

Age, Industry, and Unemployment Risk During a Pandemic Lockdown

James Graham*
University of Sydney

Murat Ozbilgin
Reserve Bank of New Zealand

May 26, 2021

Abstract

This paper models the macroeconomic and distributional consequences of lockdown shocks during the COVID-19 pandemic. The model features heterogeneous life-cycle households, labor market search frictions, and multiple industries of employment. We calibrate the model to data from New Zealand, where the health effects of the pandemic were especially mild. In this context, we model lockdowns as supply shocks, ignoring the demand shocks associated with health concerns about the virus. We then study the impact of a large-scale wage subsidy scheme implemented during the lockdown. The policy prevents job losses equivalent to 6.8% of steady state employment. Moreover, we find significant heterogeneity in its impact. The subsidy saves 17% of jobs for workers under the age of 30, but just 3% of jobs for those over 50. Nevertheless, our welfare analysis of fiscal alternatives shows that the young prefer increases in unemployment transfers as this enables greater consumption smoothing across employment states.

Keywords: COVID-19, Pandemic, Lockdowns, Unemployment Risk, Wage Subsidy, Life-Cycle, Consumption

JEL: E21, E24, E32, E62, J64

*Email: james.a.graham@sydney.edu.au and Murat.Ozbilgin@rbnz.govt.nz. For their many helpful comments and suggestions, we would like to thank Nicholas Sander, Sebastian Graves, William Gamber, Victoria Gregory, Robert Kirkby, Punnoose (Reuben) Jacob, Adam Richardson, Karsten Chipeniuk, Gael Price, Andrew Binning, and seminar participants – in person and virtual – at the University of Canterbury, Jinan University, McMaster University, Monash University, the Reserve Bank of New Zealand, the Reserve Bank of Australia, the Reserve Bank of South Africa, the New Zealand Treasury, The New Zealand Ministry of Social Development, the Virtual East Asian Macroeconomics Seminar, and the Virtual Australian Macroeconomics Seminar. All conclusions and opinions are those of the authors and do not necessarily reflect the views of the Reserve Bank of New Zealand.

1. Introduction

The COVID-19 pandemic of 2020 was the cause of enormous macroeconomic disruption around the world. The primary economic effects of the pandemic were the result of restrictions on producer and consumer activity (i.e. lockdowns), as well as decreased consumer demand due to health concerns associated with the virus. As stimulus measures, governments quickly implemented large fiscal interventions in the form of employment subsidies, loans to firms, and direct transfers to households. While the fight against the virus continues, many important economic questions remain: what are the macroeconomic costs of the lockdowns used to contain COVID? To what extent did fiscal interventions help to offset these costs? And how important is household heterogeneity in helping to understand the effects of these lockdown and fiscal policies?

This paper studies the macroeconomic and distributional consequences of lockdowns imposed during the pandemic. In general, pandemic recessions are the result of both supply and demand shocks, which complicates the study of the economic impact of lockdowns on their own. For this reason, we build a model of a lockdown calibrated to data from New Zealand, whose experience with the pandemic differed markedly from many other countries. Because New Zealand is a remote island nation, and because its government quickly closed international borders and imposed a strict national lockdown, the virus was effectively eliminated in the community by June 2020.¹ As a result, the health effects of the pandemic and their consequences for consumer demand have been limited.² This is in stark contrast to the effects of the pandemic in countries like the USA.³ Thus, the experience of New Zealand presents a useful case study for investigating the macroeconomic effects of lockdowns in isolation.

We study lockdowns in a heterogeneous agent overlapping generations model with multiple industries and labor market search frictions. Households in the model differ by age, wealth, employment status, and industry of employment. They choose how much to consume and save over their life-cycle, subject to fluctuations in employment determined by equilibrium in the labor market. A fiscal authority raises taxes, issues debt, funds transfers to unemployed workers, and provides employment subsidies during the lockdown. We solve for general equilibrium of the model and calibrate its steady state to match the distributions of wages, employment, and employment risk across age and industry using data on the New Zealand labor market prior to the pandemic.

We solve for equilibrium dynamics of the model in response to a pandemic-induced lockdown shock and the associated fiscal policy responses. We characterize a lockdown as a sequence of negative shocks to industry-level productivity (see also Guerrieri et al., 2020; Bilbiie et al., 2020;

¹For a timeline of key events in New Zealand during the pandemic, see <https://covid19.govt.nz/alert-system/history-of-the-covid-19-alert-system/>.

²For example, retail spending via debit, credit, and charge cards had largely recovered to pre-COVID levels by June 2020. See <https://www.stats.govt.nz/information-releases/electronic-card-transactions-november-2020>.

³Using a variety of micro-data sources gathered during the early phase of the pandemic, Chetty et al. (2020) showed that greater virus spread was associated with larger reductions in spending, particularly at firms in industries that required a lot of in-person interaction.

Bayer et al., 2020).⁴ In contrast to many other papers in the literature, but consistent with the limited spread of COVID-19 in New Zealand, we do not model the epidemiological aspects of the virus and we ignore the effects of demand shocks due to health concerns (see Auerbach et al., 2020; Eichenbaum et al., 2020; Krueger et al., 2020; Faria-e-Castro, 2020; Kaplan et al., 2020; Baqaee et al., 2020; Farhi et al., 2020). Instead, we assume that productivity shocks capture the entirety of the impact of a lockdown since firms cannot utilize labor resources at their previous rates. We further assume that firms in the services sector are disproportionately affected by a lockdown. This is because service workers typically need to interact with customers to carry out their jobs and as a result are much less able to work from home.⁵ In addition, service sector firms such as those in the tourism, accommodation, and travel industries have been especially affected by ongoing restrictions on international travel.

One of our main contributions to the literature is to emphasize the importance of household age in assessing the distributional effects of lockdowns and their associated fiscal policies.⁶ While the health effects of COVID-19 have been disproportionately felt by the old, young households are more likely to be affected by fluctuations in the labor market. One reason for this is that the young are much more likely to work in the service sector, which in turn is more exposed to the effects of a lockdown. In our model, we capture the effects of these age-dependent exposures through two novel features of our labor market structure. First, we incorporate life-cycle labor market search dynamics, following Lugauer (2013), De la Croix et al. (2013), and Chéron et al. (2013). Since households must search before finding a job, the young take time to settle into employment and are therefore more likely to be unemployed than the old. Second, we incorporate endogenous job separations, following Den Haan et al. (2000) and Fujita et al. (2012). In combination these model features allow us to match the life-cycle profiles of wages and job separations observed in the data. During a lockdown, larger shocks to the service sector result in young workers being laid off in much larger numbers, which disproportionately raises their unemployment risk.

Another contribution of the paper is our study of the macroeconomic and distributional effects of large-scale wage subsidy policies enacted during the pandemic. We consider the wage subsidy scheme introduced by the New Zealand government. This policy represented an exceptionally large fiscal intervention: between March and June 2020 the wage subsidy scheme paid firms approximately 50% of the median wage for each worker employed and the scheme supported approximately 75% of the New Zealand labor force.⁷ In order to be eligible for the subsidy, firms needed to observe a 30% decline in revenues over the previous month. To capture this feature of the policy, we introduce a revenue-dependent subsidy policy in the model and calibrate it to match the 75% of firm-worker matches that received the subsidy during the first

⁴Bloom et al. (2020) provide some direct empirical support for the supply shock hypothesis, showing that total factor productivity in the UK is likely to have been significantly lower during 2020 due to firms having to respond to various COVID containment measures.

⁵See Dingel et al. (2020) and Bartik et al. (2020).

⁶Carroll et al. (2020) also study the effect of lockdown shocks on household age through the lens of a partial equilibrium life-cycle model.

⁷See Section 2 for details.

quarter of the lockdown. Because revenues fall by more for firms in the service sector and with younger workers, these firms are significantly more likely to receive the wage subsidy. This suggests that the wage subsidy scheme was reasonably well-targeted, in that it most benefited those workers disproportionately affected by the lockdown.

In order to discipline our lockdown exercise in the model, we calibrate the sequence of lockdown shocks to match the declines in employment across the services and non-services sectors in the first two quarters of the pandemic. We then calibrate the wage subsidy policy parameters to match the size of the wage subsidy received by firms and the fraction of employees supported by the subsidy. As in the data, a lockdown in the model generates a 5.2% decline in service sector employment and a 1% decline in non-services employment in the first quarter of the pandemic. Comparing the baseline economy to a counterfactual economy absent the wage subsidy scheme, we find that the subsidy saves a large number of jobs. In aggregate, the policy preserves 6.8% of steady state employment relationships, which is equivalent to 183,600 jobs in the New Zealand labor market. In the cross-section, the subsidy saves 8.7% of service sector jobs and 5.5% of non-service sector jobs. But the largest differences are by age. The wage subsidy saves 17.1% of jobs for workers under the age of 30, but just 3.1% of jobs for workers over the age of 50.

We also use the calibrated model to study the effects of two alternative fiscal policies. We first consider a policy that raises unemployment benefits, and second we consider a policy that pays a lump-sum transfers to all households. These alternatives mimic policies that were adopted in the United States and other parts of the world during the pandemic.⁸ In order to compare policies on a dollar-for-dollar basis, we assume that each policy implies the same total fiscal transfer expenditures (i.e. the sum of unemployment benefits, wage subsidies, and lump-sum transfers). We find that although the alternative policies do not prevent unemployment during the lockdown, raising unemployment benefits enables more consumption smoothing among young households than does the wage subsidy policy. Since the young earn low wages in normal times, a large increase in unemployment benefits raises average youth income despite a large increase in the unemployment rate. This represents a large increase in unemployment insurance for those most likely to use it.⁹

Finally, we conduct a welfare analysis to study the relative merits of each fiscal policy response to a lockdown. Average welfare gains are higher for the wage subsidy policy than either of the alternatives. However, this masks significant heterogeneity in the welfare benefits of these policies. We find that young households are much more likely to favor the policy that raises unemployment benefits, and the welfare gains for these households are relatively large. Because young workers are much more likely to become unemployed than other workers during the pandemic, higher unemployment benefit payments help young households smooth consumption better than does a policy that simply preserves employment for many, but not all, young workers.

⁸For papers in the macroeconomics literature studying the effects of these policies in the US, see Kaplan et al. (2020), Fang et al. (2020), Carroll et al. (2020), Bayer et al. (2020), and Faria-e-Castro (2020).

⁹In a closely related HANK model, Graves (2020) shows that the presence of unemployment insurance helps to smooth the consumption response of low-wealth households to aggregate shocks that generate higher unemployment risk.

In contrast, older households prefer the wage subsidy policy to other fiscal policies, although the welfare gains are relatively small. Although unemployment risk for older workers remains relatively low during the lockdown, the higher wages that are earned later in life imply larger costs of job loss even if unemployment benefits are raised.

1.1. Literature Review

Many early papers in the macroeconomics literature on COVID-19 incorporated epidemiological model features in order to study the evolution of health and economic outcomes during the pandemic (see for example: Berger et al., 2020; Eichenbaum et al., 2020; Krueger et al., 2020; Acemoglu et al., 2020; Kaplan et al., 2020). These papers show both that households optimally reduce consumption in response to the health risks of COVID-19, and that country-wide lockdowns imply a strong tradeoff between health outcomes and economic activity. In the current paper, we focus on the effect of lockdowns alone, thereby concentrating our analysis on an extreme point along the tradeoff schedule discussed in the literature. We focus on the example of New Zealand, where early and strict lockdowns led to virtual elimination of the virus within the community. Here, the primary concern of economists and policymakers is not a tradeoff between health and the economy, but the macroeconomic and distributional consequences of the lockdowns themselves.

In order to capture heterogeneity in the effects of a pandemic and associated lockdowns, many papers build models with multiple sectors and differential exposure to shocks. Baqaee et al. (2020) and Farhi et al. (2020) study production-based economies in which sector-specific shocks are amplified through input-output linkages built into supply chains in the model. Farhi et al. (2020) show that pandemic shocks lead to higher unemployment that is concentrated in the most heavily affected sectors. Gregory et al. (2020) study the effect of lockdowns on the labor market in a directed search model with multiple industries. They argue that service sector workers experience greater job risk, which leads to a much slower recovery in the service sector as those workers take longer to find stable employment. Kaplan et al. (2020) model workers with different industries and occupations as having differential exposure to pandemic shocks via their ability to work from home. In our paper, we capture these differential household exposures to the pandemic through industry of employment and worker age. We show that younger workers are much more likely to be employed at firms in the services sector, which in turn were much more exposed to the effects of lockdowns than were other firms.

Several papers in the literature build structural macroeconomic models to assess the effects of various fiscal policy responses to the pandemic. Carroll et al. (2020) build a partial-equilibrium heterogeneous agent life-cycle model to study the effects on consumption of higher unemployment insurance payments and direct stimulus checks under the US CARES Act of 2020. They show that direct transfers help stabilize consumption expenditure in the short term, but that increases in unemployment insurance are more effective if the employment effects of the pandemic are likely to persist. Bayer et al. (2020) builds a general equilibrium HANK model to study the same policies, showing that higher unemployment benefits generate larger fiscal multipliers than transfers because they offset the effects of higher unemployment risk. Faria-e-Castro

(2020) studies a DSGE model with borrowers and savers, and shows that borrowers value fiscal interventions that most resemble direct cash payments: either lump-sum transfers or higher unemployment insurance payments. In the current paper, we show that New Zealand’s wage subsidy scheme was very effective in preserving employment relationships during the pandemic which in turn helps to stabilize aggregate consumption. However, as in the above-cited literature, we also find that young households, who tend to be less wealthy, benefit more from a policy that raises unemployment benefits. This is because higher unemployment benefits raise the insurance value of unemployment, whereas wage subsidies only preserve incomes conditional on remaining employed.

Finally, we follow a recent literature that studies the effect of rising unemployment risk in models with household heterogeneity and frictional labor markets. While our model is closely related to the model in Graves (2020), we eschew the two-asset and New Keynesian features and incorporate household age, multiple industries of employment, and endogenous job separations. As in Gornemann et al. (2016), Ravn et al. (2017), and Graves (2020), we find that low wealth households are most affected by increases in employment risk as they are less able to insure against unemployment. For this reason, young and poor households gain most from larger unemployment payments during the pandemic as they benefit from the greater insurance value of unemployment (see also Graves, 2020).

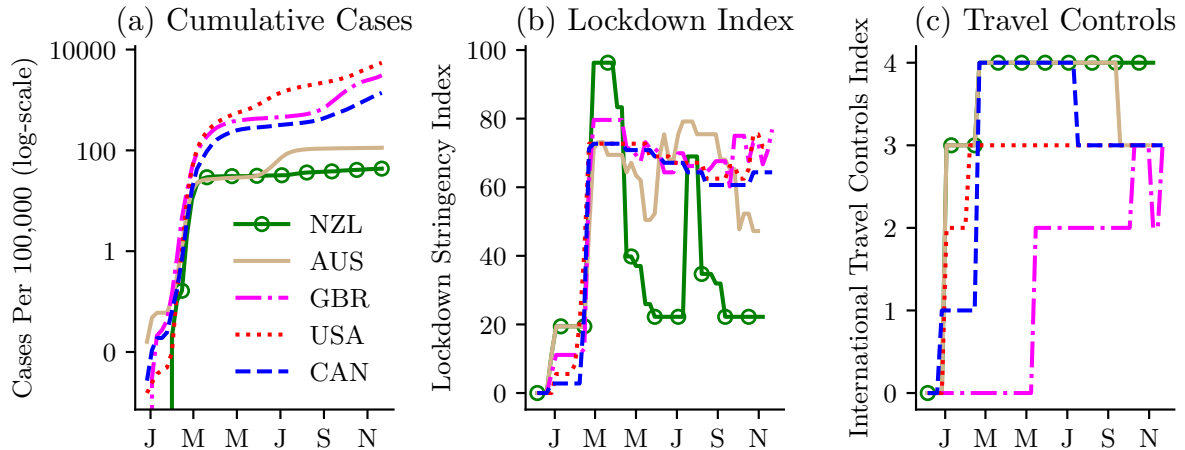
2. Motivating Evidence From New Zealand

The global outbreak of COVID-19 in 2020 was associated with significant macroeconomic disruption. However, the evolution of the pandemic in New Zealand presents an interesting case-study due to its relatively swift and stringent lockdowns, and the limited spread of the virus within its borders. The strict lockdowns imposed in New Zealand imply potentially large declines in economic output due to restrictions on productive activity, while the small COVID case-load suggests a very limited role for declining domestic demand due to health concerns. This is in stark contrast to the experience of countries such as the US, where reductions in activity due to fear of the virus seems to have dominated the effects of the lockdowns themselves (Chetty et al., 2020).

Figure 1 tracks the evolution of the virus and the imposition of lockdowns across several countries, using data collated by authors at Oxford University (Hale et al., 2020).¹⁰ Figure 1(a) shows that New Zealand experienced an outbreak of COVID cases along with other countries at the beginning of 2020. However, case numbers in New Zealand stabilized quickly. As at December 2020, New Zealand had just 43 cumulative cases per 100,000 people. This compared favourably to case numbers per 100,000 in countries such as Australia (112), Canada (1133), Great Britain (2600), and the US (4553). Rapid and extensive lockdowns in March 2020 are one reason cited for New Zealand’s success in limiting the spread of the virus. Figure 1(b) compares the timing and stringency of restrictions on social and economic activities in response to the

¹⁰COVID-related data from <https://coronavirus.jhu.edu/>, with population data from <https://data.oecd.org/pop/population.htm>.

Figure 1: COVID-related Lockdowns by Country



Notes: The overall lockdown stringency index includes information about school, workplace, and public transport closures, restrictions on social gatherings, stay-at-home requirements, and restrictions on internal and international travel.

Source: Corona Virus Government Response Tracker from Oxford University (Hale et al., 2020). Population data from the OECD

pandemic. The data shows that New Zealand imposed some of the strictest lockdown measures of any country in early March.¹¹ By May, as the number of new and active cases fell, restrictions on activity in New Zealand were gradually lifted. This contrasted with many other countries, where various restrictions remained in place throughout the year.¹²

Another important factor in accounting for New Zealand’s success in dealing with the pandemic is that it is an isolated island nation. This enabled New Zealand to exercise strict control over international border crossings. This is reflected in Figure 1(c) which shows that New Zealand quickly implemented international travel restrictions, which it has maintained throughout 2020. The large reduction in international travel, in conjunction with mandatory quarantine for the few travelers entering the country, has significantly reduced the number of imported COVID cases that can then be spread throughout the population.


The limited health effects of COVID-19 in New Zealand suggest that there was a relatively small role for health-related declines in demand for goods and services. However, the strong lockdowns imposed in New Zealand suggest a potentially large impact on economic activity. Additionally, while domestic lockdowns in New Zealand were short-lived, international travel restrictions are ongoing. This suggests differential economic impacts across sectors more or less exposed to international tourism and travel.

To illustrate the economic effects of these lockdowns in New Zealand, Figure 2 shows annual employment growth across industries during the first three quarters of 2020 relative to historical

¹¹For details of the various rules governing lockdowns in New Zealand, see <https://covid19.govt.nz/alert-system/about-the-alert-system/>.

¹²In August 2020 New Zealand’s largest city, Auckland, experienced a brief resurgence in COVID cases. This prompted the government to impose a short-lived lockdown in the city, as is reflected in the jump in lockdown stringency for New Zealand shown in Figure 1(a).

Figure 2: Changes in Employment by Industry During the Pandemic



Notes: Historical average of annual employment growth rates computed using data from 2009:Q1 to 2019:Q4.
Source: Authors' calculations using the Household labor Force Survey (2018)

averages. We split industries into service sector and non-service sector groups, where definitions and the employment composition of each are reported in Table A.7 in Appendix A.¹³ We consider these broad groups of industries for two reasons. First, as noted in the recent literature, service sector workers are less likely to be able to work from home than other workers, differentially affecting economic activity during the lockdown (see Dingel et al., 2020; Bartik et al., 2020). Second, many industries in services are heavily dependent on international travelers. For example, in New Zealand in 2019 purchases by international tourists comprised 95% of accommodation services, 42% of food and beverage services, and 25% of arts and recreation services (Zealand, 2020). These industries are especially adversely affected by ongoing international border closures.

Figure 2(a) shows that service sector employment fell by 3% in the second quarter, and by over 4% in the third quarter. In contrast, Figure 2(b) shows that non-services employment growth was positive at 1% in the second quarter, and fell by just -0.1% in the third quarter.¹⁴ Insofar as employment growth reflects the growth of economic activity, the data suggests that service sector industries experienced a far larger contraction during the pandemic than did non-service industries. We take this as evidence of significant heterogeneity in the effects of lockdowns across industries.

In anticipation of the economic effects of the pandemic, the New Zealand government implemented a broad-based wage subsidy scheme.¹⁵ This subsidy was similar to policies adopted in other countries at this time.¹⁶ In New Zealand, firms and self-employed workers could apply for a wage subsidy from 1 March 2020 that paid a flat rate of NZ\$585.80 per full-time employee


¹³Our definition of service sector industries excludes the following industries: education and training; health care and social assistance; professional, scientific, and technical services; public administration and safety. We make these exclusions because the heavy involvement of the government in these industries in New Zealand limits their exposure to the market-based effects of pandemic.

¹⁴Note that New Zealand's first COVID case was not reported until February 28, and the government's response to the pandemic began in March (Hale et al., 2020). The first economic effects of the pandemic are largely not visible in quarterly data until 2020:Q2.

¹⁵For details, see <https://www.workandincome.govt.nz/covid-19/wage-subsidy/index.html>.

¹⁶For example, see Bishop et al. (2020) for an empirical analysis of the Job Keeper scheme adopted in Australia.

Figure 3: New Zealand Wage Subsidy and Income Relief Payment Schemes



Notes: Number of wage subsidy and COVID income relief recipients divided by total number of employees. Number of employees measured using Monthly Employment Indicators data.

Source: Statistics New Zealand and Ministry of Social Development.

per week and was available for eight weeks.¹⁷ The subsidy was equivalent to approximately 50% of median weekly earnings for full-time workers in 2020.¹⁸ Receipt of the wage subsidy was subject to several conditions: firms expected a 30% drop in revenue over the previous month due to the pandemic; firms must retain subsidized employees for the duration of the subsidy; and firms must continue to pay employees at least 80% of their usual wages for the duration of the subsidy.

Figure 3 shows that around 70% of all employees in New Zealand were supported by the subsidy between March and June 2020. Limited extensions of the wage subsidy were available between June and September, but many fewer firms received these payments. Due to broad coverage and high take-up, the subsidy was an expensive fiscal intervention with NZ\$14 billion paid out in 2020, equivalent to 40% of total government expenditure on social security in 2019.¹⁹ The New Zealand government also provided other fiscal interventions in response to the pandemic, but these were much smaller in scope. For example, unemployment benefit rates were effectively doubled for workers that lost jobs as a result of the pandemic between March and November 2020.²⁰ However, as Figure 3 shows, the number of workers supported by these additional unemployment payments was dwarfed by the number of workers receiving the wage subsidy.

Finally, we consider the likely distributional impacts of the pandemic by worker age. Figure 4 illustrates industry composition, job separation rates, and earnings by industry and age using data from the Linked Employer-Employee Database (LEED) for 2018.²¹ The data reveal distinct

¹⁷Employers of part-time workers were eligible for NZ\$350 per worker per week.


¹⁸Household income statistics from the Household labor Force Survey are available from Statistics New Zealand at <http://nzdotstat.stats.govt.nz/>.

¹⁹The cost of wage subsidy scheme is reported by the Ministry of Social Development at <https://www.msd.govt.nz/about-msd-and-our-work/publications-resources/statistics/covid-19/index.html>. Data on fiscal expenditure is reported by the New Zealand Treasury at <https://www.treasury.govt.nz/information-and-services/financial-management-and-advice/revenue-and-expenditure>.

²⁰Information on the COVID income relief payment is reported by the Ministry of Social development at <https://www.workandincome.govt.nz/covid-19/income-relief-payment/index.html>.

²¹Statistics New Zealand maintains LEED, which is compiled from administrative data associated with New Zealand’s tax system. This covers all workers at businesses and organizations producing goods and services in

Figure 4: Employment and Job Separation Rates by Industry Group and Age



Source: Authors' calculations using Statistics New Zealand's Linked Employer-Employee Data.

life-cycle patterns in employment composition, employment risk, and earnings. Figure 4(a) shows that young workers are significantly more likely to find themselves employed in the service sector than are older workers. Figure 4(b) shows that both young and service sector workers face higher rates of job separation than older and non-service sector workers, and are thus far more likely to leave their current place of employment in any given period.²² Figure 4(c) shows that young workers earn less than older workers in both sectors, while service sector workers earn less than non-service sector workers at all ages.

Taken together, these facts suggest that younger workers are likely to be disproportionately affected by the effects of the pandemic. Because the pandemic disproportionately affected the service sector, as illustrated in Figure 2, younger workers are more likely to lose their jobs. Moreover, these are workers that already face higher employment risk and are compensated with lower earnings. Young workers are thus more likely to be affected by the pandemic, but also the least able to weather its effects. For this reason, it is important to understand the distributional consequences of the fiscal interventions undertaken during the COVID-19 pandemic.

3. Model

We build a structural macroeconomic model in order to study the aggregate and distributional consequences of a pandemic-induced lockdown in New Zealand. The main model ingredients are life-cycle heterogeneous agents, multiple industries, frictional labor markets, endogenous employment risk, and exogenous productivity shocks. Our model features are motivated by three main considerations. First, the strict lockdowns and limited health effects experienced by New Zealand suggest a much stronger role for supply shocks than for demand shocks. For this reason, we model a lockdown as a sequence of productivity shocks reducing the ability of firms to produce using a given level of labor inputs. Second, there is significant heterogeneity in the effect

New Zealand with revenue greater than \$30,000.

²²Note that job separations data encompass both employee layoffs and quits, which means that separations may result in either unemployment or a job-to-job transition.

of lockdowns on production across industries. For this reason, we model firms as being more or less exposed to the negative productivity shocks accompanying lockdowns. Third, the effects of a lockdown are likely to differ by age. Our model replicates the life-cycle profile of employment across industries, which captures the fact that young workers are disproportionately exposed to a lockdown through higher rates of employment in the services sector. Additionally, the model mimics the age distribution of employment risk so that young workers are more likely to be laid off following a shock than are old workers. We believe these features provide a plausible basis for analyzing the consequences of lockdown shocks and the fiscal responses to them.

3.1. Households

Households live for a finite number of periods, where their age is indexed by $k \in [1, \dots, K]$. At age $K + 1$ households retire and consume their remaining networth. The problem of a household of age k is:

$$\begin{aligned}
V_k(b, e, i) &= \max_{c, b'} u(c) + \beta \mathbb{E} [V_{k+1}(b', e', i')] \\
\text{s.t. } c + b' &= (1 - \tau_y) [ew_{i,k} + (1 - e)\omega_u] + b(1 + (1 - \tau_y)r) \\
b' &\geq 0 \\
e', i' &\sim \Gamma_{k,e,i}
\end{aligned} \tag{1}$$

Above, e denotes the employment status of the household, where $e \in \{0, 1\}$ reflects unemployment and employment, respectively. The industry of employment is indexed by $i \in \{1, 2\}$, where $i = 1$ indicates the services sector, and $i = 2$ is all other industries. The household chooses consumption c , and liquid asset holdings b , where the latter is subject to a no-borrowing constraint. In retirement, households consume all remaining networth b , where final period utility is weighted by ψ_w :

$$\begin{aligned}
V_{K+1}(b) &= \psi_w u(c) \\
c &= b
\end{aligned} \tag{2}$$

Household income depends on age, employment status, and industry of employment. If employed, the household receives a wage $w_{k,i}$, which differs by age and industry. Unemployed households receive a constant unemployment benefit ω_u . All household income, including interest income on assets, is taxed at rate τ_y . Income is subject to idiosyncratic risk due to changes in employment status and industry of employment. This risk is characterized by a Markov chain $\Gamma_{k,e,i}$, which is exogenous from the perspective of the household. However, the parameters of the Markov chain are governed by the outcomes of a labor market search processes, described in more detail in Sections 3.3.

The final consumption good c is a composite of services and other goods consumption via an Armington aggregator:

$$c = \left[\chi \frac{1}{\varepsilon_c} c_1^{\frac{\varepsilon_c - 1}{\varepsilon_c}} + (1 - \chi) \frac{1}{\varepsilon_c} c_2^{\frac{\varepsilon_c - 1}{\varepsilon_c}} \right]^{\frac{\varepsilon_c}{\varepsilon_c - 1}} \tag{3}$$

Above, ε_c gives the elasticity of substitution between the two types of consumption. The solution to the expenditure minimization problem implies the following consumption demand functions

$$c_1 = \chi \left(\frac{p_1}{p} \right)^{-\varepsilon_c} c, \quad c_2 = (1 - \chi) \left(\frac{p_2}{p} \right)^{-\varepsilon_c} c, \quad (4)$$

where $\frac{p_i}{p}$ is the relative price of consumption in industry i , and the aggregate price index is given by

$$p = [\chi p_1^{1-\varepsilon_c} + (1 - \chi) p_2^{1-\varepsilon_c}]^{\frac{1}{1-\varepsilon_c}}. \quad (5)$$

3.2. Production Firms

Competitive firms in each industry produce output Y_i via linear production technologies $Y_i = Z_i \tilde{N}_i$, where Z_i is industry-specific productivity, and \tilde{N}_i is efficiency-adjusted units of labor. Firms maximize real profits

$$\max_{\tilde{N}_i} \frac{p_i}{p} Y_i - h_i \tilde{N}_i \quad (6)$$

where $\frac{p_i}{p}$ is the relative price of output, and h_i is the industry-specific real wage rate per efficiency unit of labor.

3.3. Labor Markets

Labor markets in each industry feature search frictions. Unemployed households search for work in their current industry i , while labor market entrepreneurs in each industry post vacancies to attract workers in that industry. Let v_i be the number of vacancies posted by entrepreneurs in industry i , and u_i is the total mass of workers searching for work in industry i . The matching technology is a Cobb-Douglas function given by


$$m(u_i, v_i) = M_i u_i^\alpha v_i^{1-\alpha},$$

where M_i is industry-specific match productivity, and α is the matching elasticity. The rate at which entrepreneurs fill vacancies is defined as $q_i^v = \frac{m(u_i, v_i)}{v_i}$. The rate at which unemployed households find jobs is defined as $q_i^u = \frac{m(u_i, v_i)}{u_i}$.

We assume that labor market entrepreneurs cannot age-discriminate between potential workers when posting job vacancies.²³ This means that entrepreneurs may be matched with workers of any age. Because households retire at age $K + 1$, any existing employment relationships with age- K workers are destroyed with certainty in the following period. Entrepreneurs matched to a worker receive the industry-specific production wage h_i . But each match is subject to an idiosyncratic productivity shock x each period, where x is drawn from a log-normal distribution $\log x \sim \mathcal{N}(\mu_x, \sigma_{x,k})$ with age-dependent standard deviation $\sigma_{x,k}$. The entrepreneur earns total revenue $h_i x$ and in turn pays workers an industry- and age-specific wage $w_{i,k}$.

²³This assumption follows Lugauer (2013), who builds a model with undirected search and life-cycle workers to study the effects of demographic change.

Figure 5: Endogenous Job Separations and the Distribution of Match Productivity



The value of a filled job for an entrepreneur in industry i , matched to a worker age k , and with match-specific productivity x is

$$J_{i,k}(x) = h_i x - w_{i,k} + \beta \mathbb{E} [\max\{0, J_{i,k+1}(x')\}] \quad (7)$$

where expectations are taken over x' . As in Fujita et al. (2012) we allow for endogenous job separations. Separations follow from endogenous firm shut-down decisions. Note that a shut down need not occur when current flow profits are negative, as entrepreneurs also consider the present discounted value of future profits. Rather, an entrepreneur ceases to operate and a job separation occurs when $J_{i,k}(x) \leq 0$. This shut down condition defines a threshold productivity level

$$\underline{x}_{i,k} = \frac{w_{i,k} - \beta \mathbb{E} [\max\{0, J_{i,k+1}(x')\}]}{h_i}, \quad (8)$$

which depends on industry-level productivity through the production wage h_i . All job matches with idiosyncratic productivity below this threshold are destroyed. Note that separations apply to both existing employment arrangements as well as newly matched employers and employees. This means that a worker may match with and separate from an entrepreneur in the same period before production takes place.

Figure 5 illustrates the log-normal distribution of match-specific productivity. The black vertical line indicates the threshold productivity $\underline{x}_{i,k}$ such that all matches with lower productivity are destroyed. When industry-level productivity Z_i falls, the production wage h_i also falls, which increases the threshold to $\underline{x}'_{i,k}$ as indicated by the red vertical line. The fall in aggregate productivity results in an increase in the fraction of matches that are unprofitable and which result in separations. Thus, a fall in aggregate productivity results in a larger number of job losses through employment separations.

Thus, we can derive expressions for job separation rates $\rho_{i,k}$ using the distribution of match-specific productivity:

$$\rho_{i,k} = \int^{\underline{x}_{i,k}} \phi_k(x) dx = \Phi_k(\underline{x}_{i,k}), \quad (9)$$

where $\phi_k(\cdot)$ and $\Phi_k(\cdot)$ denote log-normal PDFs and CDFs with age-dependent standard deviations $\sigma_{x,k}$. As discussed with reference to Figure 5 above, the job separation rates $\rho_{i,k}$ are a function

of aggregate productivities Z_i through the production wage h_i in the denominator of Equation (8).

Labor market entrepreneurs post vacancies up to the point at which marginal benefit equals marginal cost. The real vacancy posting cost in each industry is κ_i . The marginal benefit of a vacancy is the expected value of filling a job. The job filling value function for a worker of age k is described in Equation 7. However, the expected value is governed by the probability of finding a worker q_i^v , the probability that a given worker is age k , and the probability that the employment match does not separate before production takes place. The probability that a searching worker chosen at random is aged k is $\frac{u_{i,k}}{u_i}$, where $u_i = \sum_{k=1}^K u_{i,k}$. The assumption of free-entry thus implies that

$$\kappa_i = q_i^v \sum_{k=1}^K \frac{u_{i,k}}{u_i} \mathbb{E}_x [\max\{0, J_{i,k}(x)\}], \quad (10)$$

where the expectation is taken over idiosyncratic match productivity x . The maximum operator within the expectation takes into account the probability that a separation occurs before production takes place.

Finally, we assume that entrepreneurs are risk-neutral, own their own firms, and that all profits earned each period are immediately consumed.

3.3.1. Wage Determination

Because Nash bargaining is complicated by the two-sided heterogeneity of workers and firms, we follow several recent papers in the heterogeneous agents literature by assuming a simple wage setting rule.²⁴ In steady state, we assume that wages are determined by the marginal labor market entrepreneur. That is, wages are set by the firm that is indifferent between continuing and shutting down. For such a firm the value of a filled job is nil, and so Equations (7) and (9) yield

$$w_{i,k} = h_i \Phi_k^{-1}(\rho_{i,k}) + \beta \mathbb{E} [\max\{0, J_{i,k+1}(x')\}], \quad (11)$$

where Φ_k is the CDF over match-specific productivity x .

Outside of steady state, we assume that worker wages are sticky and respond to production wages with a constant elasticity ε_w :

$$w_{i,k,t} = w_{i,k} \left(\frac{h_{i,t}}{h_i} \right)^{\varepsilon_w}. \quad (12)$$

3.3.2. Employment and Industry Transitions

Given the labor market search environment, we are now in a position to characterize the household employment and industry transitions described by the Markov chain $\Gamma_{k,e,i}$.

Consider a household that is currently employed ($e = 1$) and working in industry i . At the end of the period, employees are subject to endogenous job separations at rate $\rho_{i,k+1}$ (i.e. separations


²⁴See, for example, Gornemann et al. (2016) and Graves (2020) for other examples of simple wage setting rules in models with labor search and heterogeneous agents.

occur immediately prior to age $k + 1$). Separated workers that were employed in industry i may switch industries with exogenous probability $1 - \pi_{ii}$. After industry transitions take place, workers search for jobs in their new industry i' and match with a potential employer with probability $q_{i'}^u$. However, because all employment relationships are subject to match productivity shocks at the beginning of the period, even new matches may separate before production takes place. Thus, the effective job finding rate is $q_{i'}^u(1 - \rho_{i',k+1})$: the probability of finding a job and surviving the initial match productivity shock.

Now consider a household that is currently unemployed ($e = 0$) and has previously worked in industry i . Before searching for work, the household may switch industries with exogenous probability $1 - \pi_{ii}$. The household then searches for jobs in the new industry i' and matches with a potential employer with probability $q_{i'}^u$. As above, match productivity shocks take place which can disrupt employment before it begins, and so the effective job finding rate is again $q_{i'}^u(1 - \rho_{i',k+1})$.

Figure 6 shows an event tree illustrating the employment and industry transitions faced by age- k worker that is currently employed in industry i . The initial state vector of the worker is $(k, e = 1, i)$, and the transition to outcome i' indicates that a worker has switched from industry i to another industry i' . Along the upper branch, the worker is not separated and remains employed and in the same industry i . Along the lower branch, the worker is separated from their current job, may switch industries, searches for work, and may end up employed or unemployed in either the original industry i or the new industry i' .

Figure 6: Employment and Industry Transitions Event Tree



The Markov chain $\Gamma_{k,e,i}$ can be separated into two sub-Markov chains: $\Gamma_{e,k}^{i,i}$ describes employment transitions for workers that remain in the same industry i ; and $\Gamma_{e,k}^{i,i'}$ describes employment transitions for households that switch from industry i to i' . Let $[1 - e, e]'$ be the state vector describing employment status. Recall that unemployment corresponds to $e = 0$ and employment corresponds to $e = 1$. Then for each transition matrix, the upper row reports

probabilities for transitions from unemployment, and the bottom row reports probabilities for transitions from employment. For households remaining in the same industry i , the transition matrix is given by

$$\Gamma_{e,k}^{i,i} = \begin{bmatrix} \pi_{ii}(1 - q_i^u(1 - \rho_{i,k+1})) & \pi_{ii}q_i^u(1 - \rho_{i,k+1}) \\ \rho_{i,k+1}\pi_{ii}(1 - q_i^u(1 - \rho_{i,k+1})) & 1 - \rho_{i,k+1} + \rho_{i,k+1}\pi_{ii}q_i^u(1 - \rho_{i,k+1}) \end{bmatrix},$$

For households that switch from industry i to i' , the transition matrix is given by

$$\Gamma_{e,k}^{i,i'} = \begin{bmatrix} (1 - \pi_{ii})(1 - q_{i'}^u(1 - \rho_{i',k+1})) & (1 - \pi_{ii})q_{i'}^u(1 - \rho_{i',k+1}) \\ \rho_{i,k+1}(1 - \pi_{ii})(1 - q_{i'}^u(1 - \rho_{i',k+1})) & \rho_{i,k+1}(1 - \pi_{ii})q_{i'}^u(1 - \rho_{i',k+1}) \end{bmatrix}$$

Finally, the full Markov transition matrix for employment and industry transitions is

$$\Gamma_{k,e,i} = \begin{bmatrix} \Gamma_{e,k}^{1,1} & \Gamma_{e,k}^{1,2} \\ \Gamma_{e,k}^{2,1} & \Gamma_{e,k}^{2,2} \end{bmatrix} \quad (13)$$

Hence, $\Gamma_{k,e,i}$ is a 4×4 matrix that multiplies the combined state vector over employment status and industry: $[(1 - e, 1), (e, 1), (1 - e, 2), (e, 2)]'$. Each row of the transition matrix corresponds to the current labor market status and industry of a household, while each column, in turn, corresponds to the states into which workers transition next period.

3.3.3. Labor Market Flows and Aggregate labor Supply

Let the number of age- k households and employees in industry i be denoted by $I_{i,k}$ and $N_{i,k}$, respectively. The number of job searchers in industry i of age- k evolves according to the following law of motion:

$$u_{i,k} = \pi_{ii}(I_{i,k-1} - (1 - \rho_{i,k})N_{i,k-1}) + (1 - \pi_{i'i'}) (I_{i',k-1} - (1 - \rho_{i',k})N_{i',k-1}) \quad (14)$$

Above, the first additive term corresponds to households in industry i last period whose jobs were destroyed but who stayed in industry i , as well as households who were unemployed in industry i , stayed in the same industry, but could not find a job. The second additive term has a similar interpretation, but tracks households in the other industry i' who were separated or were searching but but switched industries from i' to i .

The number of age- k workers employed in industry i is given by the sum of age- $k - 1$ workers who kept their job this period, and the number of job searchers who found a job and survived endogenous separations:

$$N_{i,k} = (1 - \rho_{i,k})N_{i,k-1} + u_{i,k}q_i^u(1 - \rho_{i,k}) \quad (15)$$

Finally, aggregate labor supply in each industry is given by total efficiency units of labor provided by working households. This consists of the total number of workers in each age group scaled by average match productivity in that age group:

$$\tilde{N}_i = \sum_{k=1}^K N_{i,k}\bar{x}_{i,k}$$

where $\bar{x}_{i,k} = \int^{\underline{x}_{i,k}} x\phi_k(x)dx$.

3.4. Government

There is a government that manages fiscal policy along any transition path. Government expenditures consist of non-valued government spending G_t , unemployment benefits UB_t , wage subsidies WS_t , and lump-sum transfers TR_t . The government then collects taxes T_t , and issues debt B_t^g at interest rate r . The government budget constraint is

$$G_t + (1 + r)B_t^g + UB_t + WS_t + TR_t = B_{t+1}^g + T_t. \quad (16)$$

We assume that government spending G_t is allocated across services and other goods in the same way that household consumption expenditure is allocated, according to Equation (3).²⁵ Unemployment benefits UB_t and tax revenues T_t are aggregated across agents in the economy:

$$UB_t = \sum_{k=1}^K \left(\int (1 - e)\omega_u d\lambda_{k,t} \right)$$

$$T_t = \tau_y \left(\sum_{k=1}^K \int (1 - e)\omega_u d\lambda_{k,t} + \sum_{k=1}^K \int ew_{i,k} d\lambda_{k,t} + \sum_{k=1}^{K+1} \int rbd\lambda_{k,t} + \Pi_t \right)$$

where $\lambda_{k,t}$ is the distribution over households of age k at time t , and Π_t are labor market entrepreneur profits. Wage subsidies WS_t and lump sum transfers TR_t are only distributed following the lockdown shocks discussed in Section 5.

In all of the experiments that follow, we assume that government debt B_{t+1}^g adjusts to ensure that the fiscal budget constraint holds in each period t along a transition path. Government spending G_t is assumed to remain constant. We then adjust the tax rate τ_y to ensure that government debt returns to steady state in the long run. We discuss the details of this tax policy in Section 5.

3.5. Equilibrium

A formal definition of the recursive competitive equilibrium is stated in Appendix B.2. In the experiments conducted in Section 5, we solve for equilibrium along a transition path following a series of ex-ante unexpected shocks. Along the transition, we maintain the assumption of a small open economy. That is, the interest rate r is held constant, and the domestic economy may borrow from or save with the rest of the world. This is reflected in fluctuations in net exports in the aggregate real resource constraint:

$$\frac{p_{1,t}}{p_t} Y_{1,t} + \frac{p_{2,t}}{p_t} Y_{2,t} = C_t + \Pi_t + G_t + \kappa_1 v_{1,t} + \kappa_2 v_{2,t} + NX_t,$$

where the left side of the equation is total output, C_t is aggregate household consumption, Π_t is firm profits which are consumed by entrepreneurs during the period in which they are earned, $\kappa_i v_{i,t}$ are total vacancy posting costs for industry i , and NX_t are net exports. As with government spending, we assume that net exports are allocated across services and other goods in the same way that household consumption expenditure is allocated, according to Equation (3).²⁶

²⁵See Appendix B.1 for details.

²⁶See Appendix B.1 for details.

Table 1: Externally Calibrated Model Parameters

Description	Parameter	Value	Source
Risk Aversion	γ	2.000	Standard
Real Interest Rate	r	0.023	Reserve Bank of New Zealand
Unemployment Benefit Rate	ω_u	0.252	Work and Income New Zealand
Services Consumption Share	χ	0.519	Statistics New Zealand
Elasticity of Substitution	ε_c	2.000	Hobijn et al. (2019)
Matching Function Elasticity	ε_m	0.750	Razzak (2009)
Real Wage Elasticity	ε_w	0.100	Authors' Estimates using LEED
Vacancy Filling Rate	q^v	0.700	Christoffel et al. (2009)

4. Model Calibration

We calibrate the model so that the steady state matches a range of household and labor market characteristics in New Zealand prior to the onset of the pandemic in 2020. Importantly, the calibrated model reproduces the observed industry and age distributions of employment, employment risk, and wages. This ensures that the model captures ex-ante household exposures to lockdown shocks.

The model is calibrated at a quarterly frequency. Households are born at age 20 and work for 180 periods until retirement at age 65. Table 1 reports the other externally calibrated parameters. The risk aversion parameter γ is set to 2. The annual real interest rate is 2.3 percent, the average for the years 2000 to 2019. The unemployment benefit rate ω_w is set at 20% of the median wage, consistent with the Jobseeker Support benefit rate in 2019 for adults aged over 25. The consumption share of services χ is set to 0.519, which is the average for the years 2000 to 2019. The elasticity of substitution between services and non-services ε_c is set to 2, consistent with recent empirical evidence on substitutability between goods and services in Hobijn et al. (2019).

The matching function elasticity ε_m is set to 0.75, following empirical estimates from New Zealand in Razzak (2009). We estimate the wage elasticity ε_w using quarterly, industry-by-region LEED data from 1999-2019. Combining the wage dynamics equation Equation (12) with the first order condition for the production firms ($h_i = \frac{p_i Y_i}{p N_i}$) yields an estimable relationship between wages and employment: $\Delta \log w_{i,k,t} = -\varepsilon_w \Delta \log N_{i,t} + \varepsilon_w \Delta \log \left(\frac{p_i Y_i}{p} \right)$. We regress the log-difference in wages on the log-difference in employment, controlling for industry-year and region-year fixed effects. Because employment is endogenous to output which is unobserved, we run an instrumental variables regression, constructing a Bartik instrument from industry shares within each region and national employment growth rates in each industry. The IV results yield an elasticity of 0.121 (S.E. 0.015), so we set ε_w to 0.10. The vacancy filling rate q^v is assumed to be the same across industries, and is set to 0.70 (see Christoffel et al., 2009).

Table 2: Model Parameters and Moments

Parameter		Value	Moment	Model	Data	Source
<i>A. Miscellaneous Parameters</i>						
Discount factor	β	0.994	Networth/GDP	0.767	0.769	RBNZ, 2019
Retirement weight	ψ	140.000	Networth: 55-64/25-34	12.068	12.099	HES, 2017
Relative productivity	Z_1/Z_2	0.385	Earnings: Non-Services/Services	1.301	1.301	LEED, 2019
Tax rate	τ_y	0.212	Government Spending/GDP	0.188	0.188	RBNZ, 2017
<i>B. Job Separation Process Parameters</i>						
Mean	μ_ρ	0.026	Separation rates by age	See Figure 7(a)		LEED, 2019
Persistence	ρ_ρ	0.978	Separation rates by age	See Figure 7(a)		LEED, 2019
Initial value, services	$\rho_{1,k=0}$	0.091	Separation rates by age	See Figure 7(a)		LEED, 2019
Initial value, other	$\rho_{2,k=0}$	0.068	Separation rates by age	See Figure 7(a)		LEED, 2019
<i>C. Idiosyncratic Match Productivity Process Parameters</i>						
Mean	μ_σ	1.250	Earnings: 30-34/20-24	1.577	1.580	LEED, 2019
Persistence	ρ_σ	0.963	Earnings: 40-44/30-34	1.171	1.170	LEED, 2019
Initial standard deviation	$\sigma_{k=0}$	0.100	Vacancy costs/GDP	0.032	0.020	Standard
<i>D. Industry Labour Market Parameters</i>						
Industry transition, services	$1 - \pi_{11}$	0.121	Industry-age composition	See Figure 7(c)		LEED, 2019
Industry transition, other	$1 - \pi_{22}$	0.064	Industry-age composition	See Figure 7(c)		LEED, 2019
Employment probability, $k = 0$	$\lambda_{e=1,k=0}$	0.999	Industry-age composition	See Figure 7(c)		LEED, 2019
Services probability, $k = 0$	$\lambda_{i=1,k=0}$	0.584	Jobseeker support by age	See Figure 7(d)		MSD, 2019
Job finding rate	q^u	0.601	Jobseeker support by age	See Figure 7(d)		MSD, 2019

Source: Authors' calculations using data from the Household Economic Survey, the Ministry of Social Development, the Linked Employer-Employee Database, and the Reserve Bank of New Zealand.

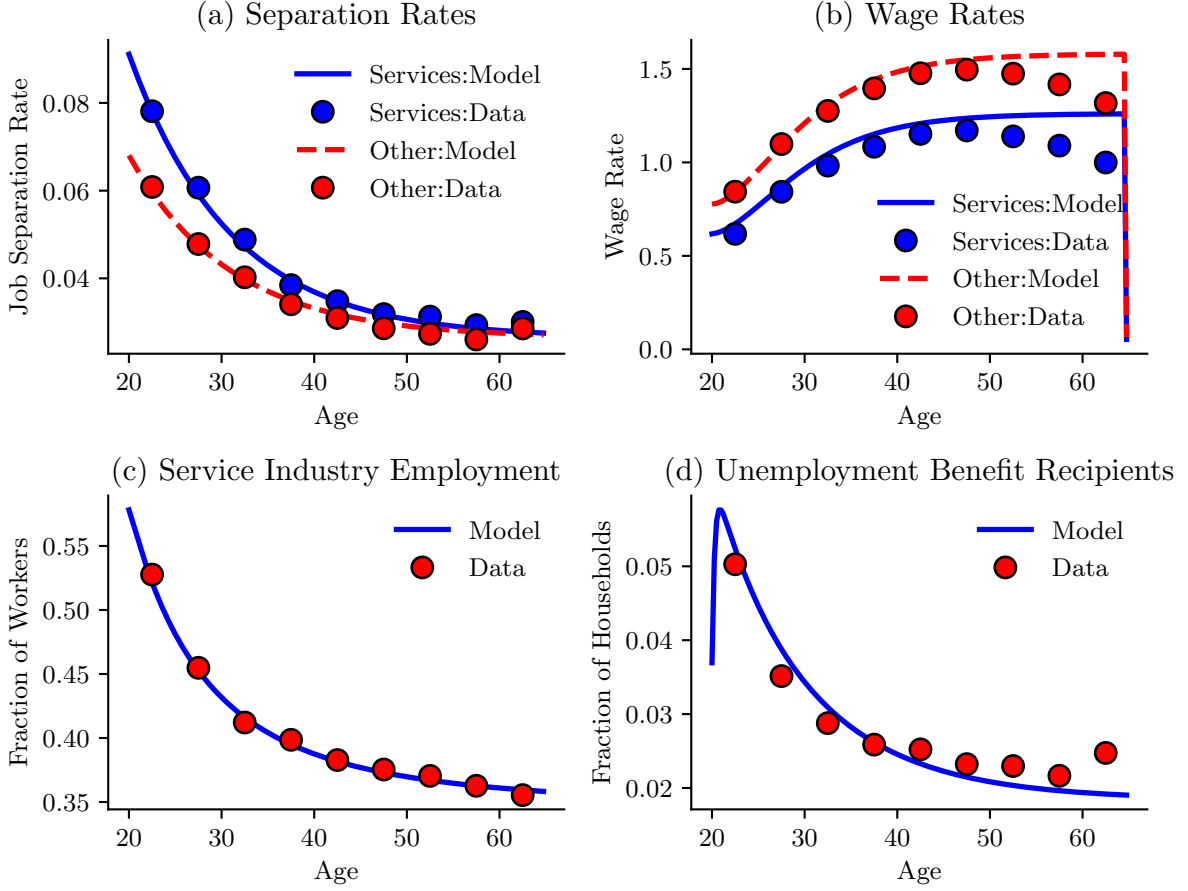
Table 2 reports internally calibrated parameters, and the moments governed by those parameters in both the model and data. Panel A reports parameters governing wealth accumulation, cross-industry earnings, and taxes. The discount factor β is set to match the ratio of networth to quarterly GDP. The utility weight on retirement ψ is set to match networth held by households in the decade prior to retirement relative to the networth held by households aged 25 to 34. Productivity in the non-services sector Z_2 is normalized to one, and the relative size of productivity in the services sector Z_1/Z_2 is set to match the ratio of average labor earnings across industries. We set the tax rate τ_y to match the ratio of government spending and GDP, since government spending is determined residually from the government budget constraint.

Panel B of Table 2 reports parameters that determine job separations. New Zealand’s LEED data provides job separation rates by industry and by household age in 5-year age groups. Because we solve the model at a quarterly frequency, we generate job separation rates for every age using a simple auto-regressive process, $\rho_{i,k} = (1 - \rho_\rho)\mu_\rho + \rho_\rho\rho_{i,k-1}$, where μ_ρ is the average separation rate, ρ_ρ is the persistence of the process, and there is a different initial condition for each industry $\rho_{i,k=0}$. We set the parameters of the auto-regressive process so that the average separation rates by industry and age group in the model match those observed in the LEED data. Figure 7(a) shows that separation rates in the model provide an extremely close fit to the data.

Panel C of Table 2 reports the parameters of the idiosyncratic match productivity process. Recall that match productivity is distributed according to a log-normal distribution with a constant mean μ_x and age-dependent standard deviations $\sigma_{x,k}$. We normalize μ_x so that $\mathbb{E}(x|k = 1) = 1$. For parsimony, we model the age-dependent standard deviation as an auto-regressive process with mean μ_σ , persistence ρ_σ , and initial condition $\sigma_{k=0}$. Conditional on job separation rates, the match productivity process governs steady state wages $w_{i,k}$ via Equation (11). Match productivity is also related to vacancy costs through the vacancy posting Equation (10). Thus, we set the match productivity parameters to target the slope of life-cycle earnings and the vacancy costs-to-GDP ratio. Specifically, we target the earnings ratio for 30-to-34 year-old workers to 20-24 year old workers, and the ratio for 40-to-44 year-old workers to 30-to-34 year-old workers. Figure 7(b) shows the life-cycle profile of earnings for each industry in the model and the data. The model does a good job of matching the early life-cycle earnings profile. The model fails to capture the hump-shape in life-cycle earnings because the simple auto-regressive process generates a monotonically increasing variance of match productivity over age. The rising variance of match productivity is associated with a monotonically increasing wage profile, rather than the declining wage profile prior to retirement observed in the data. It is possible to generate the hump-shape wage profile, but at the cost of a more complicated parameterization of the match productivity process. Finally, note that sudden drop in wages for the oldest workers is because the continuation value of a filled job for workers immediately prior to retirement is zero, which results in the wage decline via Equation (11).

Panel D of Table 2 reports parameters for industry employment transitions and the job finding rate. These parameters include the industry transition rates $1 - \pi_{11}, 1 - \pi_{22}$, initial probability of employment $\lambda_{e=1,k=0}$, initial probability of working in the services sector $\lambda_{i=1,k=0}$,

Figure 7: Job Separation and Wage Rates by Industry and Age



Notes: Authors' calculations and data from HLFS, LEED, and MSD

and the job finding rate q^u . Note that we allow industry transition rates to differ across industries, but assume that the job finding rate q^u is the same across industries in the steady state. We use a simulated method of moments procedure to choose these parameters by matching the life-cycle profiles of industry employment composition and the rate of workers on the unemployment benefit.²⁷ We choose to match the rate of unemployment benefit receipts, rather than the unemployment rate, because our measure of job separations taken from LEED data includes both separations into unemployment and job-to-job transitions. This means that the model-implied unemployment rate is too high relative to the observed unemployment rate. However, the rate at which workers in the model receive unemployment benefits corresponds to the same rate in the data. Figure 7(c) shows that the model closely matches the rate of employment in the services sector by age. Figure 7(d) shows that the model fits the rate of unemployment benefit receipts on average across the life-cycle, where younger workers are more likely to find themselves without work than are older workers.

²⁷Note that industry composition and unemployment benefit receipts are easily computed from the Markov chain transition matrix $\Gamma_{k,e,i}$ in Equation (13). This does not require us to solve the entire model for this part of the calibration.

5. Effects of a Lockdown

5.1. Calibration of Lockdown Shocks and Policy Responses

First, we introduce a sequence of industry-specific productivity shocks $Z_{1,t}$ and $Z_{2,t}$. These shocks reflect the effects of a strict domestic lock-down in the first quarter of the pandemic, and international border closures that persist for an additional three quarters. We assume that service sector firms are adversely affected by both the domestic lock-down and the international border closures. We allow the effects of the domestic lock-down and international border to differ, so that the size of the services sector productivity shock in the first quarter of the lockdown differs from the shocks that occur in the following three quarters. We then assume that other industry firms are affected by the initial domestic lock-down only. Thus, non-service firm productivity declines in the first quarter of the lockdown only.

Second, we introduce a conditional wage subsidy policy indexed by policy parameters τ_w and Θ . As discussed in Section 2, a conditional wage subsidy was the primary fiscal policy response to the lockdown in New Zealand. We characterize the policy as a lump-sum payment τ_w , and a condition under which a firm becomes eligible to receive the payment. We assume the subsidy is available in the first quarter of the lockdown only, which approximates the sharp fall in subsidy receipts from July onward, as shown in Figure 3. Firms are eligible for the subsidy if they experience a sufficiently large fall in revenues relative to steady state:

$$\frac{h_{i,t}x - h_i x}{\mathbb{E}(J_{i,k})} \leq \Theta \quad (17)$$

We scale the decline in revenues by the steady state average value of an age k worker firm $\mathbb{E}(J_{i,k})$. We do this rather than scale by steady state revenues $h_i x$ so that the policy generates a distribution of subsidy recipients. Because the distribution over match productivity varies by age, Equation 17 implies that the probability of receiving a subsidy for a firm in industry i with a worker aged k is

$$\Pr(\text{Subsidy}|i, k) = \Pr\left(x \geq \frac{\Theta \mathbb{E}(J_{i,k})}{h_{i,t} - h_i}\right) = 1 - \Phi_k\left(\frac{\Theta \mathbb{E}(J_{i,k})}{h_{i,t} - h_i}\right) \quad (18)$$

where the first equality follows from the fact that $h_{i,t} < h_i$ during the lockdown, and Φ_k is the CDF over match productivity for age k workers.

We calibrate the lock-down shocks and policy parameters to match several key observations about the evolution of the New Zealand labor market in 2020. The details of this calibration are reported in Table 3. First, we set $Z_{1,t=1}$ and $Z_{1,t=2,3,4}$ to match the declines in service sector employment during the first two quarters of the lockdown. Second, we set $Z_{2,t=1}$ to match the decline in non-service sector employment in the first quarter of the lockdown. Figure 2 shows that relative to historical average growth rates, services employment declined by 5.2% in 2020:Q2 and by 6.7% in 2020:Q3. Non-service sector employment declined by 1.0% in 2020:Q2. To match these targets, the calibration yields a 59.5% decline in services productivity relative to steady state in the first quarter, a 6.5% decline in services productivity for the following three quarters, and a 33.2% decline in non-services productivity in the first quarter.

Third, we set τ_w to match the effective size of the wage subsidy. Payments under the wage subsidy scheme were \$585.80 per worker per week, which is approximately 50% of median weekly earnings in New Zealand. Note, however, that while costs for firms in our model consist of wages only, firms in reality also face a variety of other fixed and variable costs. In order to map the model to the data, we adjust the size of the wage subsidy to account for the proportion of total costs consisting of labor costs. Using data from the 2019 Annual Enterprise Survey, we compute that wages and salaries are 18% of firms' total expenditures on average.²⁸ Thus the model-implied wage subsidy τ_w is equivalent to $0.5 \times 18 = 9\%$ of the median wage in the model. Finally, we set Θ so that 75% of all workers receive the wage subsidy, which is the peak coverage of the subsidy as shown in Figure 3.

Table 3: Pandemic Parameters and Moments

Parameter	Value	Moment	Model	Data
$Z_{1,t=1}/Z_1$	0.405	Employment Growth, Services, 2020:Q2	-0.052	-0.052
$Z_{1,t=2,3,4}/Z_1$	0.935	Employment Growth, Services, 2020:Q3	-0.067	-0.067
$Z_{2,t=1}/Z_2$	0.668	Employment Growth, Other, 2020:Q2	-0.010	-0.010
τ_w	0.113	Subsidy Size to Median Wage	0.090	0.090
Θ	-0.124	Fraction Receiving Subsidy	0.751	0.750

Notes: Employment growth rates are computed relative to historical average growth rates.

Source: Authors' calculations using data from the Household Labour Force Survey and the Ministry of Social Development.

Finally, as noted in Section 3.4 we ensure that the fiscal budget constraint holds in each period by allowing government debt B_{t+1}^g to adjust in response to shocks to revenue and spending (i.e. unemployment benefits, and wage subsidies). We then adjust the tax rate τ_y to ensure that government debt returns to its steady state value in the long run. Specifically, we assume that the tax rate follows an AR(1)

$$\tau_{y,t} = (1 - \rho_\tau)\tau_y + \rho_\tau\tau_{y,t-1}$$

with persistence $\rho_\tau = 0.95$. At the onset of the lockdown, the tax rate remains at its steady state value. We assume that in the period immediately after the lockdown is lifted, the tax rate jumps and then follows the AR(1) process over the remainder of the transition path. We calibrate the size of the initial jump in the tax rate to ensure that government debt has returned to its steady state value after 80 model periods (i.e. 20 years).

²⁸The Annual Enterprise Survey is provided by Statistics New Zealand and reports financial statistics from a survey of over 500,000 businesses.

5.2. A Lockdown With and Without a Fiscal Policy Response

To study the effects of the pandemic induced lockdown and fiscal policy response, we compare the evolution of the model economy to a counterfactual economy with identical productivity shocks but absent the wage subsidy policy. In Figures 8 to 11, solid blue lines illustrate the baseline economy with lockdown shocks and policy while dashed red lines illustrate the counterfactual economy absent the subsidy.

Figure 8 shows the evolution of several macroeconomic variables over the course of the lockdown. In the first quarter of the lockdown, productivity in both industries declines significantly. This leads to a large fall in output.²⁹ With falling output and sticky wages, firm profits decline which leads to a decrease in demand for labor. As a result, aggregate employment falls by around 2.7% in the first quarter, and is 3.9% below steady state in the second quarter. Lower employment is also associated with a higher number of workers out of work and receiving unemployment benefits, with this number rising by 2.6 and 3.8 percentage points in each of the first two quarters. Rising unemployment is associated with a decrease in household income, which leads to a 1.82% decline in household consumption in the first period after the shock. Over the course of the lockdown, government debt rises 5.9% in order to finance both unemployment benefits and wage subsidies. When the lockdown is over, the tax rate jumps by one percentage point in order to ensure debt sustainability over the long run.

Figure 8 also illustrates that the wage subsidy policy acts to dampen the macroeconomic contraction induced by the lockdown. In the counterfactual economy with no policy intervention, on impact employment would have fallen a further 5%, unemployment would have increased a further 5%, and aggregate consumption expenditure would have fallen an additional 1%. Notice that although the wage subsidy is only in place for the first quarter of the lockdown, it has a persistent effect on macroeconomic aggregates. In particular, employment and consumption are higher for the entire duration of the lockdown with the wage subsidy policy in place. Government debt and taxes after the lockdown are only marginally higher under the wage subsidy policy than the counterfactual with no such policy in place. This is because the number of unemployment benefits paid out are larger and tax revenues from incomes are lower in the absence of the wage subsidy. As a result, government debt rises in any case, and higher future taxes will be required to repay this debt.

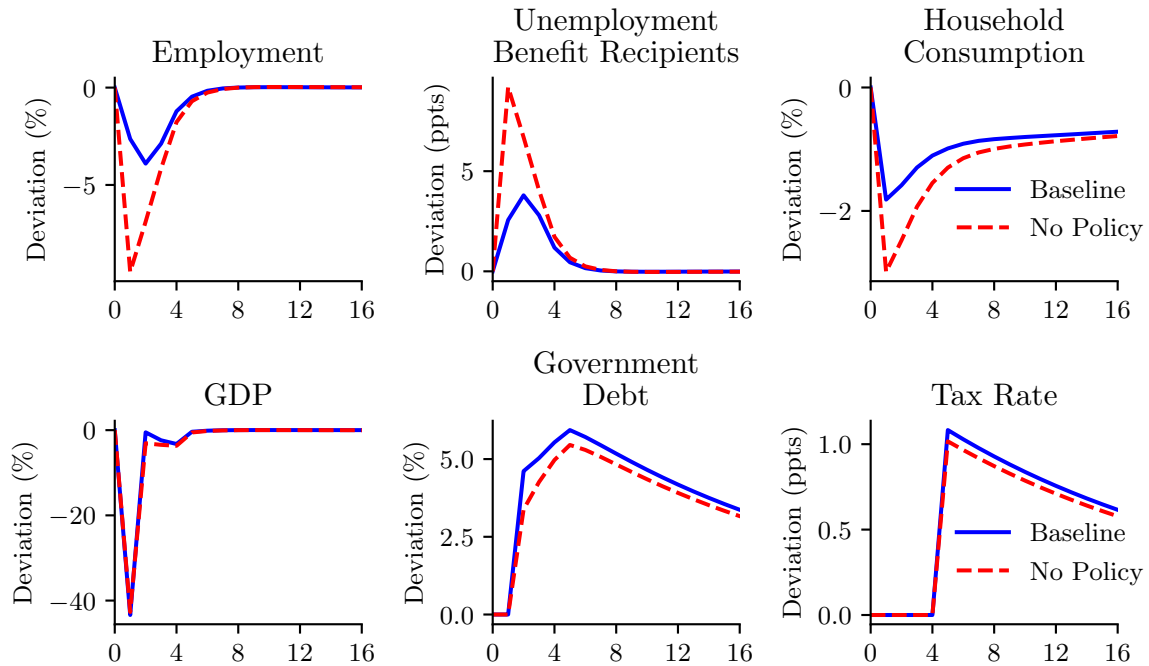
Note that the model predicts an exceptionally large decline in GDP. This is in contrast to the 11% quarterly decline in GDP observed in New Zealand in June 2020.³⁰ The reason for this discrepancy is that in our model output is entirely supply determined.³¹ The size of the productivity shocks $\{Z_{1,t}, Z_{2,t}\}$ that drive output are then calibrated to match the observed declines in employment conditional on the size and scope of the of wage subsidy policy implemented in New Zealand. With 75% of New Zealand workers receiving the subsidy, we require particularly large productivity shocks to squeeze entrepreneur profits in order to induce enough job separation to reduce employment. Note, however, that because output is

²⁹Note, aggregate real GDP is defined as $\frac{p_1}{p}y_1 + \frac{p_2}{p}y_2$.

³⁰See <https://www.stats.govt.nz/indicators/gross-domestic-product-gdp>.

³¹See Guerrieri et al. (2020) for a model of the pandemic in which aggregate output is partly demand-determined.

Figure 8: Aggregate Macroeconomic Variables During the Lockdown

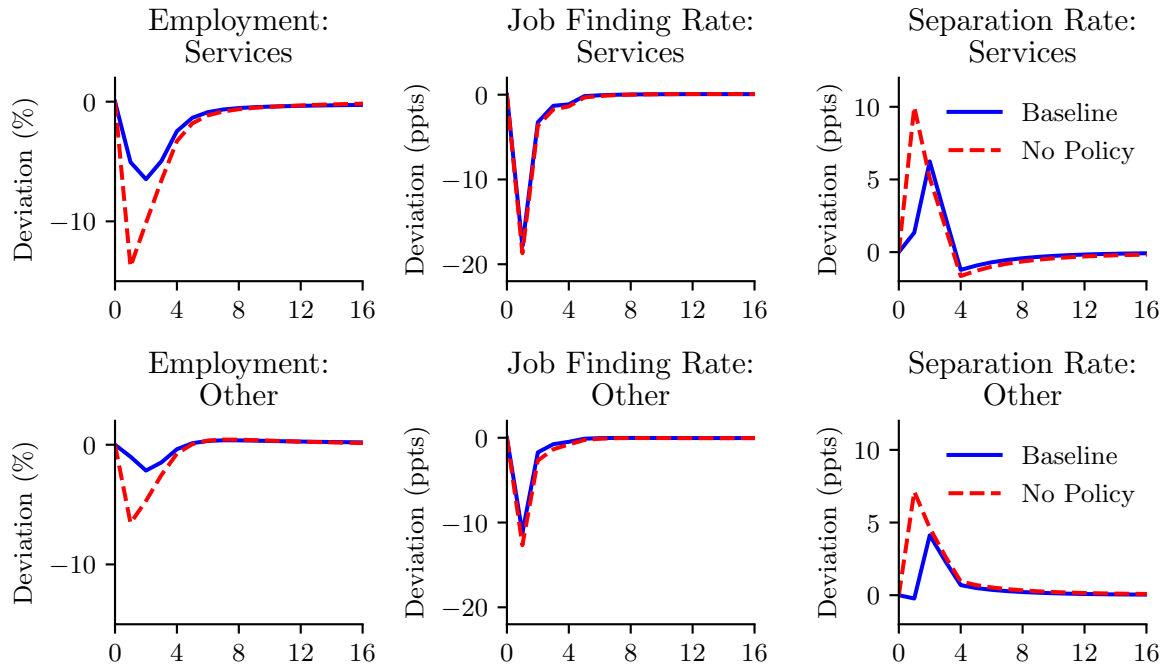


supply-determined Figure 8 shows essentially no difference in the evolution of GDP across economies with and without the wage subsidy policy. Rather, output is entirely determined by the size and scope of the lockdown, while the wage subsidy policy has a large impact on the distribution of employment, incomes, and consumption.

Figure 9 illustrates labor market outcomes across industries during the lockdown. The service sector is more adversely affected by the lockdown than the non-service sector because it is affected by a larger initial shock and is subject to the ongoing effects of border closures. As is observed for aggregate variables, the wage subsidy significantly dampens the effects of the lockdown shocks across industries. Absent the wage subsidy, services employment would have fallen an additional 10 percent below steady state, and non-service sector employment would have fallen by a further 5 percent. The third column in Figure 9 shows that average job separation rates change very little with the wage subsidy in place, whereas job separation rates rise significantly in the absence of the wage subsidy. In contrast, job finding rates are little changed by the presence of the wage subsidy. Thus, the wage subsidy largely acts through preventing job losses, rather than by promoting firm hiring.

In addition, Figure B.16 in Appendix B.3 illustrates changes in the industry composition of workers during the lockdown. While industry switching probabilities π_{ii} are exogenous, industry composition endogenously adjusts along a transition path because households only switch industries after a job separation. Since job separation rates in the services industry rise by more than separation rates in the other industry, more service workers are exposed to exogenous industry switching events. As a result, the services share of employment shrinks by around one percent and recovers slowly following the lockdown. Note that these industry transitions somewhat alleviate pressure on workers that remain in the service sector following the shocks. If,

Figure 9: Labor Market Outcomes During the Lockdown By Industry



in contrast, there was no industry switching during the lockdown, the higher number of workers searching for a service sector job would further reduce job finding rates and the employment prospects of service sector workers.

The effects of the lockdown are also unequally distributed across households. Figure 10 illustrates average job separation rates for workers aged 20 to 25, 25 to 30, and 30 to 35. We focus here on the youngest workers since they are most affected by the lockdown.³² We find that job separations rates rise by much more for the youngest workers. This is for two reasons. First, young workers are more likely to be employed in the service sector which faces larger lockdown shocks. Second, younger workers are associated with lower average current productivity and higher expected future productivity and profitability because $\sigma_{\sigma,k}$ is increasing with age. This means that wages for young workers are high relative to their current marginal product of labor, and with sticky wages the profitability of these workers is very sensitive to shocks. And so although industry-level productivity shocks are the same across workers of all ages, the job separation rates of the young are much more sensitive to these shocks than are the job separation rates of the old.

Figure 11 shows the age distribution of consumption expenditures for the youngest households over the lockdown.³³ As unemployment rises, and disproportionately among young workers, consumption falls by for young than older households. This is both because average income falls further for the young, but also because young households have lower stocks of savings with which to insure against employment shocks. Although unemployment also rises for older households, their large savings buffers mean that their expenditures are virtually unchanged.

³²See Figure B.17 in the Appendix for job separation rates for workers in all age groups.

³³See Figure B.19 in the Appendix for the consumption patterns of households in all age groups.

Figure 10: Job Separation Rates by Age




Figure 11: Household Consumption by Age




Figure 12: Wage Subsidy Receipts by Industry and Worker Age




Figure 12 shows that young workers and those in the service sector are most likely to receive the wage subsidy. We report the fraction of continuing employment relationships supported by the wage subsidy in the first quarter of the lockdown. Younger and service sector workers are more likely to receive the subsidy for two reasons. First, larger productivity shocks in the service sector lead to larger declines in revenues so that these firms are more likely to satisfy Equation 17. Second, firms with younger workers and lower average productivity experience larger relative declines in revenue for a given aggregate productivity shock and so are also more likely to receive the wage subsidy.

In order to assess who benefits the most from the wage subsidy, Table 4 reports the number of jobs saved by the wage subsidy in the first quarter of the lockdown for different groups of workers. The table shows the number of jobs saved as a fraction of steady state employment

Table 4: Jobs Saved by the Wage Subsidy Scheme

	Sector of Employment			Worker Age			
	Total	Services	Non-Services	20-29	30-39	40-49	50-65
Jobs Saved	0.068	0.087	0.055	0.171	0.055	0.037	0.031

Notes: The measure of jobs saved reported for each group is the difference between employment in the baseline model with the wage subsidy and employment in the model absent the wage subsidy. Jobs saved computed during the first quarter of the lockdown. All values are computed as fractions of steady state employment.

within each group. The first column shows that the wage subsidy scheme preserved 6.8% of aggregate steady state employment relationships, which is equivalent to 183,600 jobs in the New Zealand labor market.³⁴ This model-based estimate is similar to empirical estimates of the effect of wage subsidy schemes from similar countries. For example, Bishop et al. (2020) estimate that the Australian Job Keeper wage subsidy scheme saved around 700,000 jobs, or around 5.4% of the Australian labor force.³⁵ The second and third columns of Table 4 show that the service sector disproportionately benefits from the wage subsidy scheme: the subsidy saves 8.7% of service sector jobs and 5.5% of non-service sector jobs. The remaining columns of Table 4 show that the subsidy saves significantly more jobs among young workers than among old workers. For workers aged 20 to 29, 17.1% of jobs would have been lost absent the subsidy. In contrast, for workers aged 30 to 39, 5.5% of jobs are saved; for workers aged 40 to 49, 3.7% of jobs are saved; and for workers over the age of 50, 3.1% of jobs are saved. These results indicate that the wage subsidy scheme had large distributional consequences for employment outcomes during the COVID-19 pandemic.

5.3. Evolution of the Lockdown Under Alternative Policies

While the wage subsidy scheme prevents much job loss during a lockdown, there are many opportunity costs of committing such large fiscal expenditures to this purpose. We now study two alternative policies to a wage subsidy: a one-time lump-sum transfer, and a one-period increase in the transfer paid to unemployed workers. We consider these policies, in particular, because they were two of the most common policy responses to the pandemic in other countries. We model the effect of each policy alternative assuming the same sequence of lockdown shocks $\{Z_{1,t}, Z_{2,t}\}$. We also assume equal costs of fiscal transfers across policies. That is, each policy alternative

³⁴The Household labor Force Survey reports that there were approximately 2.7 million employees in New Zealand in 2019.

³⁵The Australian Bureau of Statistics reports that there were approximately 12,860,700 employees in Australia in 2020. See <https://www.abs.gov.au/statistics/labor/employment-and-unemployment/labor-force-australia/latest-release>.

has the same total cost of unemployment benefits, wage subsidies, and direct transfers.³⁶ This ensures that comparisons across policies are made on a dollar-for-dollar basis.

In order to understand how much is being redistributed from the wage subsidy to these alternative policies, we report the size of payments under each policy in Table 5. The table shows the size of the lump-sum transfer and unemployment benefit, respectively, relative to measures of average incomes for different groups of households. Each column in turn reports average incomes for: the unemployed, all households, households aged 20 to 29, households aged 30 to 39, households aged 40 to 49, and households aged 50 to 65. Under the lump-sum transfer every household receives a payment equivalent to 5.1% of average household income, but the payment is equivalent to 25% of unemployed incomes. Under the policy that raises unemployment benefits, the size of the new transfer is more than three times the size of the unemployed benefit in steady state, and is 63% of average household income. For young households, the unemployed benefit is equivalent to 93% of average incomes, while for older households the benefit is a little over half of average incomes.

Table 5: Payments Under Alternative Pandemic Policies Relative to Average Incomes

	Unemployed	All Ages	Age 20-29	Age 30-39	Age 40-49	Age 50-65
Transfers	0.253	0.051	0.075	0.052	0.046	0.045
Unemployment Benefits	3.115	0.632	0.927	0.637	0.565	0.555

Notes: The size of the lump-transfer and unemployment benefit of each policy is determined in general equilibrium subject to the constraint that the fiscal cost of redistribution under each policy alternative is equal to the cost under the baseline wage subsidy scheme. All values are computed as fractions of averages incomes for each group of households.

We now compare and study equilibrium outcomes under each of the policy alternatives over the course of the lockdown shocks. Figure B.20 in Appendix B.3 shows the evolution of employment, consumption, and GDP under each policy. While the wage subsidy policy directly supports the labor market thereby preventing unemployment, the lump-sum transfer and unemployment benefits policies do not so employment falls significantly more under these policies. The paths of aggregate consumption are very similar under the wage subsidy and the policy with higher unemployment benefits, but consumption falls by more under the lump-sum transfer policy. The evolution of GDP is essentially invariant to the form of policy intervention, since output is largely determined by the size of the productivity shocks that occur during the lockdown.

Figure 13 shows the response of household consumption by age group under the three policy alternatives. Figure B.21 in the Appendix shows consumption patterns for households in all age groups. For the youngest group of households, the smallest consumption decline occurs

³⁶Recall that government expenditures G_t adjust to ensure the fiscal budget constraint holds in each period, since government debt is held constant.

Figure 13: Consumption By Age Under Alternative Policies





Figure 14: The Unemployment-Income Curve Under Alternative Policies



Notes: Each figure plots the rate of unemployment benefit recipients against average incomes for each age group and across alternative policies. Points on each unemployment-income curve are labeled by worker age.

under the unemployment benefit policy. Recall that young workers earn around half as much as middle-aged workers (see Figure 7). This implies a smaller spread between the unemployment benefit and wages for younger workers than older workers. Thus, an increase in unemployment benefits represents a significant increase in expected income for the youngest households who are also the most likely to become unemployed. Indeed, recall from Table 5 that the higher unemployment benefit paid out is nearly 100% of average income for young workers. For older workers with higher earnings, a reduction in unemployment risk represents a larger increase in expected income than does an increase in unemployment benefits. Finally, consumption for the oldest workers hardly varies across policy interventions since their larger stock of savings insures them against income fluctuations. Note that the lump-sum transfer policy has a much smaller effect on the evolution of consumption as the payment is relatively small and much of it is directed towards older households who can comfortably smooth consumption without such payments.

Figure 14 provides an alternative illustration of the relative costs and benefits of the various policy responses for different groups of households. The figure plots points along an unemployment-income curve for workers of different ages. Points in the upper left show outcomes for young workers, and points to the lower right show outcomes for older workers. The solid blue lines represent unemployment-income outcomes under the baseline wage subsidy policy,

while the dashed red and green lines show unemployment-income outcome under the higher unemployment benefits and lump-sum transfer policies, respectively. First, consider a shift from the baseline wage subsidy policy to the higher unemployment benefits policy (left panel). For young workers, this change represents an inward shift of the unemployment-income curve: while young workers are much more likely to become unemployed under this policy, average incomes also rise significantly. For older workers, this change represents a direct upward shift of the unemployment-income curve: older workers face higher unemployment risk, with virtually no change in their expected income. Next, consider a shift from the baseline wage subsidy policy to the lump-sum transfers policy (right panel). For young workers, this change represents a shift of the unemployment-output curve towards the upper left: as unemployment for the young increases, the transfer is not large enough to offset the fall in average incomes. For older workers, the change represents a small inward shift of the curve: older workers experience both a small increase in unemployment and a small increase in average income inclusive of transfers.

In combination, Figures 14 and 13 suggest that there were large differences in the relative merits of the alternative policies implemented during the lockdown for households across the age distribution. Since it is unlikely that any one policy would be preferred by all groups of households, we next consider both the aggregate and cross-sectional welfare benefits of the various policy alternatives studied in this section.

5.4. Household Welfare During the Lockdown Under Alternative Policies

Finally, we study the welfare implications of the three policy interventions during the lockdown. We compute welfare gains for each policy intervention via the Consumption Equivalent Value (CEV) relative to the economy with no policy intervention. The CEV can be interpreted as the percentage gain in life-time consumption enjoyed under a particular policy intervention. We study these gains for different groups from the perspective of households in the period immediately prior to the lockdown, but who have no ability to change their behavior before the shocks occur. These households know the value of their own state variables on the eve of the lockdown, they know the effects of the shocks and each policy intervention, and they know the distribution over possible outcomes although they do not know their employment status at the onset of the lockdown. Table 6 reports the average CEV across all households and the fraction of households with a positive CEV under each of the three policy alternatives. We find that 65% of households are in favor of the wage subsidy policy with an average welfare gain of 0.21%, 100% of households are in favor of the lump-sum transfers policy with an average gain of 0.16%, and 41% of households are in favor of the policy with higher unemployment benefits with an average gain of 0.11%.

However, there is significant variation in welfare gains across different groups of households. Figure 15(a) shows that the youngest households enjoy the largest welfare gains overall, but also most prefer the policy that raises unemployment benefits during the lockdown. This is directly related to Figure 13, which shows much smaller consumption fluctuations among younger households under the unemployment benefits policy. This policy raises the insurance value of unemployment, which is particularly valuable for young households that face low wages to

Table 6: Welfare Gains from Policy Interventions During the Pandemic

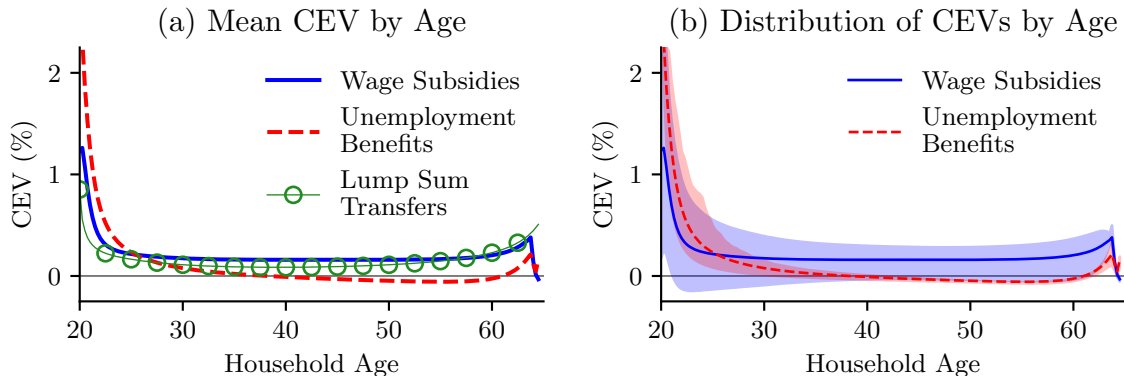
	Wage Subsidies	Lump Sum Transfers	Unemployment Benefits
Mean CEV (%)	0.211	0.158	0.109
Fraction CEV ≥ 0	0.649	1.000	0.412

Notes: Welfare gains and losses are computed using expected value functions from the perspective of households prior to the pandemic. Welfare statistics are computed for households drawn from the pre-pandemic steady state distribution.

begin with and who are more likely to become unemployed during the lockdown. The welfare benefits of all three policies fall as households age because the larger stock of assets held by older households enables them to better smooth consumption during the lockdown recession. The greater ability to smooth consumption reduces the value of policies that prevent unemployment or provide relatively small payments to these households.

Finally, 15(b) shows the 5th-to-95th percentiles of CEVs across age for the wage subsidy and unemployment benefits policy. A large proportion of young households under the unemployment benefit policy enjoy larger welfare gains than the average gain under the wage subsidy. Additionally, a small proportion of young households under the wage subsidy experience welfare losses relative to the world with no policy intervention. Nearly the entire distribution of households over the age of 40 receive larger CEVs under the wage subsidy policy than under the unemployment benefits policy. These results again highlight the dispersed welfare gains of different policy responses to the lockdown. Many more young households prefer policies with direct payments, while the majority of older households prefer policies that maintain their employment relationships as is the case under a wage subsidy scheme.

Figure 15: Welfare Gains from Policy Interventions During the Lockdown by Household Age



Notes: Panel (a) shows the mean CEV for households of each age group under each of the three policies. In Panel (b), the shaded areas show the 5th-to-95th CEV percentiles for households in each age group under the wage subsidy policy and the unemployment benefits policy. The solid and shaded lines show the mean CEVs, as in Panel (a).

6. Conclusion

In this paper, we study the macroeconomic and distributional consequences of the lockdowns implemented in response to the COVID-19 pandemic of 2020. In order to isolate the effects of lockdowns we consider the case of New Zealand, which imposed strict domestic lockdowns and international border closures but which endured very little exposure to the health effects of the virus itself. We build a life-cycle heterogeneous agents model with labor market frictions and multiple industries, calibrated to the New Zealand economy. We then study the dynamic response of the model to a series of industry-level negative productivity shocks, which mimic the effects of a lockdown on production and the demand for labor. The calibrated lockdown shocks generate the same declines in model employment as observed in New Zealand data over the course of 2020, suggesting an appropriate characterization of the effect of these shocks. Finally, we model the large wage subsidies paid to firms in New Zealand during the lockdown and study both the aggregate and cross-sectional effects of this unprecedented fiscal intervention.

We find that pandemic-induced lockdowns disproportionately affect both service sector workers and younger workers. Larger declines in service sector productivity result in larger increases in job separations in that industry, where young workers are more likely to be employed. Additionally, lower initial productivity among the youth in general leads employers to shed young workers more quickly than older workers. In the presence of a no-borrowing constraint and with young workers having had less time to accumulate savings, consumption is much more volatile for young households during the lockdown than it is for older households with the resources to self-insure against employment shocks.

Using counterfactual model outcomes, we show that the wage subsidies offered by the New Zealand government prevented a large number of job losses during the lockdown. Moreover, we find that this subsidy disproportionately benefited service sector and young workers; those who were most affected by lockdowns in the first place. The varied cross-sectional effects of the

subsidy are due to the conditional nature of its implementation: firms were only eligible for the wage subsidy if they experienced a sufficiently large decline in revenue during the initial lockdown period. In the model, firms in the service-sector and those that employ young workers experience a disproportionately large decline in revenues, which leaves them more likely to be eligible for the wage subsidy.

Finally, we study the welfare consequences of the wage subsidy policy in comparison to other possible fiscal interventions. We consider policies with the same fiscal outlays as the wage subsidy policy. The first policy pays a lump sum transfer to all households, while the second policy raises unemployment benefits. We find that young workers are better off under a policy that raises unemployment benefits. This is because the higher insurance value of unemployment for those who are most likely to become unemployed during a lockdown is more valuable than the increased probability of retaining a job under the wage subsidy policy.

While our model captures many rich features of the household life-cycle and exposure to the labor market, we eschew other model features that may also be important for understanding the impact of the COVID-19 Pandemic. First, our search model of the labor market does not allow for long-term labor market scarring. The effects of labor market scarring may increase the benefits of wage subsidy policies that prevent the onset of long unemployment spells. Second, for tractability we assume that households only have access to a liquid asset to facilitate saving. However, the large literature studying HANK models suggests that even wealthy households are sensitive to an increase in labor market risk when much of their wealth is held in illiquid assets (see Graves, 2020). Introducing a second illiquid asset would also likely amplify the welfare benefits of fiscal interventions during the lockdown, since even wealthy households would stand to benefit from these income-smoothing policies.

References

- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston, *Optimal Targeted Lockdowns in a Multi-Group SIR Model*, Working Paper, National Bureau of Economic Research, 2020.
- Auerbach, Alan J, Yuriy Gorodnichenko, and Daniel Murphy, *Inequality, fiscal policy and COVID-19 restrictions in a demand-determined economy*, tech. rep., National Bureau of Economic Research, 2020.
- Baqae, David and Emmanuel Farhi, *Nonlinear production networks with an application to the covid-19 crisis*, tech. rep., National Bureau of Economic Research, 2020.
- Bartik, Alexander W, Zoe B Cullen, Edward L Glaeser, Michael Luca, and Christopher T Stanton, *What jobs are being done at home during the COVID-19 crisis? Evidence from firm-level surveys*, tech. rep., National Bureau of Economic Research, 2020.
- Bayer, Christian, Benjamin Born, Ralph Luetticke, and Gernot Müller, *The Coronavirus Stimulus Package: How large is the transfer multiplier?*, CEPR Discussion Papers, CEPR Discussion Papers, 2020.
- Berger, David, Kyle Herkenhoff, Chengdai Huang, and Simon Mongey, “Testing and reopening in an SEIR model”, *Review of Economic Dynamics* (2020).
- Bilbiie, Florin O and Marc J Melitz, *Aggregate-Demand Amplification of Supply Disruptions: The Entry-Exit Multiplier*, Working Paper, National Bureau of Economic Research, 2020.
- Bishop, James, Iris Day, et al., *How many jobs did JobKeeper keep?*, tech. rep., Reserve Bank of Australia, 2020.
- Bloom, Nicholas, Philip Bunn, Paul Mizen, Pawel Smietanka, Gregory Thwaites, et al., *The impact of Covid-19 on productivity*, tech. rep., National Bureau of Economic Research, 2020.
- Carroll, Christopher D, Edmund Crawley, Jiri Slacalek, and Matthew N White, *Modeling the Consumption Response to the CARES Act*, Working Paper, National Bureau of Economic Research, 2020.
- Chéron, Arnaud, Jean-Olivier Hairault, and François Langot, “Life-cycle equilibrium unemployment”, *Journal of Labor Economics* 31 (2013), 843–882.
- Chetty, Raj, John N Friedman, Nathaniel Hendren, and Michael Stepner, “the Opportunity Insights Team (2020). “How did COVID-19 and stabilization policies affect spending and employment? A new real-time economic tracker based on private sector data.””, *NBER Working Paper* 27431 (2020).
- Christoffel, Kai, Keith Kuester, and Tobias Linzert, “The role of labor markets for euro area monetary policy”, *European Economic Review* 53 (2009), 908–936.
- De la Croix, David, Olivier Pierrard, and Henri R Sneessens, “Aging and pensions in general equilibrium: Labor market imperfections matter”, *Journal of Economic Dynamics and Control* 37 (2013), 104–124.
- Den Haan, Wouter J, Garey Ramey, and Joel Watson, “Job destruction and propagation of shocks”, *American economic review* 90 (2000), 482–498.

- Dingel, Jonathan I and Brent Neiman, *How many jobs can be done at home?*, tech. rep., National Bureau of Economic Research, 2020.
- Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt, *The macroeconomics of epidemics*, Working Paper, National Bureau of Economic Research, 2020.
- Fang, Lei, Jun Nie, and Zoe Xie, “Unemployment insurance during a pandemic”, *Federal Reserve Bank of Kansas City Working Paper* (2020).
- Farhi, Emmanuel and David Rezza Baqaee, *Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis*, tech. rep., National Bureau of Economic Research, 2020.
- Faria-e-Castro, Miguel, “Fiscal policy during a pandemic”, *FRB St. Louis Working Paper* (2020).
- Fujita, Shigeru and Garey Ramey, “Exogenous versus endogenous separation”, *American Economic Journal: Macroeconomics* 4 (2012), 68–93.
- Gornemann, Nils, Keith Kuester, and Makoto Nakajima, “Doves for the rich, hawks for the poor? Distributional consequences of monetary policy” (2016).
- Graves, Sebastian, “Does Unemployment Risk Affect Business Cycle Dynamics?”, *FRB International Finance Discussion Paper* (2020).
- Gregory, Victoria, Guido Menzio, and David Wiczer, “Pandemic recession: L-shaped or V-shaped?”, *Covid Economics* (2020), 88.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, *Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?*, tech. rep., National Bureau of Economic Research, 2020.
- Hale, Thomas, Sam Webster, Anna Petherick, Toby Phillips, and Beatriz Kira, *Oxford COVID-19 Government Response Tracker*, tech. rep., Blavatnik School of Government, Oxford University, 2020.
- Hobijn, Bart and Fernanda Nechio, “Sticker shocks: using VAT changes to estimate upper-level elasticities of substitution”, *Journal of the European Economic Association* 17 (2019), 799–833.
- Kaplan, Greg, Benjamin Moll, and Giovanni L Violante, *The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the US*, tech. rep., National Bureau of Economic Research, 2020.
- Krueger, Dirk, Harald Uhlig, and Taojun Xie, *Macroeconomic Dynamics and Reallocation in an Epidemic: Evaluating the “Swedish Solution”*, Working Paper, National Bureau of Economic Research, 2020.
- Lugauer, Steven, “Demographic change and the great moderation in an overlapping generations model with matching frictions”, *Available at SSRN 2334344* (2013).
- Ravn, Morten O and Vincent Sterk, “Job uncertainty and deep recessions”, *Journal of Monetary Economics* 90 (2017), 125–141.
- Razzak, Weshah A., “On the Dynamic of Search, Matching and Productivity in New Zealand and Australia”, *International Journal of Applied Economics* 6 (2009), 90–118.

Zealand, Statistics New, *Tourism Satellite Account: Year Ended March 2020*, tech. rep., Statistics New Zealand, 2020, URL: <https://www.stats.govt.nz/information-releases/tourism-satellite-account-year-ended-march-2020>.

Appendix

A. Data

- The Blavatnik School of Government at Oxford University compiles government responses to the pandemic for virtually every country in the world. Data from their Government Response Tracker is available for download at <https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker>.
- Data from New Zealand's Wage subsidy scheme was updated weekly by the Ministry of Social Development throughout 2020. Data is available at <https://www.msd.govt.nz/about-msd-and-our-work/publications-resources/statistics/covid-19/index.html>.
- Statistics for industry composition, job separation rates, and earnings by age are drawn from the 2018 Linked Employer-Employee Database (LEED). This data is available from Statistics New Zealand at <http://nzdotstat.stats.govt.nz/wbos/index.aspx>.
- Unemployment benefit receipts by age are available from the Ministry of Social Development at <https://www.msd.govt.nz/about-msd-and-our-work/publications-resources/statistics/benefit/>.
- Financial statistics for a sample of New Zealand businesses is available in the Annual Enterprise Survey from Statistics New Zealand at <https://www.stats.govt.nz/information-releases/annual-enterprise-survey-2019-financial-year-provisional>
- Data on employment by ANZSIC06 industry code is available in the Household Labour Force Survey from Statistics New Zealand at <http://infoshare.stats.govt.nz>.
- Data on monthly employment are available in the Monthly Employment Indicators from Statistics New Zealand at <http://infoshare.stats.govt.nz>
- Data on household network are available in the Household Economic Survey from Statistics New Zealand at <http://nzdotstat.stats.govt.nz/wbos/index.aspx>.
- Data on GDP and interest rates are available from the Reserve Bank of New Zealand at <https://www.rbnz.govt.nz/statistics>

Table A.7: Employment Shares by Industry and Sector

Service Sector		Non-Service Sector	
Industry	Share	Industry	Share
Arts, Recreation and Other Services	0.060	Agriculture, Forestry and Fishing	0.049
Financial and Insurance Services	0.032	Construction	0.087
Information Media and Telecommunications	0.016	Education and Training	0.086
Rental, Hiring and Real Estate Services	0.020	Electricity, Gas, Water and Waste Services	0.009
Retail Trade and Accommodation	0.147	Health Care and Social Assistance	0.109
Transport, Postal and Warehousing	0.046	Manufacturing	0.092
Wholesale Trade	0.040	Mining	0.002
		Not Specified	0.016
		Professional, Scientific, Technical, Administrative and Support Services	0.128
		Public Administration and Safety	0.062

Notes: Employment shares computed from both paid and self-employed workers across ANZSIC06 industries for 2019. Sectoral allocation chosen by authors for the purposes of the current study.

Sources: Author's calculations using data from the Household Labour Force Survey.

B. Model

B.1. Additional Model Details

Profits are earned by firms involved in the production of goods as well as entrepreneurs in the labour market. The owners of all of these firms must also pay vacancy hiring costs. Total real firm profits are thus defined as:

$$\Pi_t = \sum_{i \in \{1,2\}} \left\{ \left(\frac{p_{i,t}}{p_t} Y_{i,t} - h_{i,t} \sum_{k=1}^K N_{i,k,t} \bar{x}_{i,k,t} \right) + \left(h_{i,t} \sum_{k=1}^K N_{i,k,t} \bar{x}_{i,k,t} - \sum_{k=1}^K N_{i,k,t} w_{i,k,t} \right) - \kappa_i v_{i,t} \right\}$$

As noted in the text, firm profits are consumed within the period in which they are earned.

We assume that expenditures due to profits Π_t , government spending G_t , and net exports NX_t are allocated across services and other consumption goods in the same way as household expenditure in the CES aggregator in Equation (3). Final expenditures are given by

$$\begin{aligned} \Pi_t &= \left[\chi^{\frac{1}{\varepsilon_c}} \Pi_{1,t}^{\frac{\varepsilon_c-1}{\varepsilon_c}} + (1-\chi)^{\frac{1}{\varepsilon_c}} \Pi_{2,t}^{\frac{\varepsilon_c-1}{\varepsilon_c}} \right]^{\frac{\varepsilon_c}{\varepsilon_c-1}} \\ G_t &= \left[\chi^{\frac{1}{\varepsilon_c}} G_{1,t}^{\frac{\varepsilon_c-1}{\varepsilon_c}} + (1-\chi)^{\frac{1}{\varepsilon_c}} G_{2,t}^{\frac{\varepsilon_c-1}{\varepsilon_c}} \right]^{\frac{\varepsilon_c}{\varepsilon_c-1}} \\ NX_t &= \left[\chi^{\frac{1}{\varepsilon_c}} NX_{1,t}^{\frac{\varepsilon_c-1}{\varepsilon_c}} + (1-\chi)^{\frac{1}{\varepsilon_c}} NX_{2,t}^{\frac{\varepsilon_c-1}{\varepsilon_c}} \right]^{\frac{\varepsilon_c}{\varepsilon_c-1}} \end{aligned}$$

The solutions to the expenditure minimization problems imply the following demand functions

$$\begin{aligned} \Pi_{1,t} &= \chi \left(\frac{p_{1,t}}{p_t} \right)^{-\varepsilon_c} \Pi_t, & \Pi_{2,t} &= (1-\chi) \left(\frac{p_{2,t}}{p_t} \right)^{-\varepsilon_c} \Pi_t \\ G_{1,t} &= \chi \left(\frac{p_{1,t}}{p_t} \right)^{-\varepsilon_c} G_t, & G_{2,t} &= (1-\chi) \left(\frac{p_{2,t}}{p_t} \right)^{-\varepsilon_c} G_t \\ NX_{1,t} &= \chi \left(\frac{p_{1,t}}{p_t} \right)^{-\varepsilon_c} NX_t, & NX_{2,t} &= (1-\chi) \left(\frac{p_{2,t}}{p_t} \right)^{-\varepsilon_c} NX_t, \end{aligned}$$

where $\frac{p_{i,t}}{p_t}$ is the relative price in industry i .

B.2. Formal Definition of Recursive Competitive Equilibrium

Let $s \equiv (k, b, e, i)$ define the household state vector, let $\lambda_k(s)$ be the measure of households at age k with $\sum_{k=1}^K \int \lambda_k(s) ds = 1$, and let $\mathcal{Z} \equiv (Z_1, Z_2, \tau_w, \Theta)$ denote a vector of exogenous lockdown and fiscal policy shocks. To ease notation, we write the aggregate state vector as $\Omega = (\mathcal{Z}, \lambda)$.

A recursive competitive equilibrium of the model consists of household decision rules $\{c(s; \Omega), b(s; \Omega)\}$, labor demands of production firms $\{\tilde{N}_i(\Omega)\}_{i=1,2}$, labor market entrepreneur shutdown decisions characterized by the thresholds $\{\underline{x}_{i,k}(\Omega)\}_{i=1,2,k=1,\dots,K}$, labor demands of entrepreneurs $\{N_{i,k}(\Omega)\}_{i=1,2,k=1,\dots,K}$, the vacancy positing decisions of labor market entrepreneurs $\{v_i(\Omega)\}_{i=1,2}$, the mass of workers searching for work $\{u_{i,k}(\Omega)\}_{i=1,2,k=1,\dots,K}$, production wages $\{h_i(\Omega)\}_{i=1,2}$, worker wages $\{w_{i,k}(\Omega)\}_{i=1,2,k=1,\dots,K}$, relative prices $\{\frac{p_i}{p}(\Omega)\}_{i=1,2}$, job finding and job filling rates $\{q_i^u(\Omega), q_i^v(\Omega)\}_{i=1,2}$, job separation rates $\{\rho_{i,k}(\Omega)\}_{i=1,2,k=1,\dots,K}$,

fiscal policies $\{G(\Omega), UB(\Omega), WS(\Omega), TR(\Omega), T(\Omega), B^g(\Omega)\}_{t=0}^{\infty}$, net exports $NX(\Omega)$, and transition functions $Q_{k,k+1}(\Omega)$ that take the distribution of households over states s at age k to future states s' at age $k+1$, such that:

1. Households optimize by solving the problems (1) and (2), with associated decision rules $\{c(s; \Omega), b(s; \Omega)\}$
2. Firms in the production sector maximize profits by solving (6), with associated labor demands $\{\tilde{N}_i(\Omega)\}_{i=1,2}$
3. Labor market entrepreneurs make shutdown decisions to maximize profits, by solving the problem (7), which is associated with shutdown thresholds $\{x_{i,k}(\Omega)\}_{i=1,2,k=1,\dots,K}$ and labor demands $\{N_{i,k}(\Omega)\}_{i=1,2,k=1,\dots,K}$
4. Given shutdown decisions, job separation rates are determined by (9)
5. Potential labor market entrepreneurs post vacancies $\{v_i(\Omega)\}_{i=1,2}$ subject to the entry condition (10)
6. The number of job searchers at each age and in each industry satisfy the laws of motion (14)
7. Given wages $\{h_i(\Omega)\}_{i=1,2}$, markets for production workers clear:

$$\tilde{N}_i(\Omega) = \sum_{k=1}^K N_{i,k}(\Omega) \bar{x}_{i,k}(\Omega), \quad i = 1, 2$$

8. Given relative prices $\{\frac{p_i}{p}(\Omega)\}_{i=1,2}$, the individual goods markets clear:

$$Y_i(\Omega) = \chi \left(\frac{p_1}{p}(\Omega) \right)^{-\varepsilon_c} (C(\Omega) + \Pi(\Omega) + G(\Omega) + NX(\Omega))$$

where $C(\Omega)$ is aggregate household consumption.

9. The aggregate resource constraint holds

$$\frac{p_2}{p}(\Omega) Y_2(\Omega) + \frac{p_2}{p}(\Omega) Y_2(\Omega) = C(\Omega) + \Pi(\Omega) + G(\Omega) + \kappa_1 v_1(\Omega) + \kappa_2 v_2(\Omega) + NX(\Omega)$$

10. Government debt $B^g(\Omega)$ ensures that the government budget constraint in Equation (16) holds for all Ω
11. The transition functions $Q_{k,k+1}(\Omega)$ are induced by the exogenous shocks and policies, the processes for idiosyncratic risk, and all of the decision rules. As a result, these transition functions are consistent with individual behavior.

B.3. Additional Model Figures

Figure B.16: Industry Composition




Figure B.17: Job Separation Rates by Age

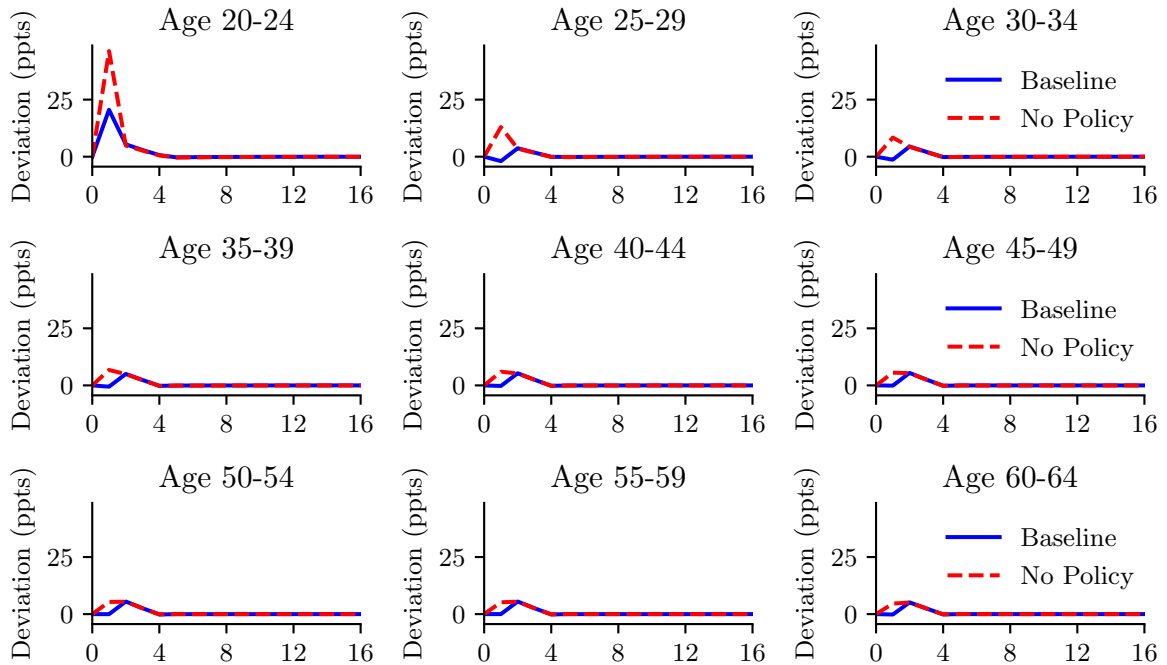


Figure B.18: Unemployed Searchers by Age

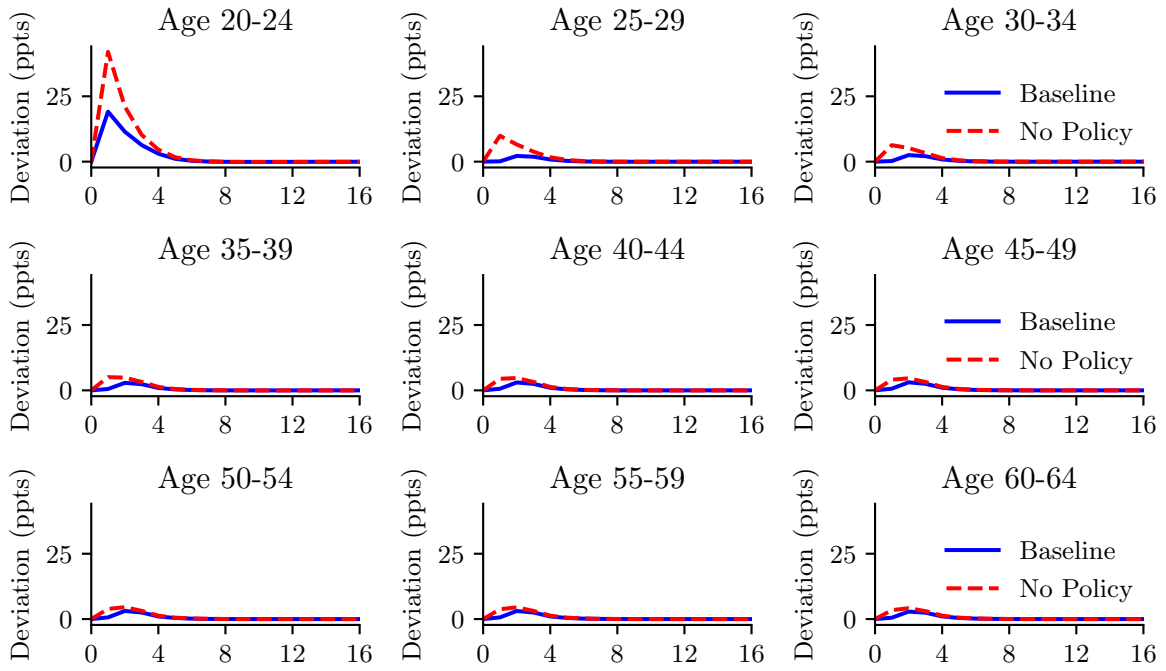


Figure B.19: Household Consumption by Age

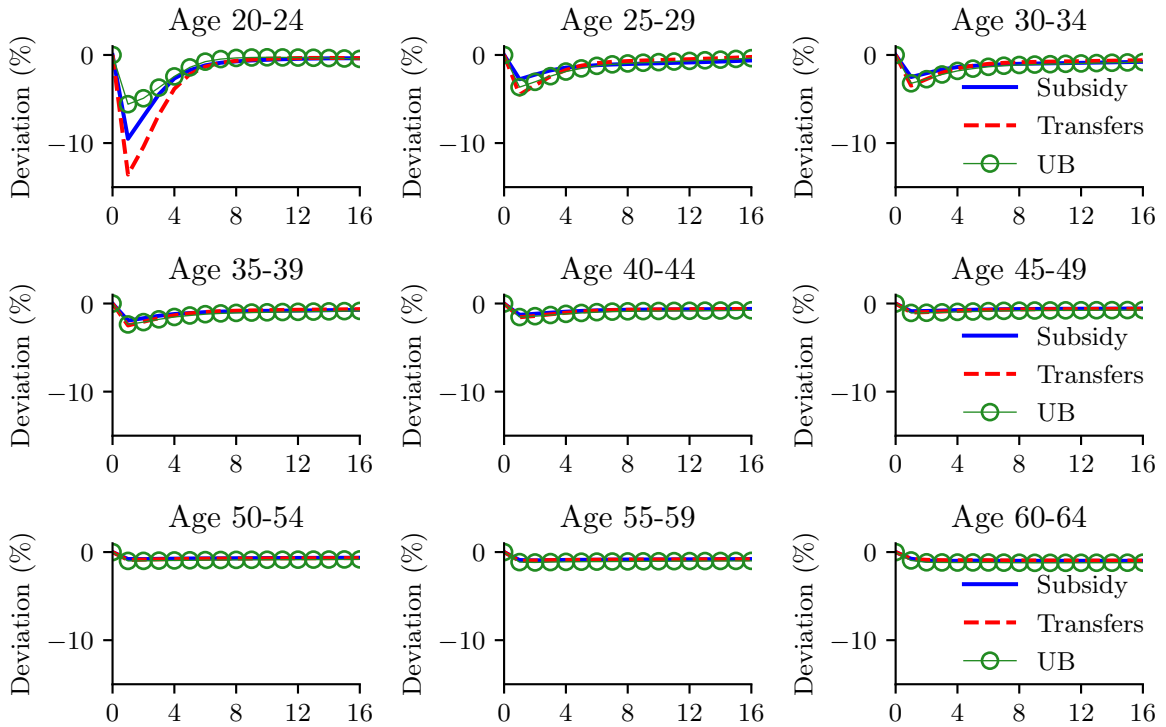


Figure B.20: Aggregate Variables Under Alternative Policies




Figure B.21: Consumption by Age Under Alternative Policies

