products, (4) Household items (TV, electronics, furniture, appliances), (5) Recreational goods (sports and fitness equipment, bicycles, toys, games), (6) Rent, (7) Mortgage (scheduled or monthly), (8) Utilities and telecommunications (natural gas, electricity, cable, internet, cellphone), (9) Vehicle payments (scheduled or monthly), (10) Paying down credit card, student loans, or other debts, (11) Charitable donations or giving to family members, (12) Savings or investments, (13) Other, specify

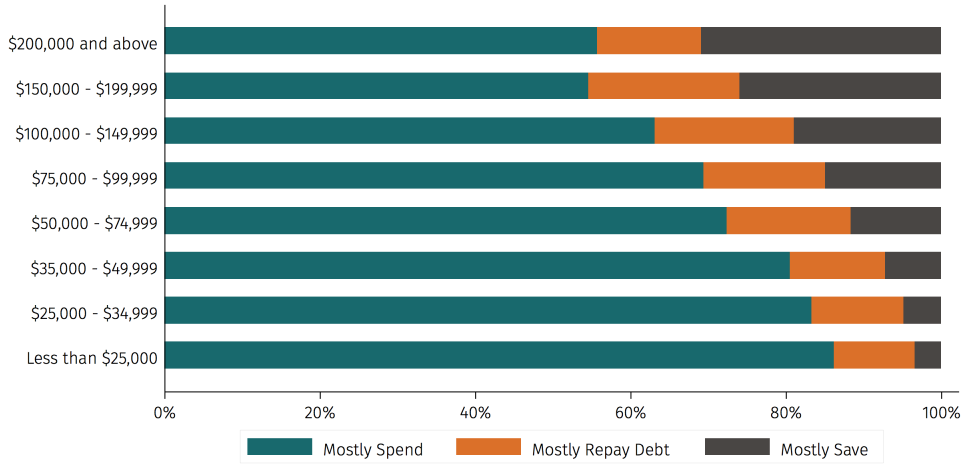
Households selected a single answer to Question 3, whose three options are labeled mostly spend, mostly repay debt and mostly save, and selected every category of Question 4 that applied to them. On average, households selected three categories, and 80% selected six or fewer. Question 3 differs from the questions that were asked about the 2001 and 2008 rebates. The Michigan Survey of Consumers asked whether the rebate would lead the household "mostly to increase spending, mostly to increase saving, or mostly to pay off debt," which is a question about the change in behavior, and the same question was asked again for the 2020 payments (Sahm, Shapiro and Slemrod, 2020). The Nielsen panel asked how the payment was used, with the same three options. Since the Pulse Survey asks only how the payment was used, its shares are not marginal propensities, and I do not compare them with the shares from the earlier surveys.

C.2 Use of the First Payment

Three quarters of recipient households reported using the first payment mostly to pay for expenses, while 14% reported mostly repaying debt and 11% mostly saving. Figure A.1 plots these three shares by income bracket. The spending share is decreasing in income, from 86% of households earning less than $25,000 to about 55% of households earning $200,000 or more, and nearly all of this difference reflects households switching from spending to saving, since 3% of the lowest bracket reported mostly saving compared to almost 30% of the highest. The debt-repayment share is flatter across the distribution,

rising from about 10% to 20%. This income gradient is the same one that the earlier rebate literature documents for 2001 and 2008 (Shapiro and Slemrod, 2003; Johnson et al., 2006; Agarwal, Liu and Souleles, 2007; Parker, 2011), and I show in Section C.5 that it survives the demographic controls of Table A.7.

Figure A.1: Majority use of the first payment, by income bracket



Notes: Household Pulse Survey, week 12. Share of recipient households in each bracket of annual household income that reported using the payment mostly to pay for expenses, mostly to pay off debt, or mostly to add to savings.

Table A.3 crosses the categories of Question 4 with the majority use reported in Question 3, and the two answers are consistent with each other. Households that reported mostly spending selected 3.7 categories on average, led by food (83%), utilities (64%) and household supplies (62%). Among these households, 17% also selected the debt category, and 29% and 28% selected mortgage and vehicle payments respectively, which repay debt while keeping the good that secures it, while 6% selected saving. Households that reported mostly repaying debt selected 2.9 categories on average, with 68% selecting the debt category, 20% mortgage payments and 31% vehicle payments, although a large share also selected food (44%), supplies (32%) and utilities (33%). Households that reported mostly saving selected 1.8 categories on average, and 71% of them selected saving, with food, the next most frequent category, selected by 19%.

C.3 Use by Labor Status

The survey makes it possible to compare directly the use of the payment by households that had lost income and by households that had not, since the payments were made to both groups. This comparison speaks to the case for a narrower program, which rested on the presumption that households that had not lost income would save the payment or repay debt (International Monetary Fund, 2020). In the first round, 63% of employed households reported mostly spending the payment. The corresponding share is 81% for the anticipating group and 84% for households that had lost income. Mostly repaying debt was reported by 20% of employed households and by 11% of each of the other two

Table A.3: Categories selected, by reported majority use of the first payment

		Use EIP mostly to…		
		Pay for expenses	Pay off debt	Add to savings
1	Food (groceries, eating out, take out)	0.832	0.444	0.194
2	Clothing (clothing, accessories, shoes)	0.228	0.126	0.061
3	Household supplies and personal care products	0.624	0.316	0.123
4	Household items	0.062	0.063	0.072
5	Recreational goods	0.024	0.028	0.038
6	Rent	0.358	0.195	0.068
7	Mortgage (scheduled or monthly)	0.294	0.203	0.078
8	Utilities and telecommunications	0.639	0.330	0.097
9	Vehicle payments (scheduled or monthly)	0.276	0.308	0.058
10	Paying down credit card, student loans, or other debts	0.172	0.655	0.097
11	Charitable donations or giving to family members	0.037	0.049	0.133
12	Savings or investments	0.056	0.113	0.711
13	Other, specify	0.050	0.068	0.065

Notes: Each column is one answer to the majority-use question and each row is the share of households in that column that selected the category. Columns sum to more than one because households may select multiple categories.

groups, and mostly saving by 19% of employed households, compared to 9% and 5%. The majority of every group therefore reported mostly spending. In Table A.4, I estimate these gaps with controls using the specification

$$Use_{it} = \gamma_t + \sum_g \alpha_g \, Status_{git} + \beta X_{it} + u_{it}, \tag{33}$$

estimated by ordinary least squares on recipient households in weeks 7 to 12. The dependent variable is an indicator for each of the three uses and the status indicators mark the anticipating and lost-income groups. I include week fixed effects, γ_t, and a vector of household characteristics, X_{it}, consisting of income bracket, sex, Hispanic origin, race, education, household size and the number of children. Households that had lost income were 19 percentage points more likely to report mostly spending, 7 points less likely to report mostly repaying debt and 12 points less likely to report mostly saving. The corresponding gaps for the anticipating group are 14, 4 and 10 points. Since the three uses are exhaustive, each row of the table sums to zero.

The anticipating group accounts for between 2.5% and 4.5% of households in every income bracket, and its behavior is that of the group that had already lost income. Its spending share is 81%, compared to 84% for the lost-income group and 63% for the employed. Further, its category selections in Table A.5 track those of the lost-income group almost line by line. A screen on a realized loss therefore misses households that are about to lose income and that already spend as though they had lost it. This is a cost of the loss criterion which its precision does not record. In the remainder of this appendix, I pool the two groups.

Figure A.2 plots the first-round usage shares by income bracket and by loss status.

Table A.4: Majority use of the first payment, by labor status

	Use EIP mostly to…		
	Pay for expenses	Pay off debt	Add to savings
Base Group: Employed			
Anticipating Unemployment	0.141	-0.037	-0.099
	(0.004)	(0.003)	(0.003)
Unemployed	0.189	-0.066	-0.121
	(0.002)	(0.001)	(0.001)
Observations	395,695	395,695	395,695
R^2	0.11	0.02	0.09

Notes: Each column is a linear probability models estimated on recipient households pooled over weeks 7 to 12 of the Household Pulse Survey. The omitted group is households that neither lost employment income nor expected to. Standard errors clustered by household in parentheses. Each row sums to zero across the three columns because the three uses are exhaustive.

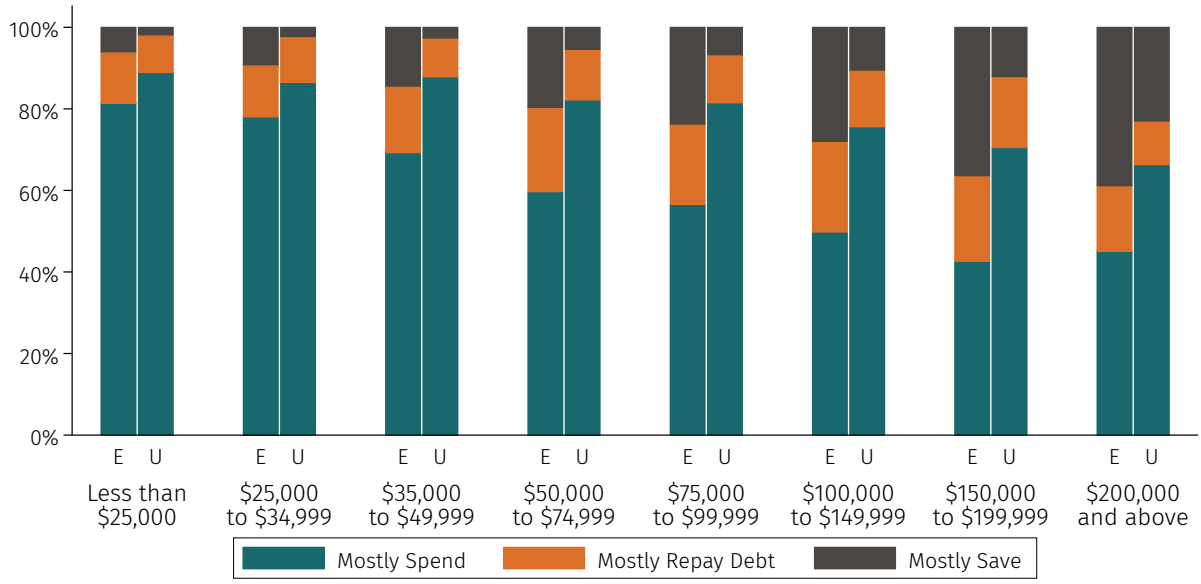
The share of households that lost income falls with income, from 60% below $25,000 to 47% at $75,000 to $99,999 and 33% above $200,000. Within every bracket, households that lost income are more likely to report spending than households that did not, and the gap widens with income. Below $25,000, 89% of households that lost income report mostly spending, and 81% of households that did not. Above $200,000, on the other hand, the shares are 66% and 45%, respectively.[21] As such, income is a poor proxy for loss, and it is the loss gradient within income brackets that the precision measure of Section 2.1 picks up. Table A.8 in Section C.5 estimates the gap bracket by bracket.

Table A.5 crosses the categories of Question 4 with labor status. On consumer goods, categories 1 to 5, the affected groups selected at higher rates. For example, food was selected by 79% of the lost-income group compared to 60% of employed households, while household supplies were selected by 58% compared to 45%, and the anticipating group selected both at nearly the same rates as the lost-income group, 77% and 57%. Clothing, at just under 20%, and the small categories of household items and recreational goods were selected at similar rates by all three groups. On recurring commitments, categories 6 to 9, the differences are stark. In particular, rent was selected by 39% of the lost-income group and 35% of the anticipating group, compared to 19% of employed households, while mortgage payments were selected by 30% and 36%, compared to 19%. Utilities were selected by 41% of employed households compared to 65% and 63%, and vehicle payments by 17% compared to 31% and 32%. This is consistent with the payment being used to keep up commitments that were entered into before the lockdown. On wealth, categories 10 to 12, the three groups selected debt repayment at similar rates, between 22% and 25%. Saving was selected by 20% of employed households compared to 11% and 9%, and giving by 8% compared to 4% and 3%. As such, employed households did not, on the whole, direct the payment to their balance sheets.

Table 2 in the main text extends this comparison to the second and third rounds. I

[21] At $50,000 to $74,999 the shares are 82% and 60%. The difference is made up by saving, which rises from 2% and 6% in the lowest bracket to 23% and 39% in the highest, so that at the bottom of the distribution

Figure A.2: Majority use of the first payment, by income bracket and loss of employment income



Notes: Household Pulse Survey, week 12. "E" bars are households reporting no loss of employment income since March 13, 2020 and no expected loss; "U" bars are households that lost employment income or expected to. Income brackets are annual household income.

Table A.5: Categories selected, by labor status, first payment

		Employed	Anticipating Unemp.	Unemp.
1	Food (groceries, eating out, take out)	0.598	0.766	0.791
2	Clothing (clothing, accessories, shoes)	0.186	0.200	0.203
3	Household supplies and personal care products	0.455	0.572	0.579
4	Household items	0.076	0.055	0.054
5	Recreational goods	0.032	0.011	0.022
6	Rent	0.186	0.345	0.394
7	Mortgage (scheduled or monthly)	0.189	0.357	0.305
8	Utilities and telecommunications	0.407	0.649	0.632
9	Vehicle payments (scheduled or monthly)	0.174	0.313	0.320
10	Paying down credit card, student loans, or other debts	0.240	0.246	0.225
11	Charitable donations or giving to family members	0.075	0.038	0.030
12	Savings or investments	0.198	0.106	0.089
13	Other, specify	0.065	0.047	0.045

Notes: Household Pulse Survey, week 12, person weights. The employed are households with no loss of employment income since March 13, 2020 and none expected in the next four weeks. The anticipating group had no loss yet but expected one, and the lost-income group had a loss since March 13. Each entry is the share of the group that selected the category, and columns sum to more than one.

estimate it by weighted least squares on the pooled sample of recipients in weeks 7 to 12, 22 to 27 and 28 to 33, using household weights. The dependent variable is again an indicator for each of the three uses. The regressors are an indicator for having lost employment income, its interactions with indicators for the second and third rounds, week fixed effects, and controls for income bracket, education, sex, Hispanic origin, race, marital status, household size, the number of children, age and age squared. The gap in round one is given by the coefficient on the loss indicator, the gap in a later round by that coefficient plus the round interaction, and the interactions therefore provide the tests of the reversal. Standard errors are estimated using the Census Bureau's successive-difference replication over its 80 replicate weights. Estimating each round separately yields gaps within a percentage point of those reported, and the round-one gap under this specification, 19.6 points, matches Table A.4.

Since all three rounds screened on the same prior-year record, one might expect the record to have grown staler with time, so that precision on the loss criterion would have fallen across rounds. In fact, I find the opposite. Under household weights, with the population loss rate computed on the same sample as the bracket rates, precision is 0.038 in the first round and 0.048 in the second, and it rises under every consistent choice of weight and sample. An apparent decline between the rounds appears only when person weights are used for the first round and household weights for the second, a combination which the first-phase public files invite, given that the household weight is published in a companion file. The third round is not comparable on this dimension, however, since from week 28 onward the survey shortens the reference period for loss of employment income from the period since March 2020 to the last four weeks, which changes the loss criterion itself.

C.4 The Share Spent, from the Categories Selected

Table A.6 reports the share of the payment spent. I construct it from the category question, since the majority-use question cannot deliver a spending share on its own.[22] Each category is classified as spending, debt repayment or saving, the "other" category, which is under 5% of selections, is dropped, and each household's payment is allocated across the categories that it selected. Thirteen percent of households selected only saving or debt categories, 21% of the employed and 8% of the affected, and their share is zero. Thirty-eight percent selected no wealth category at all, 39% of the employed and 36% of the affected, and their share is one. The rest selected a mixture, and the table reports the aggregate under the assumptions that resolve the mixtures.

C.5 Additional Regressions

Table A.7 reports the linear probability models behind Figure A.1. The models include week fixed effects and the full set of household characteristics, where the base group of

both groups spend, whereas at the top the employed save and the affected spend.

[22] As Shapiro and Slemrod (2003) note, if mostly spending meant a share of 0.51 and the other two answers a share of 0.49, the aggregate would be near one half whatever the distribution of answers.

Table A.6: The share of the first payment spent, implied by the categories selected

	secured payments as debt			secured payments as spending		
	all	no loss	lost	all	no loss	lost
equal split across categories	0.69	0.65	0.73	0.84	0.75	0.90
share to wealth $p = 0.25$	0.74	0.69	0.78	0.84	0.76	0.89
share to wealth $p = 0.50$	0.62	0.59	0.64	0.78	0.70	0.83
share to wealth $p = 0.90$	0.43	0.43	0.42	0.68	0.61	0.73

Notes: Household Pulse Survey, week 12. Each household's payment is allocated across the categories it selected in the category question, and the "other" category is dropped. Categories are classified as spending, debt repayment or saving. Mortgage and vehicle payments are counted as debt repayment in the left panel and as spending in the right panel. The first row splits the payment equally across the categories selected. The remaining rows pool debt repayment and saving into a single wealth category and assign it the share p of the payment whenever the household selected it alongside at least one spending category. In this table, the lost-income group also includes households that expected to lose employment income.

each characteristic is the omitted category, and each row sums to zero across the three uses. The income gradient survives these controls. Relative to households earning less than $25,000, households earning $200,000 or more are 13 points less likely to report mostly spending and 11 points more likely to report mostly saving. Further, households with more children spend more. Sex, Hispanic origin and race make little difference, while education lowers the spending share by five or six points at the associate's degree and above. In Table A.8, I re-estimate the specification of Table A.4 within each income bracket, pooling the two affected groups. The loss gap in the spending share rises from 9 points below $25,000 to 22 points at $200,000 and above, which is the widening visible in Figure A.2.

D Additional Results

D.1 The Targeting–Timing Surface

Figure A.3 plots the population's loss over the whole plane of precision and delay, with the frontier of Section 2 on top of it; Figure 5 in the text is the loss along that line. Table A.9 prices the frontier on the alternative statistics and endpoints discussed in Section 5.3.2.

D.2 The Crisis-Specific Surface

Figure A.4 plots the difference between the lockdown surface of Figure A.3 and the surface of the identical transfer in the economy with no shock. At any fixed delay, the difference across precision is the crisis-specific return to precision of Table 8, and it is nearly linear in θ, which is why the corner of Section 5.3.1 is robust.

Table A.7: Majority use of the first payment and household characteristics

	Use EIP mostly to…		
	Pay for expenses	Pay off debt	Add to savings
Income (base: less than $25,000)			
$25,000 - $34,999	-0.020	0.014	0.006
	(0.008)	(0.007)	(0.005)
$35,000 - $49,999	-0.043	0.024	0.019
	(0.010)	(0.009)	(0.006)
$50,000 - $74,999	-0.069	0.034	0.036
	(0.011)	(0.010)	(0.008)
$75,000 - $99,999	-0.088	0.043	0.045
	(0.013)	(0.011)	(0.009)
$100,000 - $149,999	-0.112	0.047	0.064
	(0.015)	(0.012)	(0.011)
$150,000 - $199,999	-0.134	0.050	0.083
	(0.017)	(0.014)	(0.013)
$200,000 and above	-0.134	0.023	0.110
	(0.024)	(0.020)	(0.019)
Sex (base: male)			
Female	-0.012	0.006	0.006
	(0.011)	(0.009)	(0.009)
Hispanic origin (base: not Hispanic)			
Hispanic	-0.009	0.024	-0.015
	(0.014)	(0.012)	(0.009)
Race (base: white alone)			
Black, Alone	0.028	0.008	-0.036
	(0.022)	(0.018)	(0.015)
Asian, Alone	0.004	-0.008	0.004
	(0.024)	(0.023)	(0.017)
Any other race alone, or race in combination	0.010	0.003	-0.013
	(0.012)	(0.010)	(0.008)
Education (base: no high school)			
Some high school	-0.008	0.022	-0.014
	(0.026)	(0.024)	(0.016)
High school graduate or equivalent (for example GED)	0.002	0.020	-0.022
	(0.027)	(0.025)	(0.017)
Some college, but degree not received or is in progress	-0.041	0.048	-0.007
	(0.028)	(0.026)	(0.018)
Associate's degree (for example AA, AS)	-0.063	0.064	-0.001
	(0.030)	(0.027)	(0.020)
Bachelor's degree (for example BA, BS, AB)	-0.052	0.044	0.008
	(0.030)	(0.028)	(0.020)
Graduate degree (for example master's, professional, doctorate)	-0.066	0.047	0.019
	(0.031)	(0.028)	(0.021)
Household size	-0.001	0.001	-0.000
	(0.001)	(0.001)	(0.001)
Children in household	0.019	-0.002	-0.017
	(0.005)	(0.004)	(0.003)
Constant	0.825	0.065	0.110
	(0.030)	(0.027)	(0.020)
Observations	396,418	396,418	396,418
R^2	0.05	0.01	0.05

Notes: Each column is a linear probability model estimated with recipient households pooled over weeks 7 to 12 of the Household Pulse Survey. The base group of each characteristic is the omitted category. Standard errors clustered by household in parentheses.

Table A.8: The loss gap in majority use, by income bracket, first payment

	Use EIP mostly to…		
	Pay for expenses	Pay off debt	Add to savings
Less than $25,000	0.086	-0.035	-0.050
	(0.004)	(0.003)	(0.002)
$25,000 - $34,999	0.138	-0.052	-0.085
	(0.004)	(0.003)	(0.003)
$35,000 - $49,999	0.182	-0.073	-0.107
	(0.004)	(0.003)	(0.003)
$50,000 - $74,999	0.206	-0.074	-0.131
	(0.003)	(0.003)	(0.003)
$75,000 - $99,999	0.220	-0.076	-0.141
	(0.004)	(0.003)	(0.003)
$100,000 - $149,999	0.212	-0.066	-0.144
	(0.004)	(0.003)	(0.003)
$150,000 - $199,999	0.215	-0.049	-0.163
	(0.006)	(0.005)	(0.005)
$200,000 and above	0.222	-0.047	-0.176
	(0.010)	(0.007)	(0.009)

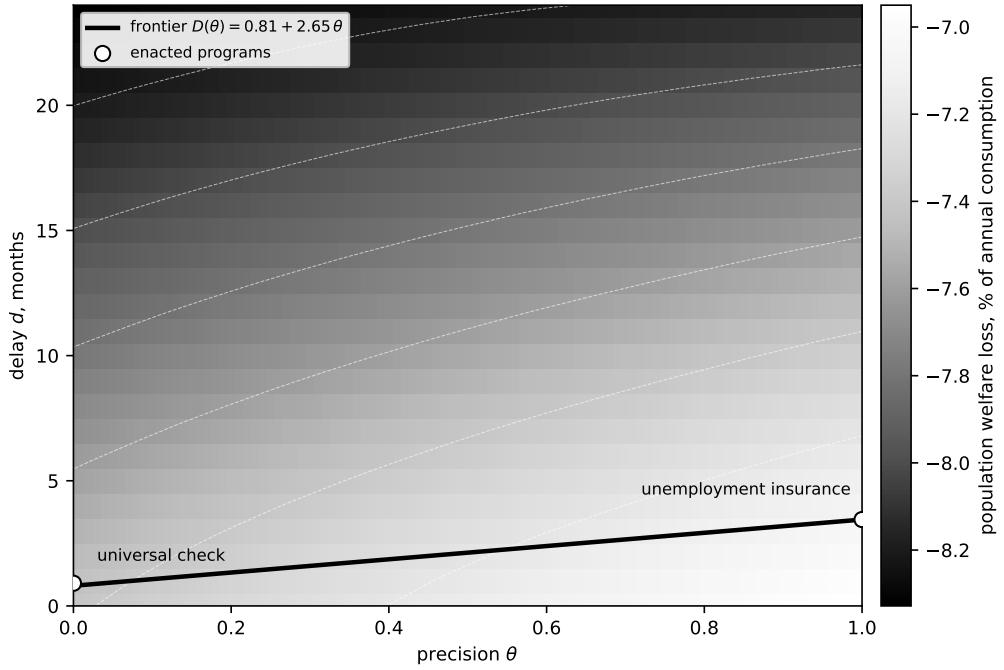
Notes: Each entry is the difference (within the income bracket) between households that lost or expected to lose employment income and households that did neither. Standard errors clustered by household in parentheses.

Table A.9: The price of the screen: frontier statistics and endpoint alternatives

	$D(1)$	Net	Forfeited	Share	θ^*
Return to Full Targeting, Both Paid at Once: +0.47 pp					
Frontier Statistic					
Mean (Priced)	**3.45**	**+0.38**	**0.09**	**19.3%**	**1.00**
90th Percentile	6.10	+0.28	0.19	40.9%	1.00
Endpoint Tier					
FPUC ($600 Supplement)	2.69	+0.41	0.06	13.3%	1.00
PUA	4.67	+0.33	0.14	29.5%	1.00
PEUC (Extended Benefits)	6.58	+0.25	0.22	46.9%	1.00

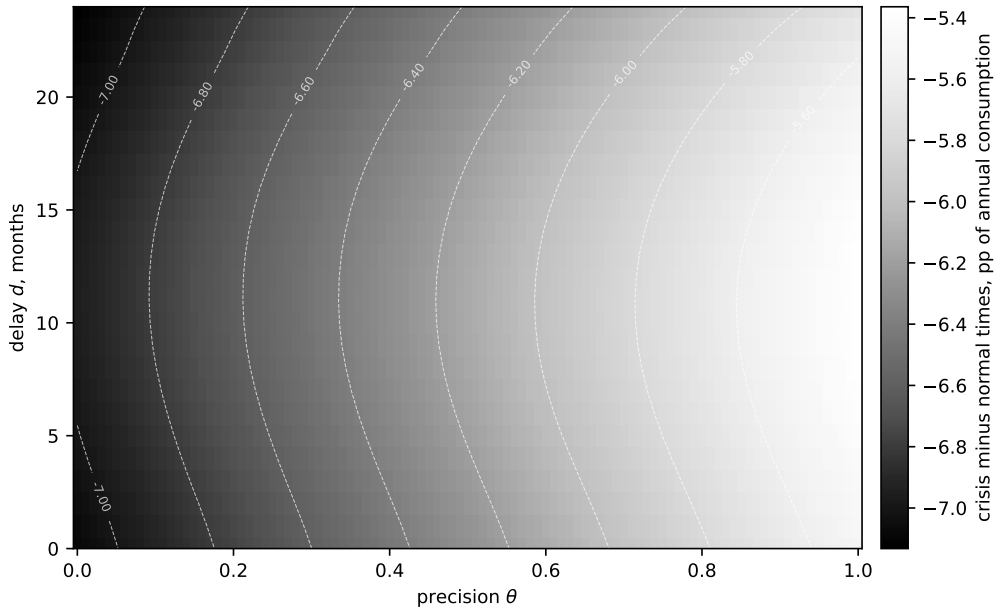
Notes: Population welfare is in percentage points of annual consumption, on the lockdown surface with W interpolated linearly in d at each θ. The column $D(1)$ is the delay of the fully targeted design, in months. The net column compares the fully targeted design paid when its screen allows with the universal check paid when its rail allows. The forfeited column is the return with both designs paid at once less the net return, and the share column divides the two (the return with both paid at once is $W(1,0) - W(0,0) = +0.47$). The last column is the precision that minimizes the loss along the frontier.

Figure A.3: The targeting–timing surface and the frontier



Notes: Population welfare loss from the lockdown, as a percentage of annual consumption, for every program (θ, d) on the grid; lighter is a smaller loss. The line is the frontier of Section 2, $d = D(\theta)$, and the dots are the two enacted designs at its endpoints.

Figure A.4: The crisis-minus-normal-times surface

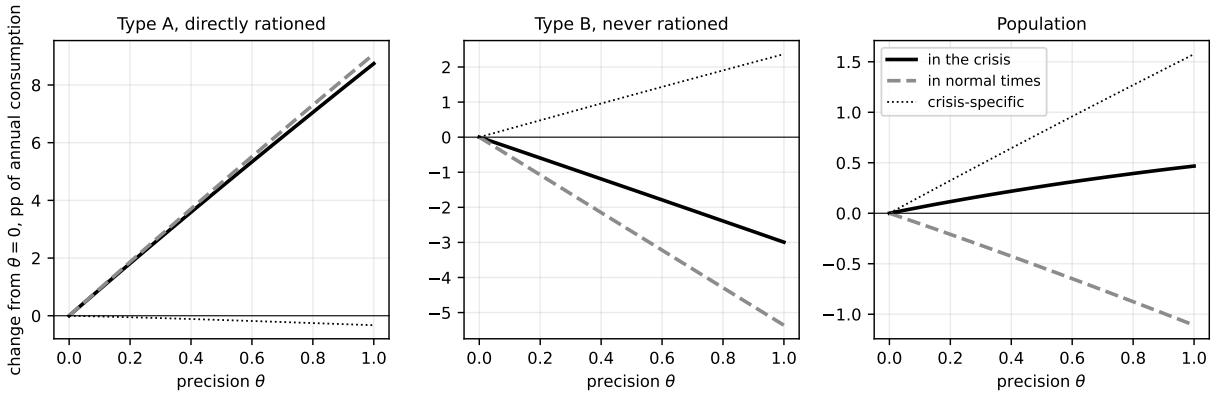


Notes: Population welfare in the lockdown minus the welfare of the identical transfer in the economy with no shock, in percentage points of annual consumption, for every (θ, d). At a fixed d, the difference across θ is the crisis-specific return to precision reported in Table 8. Since the no-shock welfare falls with delay faster than the crisis welfare over most of the grid, the difference rises with delay, and no timing conclusion is drawn from it.

D.3 The Return to Precision by Type

Figure A.5 plots the objects reported in Table 8 across the whole range of θ. For each group, the figure plots the change in welfare relative to the universal check as the transfer is concentrated on the shut sector, with the transfer paid at the onset, both in the crisis and in normal times.

Figure A.5: The return to precision with free screening, by type



Notes: Change in the consumption-equivalent welfare of each group relative to the universal check, in percentage points of annual consumption, as the transfer is concentrated on the affected sector, with the transfer paid at $d = 0$. The solid line is the lockdown; the dashed line is the identical transfer in the economy with no shock; the dotted line is their difference.

D.4 The Cost of Delay by Type, and the Planner's Weight

Table A.10 reports the cost of a month of delay for the fully targeted transfer by type, computed both over the range that the frontier spans and over the whole grid, in the crisis and in normal times. Figure A.6 then plots the precision that minimizes the loss along the frontier as the planner's weight on the shut sector varies.

D.5 Where a Transfer Dollar Goes

Table A.11 reports the cumulative twelve-month responses per unit of transfer for the two polar designs, each solved with its full path, in the crisis and in normal times. These responses are the evidence on which the one-way linkage of Section 5.3.1 rests.

D.6 Welfare Overlays

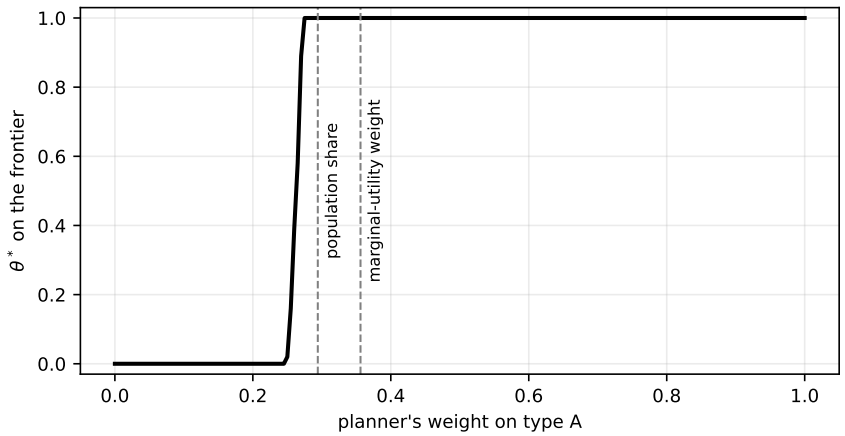
Table A.12 recomputes the population loss for the crisis with no transfer and for the two polar designs paid at the onset under the two welfare-only overlays of Section 3.9. Crediting type A with the full leisure value of its rationed hours reduces every loss by about 2.5 pp, while deducting the variety wedge of exclusion rationing raises every loss by between 0.45 and 1.4 pp, depending on the elasticity of substitution across varieties. The return to

Table A.10: The cost of a month of delay, fully targeted transfer

	$d \leq 4$		$d \leq 24$	
	Crisis	Normal Times	Crisis	Normal Times
Type A (Recipient)	0.080	0.043	0.158	0.052
Type B (Never Paid)	0.013	0.056	0.005	0.038
Population	0.033	0.052	0.051	0.044

Notes: Consumption-equivalent welfare lost per month that the fully targeted transfer ($\theta = 1$) is delayed, in percentage points of annual consumption, measured by the slope of $W(1, d)$ in d. The first two columns use $d \leq 4$ months, the range that the frontier spans, and the last two use the whole 24-month grid. The population entries are regression slopes, and the type entries over the whole grid are the average slope, $[W(1, 0) - W(1, 24)]/24$. *In the crisis* is the lockdown surface and *in normal times* the identical transfer with no shock. Fitting $W(1, d) = a - bd - cd^2$ over the grid in the crisis gives $b = 0.0349$ and $c = 0.00066$ per month. The corresponding Jensen error from collapsing a program with a two- or four-month dispersion of payments to its mean date is 0.003 or 0.011 pp.

Figure A.6: The optimum on the frontier as the planner's weight on the affected sector varies



Notes: Precision that minimizes the weighted loss along the frontier, $\omega\,\lambda_A + (1 - \omega)\,\lambda_B$ with W interpolated in d, as the weight ω on type A varies. The dashed lines mark the population share and the marginal-utility weight.

Table A.11: Where a transfer dollar goes: cumulative responses in the crisis and in normal times

	Fully Targeted, $\theta = 1$		Universal Check, $\theta = 0$	
	Crisis	Normal Times	Crisis	Normal Times
Consumption, A	1.021	1.065	0.573	0.688
Rationed-Good Purchases, A	1.008	1.223	0.456	0.737
Rationed-Good Purchases, B	-0.160	-0.224	0.232	0.267
Output, Sector 2	0.047	0.051	-0.064	-0.011
Wage, Sector 2	0.044	-0.238	0.163	0.048
Consumption, B	-0.142	-0.429	0.299	0.239

Notes: Each entry is a cumulative deviation over the first twelve months after a transfer paid at the onset, relative to the same economy with no transfer. Entries are scaled by the transfer received by the type in question, so that type-A rows are per unit of S_A, while type-B and sector-2 rows are per unit of S_B under the universal check and per unit of S_A under the fully targeted design. Only type A is paid under the fully targeted design. *Crisis* is the lockdown, and *normal times* is the identical transfer in the economy with no shock.

full precision is 0.46 pp under the consumption-only measure, compared to between 0.46 and 0.47 pp under the overlays.

Table A.12: Welfare overlays: population loss under alternative measures

	Consumption Only	Leisure, $\xi = 1$	Variety, $\varepsilon = 4$	Variety, $\varepsilon = 7$	Variety, $\varepsilon = 10$
No Transfer	-7.71	-5.24	-9.09	-8.39	-8.16
Universal Check ($\theta = 0$)	-7.42	-4.95	-8.79	-8.10	-7.87
Fully Targeted ($\theta = 1$)	-6.96	-4.47	-8.34	-7.64	-7.41
Return to Precision, pp	+0.46	+0.47	+0.46	+0.46	+0.46

Notes: Population consumption-equivalent welfare change from the lockdown, as a percentage of annual consumption, for the crisis with no transfer and for the two polar transfer designs, both paid at the onset ($d = 0$), under each welfare measure. *Consumption only* is the measure used throughout the paper. *With leisure* credits type A with the leisure value of its rationed hours at the labor disutility in the household problem, $\xi = 1$, which is the upper bound of Section 3.9. *With variety* deducts the loss of variety under exclusion rationing (Appendix A.5) at the stated elasticity of substitution across varieties. Each overlay is added to lifetime utility before the consumption-equivalent transform. Since the overlays are type-level objects, this table aggregates over the two types, whereas the text aggregates over the eight cohorts, and the two aggregations give the consumption-only figures of Tables 7 and 8 to within 0.001 pp. The last row is the difference between the two designs under each measure.