

Pandemics through the lens of occupations

Anand Chopra

Indian Institute of Technology Kanpur

Michael B. Devereux

University of British Columbia, NBER

Amartya Lahiri

University of British Columbia

Abstract. We outline a macro-pandemic model where individuals can select into working from home or in the market. Market work increases the risk of infection. Occupations differ in the ease of substitution between market and home work and in the risk of infection. We examine the evolution of a pandemic in the model as well as its macroeconomic and distributional consequences. The model is calibrated to British Columbian data to examine the implications of shutting down different industries by linking industries to occupations. We find that endogenous choice to self-isolate is key: it reduces the peak weekly infection rate by two percentage points but reduces the trough consumption level by four percentage points, even without policy-mandated lockdowns. The model also produces widening consumption inequality, a fact that has characterized COVID-19.

Résumé. *Les pandémies envisagées à travers le prisme des professions.* Nous décrivons un modèle de macro-pandémie où les individus peuvent choisir de travailler à domicile ou sur le marché. Le travail sur le marché augmente le risque d'infection. Les professions diffèrent par la facilité de substitution entre le marché du travail et le travail à domicile, et par le risque d'infection. Nous examinons l'évolution d'une pandémie dans le modèle ainsi que ses conséquences macroéconomiques et distributionnelles. Le modèle est calibré sur les données de la Colombie-Britannique pour examiner les implications de la fermeture de différentes industries en reliant les industries aux professions. Nous constatons que le choix endogène de s'auto-isoler est essentiel : il réduit le taux d'infection hebdomadaire maximal de deux points de pourcentage, mais réduit le niveau de consommation minimal de quatre points de pourcentage, même sans les verrouillages imposés par la politique. Le modèle produit également une inégalité de consommation croissante, un fait qui a caractérisé le COVID-19.

JEL classification: E00, I00

Corresponding author: Michael Devereux, Michael.Devereux@ubc.ca

We are grateful for the comments from the two anonymous referees, Chris Telmer, Giovanni Gallipoli and seminar participants at the Bank of Canada, the Government of British Columbia and the Vancouver School of Economics. Devereux and Lahiri thank SSHRC for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Canadian Journal of Economics / *Revue canadienne d'économie* 2022 55(S1)
February 2022. Printed in Canada / *Février 2022. Imprimé au Canada*

ISSN: 0008-4085 / 22 / pp. 540–580 / DOI: 10.1111/caje.12547

© Canadian Economics Association

1. Introduction

AS THE CORONAVIRUS and the disease it spawns, COVID-19, has spread across the globe, it has unleashed fearsome challenges for epidemiologists, public health professionals, economists and public policy officials to both understand and devise methods to arrest its spread and attenuate its effects. The particular challenge in managing the crisis is the tradeoff between the public health and economic costs of the disease. In this paper, we formalize a macroeconomic model and blend it with the standard Susceptible–Infectious–Recovered (SIR) model for pandemics. Relative to standard macroeconomic models, our structure has two main innovations. First, we conceptualize aggregate output as being the outcome of the labour effort of different occupations. We believe this focus is important because a significant source of contagion of infections is the interaction of people while working. However, labour is supplied through many different occupations with different degrees of required social proximity and interaction. Hence, the economic costs of pandemics are likely to vary significantly with different occupation mixes in society.

A second innovation of the model is that it allows individuals to choose whether to work from the market or work from home. The ease with which individuals can substitute into working from home varies across occupations. Our model and empirical implementation incorporates this variation. Allowing individuals to choose the location of their work allows them to self-insure against infection risk. This appears to be a key decision facing private agents and policy-makers alike as they decide on the methods of dealing with the crisis.

Our model is calibrated to British Columbian occupation and labour force data from the 2016 Canadian census. We show that the endogeneity of market participation is quantitatively important with the peak weekly infection rate lower by two percentage points. Risk aversion is the key driver for this result. Moreover, the effect of lockdowns and their relaxation is also significantly impacted by the endogeneity of market participation. Lockdowns that are relaxed while infections are still growing result in slower recoveries as individuals endogenously choose to stay away from market work, even without coercion.

We conduct a number of robustness checks on our results. We uncover three important insights from these exercises. First, our baseline model assumes a common infection risk from market work for all occupations. Relaxing this assumption and allowing for occupation-specific infection risk from market work does not change the baseline results by much. Second, the degree of risk aversion is very important. With risk neutrality, the difference between endogenous and exogenous choices of work location almost disappears. Third, the productivity of individuals contingent on becoming infected is important. The lower the productivity of the infected, the lower is the peak of the overall infection rate.

Since the onset of the pandemic, a vast literature (some of which is discussed below) has grown, focusing on many aspects of human interaction that gives rise to infection risk. The model we develop below allows for multiple

sources of risk, but as argued above, our main focus is on the risk associated with market work and the consequent choices that individuals make related to this risk. That is not to argue that other activities, such as shopping, leisure, sports or other venues for social interaction are not equally as important. However, they are not the main margins of choice within our model.

The structure that we formalize also has distributional implications. Because individual skills are occupation-specific and occupations have different degrees of substitutability between market and home work, a pandemic induces heterogeneity in infection rates and consumption across occupations. Occupations that are easy to perform from home without significant loss of income allow individuals to insulate themselves from infection risk without sacrificing consumption. Contrarily, occupations that are harder to perform from home see higher infections and lower consumption.

The model matches the data fact that the higher the income group, the greater the initial drop in consumption and the slower the consumption recovery. The model also predicts that the trough of consumption of the lowest income groups comes much later than for relatively richer groups. Hence, we may not have seen the worst yet in terms of the distributional consequences of COVID-19.

We believe our results point to the importance of public health initiatives that prioritize building societal confidence in the safety of market work. One such measure is widespread and frequent testing of the general population. Absent such measures, economic recovery from COVID-19 may be slow and prolonged and continue until vaccines become more widely available to the working age population.

Since the beginning of the global COVID-19 pandemic, there has been an explosion of research studies analyzing the links between macroeconomic and epidemiological models. The work has been so prolific that already there are a already number of reviews of the recent literature (Brodeur et al. 2020 and Hur and Jenuwine 2020). We do not attempt here to provide a complete summary of new papers in this area. Instead, we describe the connection of our paper with other work that is closest in spirit or implementation to ours.

Our paper is closest on the economic side to Eichenbaum et al. (2021) and takes some of the calibration of disease dynamics from Atkeson (2020). Recent work in Chetty et al. (2020) documents the evolution of inequality along various margins in the aftermath of COVID-19. We use the consumption inequality facts presented in Chetty et al. (2020) as a benchmark to compare the distributional dynamics of our model.¹

1 Carroll et al. (2020) and Coibion et al. (2020) provide estimates of household consumption responses to the US Cares Act. Matching the pass-through of stimulus benefits into consumption in a model requires a richer framework (see Bayer et al. 2020 and Kaplan et al. 2020) from which we abstract in this paper.

Other papers closely related to ours are Dingel and Neiman (2020), who construct a measure of ease of home work by occupation from O*NET data. We use a similar measure to benchmark our simulations. In addition, a number of empirical studies have focused on the decision and consequences of working from home during the pandemic. Bartik et al. (2020) find that employers in industries where more people were working from home pre-COVID had lower productivity losses during COVID. Adams-Prassl et al. (2020) show using survey data for US, UK and Germany that workers who have a lower ability to work from home are more likely to lose their jobs and experience a fall in earnings. Beland et al. (2020) shows similar results for unemployment and wages in Canada until April 2020 using the Labour Force Survey (LFS) data. Consistent with our mechanism, a survey by Brynjolfsson et al. (2020) indicates a relationship between increasing COVID-19 cases and a shift in the number of respondents working more from home across US states.

An important theme in our study is the distinction between changes in activity due to lockdowns versus voluntary behaviour. Using high-frequency cellphone data, Brzezinski et al. (2020), Chen et al. (2020), Chernozhukov et al. (2021), Goolsbee and Syverson (2021) and Gupta et al. (2020) find a modest to large role of voluntary behaviour in the reduction in mobility relative to stay-at-home and schools and business closure orders in the US. Maloney and Taskin (2020) find a larger role for voluntary behaviour than government orders in advanced economies in reducing mobility. This finding is supported by Caselli et al. (2020) for both mobility and job postings.²

Besides Eichenbaum et al. (2021), other theoretical models that have some relation to ours, are Krueger et al. (2020), Jones et al. (2020), Aum et al. (2021) and Bodenstein et al. (2020). Work that examines optimal lockdown policies in richly parameterized SIR models can be found in Acemoglu et al. (2020), Alvarez et al. (2020) and Jones et al. (2020).

There is also a burgeoning literature examining the heterogeneous effects of the COVID-19 shock. For example, Baqaee and Farhi (2020) and Woodford (2020) study fiscal and monetary stabilization policy with multiple sectors as the COVID-19 pandemic shock can affect the various sectors differently. In contrast to these papers, we study how *ex ante* heterogeneous individuals across various occupations respond to the same shock.

The next section presents the model, while section 3 describes the dynamics of a pandemic and some theoretical results. Section 4 describes the calibration of the model and the quantitative results. Section 5 explores the robustness of the quantitative results with respect to key parameters, while section 6 compares the distributional predictions of the model with those in the data. The last section concludes.

2 Ding et al. (2020) highlight the importance of community engagement and individual willingness to follow social and public objectives as important factors in an individual's response to social distancing policies.

2. Model

We study a closed economy with a continuum of occupations with measure one. The economy consumes a final good that is produced by multiple sectors. Sectoral goods are produced by combining various occupations. Within each occupation i , there is a population $L_t(i)$ of individuals at date t . Hence, the total population is $L_t = \int_0^1 L_t(i) di$.

Infinitely lived individuals in the economy have one unit of labour time that they supply inelastically to work every period. Individuals have occupation-specific skills. These skills can be used to supply work effort from home or the market. A representative individual within occupation i receives wage $w^h(i)$ from working at home and wage $w^m(i)$ from working in the market. Labour markets are competitive so that wages reflect the marginal products of labour from home and market work in each occupation. On the other hand, the value of work from home may be extremely small or close to zero for some occupations. Our model allows for arbitrarily small values of $w^h(i)$ for some occupations.³

We assume that occupations are arranged in order of the wage from working in the market relative to working from home. Thus $\frac{w^m(j)}{w^h(j)} > \frac{w^m(i)}{w^h(i)}$ for $j > i$. An individual within each occupation will work in the marketplace if the value of market work $V^m(i)$ exceeds the value of home work $V^h(i)$. So if occupation i decides to work in the marketplace, then occupation $j > i$ will also work in the marketplace. We denote θ as the share of occupations working in the market rather than home. Thus, defining \bar{i} as the marginal occupation that is indifferent between home and market work, $\theta = 1 - \bar{i}$.

Individuals maximize the present discounted value of lifetime utility. Thus, an individual j in occupation i maximizes

$$V^j(i, t) = \sum_t \beta \ln c_t^j(i).$$

Individuals are hand-to-mouth consumers so that every period they consume their current period earnings completely: $c_t^j(i) = w_t(i)$.

The epidemiological structure of the model is very similar to recent papers noted above. At every date, within each occupation, there are potentially three types of agents in the economy: Susceptible (S), Infected (I) or Recovered (R) with the total population being $L_t = S_t + I_t + R_t$. The measure of agents in each of these groups changes over time as a function of the cumulated decisions made by individuals regarding home versus market work as well as through infections occurring through consumption activities as well as random infections, as described more fully below. In the following, we shall

3 In this context, Bloom et al. (2015) show that switching to working from home increased the productivity of call centre employees.

use the notation $K_t(i)$ to denote the measure of people of type $K = S, I, R$ in occupation i at date t .

For an individual, working in the market entails the risk of becoming infected by a virus. This risk is a function of the interaction between infected people and susceptible individuals working in the market. Once infected, individuals either continue to remain infected or recover with probability π_R or die with probability π_D .

2.1. Production

The production structure of the economy consists of a final good that is produced by combining a set of J -intermediate goods corresponding to different sectors. The final good is the numeraire for the model. The final goods technology is

$$Y_t = \left[\sum_{j=1}^J \eta^{j(1-\rho)} (Y_t^j)^\rho \right]^{\frac{1}{\rho}}, \tag{1}$$

where $\rho \leq 1$ controls the elasticity of substitution between sectoral goods and $\sum_j \eta^j = 1$ is assumed. Y_t^j denotes output of sector j used by the final goods sector at date t . Using P_t^j to denote the price of sector j goods, we can derive the demand function for sector j goods as

$$Y_t^j = Y_t \eta^j P_t^{j \frac{1}{\rho-1}}. \tag{2}$$

Equations (1) and (2) can be used to derive the price of the basket of final goods, which is unity, as it is the numeraire:

$$\sum_{j=1}^J \eta^j P_t^{j \frac{\rho}{\rho-1}} = 1 \tag{3}$$

In the special case where sectoral goods are perfect substitutes in producing the final good, or $\rho = 1$, all sectoral prices must be equal and therefore $P_t^j = 1$.

Output for each sector is produced by combining the effective labour supply by the continuum of occupations. This labour supply is adjusted for occupation-specific productivity, as we shall describe below. The sector j technology is given by

$$Y_t^j = \left[\int_{i=0}^1 \delta^j(i)^{(1-\rho^j)} (L_t^j(i))^{\rho^j} di \right]^{\frac{1}{\rho^j}}, \quad \rho^j \leq 1. \tag{4}$$

Here, labour usage $L_t^j(i)$ is defined over a continuum of occupations i , and $\rho^j \leq 1$ controls the elasticity of substitution between occupations in sector j . $\rho^j = 1$ is the case of perfect substitutability, while $\rho^j = -\infty$ is the Leontief

case of no substitutability between occupations. $\delta^j(i)^{(1-\rho^j)}$ denotes the input weight of occupation i in sector j , and we assume that $\int_0^1 \delta^j(i) di = 1$. Because both ρ^j and $\delta^j(i)^{(1-\rho^j)}$ are indexed by j , our specification allows the occupational intensity of labour to vary across sectors.

The demand by sector j for labour of occupation i is

$$L_t^j(i) = Y_t^j \delta^j(i) \left(\frac{P_t(i)}{P_t^j} \right)^{\frac{1}{\rho^j-1}}, \tag{5}$$

where $P_t(i)$ represents the price of occupation i good. It is straightforward to show using equations (4) and (5) that the price of sector j output is

$$P_t^j = \left[\int_{i=0}^1 \delta^j(i) P_t(i)^{\frac{\rho^j}{\rho^j-1}} di \right]^{\frac{\rho^j-1}{\rho^j}}.$$

Labour of occupation i is produced by using the productivity adjusted linear aggregator:

$$L_t(i) = A_i^m \{ S_t^m(i) + R_t^m(i) + \gamma_I^m I_t^m(i) \}, \tag{6}$$

$$+ A_i^h \{ S_t^h(i) + R_t^h(i) + \gamma_I^h I_t^h(i) \},$$

where $S^k(i)$, $R^k(i)$ and $I^k(i)$ denote the measures of susceptible, recovered and infected individuals, respectively, in occupation i who are working in $k = h, m$. Note that h denotes home and m denotes market. The productivity of occupation i working from home or market is denoted by A_i^k . We assume throughout that $\gamma_I^k < 1$ for $k = h, m$, so that the productivity of infected individuals is strictly less than non-infected workers in every occupation. The productivity terms are not indexed by the time subscript t indicating the assumption that productivity is constant.⁴

In addition, we allow for a tax rate on market work equal to τ_m . This tax may be rebated to each affected occupation as a lump sum rebate. We think of this parameter as controlling a policy of health-driven shutdowns in market activity.

Competitive labour markets then imply that wage rates for market and home work are given by

$$w^m(i, t) = (1 - \tau^m) P_t(i) A_i^m, \tag{7}$$

$$w^h(i, t) = P_t(i) A_i^h, \tag{8}$$

$$w_I^m(i, t) = (1 - \tau^m) P_t(i) \gamma_I^m A_i^m, \tag{9}$$

$$w_I^h(i, t) = P_t(i) \gamma_I^h A_i^h. \tag{10}$$

4 It is straightforward to allow for exogenous productivity growth in the model.

These conditions determine the after-tax wages in each occupation i for all groups of individuals. Wages in every occupation are constant over time here because both occupational productivity and prices are constant, and assuming also that taxes are constant.

2.2. The pre-infection economy

Take an initial situation where there is no risk of being infected in the marketplace. We also assume that there is no tax on market labour during this period. Then

$$V^h(i, t) = \ln w^h(i) + \beta V^h(i, t + 1) = \frac{\ln w^h(i)}{1 - \beta}$$

$$V^m(i, t) = \ln w^m(i) + \beta V^m(i, t + 1) = \frac{\ln w^m(i)}{1 - \beta},$$

where the second equality in both equations above follow from the fact that, because there are no shocks in this economy and there are no sticky variables that evolve gradually over time, wages have to be constant over time. With all other exogenous variables being constant, in each occupation the choice between home and market work depends only on the wage comparison $w^m(i) - w^h(i)$. Because this is constant, the choice between home and market work remains unchanged over time in every occupation i . Thus, an occupation that is performed from home today yields a lifetime utility of $\frac{\ln w^k(i)}{1 - \beta}$ for $k = h, m$.

There is a cutoff occupation \bar{i} that is determined by the condition $V^h(\bar{i}) = V^m(\bar{i})$. From the expressions for the value functions derived above, the cutoff condition reduces to $w^h(\bar{i}) = w^m(\bar{i})$. All occupations for which $i > \bar{i}$ will be performed through market work, while those below the threshold occupation will work from home. The threshold then determines the total number of individuals working in the marketplace as $L^m = L - \int_0^{\bar{i}} L(i) di$. The cutoff and the distribution of market versus homework remains unchanged over time.

Occupational output produced by occupation i is given by

$$L(i) = \begin{cases} L_0(i)A_i^h, & \text{if } i < \bar{i} \\ L_0(i)A_i^m, & \text{if } i \geq \bar{i}, \end{cases}$$

where $L_0(i)$ is the measure of individuals in occupation i , which is constant over time absent any pandemic dynamics.

Sectoral output, which is produced by combining occupations, is given by

$$Y^j = \left[\int_{i=0}^{\bar{i}} \delta^j(i)^{1-\rho^j} L^{j,h}(i)^{\rho^j} di + \int_{i=\bar{i}}^1 \delta^j(i)^{1-\rho^j} L^{j,m}(i)^{\rho^j} di \right]^{\frac{1}{\rho^j}}, \quad j = 1, \dots, J.$$

Note that absent a pandemic, all sectoral outputs remain constant over time because occupational goods are constant over time.

Given sectoral outputs, output of the final good in the uninfected economy is given by equation (1). This is a stationary economy with the occupations choices, population distribution across occupations, sectoral output, the final good production as well as all prices remaining constant over time.

2.3. A special case

We shall focus on a special case of the model in the analysis below. Specifically, consider the case where the following condition holds:

CONDITION 1.

$$\rho = \rho^j = 1 \text{ for all } j.$$

Under condition 1, equations (2) and (5) imply that $P_t(i) = P_t^j = 1$. Hence, all prices are constant and equal to one. This directly implies that without taxes, we must have $w_t^k(i) = w^k(i) = A_i^k$ and $w_t^l(i) = \gamma_l^k w^k(i)$. Thus, all occupational wages are also constant.

It is straightforward to verify that, under condition 1, we must have

$$Y_t = \int_{i=0}^1 L_t(i) di.$$

3. A pandemic

We shall now study the outbreak of a pandemic in the context of the special case of the model outlined in section 2.3.

Suppose that starting at some date $t = 0$, a fraction ε_0 of the population becomes infected with the virus. Hence, the measure $I_0 = \varepsilon_0 L_0$ of individuals becomes infected. This information is public knowledge. Infection can spread in multiple ways. Our focus is on infection through the workplace, but we also allow for infection through random association (e.g. through leisure activities) and through consumption activities, such as retail. We assume that these second types of infection cannot be avoided, but individuals can reduce risk by choosing to work at home.⁵

In the following, we shall denote the lifetime utility value of susceptible, infected and recovered individuals in occupation i at date t with $V^S(i, t)$, $V^I(i, t)$ and $V^R(i, t)$, respectively. It will prove useful to first describe the choices of

5 Endogenous choice of activities with respect to leisure and shopping activity are of course important, but we abstract from them here as our focus is essentially on the implications of the pandemic for the labour market and employment differences across occupations.

recovered agents, then those of infected agents and finally the choices of susceptible agents.

3.1. Recovered individuals

Recovered individuals are those who contracted the disease and survived. They are as productive in the labour market as susceptible individuals. The choices of a recovered agent i at date t working in location $j = h, m$ are determined by

$$\begin{aligned} V^{Rh}(i, t) &= \ln w^h(i) + \beta \text{Max}\{V^{Rh}(i, t + 1), V^{Rm}(i, t + 1)\}, \\ V^{Rm}(i, t) &= \ln w^m(i) + \beta \text{Max}\{V^{Rh}(i, t + 1), V^{Rm}(i, t + 1)\}. \end{aligned}$$

Because the continuation values from both home and market work (the second terms on the right-hand side of the two equations above) are the same, the threshold occupation for recovered agents is given by

$$w^h(\bar{i}^R) = w^m(\bar{i}^R).$$

This defines the measure of occupations in which recovered agents work in the market as $\theta^R = 1 - \bar{i}^R$. Crucially, the constancy of occupational wages implies that this measure of recovered agents working in the market remains constant over time.

The preceding implies that: (a) a recovered agent in occupation i chooses to always work at home or always in the market, (b) V^{Rk} , $k = h, m$ remains constant over time and (c) the threshold \bar{i}^R remains constant. We collect these results in the following:

$$\begin{aligned} V^{Rh}(i, t) &= V^{Rh}(i) = \frac{\ln w^h(i)}{1 - \beta} \\ V^{Rm}(i, t) &= V^{Rm}(i) = \frac{\ln w^m(i)}{1 - \beta} \\ (1 - \beta) V^R(i) &= \text{Max}\{\ln w^h(i), \ln w^m(i)\} \end{aligned}$$

3.2. Infected individuals

To determine the lifetime utility of an infected agent $V^I(i)$, we assume that infected individuals can earn $w_I^h(i)$ in home work and $w_I^m(i)$ in market work. As noted above, we assume that $w_I^h(i) < w^h(i)$ and $w_I^m(i) < w^m(i)$ for all i . The assumptions capture the realistic scenario of an infected worker being less productive than a healthy worker. Infected individuals recover with probability π_R and die with probability π_D . These probabilities are assumed to be time invariant in the current formulation.⁶

6 In the robustness analysis below, we allow for π_D to vary based on the degree of congestion in the capacity of the health system.

An infected individual's value functions from working either at home or in the market are defined as

$$\begin{aligned}
 V^{Ih}(i, t) &= \ln w_I^h(i) + \beta[(1 - \pi_R - \pi_D)\text{Max}\{V^{Ih}(i, t + 1), V^{Im}(i, t + 1)\} \\
 &\quad + \pi_R V^R(i, t + 1)], \\
 V^{Im}(i, t) &= \ln w_I^m(i) + \beta[(1 - \pi_R - \pi_D)\text{Max}\{V^{Ih}(i, t + 1), V^{Im}(i, t + 1)\} \\
 &\quad + \pi_R V^R(i, t + 1)].
 \end{aligned}$$

The cutoff market occupation for the set of infected individuals will be determined by $V^{Ih}(\bar{i}^I, t) = V^{Im}(\bar{i}^I, t)$. Because the continuation values from both market work and home work today are identical, the cutoff condition reduces to

$$w_I^h(\bar{i}^I) = w_I^m(\bar{i}^I).$$

The associated measure of occupations with infected people working in the marketplace is $\theta^I = 1 - \bar{i}^I$. Note that the constancy of occupational wages implies that both \bar{i}^I and θ^I are constant over time.

These results can be collected in the following expressions:

$$\begin{aligned}
 V^{Ih}(i, t) &= V^{Ih}(i) = \frac{\ln w_I^h(i) + \beta\pi_R V^R(i)}{1 - \beta(1 - \pi_R - \pi_D)} \\
 V^{Im}(i, t) &= V^{Im}(i) = \frac{\ln w_I^m(i) + \beta\pi_R V^R(i)}{1 - \beta(1 - \pi_R - \pi_D)} \\
 V^I(i) &= \text{Max}\{V^{Ih}(i), V^{Im}(i)\}
 \end{aligned}$$

In deriving these expressions, we have used the fact that the lifetime value of a recovered agent $V^R(i)$ is a constant over time. We established that above.

3.3. Susceptible individuals

Having described the choices of recovered and infected agents, we now turn to susceptible agents. For these individuals, their current choices have implications for their utility continuation values because susceptible agents are at risk of contracting the virus.

There are three ways in which an individual can get infected. The first is through interactions with infected people while working in the market, which we denote by π_t^m . The second is in the process of random interactions with infected people, through physical closeness in leisure activities or through travel, for instance. We denote this risk by π_t^l . Finally, we allow for infections through consumption activity, intended to capture the interaction of individuals through retail activity. We denote this risk as π_t^c . While π_t^l and π_t^c are common to all susceptible individuals, the risk of infection at work is specific to agents who choose to work in the market rather than at home. π_t^c , π_t^l and π_t^m

fluctuate over time as the market participation of agents changes in the economy and as the overall number of infected people changes.

When there is a risk of infection from market work, the value functions for a susceptible individual in occupation i are

$$\begin{aligned}
 V^{Sh}(i, t) &= \ln w^h(i) + \beta [(\pi_t^I + \pi_t^c) V^I(i) \\
 &\quad + (1 - \pi_t^c - \pi_t^I) \text{Max}\{V^{Sh}(i, t + 1), V^{Sm}(i, t + 1)\}], \\
 V^{Sm}(i, t) &= \ln w^m(i) + \beta [(\pi_t^m + \pi_t^I + \pi_t^c) V^I(i) \\
 &\quad + (1 - \pi_t^m - \pi_t^I - \pi_t^c) \text{Max}\{V^{Sh}(i, t + 1), V^{Sm}(i, t + 1)\}],
 \end{aligned}$$

where $V^I(i)$ is the function that we solved for above.

From the point of view of a susceptible person, the probability of being infected in the market place is equal to $\pi_t^m = \frac{\epsilon S_t^m I_t^m}{S_t^m} = \epsilon I_t^m$. Individuals take this probability of getting infected to be given exogenously when they make their decisions regarding market versus home work in their occupations.

The cutoff \bar{i}_t^S is determined by

$$\begin{aligned}
 \ln w^h(\bar{i}_t^S) &= \ln w^m(\bar{i}_t^S) + \beta \pi_t^m [V^I(\bar{i}_t^S) \\
 &\quad - \text{Max}\{V^{Sh}(\bar{i}_t^S, t + 1), V^{Sm}(\bar{i}_t^S, t + 1)\}].
 \end{aligned}$$

We now prove the following result that will be crucial in characterizing the solution:

CONDITION 2.

$$V^I(i) < \text{Max}\{V^{Sh}(i, t), V^{Sm}(i, t)\} \forall i, t.^7$$

PROPOSITION 1. *Condition 2 is necessary and sufficient for the threshold occupation \bar{i}_t^S to be increasing in the probability of infection π_t^m .*

Proof. From the expressions for $V^{Sh}(i)$ and $V^{Sm}(i)$ above, we get

$$\begin{aligned}
 V^{Sh}(i, t) - V^{Sm}(i, t) &= \ln w^h(i) - \ln w^m(i) - \beta \pi_t^m [V^I(i) \\
 &\quad - \text{Max}\{V^{Sh}(i, t + 1), V^{Sm}(i, t + 1)\}].
 \end{aligned}$$

Partially differentiating this with respect to π_t^m gives, for any occupation i ,

$$\frac{\partial [V^{Sh}(i, t) - V^{Sm}(i, t)]}{\partial \pi_t^m} = \beta [\text{Max}\{V^{Sh}(i, t + 1), V^{Sm}(i, t + 1)\} - V^I(i)].$$

7 There always exists a weakly positive wage from home or market work that will guarantee that this condition holds.

This derivative is positive under condition 2. The threshold condition is $V^{Sh}(\bar{i}^S) = V^{Sm}(\bar{i}^S)$. Because $\frac{\partial[V^{Sh}(i) - V^{Sm}(i)]}{\partial \pi_t^m} > 0$, it follows that $V^{Sh}(\bar{i}_t^S, t + 1) > V^{Sm}(\bar{i}_t^S, t + 1)$ if and only if $\pi_{t+1}^m > \pi_t^m$. Hence, the threshold $\bar{i}_{t+1}^S > \bar{i}_t^S$ when $\pi_{t+1}^m > \pi_t^m$ because $\ln w^m(i) - \ln w^h(i)$ is increasing in i . ■

COROLLARY 1. *Along paths with a falling probability of infection from market work π_t^m , the threshold occupation \bar{i}_t^S declines so that more occupations become market based over time, or $\theta_t^S = 1 - \bar{i}_t^S$ rises.*

The proposition says that as long as getting infected reduces lifetime utility sufficiently, a higher risk of infection raises the benefit of working from home. Consequently, along paths with monotonically rising π_t^m , the threshold θ_t^S falls as \bar{i}_t^S rises.

More generally, however, during the evolving dynamics of the pandemic, π_t^m will typically be non-monotonic, rising initially as more infected individuals come in contact with susceptible individuals, but then falling as both the population of infected and susceptible individuals falls. In this case, \bar{i}_t^S will move in a non-monotonic direction also. The path of \bar{i}_t^S is then determined by the peak value of π_t^m that is attained as the pandemic spreads. In this case, we can establish proposition 2.

PROPOSITION 2. *Take a sequence of infection probabilities $\pi_1^m \dots \pi_T^m$ such that $\pi_1^m < \pi_2^m < \dots < \pi_M^m$ and $\pi_M^m > \pi_{M+1}^m > \dots > \pi_T^m$.*

Then, for $\pi_t^m < \pi_M^m$, \bar{i}^S is governed by the dynamics:

$$\begin{aligned} \ln w^h(\bar{i}_t^S) &= \ln w^m(\bar{i}) + \beta \pi_t^m [V^I(\bar{i}_t^S, t + 1) - V^{Sh}(\bar{i}_t^S, t + 1)] \\ V^{Sh}(i, t) &= \ln w^h(i) + \beta [(\pi_t^I + \pi_t^c) V^I(i) + (1 - \pi_t^c - \pi_t^I) V^{Sh}(i, t + 1)] \end{aligned}$$

and for $\pi_t^m > \pi_M^m$, \bar{i}^S is governed by the dynamics:

$$\begin{aligned} \ln w^h(\bar{i}_t^S) &= \ln w^m(\bar{i}) + \beta \pi_t^m [V^I(\bar{i}_t^S, t + 1) - V^{Sm}(\bar{i}_t^S, t + 1)] \\ V^{Sm}(i, t) &= \ln w^m(i) + \beta [(\pi_t^m + \pi_t^I + \pi_t^c) V^I(i) \\ &\quad + (1 - \pi_t^m - \pi_t^I - \pi_t^c) V^{Sm}(i, t + 1)] \end{aligned}$$

Proof. This follows directly from the proof of proposition 1 when combined with the non-monotonic path of π_t^m . ■

PROPOSITION 3. *Under condition 2, the threshold occupation of susceptible individuals, \bar{i}_t^S , is greater than the threshold for recovered agents \bar{i}^R for all t . Hence, the share of occupations under market work is greater for recovered individuals than susceptible agents, or $\theta^R > \theta_t^S$.*

Proof. The cutoff condition for recovered agents can be written as $\ln w^m(\bar{i}^R) - \ln w^h(\bar{i}^R) = 0$. The threshold condition for susceptible agents is

$$\ln w^m(\bar{i}_t^S) - \ln w^h(\bar{i}_t^S) = \beta \pi_t^m \left[\text{Max} \left\{ V^{Sh}(\bar{i}_t^S, t + 1), V^{sm}(\bar{i}_t^S, t + 1) \right\} - V^I(\bar{i}) \right] > 0,$$

where the last inequality holds under condition 2. Because $\ln w^m(i) - \ln w^h(i)$ is rising in i , we must have $\bar{i}_t^S > \bar{i}^R$. The statement on the share of market occupations for types R and I follows trivially as $\theta^j = 1 - \bar{i}^j$ for $j = R, I$. ■

3.4. Occupation populations

The absolute measures of individuals working in the market in this economy at any date t is the sum of susceptible, infected and recovered individuals working in market occupations: $P_t^m = S_t^m + I_t^m + R_t^m$. The number of market workers within each group is given by

$$S_t^m = \int_{\bar{i}_t^S}^1 S_t(i) di, \tag{11}$$

$$I_t^m = \int_{\bar{i}^I}^1 I_t(i) di, \tag{12}$$

$$R_t^m = \int_{\bar{i}^R}^1 R_t(i) di. \tag{13}$$

Equation (11) is the measure of susceptible individuals engaged in market work.

Because infections are dependent on whether susceptible individuals work in the market or at home, we also need to describe the evolution of occupation-specific populations in each of the three groups S , I and R . The evolution of susceptible agents in occupation i is given by

$$S_{t+1}(i) = \begin{cases} S_t(i) - \varepsilon S_t(i) I_t^m - \pi_c C_t^S(i) C_t^I - \pi S_t(i) I_t, & \text{if } i > \bar{i}_t^S \\ S_t(i) - \pi_c C_t^S(i) C_t^I - \pi S_t(i) I_t, & \text{if } i \leq \bar{i}_t^S, \end{cases} \tag{14}$$

where ε is the matching rate that generates new infections through work and π is the probability of infection from random contact between susceptible and infected people. π_c is the exogenous probability of getting infected due to random interaction between consumption of the susceptible and infected individuals. C_t^S denotes aggregate consumption of susceptible individuals, while C_t^I is total consumption of infected individuals. We shall solve for these below. In deriving the above, we have used the fact that only susceptible agents in occupations that are provided in the market, $i > \bar{i}_t^S$, risk getting infected.

Integrating equation (14) over all occupations gives the evolution of the number of susceptible individuals as

$$S_{t+1} = S_t - \varepsilon S_t^m I_t^m - \pi_c C_t^S C_t^I - \pi S_t I_t. \tag{15}$$

Similarly, the number of infected people in occupation i evolves according to

$$I_{t+1}(i) = \begin{cases} (1 - \pi_R - \pi_D)I_t(i) + \varepsilon S_t(i)I_t^m + \pi_c C_t^S(i)C_t^I + \pi S_t(i)I_t & \text{if } i > \bar{i}_t^S \\ (1 - \pi_R - \pi_D)I_t(i) + \pi_c C_t^S(i)C_t^I + \pi S_t(i)I_t & \text{if } i \leq \bar{i}_t^S. \end{cases} \tag{16}$$

Integrating equation (16) over all occupations gives the evolution of infected individuals in the economy as

$$I_{t+1} = (1 - \pi_R - \pi_D)I_t + \varepsilon S_t^m I_t^m + \pi_c C_t^S C_t^I + \pi S_t I_t. \tag{17}$$

Lastly, the measure of recovered individuals in occupation i is given by

$$R_{t+1}(i) = R_t(i) + \pi_R I_t(i). \tag{18}$$

Integrating this over all occupations yields the evolution of recovered agents as

$$R_{t+1} = R_t + \pi_R I_t. \tag{19}$$

Noting that the total population is $L_t = S_t + I_t + R_t$, we can sum equations (15), (17) and (19) to get the evolution of the total populations as

$$L_{t+1} = L_t - \pi_d I_t.$$

We assume that the initially infected population is distributed proportionally across all occupations so that

$$S_0(i) = (1 - \varepsilon_0)L_0(i), \tag{20}$$

$$I_0(i) = \varepsilon_0 L_0(i), \tag{21}$$

$$R_0(i) = 0. \tag{22}$$

3.5. Consumption

Aggregating over consumption by all individuals in each occupation gives the occupation-specific aggregate consumption as

$$\begin{aligned} C_t^S(i) &= c_t^S(i)S_t(i), \\ C_t^I(i) &= c_t^I(i)I_t(i), \\ C_t^R(i) &= c_t^R(i)R_t(i). \end{aligned}$$

Hence, total consumption in the economy, which must be the sum of consumption by each group, is

$$C_t = \int_i \{C_t^S(i) + C_t^I(i) + C_t^R(i)\} di = C_t^S + C_t^I + C_t^R. \quad (23)$$

3.6. Government

The government maintains a balanced budget by rebating all its tax revenues through lump sum transfers to individuals. This implies that

$$G_t = g_t L_t = \tau^m \int_{i=0}^1 A_i^m \{S_t^m(i) + \gamma_I^m I_t^m(i) + R_t^m(i)\} di. \quad (24)$$

3.7. Market clearing

We conclude the description of the key relationships by describing the market clearing condition for the final good. Because the government rebates all tax revenues to individuals in the form of lump sum transfers, we have

$$C_t = Y_t. \quad (25)$$

4. Quantitative results

We now turn to quantifying the model under condition 1, where all the occupations in producing the sectoral output and sectors in the final goods technology are perfect substitutes and the sectoral weights are equal. Recall that in this special case, GDP is just a sum of all the occupational outputs. The model has a number of margins that affect outcomes. We flesh out some of these through comparative static exercises on the parameters that control these margins.

4.1. Calibration

There are 10 key parameters in the model. The parameters are mostly taken from Eichenbaum et al. (2021). They calibrate the probability parameters ε , π_C and π to jointly match the probabilities of infection from consumption and market work and assuming that by the end of the pandemic 60% of the population is infected. They estimate these probabilities by combining the estimates of Ferguson et al. (2006) with the Bureau of Labour Statistics (BLS) 2018 Time Use Survey and the 2018 BLS data on total workers. Their estimates imply that 16% of transmissions occur during consumption, 17% at work and remainder are through random interactions.

We follow Atkeson (2020) in assuming that it takes 18 days to either recover from the virus or die from it. Because our time unit is a week, we set

TABLE 1

Parameter values for calibration

Target variable	Parameters	Value
Probability of random infection	π	0.40
Probability of infection while consuming	π_C	1.5924e-07
Probability of infection through market work	ε	0.11
Probability of death	π_D	0.002
Probability of recovery	π_R	0.387
Share of work from home	\bar{i}	0.06
Weekly discount factor	β	0.999
Relative home productivity of infected	γ_I^h	0.8
Relative market productivity of infected	γ_I^m	0.8
Initial fraction infected	ε_0	0.001

$\pi_D + \pi_R = 7/18$. On the basis of South Korean data, Eichenbaum et al. (2021) estimate the mortality rate from the virus to be 0.5%. This yields individual estimates for π_R and π_D . The baseline values for the relative productivity of infected people is set to 0.8. They set this value based on the Chinese Center for Disease Control and Prevention's estimate that 80% of infected people are asymptomatic and that symptomatic individuals withdraw from work. The values of the discount factor and the initially infected are also taken from Eichenbaum et al. (2021).

We estimate the share of individuals engaged in occupations classified as fully working from home at 6%. This is based on estimates from the 2016 Canadian census along with the size of the population working in each occupation before the pandemic begins. The weekly discount factor is computed by using the annual discount factor of 0.96 and converting it to its weekly equivalent. Lastly, we start the pandemic off by assuming that there an exogenous infection of 0.1% of the population at an initial date $t = 0$. The values for these parameters that we use for the simulations are given in table 1.

4.2. Data

We implement the model using occupations mapped into the share of occupations in British Columbia (BC) using the National Occupational Classifications (NOC) using data from Statistics Canada Labour Force Survey (Statistics Canada 2021), health risk and work-from-home shares provided by Baylis et al. (2020a)⁸ and partly from the 2016 Canadian census (Statistics Canada 2016). Figure 1 shows shows the different occupations at the two-digit NOC level with their share based on the November 2019 Labour Force Survey as well as the measure of each occupation that reported some degree of working from home from the 2016 Canadian census. It shows that working from home is more prevalent in professional occupations and management

8 See Baylis et al. (2020b) for replication file.

occupations in art and finance, consistent with the findings of Hensvik et al. (2020) for the US.

Figure 2 reports the change in employment by occupation from February to April 2020 and then from April 2020 to June 2020 (based on the Labour Force Survey from Statistics Canada). Each change is weighted by the share of the occupation in the total labour force. The first two months of the pandemic saw job losses of almost 400,000 in BC. Hardest hit were occupations in service sectors such as sales representatives and service support occupations. Other occupations, such as senior managers, were almost untouched. In addition, occupations related to essential services, such as nursing occupations or occupations in front line protection services increased in employment during the first period. The June Labour Force Survey then shows a pick up of approximately 200,000 jobs from the April low point. The message from figure 2 indicates substantial heterogeneity in job losses among occupations.

The model set out above does not have an explicit role for unemployment. We assume each occupation either works at home or in the market. However, we can connect the model to the data by interpreting switches to working from home among occupations that face a major loss of income from home working as being analogous to unemployment. On the contrary, occupations such as senior management or professional occupations in business and finance can easily remain in the category of being employed while either continuing to work from home or switching to work from home following the outbreak of the pandemic.




FIGURE 1 Employment and work-from-home shares for NOC occupations




FIGURE 2 Change in employment for each occupation, February to June 2020

We should note that because switching from market to home work in our model is voluntary, the model, in a literal sense, does not generate involuntary unemployment. But to the extent that lockdown policies induce individuals to leave the market beyond their unconstrained choices (as we describe below), these policies play the role of supply restrictions leading to job losses as characterized in some of the recent COVID-19 literature (see Baqaee and Farhi 2020, Guerrieri et al. 2020 and Woodford 2020). A key message of our analysis is that such restrictions may play a secondary role relative to the endogenous choice of households in accounting for the retreat from market work and the fall in household consumption during a pandemic.

4.3. Estimating the home wage schedule

There are two complications in mapping the available data on occupation distribution and wages to the model. The first is that in the steady state of the model, occupations are characterized by either all individuals working from home or in the market. The data, however, reveal a mix of market and home work shares in each occupation. A second issue is the measurement of wages. In the model, individuals make labour supply decisions on the basis of occupation-specific wages for home versus market work. While we do have occupation-specific wages from the 2016 census, these are equilibrium wages and in most cases refer to wages paid for working in the market. We do not observe home wages separately.

We get around the first problem by adopting a threshold approach to classifying occupations into working from home or in the market. Specifically, we

classify all occupations with work-from-home shares above a certain threshold to be entirely working from home, while those below are classified as working in the market. In a 2017 survey by Regus (2017), approximately 10% of workers reported working exclusively from home. The 2016 Canadian census reports, for each occupation, the share of workers who work from home. In order to match the overall 10% share of home work, we choose the threshold share of working from home as $\hat{\lambda} = 0.194$.⁹ This represents the 90th percentile in the occupational distribution of work-from-home shares. This chosen threshold leaves four occupations (professional occupations in art, insurance, middle management and senior management) that are classified as entirely working from home in the initial steady state.

In order to map the model to the data, we also need a measure of wages from working at home in each occupation. Because our data on occupational wages reveal only equilibrium wages, the home wage schedule for workers who work exclusively in the market during normal times needs to be imputed. We make the assumption that the difference between the home and market wage schedule in any occupation is a linear, increasing function of the ease of performing the occupation-specific tasks from home. We proxy the ease with which tasks can be performed from home in an occupation by the fraction of workers in that occupation who work from home.

Using $\lambda(i)$ to denote the share of workers in an occupation that works from home, we compute the home wage schedule as

$$w^h(i) = w^m(i) [1 + (\lambda(i) - \hat{\lambda})], \quad (26)$$

where $w^m(i)$ is the median weekly wage in occupation i . Hence, occupations with home work shares below the threshold will have home wages below the market wage, while those with home work shares above the threshold will have home wages that exceed the market wage. As noted above, our data contain information on 40 two-digit occupations. Data on occupation wages come from the 2016 census.

Figure 3 shows the estimated relative home wage schedule and the shares of working from home in each occupation in the data. In the following, we shall use this estimated home wage premium schedule to determine optimal choices by individuals within each occupation.

4.4. Results

We now show the dynamics of the model starting from a steady state when 0.1% of the population gets exogenously infected with a contagious virus. We

⁹ The employment share of individuals working at home over all occupations is estimated at 10%. But the model definition of working at home instead requires that we compute the employment share of those occupations that are performed at home. This estimate from the data is 6%.



FIGURE 3 Estimated relative home wage schedule

start by showing the dynamics in the baseline case of the model with parameter values as in table 1. Figure 4 depicts the dynamic behaviour of the key variables. The time unit for our analysis is weeks and the infection and death rate statistics are expressed in weekly units.

The solid (grey) line shows the dynamics in the baseline case of the model. The infection rate peaks at slightly under 5% after around 32 weeks from the initial infection. The pandemic kills a cumulated 0.25% of the population. Moreover, 50% of the population remains susceptible at the end of the pandemic.

The baseline numbers reported in figure 4 imply that, for a country with 30 million people, 15 million people get infected at some point, while 75,000 people die. By comparison, as of April 14, 2021, total reported cases in Canada were 1.09 million, along with 23,445 fatalities. Clearly, the model underpredicts both infections and deaths. But it is widely understood that actual infection rates are much higher than those reported. The CDC in the US estimates that infection rates are up to 10 times those officially reported Mandavilli (2020).

The preceding suggests that to get closer to matching both the infection rates and death rates seen in the data, we may have to adjust death rates downwards. Using a number of alternative adjustment factors, section 4.6 provides a comparison of the infection and death rates with those in the US and Canada during the first wave of the pandemic.

Given that both medical and public health practices developed in leaps and bounds since February 2020 when the pandemic first began attracting global attention, the parameters for infection and death probabilities are likely

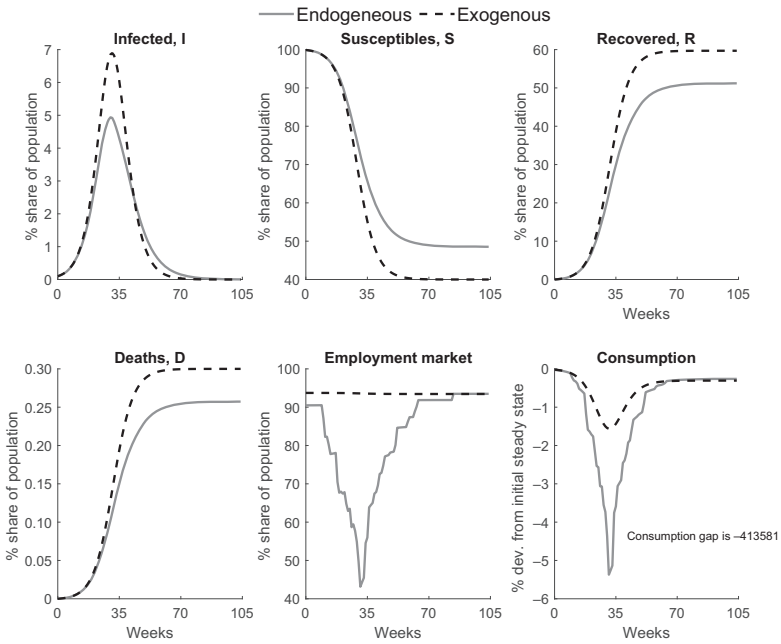


FIGURE 4 Endogenous versus exogenous choice of market and home work

to have changed significantly. Because we ignore this, our model is better interpreted as describing the initial wave of the pandemic where agents had to make the choice of market versus home work for a given institutional and technological structure of the workplace. The quantitative results should be interpreted in this light.

The economic costs of the pandemic occur because of both a decline in market employment and a fall in productivity due to infection. At the deepest point of the recession, which occurs at the same time as the peak of the infections, market employment falls by around 45 percentage points, while output and consumption decline by over 5% relative to steady state. These numbers appear to be in line with the initial estimates of job losses and output contractions reported from Canada and around the world.

A distinguishing feature of the model is that private agents endogenously select into market versus working from home. Figure 4 shows the crucial role played by these endogenous choices made by individuals. The dashed (black) line shows the dynamic responses when the work location choices of individuals are shut down completely. When work choices are constrained to remain at the pre-pandemic level, the peak infection rate and the cumulated death rate is two percentage points and 0.05 percentage points higher, respectively, than in the endogenous choice case. The higher infection and death rates occur because people continue to work in the market instead of withdrawing to working from home. Relative to the endogenous choice case, consumption

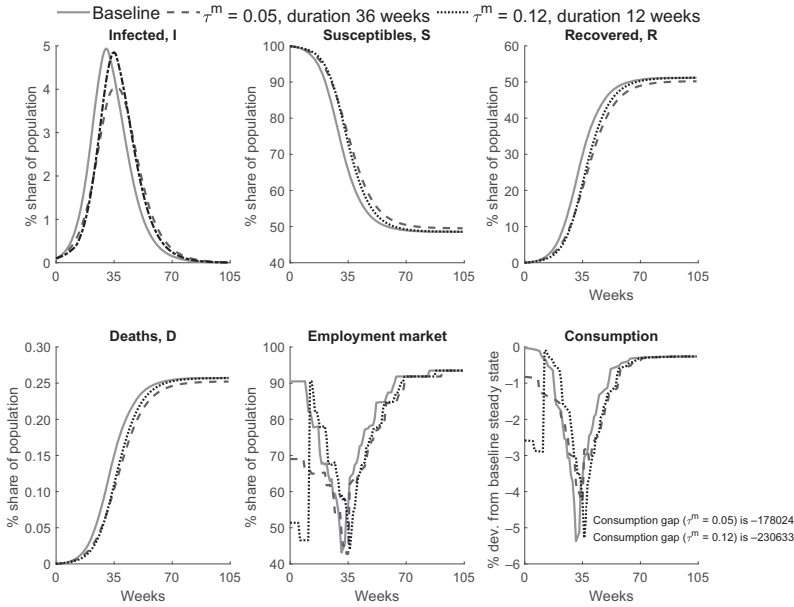


FIGURE 5 Varying the intensity of a lockdown

declines by almost four percentage points less in the exogenous choice case. Clearly, the effect of endogenous choice of work location is quantitatively substantial.

4.5. Temporary lockdowns

Countries globally have chosen various degrees of lockdowns to fight the spread of the infection. This aspect of policy behaviour can be captured in the model through taxes on market production. How effective is shutting down markets in reducing the prevalence of infections? How costly is it? Figure 5 shows the effect of two different lockdown intensities. The first is a lockdown for 36 weeks with a $\tau^m = 0.05$ (dashed black lines). The second is a shorter lockdown of 12 weeks but more intense with $\tau^m = 0.12$ (dotted black lines). The figure also shows the baseline case of $\tau^m = 0$ with the solid (grey) line. The government rebates the tax revenue through lump sum transfers to the individuals.

Figure 5 makes clear that a more intense but shorter lockdown of market work for 12 weeks does not materially alter either the peak infection rate or the cumulated number of deaths. It essentially delays the peak by the duration of the lockdown. Market employment collapses by over 40 percentage points on impact of the shock with attendant output losses of almost 3% relative to steady state. However, these effects are relatively short lived as the end of the lockdown results in a rapid recovery of market employment.

Interestingly, the economy begins a fresh contraction in market employment and output after the end of the lockdown, but this situation is due entirely to the endogenous response of individuals sheltering from infections and from the continuing rise of infections in the economy.

The other case is where the lockdown is less intense with $\tau^m = 0.05$ but lasts for 36 weeks. Figure 5 shows that this longer lockdown lowers the peak infection rate by one percentage point relative to the baseline and is less costly in terms of output. The reasons are twofold. First, the initial collapse in market work is lower under the less intense lockdown. Second, the lockdown lasts long enough to ensure that the pandemic reaches its peak before the program ends. Consequently, there is no renewed endogenous decline in market work and output after the end of the program.

An interesting feature of the temporary lockdown policies is that they have little impact on the long run outcome of the pandemic, in terms of total infections and deaths. This contrasts sharply with the impact of endogenous shifts out of market work. Comparing figure 5 with figure 4, we see that the final death rate from the pandemic is reduced from 0.3% to 0.25% of the population through endogenous responses of individuals in their occupation location choices. The additional reduction in deaths coming from lockdowns is negligible.

While each of these policies have some impact on the infection curve, neither manages to completely “flatten the curve” as described in media discussion of pandemic responses. But more severe policies can have a substantial effect. Figure 6 shows a drastic lockdown policies, involving $\tau^m = 0.20$ imposed for 52 weeks. This causes an immediate and complete shift to working from home for all occupations and significantly reduces infections and final death rates, but it comes at a huge cost in terms of overall consumption.¹⁰

4.6. Adjusted fatality rates

As we have noted above, our baseline calibration using economic and epidemiological data produces significant exit from market work and reductions in consumption but still produces infection and death rates above those in the officially measured data from Canada during the first year of the pandemic. In figure 7, we explore the impact of an alternative, lower calibration of infection rates and compare this with data on BC, Canada as a whole and the US.

10 It should be acknowledged that government-mandated lockdowns affect not just the workplace but also interaction through consumption activities and other venues for community spread of infection. A more extensive lockdown policy would then likely reduce the parameters π and π_C . But our main interest is in the decision over home or market work, and because we do not allow for endogenous choice of risk reduction in consumption activities, we do not explore this dimension of lockdowns in our baseline model.

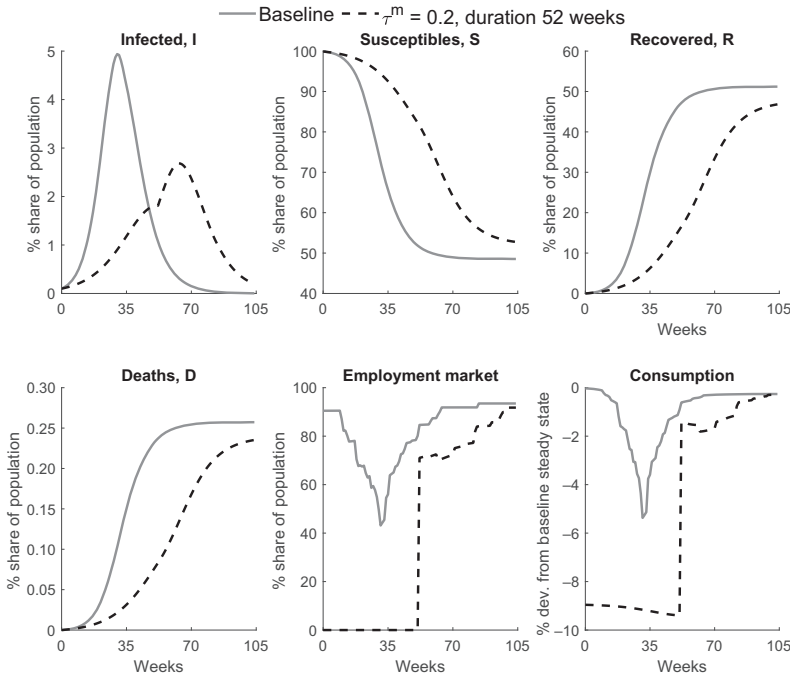


FIGURE 6 An extreme lockdown policy

Figure 7 is constructed by imposing a uniform 10% reduction in infection rates for all activities, including market work, consumption and random activities. This cuts the peak of the infection rate from 5% to 3% and reduces the long run death rate from 0.25% to 0.2%. It also leads to a much smaller reduction in market employment and consumption.

We note again, however, that measured infection rates from the disease are likely to be considerably underestimated, while death rates are more reliable. Figure 8 makes allowances for this uncertainty in measurement in comparing model and data. The left-hand panel reports the death rates as a proportion of the population in the data, the baseline model and the model with a 10% uniform reduction in infection probabilities. The right-hand side panel reports actual and imputed infection rates.

This panel is constructed as follows. From observed death rates, we impute infection rates using three different assumptions for infection fatality rates. Based on the assumption that it takes 18 days to recover or to succumb to the disease, we construct implied infection rates from daily death rates using infection fatality rates (IFR) of 0.5% (the baseline model), 0.25% and 1.0%. We then compare imputed infection rates with those from the model baseline case as well as with the model with a uniform lower infection rate. We report these estimates for Canada, the US and British Columbia.

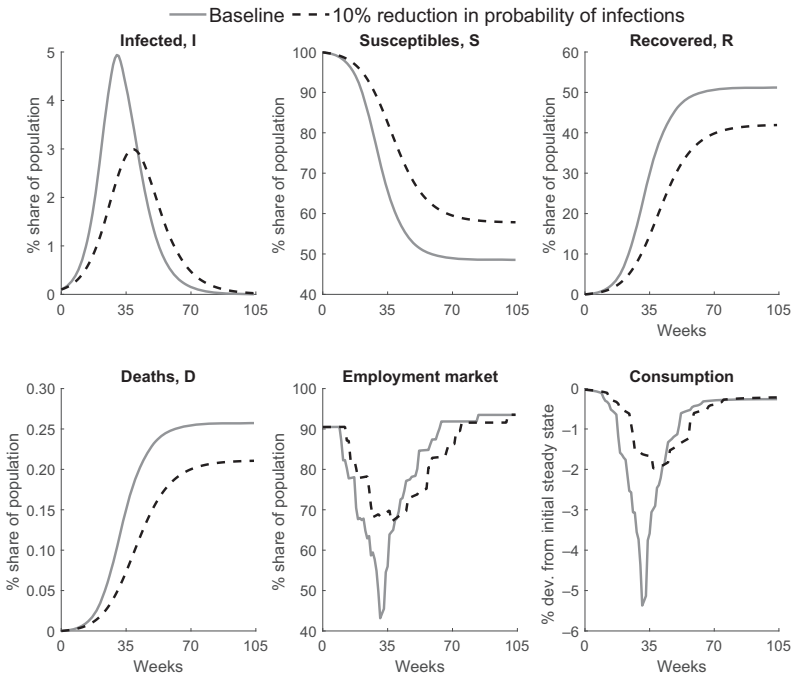



FIGURE 7 Lower probability of infections

The results indicate that, with lower infection rates and lower infection fatality rates, the model comes close to matching the observations for the US in the first year of the pandemic, but still produces numbers considerably above those of Canada and BC.

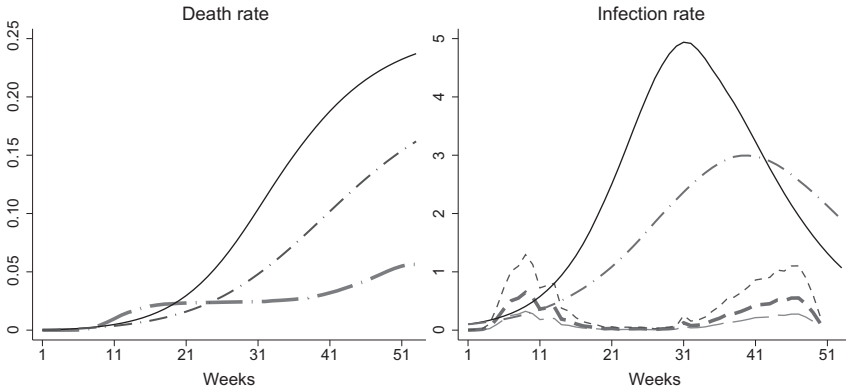
5. Robustness of results

Amongst the various assumptions underlying the baseline model, five appear to be potentially crucial. First, the model assumes that the infection risk from working in the market is common across all occupations. Second, we assumed that utility is log-linear in consumption. This leaves open the question of the sensitivity of the model to the degree of risk aversion. Third, the size of the output losses as well as the attitude to exposing oneself to the risk from market work has to be dependent on the productivity of labour contingent on becoming infected. Fourth, our baseline model did not account for either the possibility of a vaccine becoming available or the consequences of congestion in the health care system. Fifth, governments worldwide implemented various fiscal programs for pandemic relief of their citizens. How would fiscal transfers affect our results? We now examine the importance of these margins.

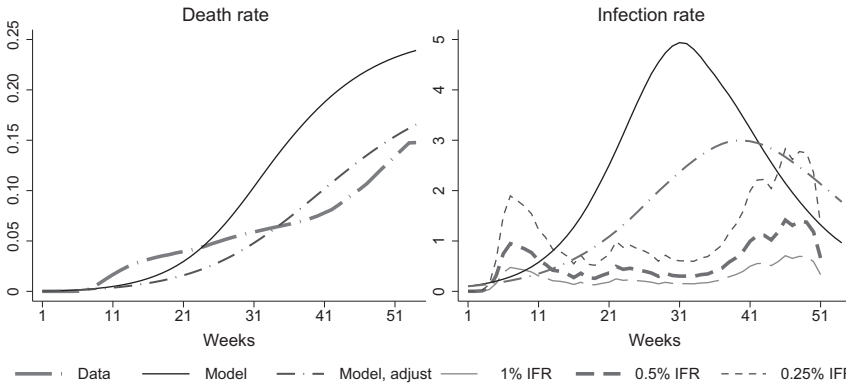
British Columbia



Canada



United States



— Data — Model - - - Model, adjust — 1% IFR — 0.5% IFR - - - 0.25% IFR

FIGURE 8 Comparing model simulations with data
NOTE: Death data for Canada and British Columbia are from the Canadian COVID-19 epidemiological data (www.canada.ca/en/public-health/services/diseases/coronavirus-disease-covid-19/epidemiological-economic-research-data.html) and from The New York Times (2021) for the US.

5.1. Occupation-specific risk

Thus far we assumed that π_t^m was independent of occupation i . This is a strong assumption. The exposure to infection risk varies across occupations because of occupation-specific characteristics like proximity to others at work, the number of people one has to interact with at work, etc. and the living conditions of the workers performing those occupations such as whether the individual lives with a health worker and the number of people the individual lives with. Baylis et al. (2020a) develop a measure of viral risk exposure across four-digit NOC occupations (henceforth referred to as VSE risk index) using a number of characteristics relating to conditions at and outside the workplace as highlighted by public health officials including the British Columbia Centre for Disease Control.

Figure 9 depicts the variation in this risk index across occupations at the two-digit level. The figure clearly indicate that different occupations experience different degrees of infection risk, quite separately from their market versus home wage premium. Health-related occupations like Professional occupations in health and nursing have the highest viral transmission risk but a much lower wage from working at home than at marketplace. On the other hand, professional occupations in finance and natural sciences have among the highest home wage premium but much lower risk scores. Does this matter?

We examine the importance of occupation-specific risk by modifying the model. Recall that the probability of getting infected through market work was $\pi_t^m = \varepsilon I_t^m$. We now modify this to

$$\pi_{it}^m = \varepsilon_i I_t^m.$$

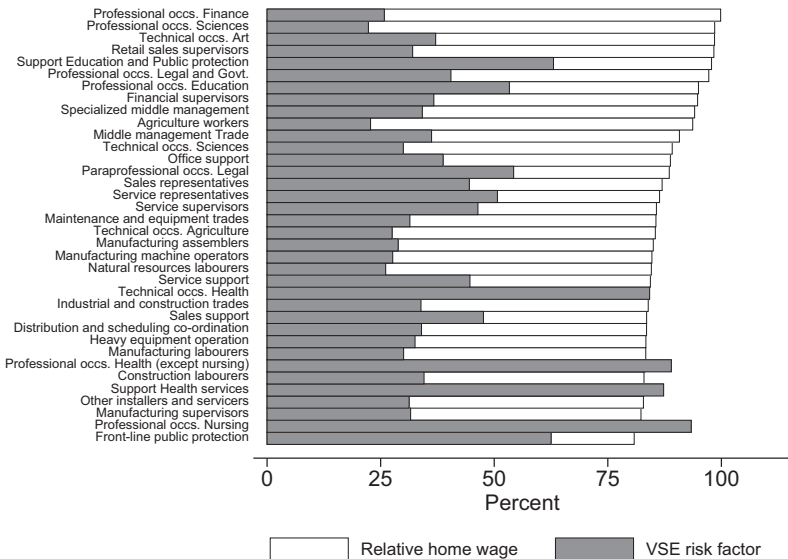


FIGURE 9 Risk indices by occupation

In order to compute the occupation-specific infection risk, we first compute the mean risk factor across all occupations based on the VSE risk tool. We then define the risk factor proportionality of each occupation as

$$z_i = \frac{\text{VSE risk index for occupation } i - \text{Mean risk index}}{\text{Mean risk index}}.$$

Using this, we compute the occupation-specific infection risk parameter ε_i as

$$\varepsilon_i = \varepsilon(1 + z_i),$$

where ε is the common risk parameter that we were using before.

Figure 10 compares the responses of the model under occupation-specific risk with those under common risk. The solid (grey) lines are the endogenous responses under common risk, while the dotted (black) lines are the endogenous responses under occupation-specific risk. The key point to note is that the differences are relatively minor in the aggregate. Put differently, the results that we obtained under common infection risk across occupations are robust to accounting for occupation-specific risk.

5.1.1. Which occupations stay home?

Our model is calibrated to the BC labour force categorized according to the National Occupational Classification. We have described how occupations differ both according to the ease of work from home and the infection risk in the labour market. It is interesting to see what the model predicts for the distribution of market exits by occupation. Figure 11 lists the occupations that exit the market and for how long, in the baseline case and in response to the two lockdown scenarios, in the case of common infection risk across occupations.

As one would expect, with common risk, the occupations exiting the market are those whose opportunity cost is lower—particularly professional occupations with less customer or client physical interaction. We see also that the lockdown policies have a relatively minor additional impact upon the occupations that exit the market and the duration of exit.

When we introduce occupation-specific risk, figure 12 shows that the type of occupations exiting the market changes quite considerably. Now individuals in “risky” occupations such as health, retail or front line protection choose to exit market work. But comparing figures 11 and 12, we see it is still true that most of the time outside the market is accounted for by occupations with a lower opportunity cost of working from home. This holds both for the baseline case and for the lockdown experiments. This observation helps to explain the result in figure 10, where we saw that the model simulations under constant risk and occupational specific risk differ by only a minor degree.

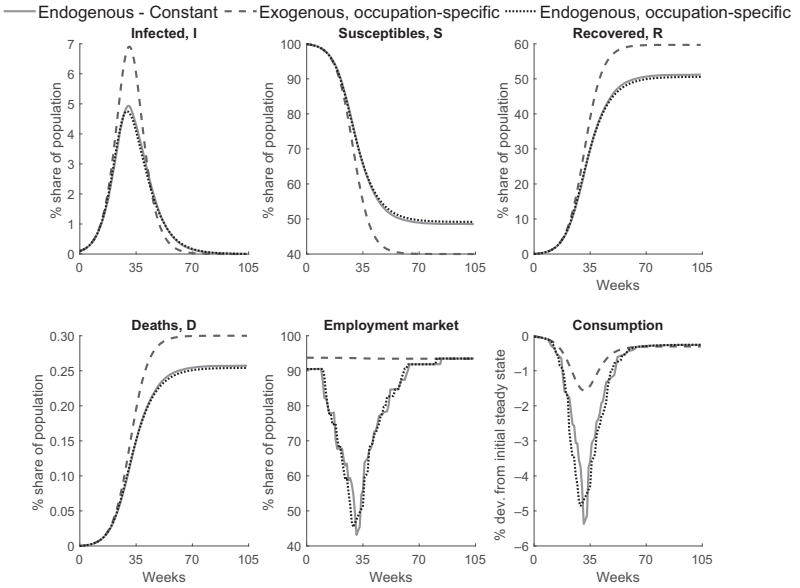


FIGURE 10 Occupation-specific infection risk




FIGURE 11 Occupations exiting the market with constant risk

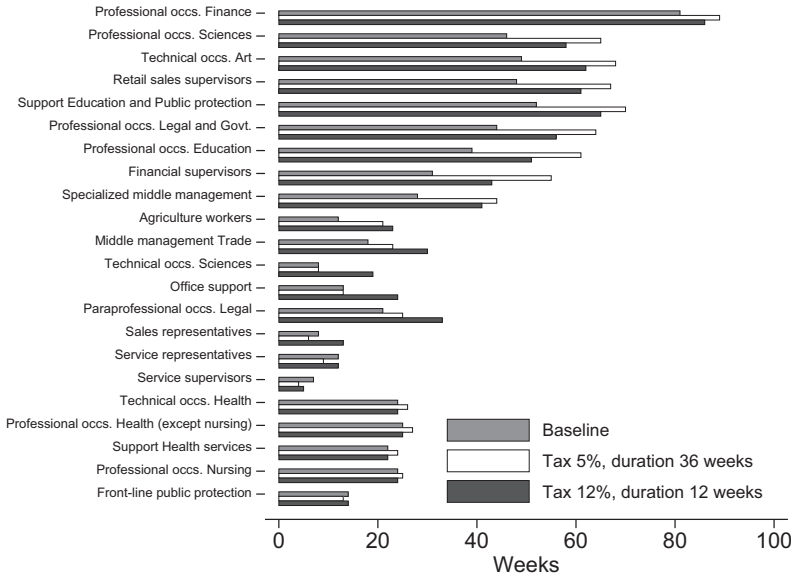


FIGURE 12 Occupations exiting the market with occupation-specific risk

5.2. Role of risk aversion

We now examine the importance of the degree of risk aversion of individuals in driving the dynamics of the pandemic. To study the importance of risk aversion, we assume that the periodic utility function of individuals is linear:

$$u(c_t) = c_t$$

Figure 13 shows the key dynamics of the economy under risk neutrality. For comparison, we also plot the responses under log-linear preferences, which was the baseline case. Figure 13 makes clear that risk aversion is a key parameter underlying the effect of endogenous choices. The responses when individuals are risk neutral are much closer to the responses under exogenous choices. This is intuitively obvious. The greater the risk aversion of the individual, the more she avoids market work when it entails the risk of exposure to infection. This leads to a reduction in the proportion of infected individuals in the economy but amplifies the economic losses of the pandemic. Thus, the tradeoff between the public health and economic costs of the disease is much more pronounced when individuals are risk averse.

5.3. Costs of infection

The main focus of our work is on the salience of the individual’s endogenous response to minimize the exposure risks posed by the pandemic. A key factor in assessing this risk is the impact of infections on earnings potential. During

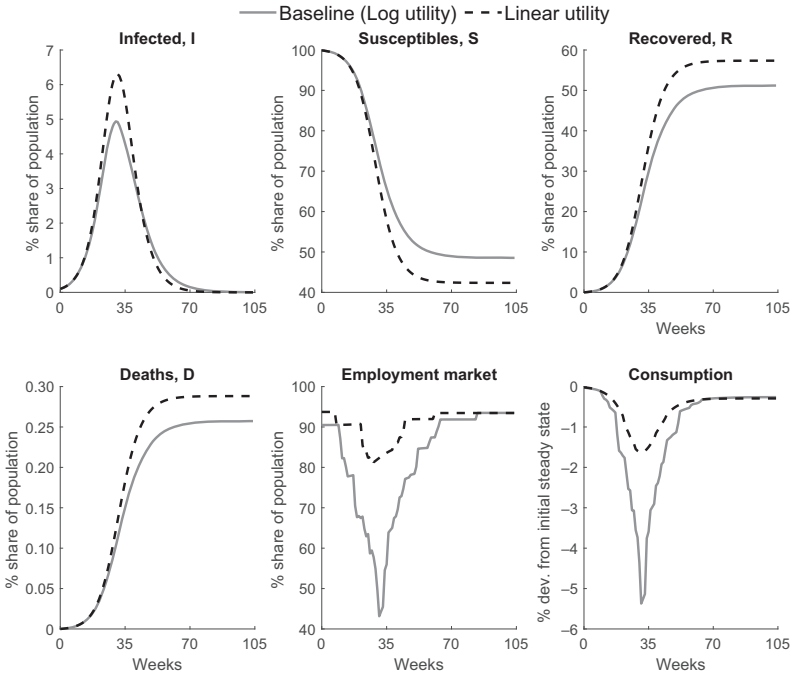


FIGURE 13 Role of risk aversion

the duration of the infection the individual’s productivity declines by $1 - \gamma_I^m$ for market work and $1 - \gamma_I^h$ for home work. As γ_I^m and γ_I^h become smaller, the cost of infection becomes higher.

In our baseline simulations, we assumed that $\gamma_I^m = \gamma_I^h = 0.8$. To examine the sensitivity of results to this parameter, we compute the model responses to two different values: $\gamma_I^m = \gamma_I^h = 0.9$ and $\gamma_I^m = \gamma_I^h = 0.6$. Figure 14 shows the responses. The figure shows that a lower γ_I^m and γ_I^h reduces the peak infection rate by a full percentage point while also delaying the peak. The infection rises more slowly and also declines less rapidly when $\gamma_I^m = \gamma_I^h = 0.6$ relative to the baseline case. As a result, the probability of infection perceived by susceptible individuals, π_t^m , rises and falls more slowly relative to the baseline. From proposition 1, we know that more occupations switch to home work as the probability of infection rises. This can be seen in the slower decline and rise of market employment when $\gamma_I^m = \gamma_I^h = 0.6$ relative to the baseline of $\gamma_I^m = \gamma_I^h = 0.8$. This pattern is mirrored in the path of overall consumption, which declines more slowly but also recovers more slowly over time.

The fact that the probability of infection falls more while the share of market employment declines less under a lower γ_I^m and γ_I^h might seem counterintuitive at first glance. The reason for this is the lower consumption of the infected. Because productivity of the infected is lower, their consumption

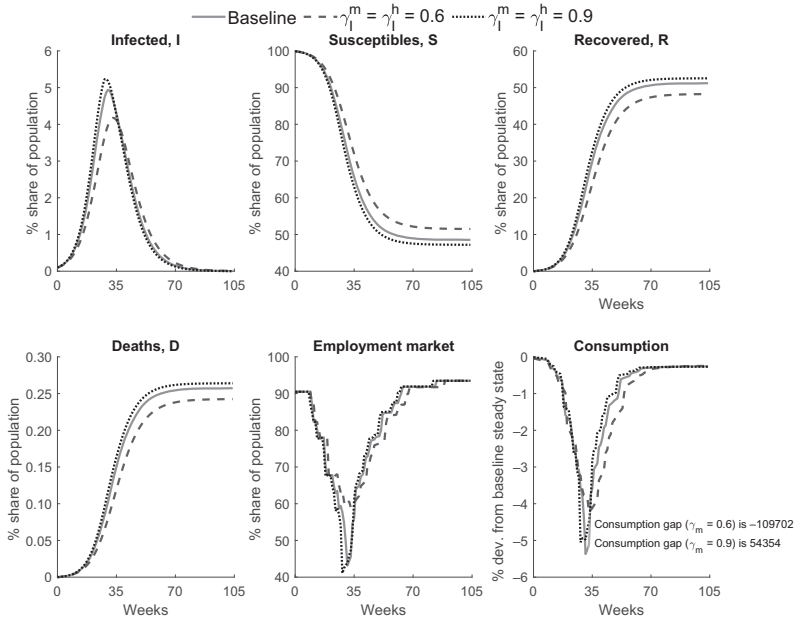


FIGURE 14 Role of productivity when infected

levels are lower now relative to the baseline case. Because the probability of susceptible individuals becoming infected while consuming is proportional to the consumption of the infected, the overall infection probability declines as the productivity of infected falls. The lower risk exposure due to lower consumption of the infected also open up greater space for risk taking by susceptible individuals who respond by withdrawing less from risky market work.

5.4. Vaccines, treatments and health care congestion

In the early stages of the pandemic, it was quite uncertain whether an effective vaccine against the disease could be developed, and how long it would take to produce such a vaccine. While our model is not ideally suited to explore in detail the impact and application of vaccines, we can explore the effect of increases in the probability of vaccine discovery on the choice of home versus market work.

Figure 15 describes the baseline calibration of our model amended to take account of a constant weekly probability of a vaccine discovery. A vaccine discovery allows all susceptible individuals to become recovered without becoming infected. We allow for a constant weekly probability δ_V that a vaccine is found. We look at two cases, one where the probability is $\delta_V = \frac{1}{26}$ (a vaccine discovery within six months) and one with the probability equal to $\delta_V = \frac{1}{13}$ (a vaccine discovery within three months). We also allow for an innovation in medical treatment. This is modelled in a similar fashion by allowing for a probability δ_c that all infected individuals switch to the recovered state. We

use the same probability assumptions for treatment as for vaccines. The simulations show only the response before the vaccine goes into effects. The results indicate that an increase in the likelihood of a vaccine arrival slightly reduces both the exit rate from market work and the consumption decline, but leaves the pandemic dynamics materially unchanged.

Figure 16 shows the impact of congestion in the health care system on the response of market versus home work. We model this by allowing for the mortality rate to be related to the number of infected individuals as follows:

$$\pi_{Dt} = \pi_D + \kappa I_t^2,$$

where κ is calibrated such that the peak mortality rate is 1%. The results show that the effect of health care congestion leads to a major exit from market work and a significant lowering of the peak infection rate, although at the same time the death rate increases quite sharply.

5.5. Fiscal responses to the pandemic

While our main objective in this paper is to explore the endogenous occupational response to pandemic risk, it is important to acknowledge that in many countries, particularly in Canada, governments made large compensatory

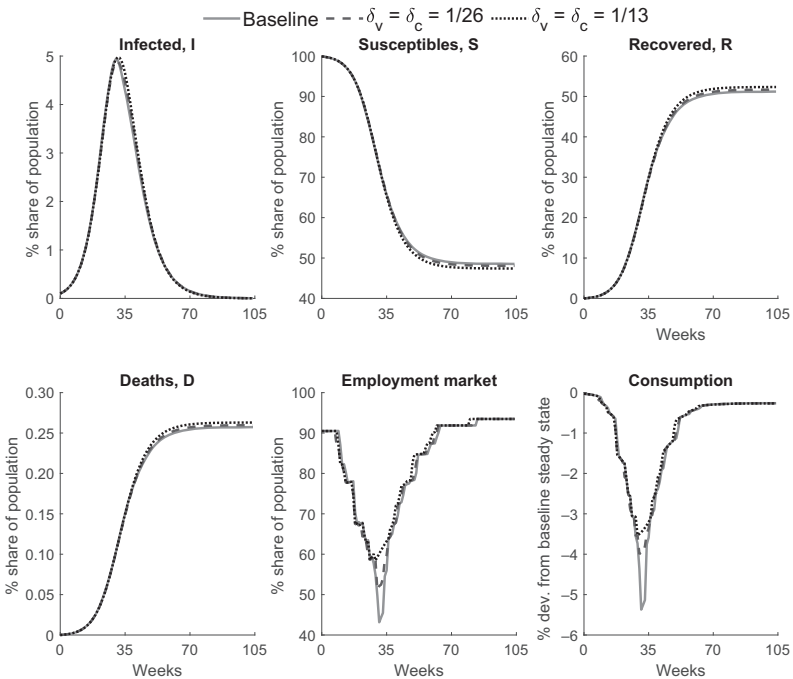


FIGURE 15 Role of vaccine and medical treatment

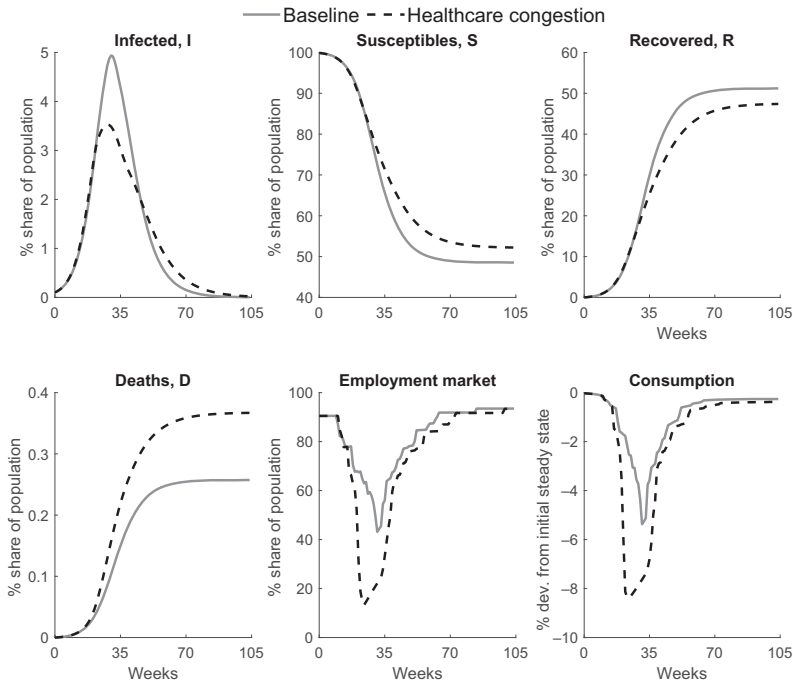


FIGURE 16 Role of congestion in health care

fiscal transfers to job losers. This clearly alleviated the consequence of loss of income due to inability to work in the market. As discussed above, our model is not ideally designed to analyze unemployment dynamics. However, we can explore the implications of fiscal transfers in affecting the choice between market work and home work. To do this, we change the structure of the model to allow for tax rebates that are directed towards occupations whose consumption is low and will lose out most from moving from the market to home work.

Figure 17 shows the impact of a 5% proportional income tax that is rebated to those who were working in the market and whose market wage is below the steady state median wage. Thus, the rebate is applied to individuals working in 18 out of the 40 occupations. The rebate amounts to a substantial increase of 19% and 8.8% in consumption for the lowest and second lowest income quartile, respectively. The tax and rebate is applied only at the onset of pandemic. The rebate reduces the cost of exiting the market for those occupations that have on average a low return to working at home. This has the effect of increasing the exit from the market and slightly reducing the peak of the infection rate as well as reducing the number of overall deaths. At the same time, the fall in total consumption is somewhat less than in the baseline. Overall, however, the fiscal response does not significantly alter the pandemic dynamics relative to our baseline.

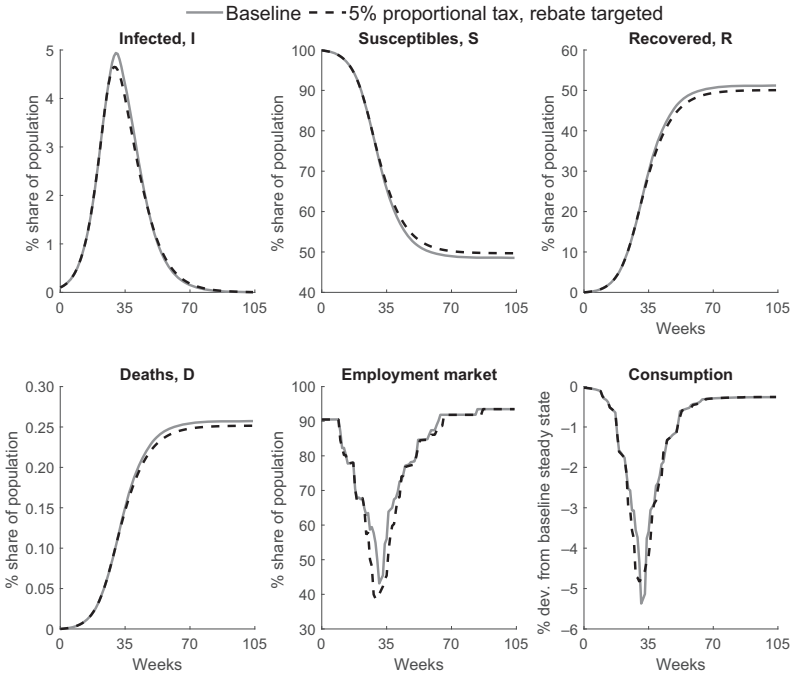


FIGURE 17 Role of progressive tax and rebate

6. Distributional consequences of the pandemic

The coronavirus epidemic has clearly disrupted economic life globally. Along with the aggregate contractionary effects, it has also had huge distributional consequences. Recent work by Chetty et al. (2020) uses credit card payment data to show that the overall consumption decline between January 2020 and June 17, 2020, was 8.9%. However, underneath this overall decline lies huge variation. This is shown in panel (a) of figure 18, which shows that during this period spending by consumers in zip codes with the top quartile of median income declined by 13.3%, while spending by those in zip codes in the bottom quartile of incomes fell by a meagre 2.8%. More broadly, the figure shows that the spending in the richest zip codes declined the most on impact of the pandemic and has shown the slowest recovery. In contrast, spending fell the least in the poorest zip codes and recovered fastest. Panel (b) of figure 18 shows aggregate consumption responses by quartiles of income for our baseline simulation.¹¹ These responses are generated from the version of the model without any lockdowns, or for $\tau_m = 0$.

11 The consumption responses displayed in panel (b) of figure 18 are aggregated over 10 occupations in each quartile, whereas the overall consumption response is aggregated over 40 occupations.

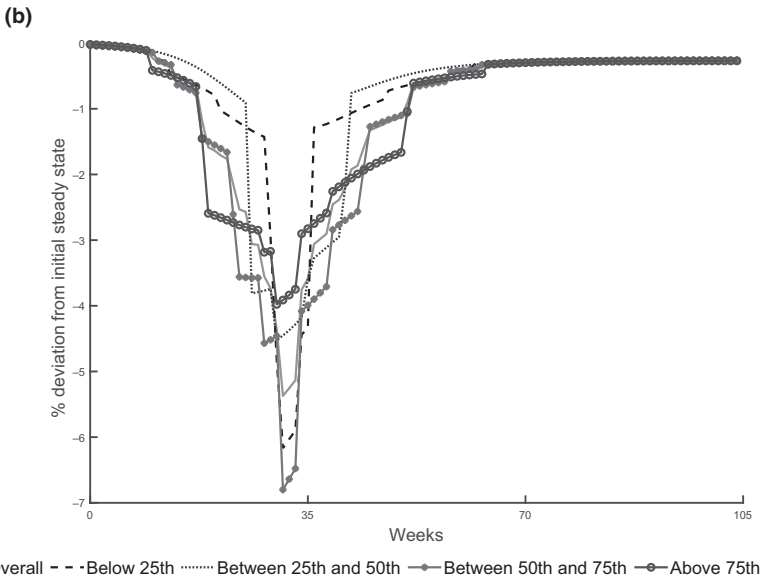
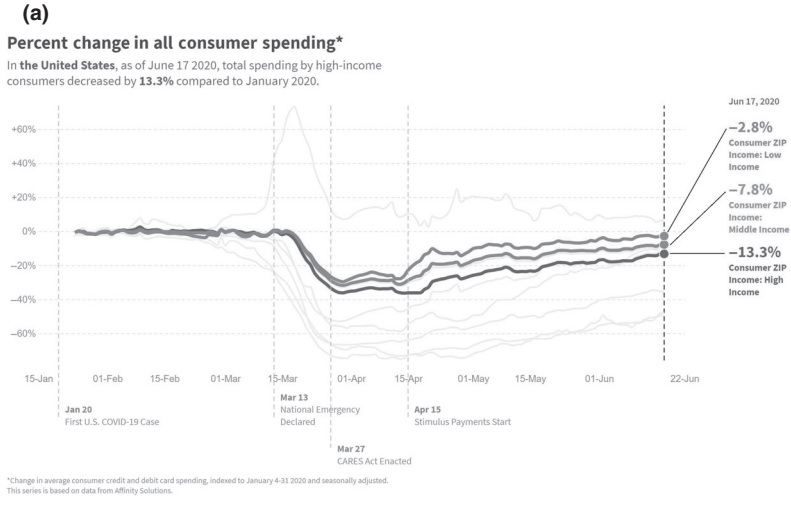


FIGURE 18 Consumption distribution by quartiles
NOTE: The figure in panel (a) can be found in Chetty et al. (2020).

A few features of panel (b) are noteworthy. First, as in the data, the consumption of the lowest income quartile in the model recovers the fastest, while the consumption recovery of the top quartile is the slowest. Second, up to 16 weeks after the start of the infection, the consumption declines are the largest for the richest quartile followed by the third quartile. The consumption declines for the two poorest quartiles are the smallest until 16 weeks though

the poorest quartile consumption decline is marginally greater than that for the second quartile. This overall pattern fits the qualitative features in panel (a). This is because the lower income groups are mostly in occupations where the home wage premium is negative and large.

7. Conclusion

We have developed a simple macro model that, in conjunction with a standard SIR model, gives rise to endogenous cycles in both infections and output. The model has the standard interaction between economic and infection outcomes. Economic market activity increases infection risks, while lockdowns reduce infections at the cost of deeper recession. The key innovation of our model is that it focuses on occupations and the endogenous choice by individuals whether to work in the market or from home.

The quantitative results from the model calibrated to Canadian data from the 2016 census suggest that the endogenous response of private agents to market risk can be significant with the weekly infection rate being two percentage points lower at peak than when private supply of market work is exogenously given. Correspondingly, the output loss at the deepest point (the trough) is four percentage points greater under the endogenous choice case.

Our results have a key policy implication. Because individuals are willing to pay to insure themselves against infection risk by withdrawing from the market, opening up the economy may not lead to a sharp recovery. Until people develop confidence about the safety of market activity, they will continue to self-isolate. While self-isolation may not be an option for people in occupations that are intensive in social interactions, their returns from market activity are often dependent on the market consumption activity by others. If a significant share of people withdraw from consuming in the market, upstream suppliers of these services and goods will suffer income and job losses as a result. Our results suggest that this effect can be quantitatively high. Dealing with this confidence problem may involve spending on intensive testing rather than direct fiscal infusions into the economy.

The model also provided interesting insights into the effects of lockdowns. We found that intense but shorter lockdowns do not materially change the infection dynamics or the cumulated number of deaths except for marginally postponing the peak of the infection. Longer and possibly less intensive lockdowns, on the other hand, may be more effective in lowering the infection peak. Moreover, the shorter but more intensive lockdown has higher cumulated output costs relative to the longer but milder lockdown. Thus, both the health and economic costs of the pandemic reduce with a milder but longer lockdown.

Supporting information

Supplementary material accompanies this article.

References

- Acemoglu, D., V. Chernozhukov, I. Werning, and M. Whinston (2020) "Optimal targeted lockdowns in a multi-group SIR model," NBER working paper no. 27102
- Adams-Prassl, A., T. Boneva, M. Golin, and C. Rauh (2020) "Inequality in the impact of the coronavirus shock: Evidence from real time surveys," *Journal of Public Economics* 189, 104245
- Alvarez, F.E., D. Argente, and F. Lippi (2020) "A simple planning problem for COVID-19 lockdown," NBER working paper no. 26981
- Atkeson, A. (2020) "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios," NBER working paper no. 26867
- Aum, S., S. Y. T. Lee, and Y. Shin (2021) "Inequality of fear and self-quarantine: Is there a trade-off between GDP and public health?," *Journal of Public Economics* 194, 104354
- Baqae, D., and E. Farhi (2020) "Supply and demand in disaggregated Keynesian economies with an application to the COVID-19 crisis," NBER working paper no. 27152
- Bartik, A.W., Z. B. Cullen, E. L. Glaeser, M. Luca, and C. T. Stanton (2020) "What jobs are being done at home during the COVID-19 crisis? Evidence from firm-level surveys," NBER working paper no. 27422
- Bayer, C., B. Born, R. Luetticke, and G. J. Müller (2020) "The Coronavirus Stimulus Package: How large is the transfer multiplier?," CEPR discussion paper no. 14600
- Baylis, P., P.-L. Beauregard, M. Connolly, N. Fortin, D. A. Green, P. Gutierrez Cubillos, S. Gyetvay, C. Haeck, T. Laura Molnar, G. Simard-Duplain, H. E. Siu, M. TeNyenhuis, and C. Warman (2020a) "The distribution of COVID-19 related risks," NBER working paper no. 27881
- (2020b) Replication file for "The distribution of COVID-19 related risks" (NBER working paper no. 27881). <https://github.com/pbaylis/vse-risk-tool>
- Beland, L.-P., A. Brodeur, D. Mikola, and T. Wright (2020) "The short-term economic consequences of COVID-19: Occupation tasks and mental health in Canada," IZA discussion paper no. 13254
- Bloom, N., J. Liang, J. Roberts, and Z. J. Ying (2015) "Does working from home work? Evidence from a Chinese experiment," *Quarterly Journal of Economics* 130(1), 165–218, ISSN 0033-5533
- Bodenstein, M., G. Corsetti, and L. Guerrieri (2020) "Social distancing and supply disruptions in a pandemic," *Covid Economics - Vetted and Real-Time Papers* (19), 1–52
- Brodeur, A., D. Gray, A. I. Suraiya, and J. Bhuiyan (2020) "A literature review of the economics of COVID-19," IZA discussion paper no. 13411
- Brynjolfsson, E., J. Horton, A. Ozimek, D. Rock, G. Sharma, and H.-Y. TuYe (2020) "COVID-19 and remote work: An early look at US data," NBER working paper no. 27344
- Brzezinski, A., V. Kecht, and D. Van Dijke (2020) "The cost of staying open: Voluntary social distancing and lockdowns in the US," Oxford discussion paper no. 910

- Carroll, C.D., E. Crawley, J. Slacale, and M. N. White (2020) “Modeling the consumption response to the CARES Act,” NBER working paper no. 27876
- Caselli, F., F. Grigoli, W. Lian, and D. Sandri (2020) “Protecting lives and livelihoods with early and tight lockdowns,” IMF working paper no. 20/234
- Chen, M.K., Y. Zhuo, M. De La Fuente, R. Rohla, and E. F. Long (2020) “Causal estimation of stay-at-home orders on SARS-CoV-2 transmission,” mimeo, UCLA Anderson School of Management
- Chernozhukov, V., H. Kasahara, and P. Schrimpf (2021) “Causal impact of masks, policies, behavior on early COVID-19 pandemic in the U.S.,” *Journal of Econometrics* 220(1), 23–62, ISSN 18726895
- Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, and 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 no. 27431
- Coibion, O., Y. Gorodnichenko, and M. Weber (2020) “How did U.S. consumers use their stimulus payments?,” NBER working paper no. 27693
- Ding, W., R. Levine, C. Lin, and W. Xie (2020) “Social distancing and social capital: Why U.S. counties respond differently to COVID-19,” NBER working paper no. 27393
- Dingel, J., and B. Neiman (2020) “How many jobs can be done at home?,” *Journal of Public Economics* 189, 104235
- Eichenbaum, M.S., S. Rebelo, and M. Trabandt (2021) “The macroeconomics of epidemics,” *Review of Financial Studies* 34(11), 5149–87
- Ferguson, N.M., D. A. T. Cummings, C. Fraser, J. C. Cajka, P. C. Cooley, and D. S. Burke (2006) “Strategies for mitigating an influenza pandemic,” *Nature* 442 (7101), 448–52
- Goolsbee, A., and C. Syverson (2021) “Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020,” *Journal of Public Economics* 193, 104311
- Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2020) “Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages?,” NBER working paper no. 26918
- Gupta, S., T. D. Nguyen, F. Lozano Rojas, S. Raman, B. Lee, A. Bento, K. I. Simon, and C. Wing (2020) “Tracking public and private responses to the COVID-19 Epidemic: Evidence from state and local government actions,” NBER working paper no. 27027
- Hensvik, L., T. Le Barbanchon, and R. Rathelot (2020) “Which jobs are done from home? Evidence from the American Time Use Survey,” CEPR discussion paper no. 14611
- Hur, S., and M. Jenuwine (2020) “Lessons on the economics of pandemics from recent research,” Federal Reserve Bank of Cleveland Economic Commentary, no. 2020-11. <https://doi.org/10.26509/frbc-ec-202011>
- Jones, C., T. Philippon, and V. Venkateswaran (2020) “Optimal mitigation policies in a pandemic: Social distancing and working from home,” NBER working paper no. 26984
- Kaplan, G., B. Moll, and G. L. Violante (2020) “The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US,” NBER working paper no. 27794

- Krueger, D., H. Uhlig, and T. Xie (2020) "Macroeconomic dynamics and reallocation in an epidemic," NBER working paper no. 27047
- Maloney, W., and T. Taskin (2020) "Determinants of social distancing and economic activity during COVID-19: A global view," *COVID Economics – Vetted and Real-Time Papers* (13), 157–77
- Mandavilli, A. (2020, June 27) "Actual coronavirus infections vastly undercounted, C.D.C. data shows," *The New York Times*. Available at www.nytimes.com/2020/06/27/health/coronavirus-antibodies-asymptomatic.html
- New York Times, The (2021) "Coronavirus (COVID-19) data in the United States [repository]," Available at <https://github.com/nytimes/covid-19-data>
- Regus (2017) "The workplace revolution," Regus. Available at www.regus.com/work-us/en-us/the-workplace-revolution/
- Statistics Canada (2016) 2016 Census of Population data tables, Statistics Canada catalogue no. 98-400-X2016304. Available at www150.statcan.gc.ca/n1/en/catalogue/98-400-X2016304
- (2021) "Labour force characteristics by occupation, monthly, unadjusted for seasonality (x 1,000)." Table 14-10-0296-01. <https://doi.org/10.25318/1410029601-eng>
- Woodford, M. (2020) "Effective demand failures and the limits of monetary stabilization policy," NBER working paper no. 27768