Department of Economics

DISCUSSION PAPER SERIES IN ECONOMICS

DP No. 32

Epidemics: A Tale of Two Workers

June 2020

Rahul Nath

https://www.ashoka.edu.in/ecodp

Plot No. 2, Rajiv Gandhi Education City,
National Capital Region
P.O.Rai,Sonepat Haryana-131029

Epidemics: A Tale of Two Workers*

Rahul Nath[†]

June, 2020

Abstract

This paper shows that the labour market opportunities available to an agent has a significant bearing on how that agent experiences the outbreak of an epidemic. I consider two types of labour (i) market labour that can only produce output in close physical proximity, and (ii) remote labour that can produce output at a distance. This paper develops a Two Agent New Keynesian model extended to include an epidemic bloc and dual feedback between economic decisions and the evolution of the epidemic. I show that an agent restricted to only supply market labour experiences higher death rates vis-à-vis their share of the population, and suffers larger declines in labour and consumption over the course of the epidemic. Post-epidemic, these agents are significantly worse off than their counterparts who have the opportunity to work from home and hence a more unequal society emerges. I then show that simple containment policies, while leading to larger losses in economic prosperity as measured by output loss, can significantly reduce death rates across the population, bring the death rates of the two groups closer together, and reduce the inequality that emerges post epidemic.

Keywords: Macroeconomic Labour, Heterogeneous Agents, New Keynesian DSGE, Epidemics, Lockdown, Covid-19

JEL Codes: E120, E240, E250, E320, E660, I140, J210, J220

*I would like to thank Judy Goh, Srijita Ghosh, Suraj Shekhar and Swagata Bhattacharjee for their useful discussions, feedback and comments.

†Ashoka University, E-mail address: rahul.rb.nath@gmail.com

A common refrain during the current Covid-19 pandemic is that it is a great leveller since the disease does not discriminate among those who contract the disease. From a purely scientific point of view it is true that the Covid-19 virus does not discriminate at the point of infection, i.e. anyone is capable of getting the disease. However, contracting the disease is a whole different ballgame since contracting the disease requires interaction, often in a social setting, and here the playing field is much less level. Over the course of the current pandemic there has been mounting evidence that some groups are over-represented among those that contract the disease. Doctors, nurses, other healthcare workers, those working in services deemed essential, minority groups and migrant workers are among the groups that make up a disproportionately higher proportion of total infected cases.

As various containment measures were initially introduced around the world, work has shifted to remote models where possible. There are many occupations, however, that cannot be easily moved to remote work. Consider a nurse in a hospital or care facility, or a fire fighter, these workers are far less likely to be able to work remotely than a lawyer, a teacher, or an economist. Many occupations have characteristics that make them difficult to carry out except in close proximity to other people. In a pandemic where disease spreads by proximity of social contact this exposes workers in such occupations to much higher risk of ultimately contracting the disease. The present paper seeks to understand why some groups are more likely to contract a given disease based on the economic opportunities available to them. Specifically, this paper focuses on whether agents may have differing experiences of a pandemic depending on the labour market opportunities available to them, i.e. whether the opportunity to work remotely is a possibility. Through the study of labour market opportunities the paper seeks to shed light on a more general question of how entrenched differences between groups that impact their economic opportunities affect the experience of the epidemic within these groups, and how such pre-epidemic differences may help to predict the level of inequality that materialises post-epidemic.

I find that when agents are restricted in their labour market opportunities, they experience the epidemic differently. Correspondingly, there is an increase in economic inequality between the two groups. Agents who can only work in market labour are worse off both in terms of death rates and economic outcomes. They experience significantly higher death rates in excess of their share of the population arising from a higher exposure to infection risk from only being able to engage in market labour. These agents also face worse economic outcomes via larger declines in labour supply and consequently lower consumption. I find that containment measures aimed at reducing the spread of the disease can reduce this inequality. All of the containment measures studied significantly reduce the death rates of all agents at the cost of slightly higher decline in output. I find that even containment measures with early exit, i.e. not reducing the infected populations to zero, can still have significant impact albeit even though they see a second wave of the epidemic. In particular I find that non-symmetric containment policies, i.e. those that treat the two groups differently, is the most preferred option from a purely economic standpoint as it leads to minimal output decline vis-*à*-vis the laissez-faire policy scenario. There are three key contributions of this paper to the literature. First is the study of how groups differing in their economic opportunities before the epidemic begins have different experiences of an epidemic. Second, I develop a simple New Keynesian framework that allows for mutual feedback between epidemics and economic behaviour. And finally, I explore how pre-epidemic group characteristics can be used to better understand the evolution of an epidemic and the design of containment measures.

I study the central question of this paper by developing a simple Two Agent New Keynesian DSGE model. The model is populated with two types of households who are identical in all respects except the labour market opportunities available to them. Both households can engage in market labour (i.e. labour requiring social proximity to others), but one household has the opportunity to also supply their labour remotely. This model is augmented with an epidemic bloc using an extension of the Susceptible-Infected-Recovered (SIR) epidemiological framework of Kermack and McKendrick (1927) that allows for dual feedback between the evolution of the epidemic and macroeconomic decisions. That is, the evolution of the epidemic affects how agents make their optimal decisions and agent decisions affect the evolution of the epidemic by altering the amount of social contact they have through the supply of market labour. This dual feedback is introduced by altering the utility function to allow agents to incorporate the current state of the epidemic into their decisions, and by altering the transmission rate of the disease to take into account social interactions.

During an epidemic, household behaviour changes to reduce engagement in economic activities that involve social proximity, out of fear of infection. The only form of social economic activity considered in this paper is market labour. The undesirability of working in market labour during the pandemic is captured via an additional epidemic factor that increases the disutility agents experience from engaging in market labour. The epidemic introduces tension into the model as some agents can substitute market labour for remote labour while others may not. All other things equal, those agents that can only supply market labour end up being more exposed to the disease if they seek to maintain their labour supply at the pre-epidemic level.

The standard SIR framework of Kermack and McKendrick (1927) does not consider how the evolution of an epidemic may respond to changing social behaviour during the epidemic. During an epidemic, as agents choose to spend less time in social economic activity they effectively reduce the transmission of the disease. The basic SIR model is extended to include this behavioural response by introducing an aggregate exposure variable that depends on how agents change their supply of market labour during the pandemic. This allows labour market decisions to impact the spread of the epidemic. The introduction of behavioural responses is not new and has been explored in various economic models that deal with disease spread. An early example of this approach is Kremer (1996) in studying the spread of AIDS, and more recently from the countless current papers studying the Covid-19 epidemic.

Since the onset of the Coivd-19 crisis there has been a very quickly growing body of work looking at the economics of epidemics. Before moving on to the main body of the paper I briefly survey the papers most relevant to this work below.

This paper is most closely related to the excellent work by Eichenbaum et al. (2020). Eichenbaum et al. (2020) consider how economic decisions affect the evolution of an epidemic, they show that there is an inevitable trade-off between controlling the evolution of an epidemic and the severity of the economic decline. The households in Eichenbaum et al. (2020) are identical pre-epidemic and are then differentiated during the epidemic by health status. This paper departs from Eichenbaum et al. (2020) by considering groups differentiated by economic opportunities that exist regardless of the presence of an epidemic. That is, I focus on entrenched differences due to the nature of the agents occupation rather than those that arise from the epidemic in order to study group differences.

Preliminary work presented by Kaplan et al. (2020) studies the macroeconomic impact of epidemics in a fully Heterogeneous Agent New Keynesian framework. They allow for remote labour within the model in a more elaborate multi-sector setting where occupations are differentiated along various dimensions. As of the date of writing no results had been circulated. This paper differs from Kaplan et al. (2020) by studying a much simplified heterogeneous agent model to allow for a clearer understanding of how epidemics evolve and the macroeconomic consequences when labour can only vary between two types. The results of Kaplan et al. (2020), when available, would provide useful insight into how differing labour along other dimensions might play a role in modifying the results of this paper.

Multiple agent types differentiated by economic opportunities are also a feature of both Bodenstein et al. (2020) and Glover et al. (2020). Glover et al. (2020) differentiate agents along age (old and young), health status and employment sector (basic and luxury), while in Bodenstein et al. (2020) agents supply labour inelastically either in a labour intensive production sector or a production sector that uses capital in production. Critically, in both papers, labour can only be supplied inelastically and in the market. In this paper I abstract away from all of these important dimensions to focus on the impact of the epidemic on the supply of labour which is differentiated in whether it can produce output without the need for close social contact necessitated by being present on-site.

The remainder of this paper is structured as follows. Section 1 describes the model, Section

2 discusses calibration and simulation methods used to solve the model, Section 3 analyses the results of the benchmark model, Section 4 considers how containment measures affect the results and Section 5 concludes.

1 The Model

The model studied in this paper is a Two Agent New Keynesian (TANK) model where the households differ in the types of labour that they may supply. The two types of labour in the economy are 'market' labour and 'remote' labour which differ in how they produce output: 'market' labour must be physically present on-site to produce output while 'remote' labour can produce output without being physically present. There is clearly a lot of heterogeneity within these two broad classes and the 'remotability' of labour is a continuous variable rather than a discrete binary variable. This paper abstracts from both of these dimensions by considering a single type of labour that can be supplied by agents that differ in how they supply this labour.

The set-up of firms in the economy follows the standard practice in the New Keynesian literature. Namely there are two types of firms in the economy - a final goods firm and a continuum of intermediate goods firms. The intermediate goods firms each produce a differentiated good which endues them with some degree of market power. These intermediate firms use both types of labour in the production of their output and face costs to their adjustment of nominal price in the spirit of Rotemberg costs. The final goods firm aggregates the output of intermediate firms and sells this aggregate output to households.

The model is closed with a central bank that controls the nominal interest rate and sets monetary policy according to a simple Taylor Rule.

There are three exogenous shocks that hit the economy - technology shocks, monetary policy shocks and epidemic shocks. The first two shocks are standard shocks in the NK literature, while the epidemic shock is introduced to study how the presence of an epidemic affects the dynamics of the economy. Epidemic shocks have a two-way impact whereby the epidemic impacts economic activity and economic activity affects the spread of the disease. The epidemic shock affects economic activity by altering household disutility from providing 'market' labour hence impacting their desire to supply 'market' labour. This then impacts the dynamics of epidemic by altering the exposure of agents to the disease.

The remainder of this section discusses the model set-up in detail and is divided into 5 sections. The first two sections describe the core macroeconomic bloc, i.e. households and firms. The third section describes the epidemic bloc of the economy. The final two sections describe the Central Bank policy rule, and aggregation and equilibrium in the model.

1.1 Households

There exist a continuum of households indexed on the unit interval $j \in [0,1]$ and make all of their decisions at the beginning of the period. These households are split into two groups depending on where on the unit interval they are indexed and they differ only in the labour supply opportunities available to them. Type-1 households fall in the interval $j = (\theta_t, 1]$ and supply both market and remote labour, i.e. they can decide whether or not to be physically present during the production of output. Type-2 households fall in the interval $j = [0, \theta_t)$ and supply only market labour, i.e. they must be physically present in the firm to produce output.

The utility of the household household of Type-j is defined over the consumption of aggregate good $(c_{t,j})$, supply of market labour $\left(n_{t,j}^M\right)$, and supply of remote labour $\left(n_{t,j}^R\right)$.[1]. Specifically I use the GHH form of Greenwood et al. (1988) which removes the presence of any Type-j level wealth effect from the labour supply decision so that labour supply only depends on the wage rate. The utility function takes general form,

$$u\left(c_{t,j}; n_{t,j}^M; n_{t,j}^R\right) = \frac{\left[c_{t,j} - \chi_j^M \Gamma_t \frac{\left(n_{t,j}^M\right)^{1+\psi}}{1+\psi} - \chi_j^R \frac{\left(n_{t,j}^R\right)^{1+\psi}}{1+\psi}\right]^{1-\sigma}}{1-\sigma} \tag{1.1}$$

where σ captures the degree of risk aversion, ψ is the inverse of the Frisch Elasticity of Labour, χ_j^M measures the disutility of providing market labour for the Type-j household and χ_j^R measures the disutility of providing remote labour for the Type-j household. The parameter, Γ_t captures the impact of the epidemic shock on the supply of market labour and it defined as,

$$\Gamma_t = 1 + \widetilde{\beta}_{t,j} \mathcal{S}_t \frac{\mathcal{I}_t}{\mathcal{N}_t} \tag{1.2}$$

where $\widetilde{\beta}_{t,j}$ is the belief of Type-j agents about the 'effective' transmission rate of the disease.[2] In this paper I assume that the belief is symmetric, i.e. both agents have the same $\widetilde{\beta}_{t,j}$. Further to the symmetric belief assumption, I also assume that $\widetilde{\beta}_{t,j} \mathcal{S}_t \frac{\mathcal{I}_t}{\mathcal{N}_t} = \beta_t^- \mathcal{S}_t^- \frac{\mathcal{I}_t^-}{\mathcal{N}_t^-}$, i.e. the agent forms a belief based on the information available to them at the beginning of the period.[3]

The general budget constraint for households is given by,

$$c_{t,j} + \frac{B_{t+1,j}}{R_t^n P_t} = \frac{W_t^M}{P_t} n_{t,j}^M + \kappa \frac{W_t^R}{P_t} n_{t,j}^R + \frac{B_{t,j}}{P_t} + D_{t,j} \tag{1.3}$$

where $B_{t,j}$ is holdings of nominal bonds, R_t^n is the nominal interest rate set by the central bank, $\frac{W_t^M}{P_t}$ is the real wage for market labour, $\frac{W_t^R}{P_t}$ is the real wage for remote labour, and $D_{t,j}$ is the

[1]Note that for Type-2 household $n_{t,j}^R = 0$

[2]The 'effective' transmission rate is defined in the section on the epidemic bloc of the model.

[3]The study of more complex belief structures is beyond the scope of the present paper and left for future work.

dividend received by virtue of household ownership of firms. Wages are determined in perfectly competitive labour markets and both households and firms take wages as given. The parameter $\kappa \leq 1$ captures the idea that working from home entails a cost in terms of lost wages. This parameter captures all aspects of the cost of working remotely (e.g. lost productivity, set-up costs, etc.) and introduces a wedge that makes market labour more desirable all other things equal.

The Type-j household therefore solves the following optimisation problem,

$$\max E_t\left[\sum_{h=0}^{\infty}\beta^h u\left(c_{t+h,j};n^M_{t+h,j};n^R_{t+h,j}\right)\right]$$

s.t.

$$c_{t+h,j}+\frac{B_{t+h+1,j}}{R^n_{t+h}P_{t+h}}=\frac{W^M_{t+h}}{P_{t+h}}n^M_{t+h,j}+\kappa\frac{W^R_{t+h}}{P_{t+h}}n^R_{t+h,j}+\frac{B_{t+h,j}}{P_{t+h}}+D_{t+h}$$

This yields the following conditions,

$$\frac{W^M_t}{P_t}=\chi^M_j\Gamma_t\left(n^M_{t,j}\right)^{\psi} \tag{1.4}$$

$$\frac{W^R_t}{P_t}=\frac{\chi^R_j}{\kappa}\left(n^R_{t,j}\right)^{\psi} \tag{1.5}$$

$$1=E_t\left[\beta\frac{u_{c,t+1,j}}{u_{c,t,j}}\frac{R^n_t}{\pi_{t+1}}\right] \tag{1.6}$$

which have the standard interpretations as labour supply conditions and consumption Euler Equations.

1.2 Firms

1.2.1 Final Goods Firms

The output of intermediate goods firms, $y_t(i)$, is bought by perfectly competitive final goods firms which costlessly aggregate the output. This aggregate output is sold to households as an aggregate consumption good.

The final goods firms aggregate output using the Dixit-Stiglitz aggregator,

$$y_t=\left[\int_0^1 y_t(i)^{\frac{\varepsilon^p-1}{\varepsilon^p}}di\right]^{\frac{\varepsilon^p}{\varepsilon^p-1}}, \tag{1.7}$$

where $\varepsilon^p>0$ measures the degree of substitutability between different goods. The final goods firms maximise their profits leading to the standard demand function for the intermediate goods firms,

$$y_t\left(i\right)=\left[\frac{P_t\left(i\right)}{P_t}\right]^{-\varepsilon^p}y_t, \tag{1.8}$$

where y_t is the aggregate demand and P_t is the aggregate price level defined as,

$$P_t=\left[\int_0^1 P_t\left(j\right)^{1-\varepsilon^p}di\right]^{\frac{1}{1-\varepsilon^p}}. \tag{1.9}$$

1.2.2 Intermediate Goods Firms

There exist a continuum of monopolistically competitive intermediate goods firms indexed on the unit interval $i\in[0,1]$ each of whom produce a differentiated good determined by their index. The intermediate goods firm produces output by employing both market and remote worker types, for which it pays market determined wages, and are subject to nominal rigidities in changing prices *à la* Rotemberg (1982). I assume that the cost of adjusting prices is an intangible cost that enters the firms optimisation problem as a form of 'disutility', i.e. it doesn't affect cash flow. The profit of the firm is given by,

$$\frac{P_{t+h,i}y_{t+h,i}}{P_{t+h}}-\frac{W^M_{t+h}}{P_{t+h}}n^M_{t+h,i}-\frac{W^R_{t+h}}{P_{t+h}}n^R_{t+h,i}-\frac{\chi^P}{2}\left(\frac{P_{t+h,i}}{P_{t+h-1,i}}-1\right)^2y_{t+h} \tag{1.10}$$

where $n^M_{t,i}$ and $n^R_{t,i}$ are the aggregate market and remote labour employed by the firm. The parameter χ^P determines the strength of the disutility arising from adjusting prices, and hence the degree of price stickiness.

The production function of the firm is given by a Cobb-Douglas production function defined over labour input only,

$$y_{t,i}=z_t\left(n^M_{t,i}\right)^{\alpha_M}\left(n^R_{t,i}\right)^{\alpha_R}; \qquad \alpha_M+\alpha_R=\alpha \tag{1.11}$$

where α_M is the income share of market labour in the production of output, and α_R is the income share of remote labour in the production of output. In order to be consistent with aggregate data the restriction $\alpha_M+\alpha_R=\alpha$ is imposed, where α is the income share of labour in the production of output.

This technology is assumed to be subject to non-idiosyncratic shocks to productivity, z_t. Productivity follows an AR(1) process in logs, i.e.

$$\ln z_{t+1}=\rho\ln z_t+\upsilon^y_t \tag{1.12}$$

where ρ measures the persistence of the shock and $\upsilon^y_t\sim N\left(0,\sigma^2_z\right)$ is a random shock. The firm is owned by both households and so uses an aggregate discount factor to discount profits given by,

$$m_{t,t+1} = (1-\theta_t)\frac{\beta\Lambda_{t+1,1}}{\Lambda_{t,1}} + \theta_t\frac{\beta\Lambda_{t+1,2}}{\Lambda_{t,2}} \tag{1.13}$$

The firm maximises profit subject to production function and demand, i.e.

$$\max E_t\left\{\sum_h^\infty m_{t,t+h}\left[\frac{P_{t+h,i}y_{t+h,i}}{P_{t+h}} - \frac{W^M_{t+h}}{P_{t+h}}n^M_{t+h,i} - \frac{W^R_{t+h}}{P_{t+h}}n^R_{t+h,i} - \frac{\chi^P}{2}\left(\frac{P_{t+h,i}}{P_{t+h-1,i}}-1\right)^2 y_{t+h}\right]\right\}$$

s.t.

$$y_{t+h,i} = \left[\frac{P_{t+h,i}}{P_{t+h}}\right]^{-\varepsilon_P} y_{t+h}$$

$$y_{t+h,i} = z_{t+h}\left(n^M_{t+h,i}\right)^{\alpha_M}\left(n^R_{t+h,i}\right)^{\alpha_R}$$

Noting that firms are ex-post identical in under Rotemberg pricing, the firms problem leads to the following two equilibrium conditions,

$$E_t\left[m_{t,t+1}\chi^P\pi_{t+1}(\pi_{t+1}-1)y_{t+1}\right] - \chi^P\pi_t(\pi_t-1)y_t = (\varepsilon_P-1)y_t - \frac{W^M_t}{P_t}\frac{\varepsilon_P y_t}{f_{M,t}} \tag{1.14}$$

$$\frac{W^M_t}{f_{M,t}} = \frac{W^R_t}{f_{R,t}} \Rightarrow \frac{W^M_t}{W^R_t} = \frac{\alpha_M}{\alpha_R}\frac{n^R_t}{n^M_t} \tag{1.15}$$

where the first is the Phillips Curve, and the second requires that firms hire each labour type until their effective marginal costs are equalised.

1.3 The Epidemic Bloc

The epidemic bloc of the model is an extension of the standard SIR epidemic model that allows for aggregate exposure to respond endogenously to economic activity, and so modify disease transmission. At any point in time an agent can be in one of five states: Susceptible ($\mathcal{S}_t$), Exposed ($\mathcal{E}_t$), Infectious ($\mathcal{I}_t$), Recovered ($\mathcal{R}_t$), and Dead ($\mathcal{D}_t$). The spread of the epidemic is described by the following system of equations,

$$\mathcal{S}_{t+1} = \mathcal{S}_t - \beta_0\mathcal{S}_t\frac{\mathcal{I}_t}{\mathcal{N}_t}\mathcal{X}_t \tag{1.16}$$

$$\mathcal{E}_{t+1} = \mathcal{E}_t - \lambda_{\mathcal{E}}\mathcal{E}_t + \beta_0\mathcal{S}_t\frac{\mathcal{I}_t}{\mathcal{N}_t}\mathcal{X}_t \tag{1.17}$$

$$\mathcal{I}_{t+1} = \mathcal{I}_t - \lambda_{\mathcal{I}}\mathcal{I}_t + \lambda_{\mathcal{E}}\mathcal{E}_t \tag{1.18}$$

$$\mathcal{R}_{t+1} = \mathcal{R}_t + (1-\gamma)\lambda_{\mathcal{I}}\mathcal{I}_t \tag{1.19}$$

$$\mathcal{D}_{t+1} = \mathcal{D}_t + \gamma\lambda_{\mathcal{I}}\mathcal{I}_t \tag{1.20}$$

$$\mathcal{N}_{t+1} = \mathcal{S}_{t+1} + \mathcal{E}_{t+1} + \mathcal{I}_{t+1} + \mathcal{R}_{t+1} \tag{1.21}$$

where β_0 is the transmission rate for the disease, $\mathcal{X}_t$ is aggregate exposure, γ is the death rate, and λ_j $j = \{E, I\}$ is the transition rate out of the respective states. The parameter β_0 is referred to as the 'basic' transmission rate, i.e. the rate of transmission were agents not to respond endogenously. Our focus is on aggregate exposure, $\mathcal{X}_t$, as this modifies the rate at which agents enter the disease states, and $\beta_t = \beta_0 \mathcal{X}_t$ is referred to as the 'effective' transmission rate. Once an agent enters the Exposed state they move mechanically, as defined by the transmission rates, through the states until they exit as either Recovered or Dead. Agents can only infect other agents while in the Infectious state.

Labour supply in market labour is the only activity that requires interacting with other agents in close proximity. Hence, the aggregate level of exposure is defined using the time spent in market labour by agents relative to their steady states, i.e.

$$\mathcal{X}_t = \sum_{j \in \mathcal{J}} w_{t,j} \frac{n_{t,j}^M}{\overline{n}_j^M}; \ \sum_{j \in \mathcal{J}} w_{t,j} = 1 \tag{1.22}$$

where $\mathcal{J}$ is the set of agent types, and $w_{t,j}$ is the weight of Type-j in the economy. In the simple two agent set-up in this paper $w_{t,1} = 1 - \theta_t$ and $w_{t,2} = \theta_t$, however this measure can be extended to any arbitrary set of agent types.

This measure of aggregate exposure captures two salient features of how economic actions affect disease spread. The first, that individual group level actions can modify overall exposure can easily be seen from (1.22). For example, either type of agent can reduce aggregate exposure by reducing their individual supply of market labour. The second feature is that these individual group level decisions have broader aggregate consequences, i.e. there are exposure externalities. To see this clearly, define group level exposure [4] as,

$$\mathcal{X}_{t,j} = \frac{w_{t,j} n_{t,j}^M}{n_t^M} \mathcal{X}_t \tag{1.23}$$

So that equilibrium labour supply decision of the other agent affects the group level exposure by changing both aggregate supply of market labour, n_t^M, and the aggregate exposure level. In the two agent case the following relationship holds for the cross partial derivative of exposure,

$$\frac{\partial \mathcal{X}_{t,j}}{\partial n_{t,k}^M} > 0 \Longleftrightarrow \frac{\overline{n}_k^M}{\overline{n}_j^M} < 1 \tag{1.24}$$

[4] This definition of group level exposure is consistent with aggregate exposure and the aggregation of market labour

such that the group level exposure will increase for the group with higher steady state market labour.[5] The calibration of the model fixes the sign of these cross-partial effects in the two-agent case.

1.4 Central Bank

The economy is closed by specifying how the central bank sets the nominal interest rate. The central bank sets the gross nominal interest rate $\mathcal{R}_t^n$ according to the Taylor Rule,

$$\frac{\mathcal{R}_t^n}{\overline{\mathcal{R}^n}} = (\pi_t)^{\rho_\pi} \eta_t, \tag{1.25}$$

where the parameter ρ_π controls the degree to which the central bank responds to price inflation in setting the nominal rate. The Taylor Rule rule is subject to uncertainty via the nominal interest rate shock η_t which follows an AR(1) process,

$$\ln \eta_{t+1} = \iota \ln \eta_t + \upsilon_t^{\mathcal{R}^n} \tag{1.26}$$

where ι is the degree of persistence of the shock and $\upsilon_t^{\mathcal{R}^n} \sim N\left(0, \sigma_{R^n}^2\right)$ is a random shock.

1.5 Aggregation and Equilibrium

Aggregate variables are given by the population weighted averages, i.e.

$$\begin{aligned}
c_t &= (1-\theta_t)\, c_{t,1} + \theta_t c_{t,2} \\
n_t &= (1-\theta_t)\left(n_{t,1}^M + n_{t,1}^R\right) + \theta_t n_{t,2}^M \\
n_t^M &= (1-\theta_t)\, n_{t,1}^M + \theta_t n_{t,2}^M \\
n_t^R &= (1-\theta_t)\, n_{t,1}^R \\
B_t &= (1-\theta_t)\, B_{t,1} + \theta_t B_{t,2}
\end{aligned}$$

In equilibrium the aggregate labour markets clear for each type of labour and hence the aggregate labour market is in equilibrium. In equilibrium $B_t = 0$, i.e. bonds are in zero net supply.

The aggregate resource constraint is derived by noting that the firm pays all profits to the households so that it makes no profit post-dividend and that the cost of adjusting prices doesn't

[5]This result is unique to the case of two agents, for more than two agents the following condition determines the sign of the externality,

$$\frac{\partial \mathcal{X}_{t,j}}{\partial n_{t,k}^M} > 0 \Longleftrightarrow \sum_{i \neq k} \left(1 - \frac{\overline{n}_k^M}{\overline{n}_i^M}\right) > 0$$

affect cash flow. Hence the dividend paid by the firm is given by,

$$D_t = \frac{P_{t,i} y_{t,i}}{P_t} - \frac{W_t^M}{P_t} n_{t,i}^M - \frac{W_t^R}{P_t} n_{t,i}^R \tag{1.27}$$

Combining with the aggregate budget constraint, and labour market clearing conditions yields the following aggregate resource constraint in equilibrium,

$$c_t + (1-\kappa) \frac{W_t^R}{P_t} n_t^R = y_t \tag{1.28}$$

this is also the goods market clearing condition for this economy. This is an intuitive relationship that requires that output be used either for aggregate consumption or to pay for the costs incurred by Type-1 households working remotely. So the inefficiency of working remotely introduces a wedge between consumption and output.

2 Calibration and Simulation

The calibration of the model uses relatively standard values from the New Keynesian DSGE literature and the model is calibrated to daily frequency as epidemics occur over days and weeks rather than quarters. The values of all calibrated parameters in the benchmark model can be found in Table: 1.

The model is solved using non-linear methods. In particular I use the Generalised Stochastic Simulation Algorithm (GSSA) of Judd et al. (2011) and Maliar and Maliar (2014). GSSA is an extension of the standard Parameterised Expectations Algorithm of den Haan and Marcet (1990) that replaces simple polynomials with more general basis functions, and replaces non-linear least squares estimation with quadrature techniques to estimate conditional expectations. In the simulation solution I employ Hermite Polynomials as basis functions and use 5 nodes in the quadrature calculations.

2.1 Calibrating Market and Remote Worker Parameters

The parameters $\left(\alpha_R, \alpha_M, \theta, \chi_1^M, \chi_1^R, \chi_2^M, \kappa\right)$ relate to the presence of market and remote workers.

Labour shares, (α_R, α_M) are calibrated using the fact that capital share of output is a well estimated parameter in the literature with a value of 0.36. So in order to be consistent with aggregate data one must have $\alpha_R + \alpha_M = 0.64$. Simple rearrangement of the first order conditions for intermediate firms requires that in steady state the following relationships hold

$$\alpha_R = \frac{\varepsilon_P}{\varepsilon_P - 1}\frac{W^R}{P}\frac{n^R}{y}$$
$$\alpha_M = \frac{\varepsilon_P}{\varepsilon_P - 1}\frac{W^R}{P}\frac{n^M}{y}$$

In order to be consistent with the aggregate data it is known that the total labour share of income equals 0.64, i.e. $\frac{\frac{W}{P}n}{y} = 0.64$. Combining these gives,

$$\alpha_R = 0.64\frac{\frac{W^R}{P}n^R}{\frac{W^R}{P}n^R + \frac{W^M}{P}n^M} \tag{2.1}$$

where the denominator comes from realising that the total wage bill can be decomposed as $\frac{W}{P}n = \frac{W^R}{P}n^R + \frac{W^M}{P}n^M$.

I use the Occupational Estimate Statistics from the Bureau of Labour Statistics (BLS) from 2009 - 2019 to estimate the wage bills. The BLS Occupational Estimate Statistics an annual statistical release that classifies employment by occupation with data on employment numbers, proportion of workforce and average annual salary estimates. The classification by occupation is crucial to the calibration as the objective is to determine the share of workers that can feasibly engage in remote work from home. I proceed by manually classifying each 6 digit Standard Occupational Classification (SOC) codes as either being able to work remotely or not. The classification of occupations is very similar to Dingel and Neiman (2020). Using this classification α_R can be calculated as follows. Let $\mathcal{O}$ be the set of occupations, and further $\mathcal{O}_R$ denote the subset of occupations that are classified as being remote-able work. Then α_R is calculated as,

$$\alpha_R = 0.64\frac{\sum_{h\in\mathcal{O}_R}\frac{W_h}{P}n_h^R}{\sum_{h\in\mathcal{O}}\frac{W_h}{P}n_h} \tag{2.2}$$

This was done for each of the surveys from 2009-2019 yielding an average value of $\alpha_R = 0.3313$, which is robust to alternative SOC classifications.

A similar approach is used to work out the share of the population in 'remote-able' occupations, this share is given by $1-\theta$. Using the BLS Occupational Estimate Statistics and the same set of 'remote-able' occupations $\mathcal{O}_R$ the $1-\theta$ share of 'remote-able' occupation households is estimated as,

$$1 - \theta = \frac{\sum_{h\in\mathcal{O}_R} n_h}{\sum_{h\in\mathcal{O}} n_h} \tag{2.3}$$

This calculation yields an average value of $\theta = 0.6144$, which again is robust to alternative SOC classifications.

The preference parameters $\left(\chi_1^M, \chi_1^R, \chi_2^M\right)$ are calibrated to ensure that workers spend a third of their time in working in steady state. Given the two labour types the 2018 BLS American Time Use Survey is used to pin down the ratio of time spent in market labour versus remote labour. The 2018 survey finds that 23.7% workers worked from home at the aggregate level, which implies that

$$\frac{n_M}{n_R} = \frac{1 - 0.237}{0.237}$$

Having pinned this ratio down one can easily find the steady state labour supply for each agent and labour type from the aggregate labour relationships. The labour supply first order conditions are then used to find the values of $\left(\chi_1^M, \chi_1^R, \chi_2^M\right)$ that ensure workers spend a third of their time working in steady state.

The final parameter to be calibrated is κ. To my knowledge there is no study that estimates this parameter so I use a benchmark value of $\kappa = 0.9$ in the analysis.

2.2 Calibrating Epidemiology Parameters

The parameters $\beta_0, \lambda_E, \lambda_I, \gamma$ are chosen to replicate the Covid-19 pandemic. The Covid-19 pandemic is a rapidly evolving real-time pandemic at the time of writing and there are large variations in the parameter estimates for the epidemiological parameters. The calibration used here attempts to use the best available estimates of these parameters.

Most epidemiological studies assume an average of 5.2 days spent in the exposed state. (Adhikari et al., 2020; Guan et al., 2020; He et al., 2020; Wang et al., 2020). There is a high degree of uncertainty surrounding the time spent in the infectious state. In this paper I assume that agents spend on average of 7 days in the infectious state. This seems a reasonable as it implies an average duration of the disease at 12 days which is roughly in accordance with the 14 day isolation/quarantine regime in most countries for those who have tested postive for Covid-19. This leads to calibration of the transition rates as $\lambda_E = \frac{1}{5.2}$ and $\lambda_I = \frac{1}{7}$. The current estimate for the case mortality rate is $\gamma = 0.02$.

In order to calibrate β_0 the concept of the basic reproduction rate, R_0, is used. The basic reproduction rate is the average number of people that an infected agent can infect before recovering and is given by the expected duration in the infected state multiplied by the transmission rate. The expected duration in the infectious state is 7 days so a single agent can infect on $R_0 = 7\beta_0$. Most studies use a value of $R_0 = 2.2$ for Covid-19 as per the mean estimates in Guan et al. (2020), this results in $\beta_0 = \frac{2.2}{7}$.

These parameters also fall within the range of values that have been used in various economic studies that incorporate some version of the SIR epidemiological model of the Covid-19 pandemic.

Table 1: Calibration for Benchmark Model

Macroeconomic Parameters (Households)			
Discount Factor	$\beta = 0.9999$	Risk Aversion	$\sigma = 2$
Frisch Elasticity	$\frac{1}{\psi} = 0.8$	Remote Work Cost	$\kappa = 0.9$
Time Spent in Remote Work	23.7%	Disutility of Market Work (Type-1)	$\chi_1^M = 4.0786$
Disutility of Remote Work (Type-1)	$\chi_1^R = 7.0189$	Disutility of Market Work (Type-2)	$\chi_2^M = 1.2384$
Share of Type-2 Households	$\theta = 0.6144$		
Macroeconomic Parameters (Firms)			
Elasticity of Substitution	$\varepsilon^P = 11$	Slope of Phillips Curve	0.1
Price Adjustment Parameter	$\chi^P = 110$	Market Labour Income Share	$\alpha_M = 0.31$
Remote Labour Income Share	$\alpha_R = 0.33$	Technology Shock Persistence	$\rho = 0.95$
Technology Shock Std. Dev.	$\sigma_z = 0.007$		
Macroeconomic Parameters (Central Bank)			
Inflation Response	$\phi_\pi = 1.5$	Monetary Shock Persistence	$\iota = 0.65$
Monetary Shock Std. Dev.	$\sigma_{R^n} = 0.0028$		
Epidemic Parameters			
Basic Transmission Rate	$\beta_0 = \frac{2.2}{7}$	Basic Reproduction Rate	$R_0 = 2.2$
Exposed Transition Rate	$\lambda_E = \frac{1}{5.2}$	Infectious Transition Rate	$\lambda_I = \frac{1}{7}$
Case Mortality Rate	$\gamma = 0.02$		

2.3 Other Parameters

The remaining set of parameters $\left(\beta, \sigma, \psi, \varepsilon_P, \chi^P, \phi_\pi, \rho, \iota, \sigma_z, \sigma_{R^n}\right)$ are chosen to match standard values in the literature. The Taylor Rule parameter $\phi_\pi = 1.5$ is standard in the New Keynesian literature.

The time discount factor β is chosen to ensure an annualised return of 4.2%, and $\sigma = 2$ is a standard value for the risk aversion parameter. A Frisch Elasticity of 0.8 is used which yields $\psi = \frac{1}{0.8}$, this is well within the standard range of values for this parameter.

The elasticity of substitution between goods, $\varepsilon_P = 11$, is chosen to ensure a steady state mark-up of 10%. Solving the Rotemberg Phillips curve forward via iterative substitution gives the slope of the forward looking Philips Curve as $\frac{\varepsilon_p}{\chi^P}$. The value of χ^P is set so that the slope of the Phillips Curve is $\frac{\varepsilon_p}{\chi^P} = 0.1$.(Schorfheide, 2008; Kaplan et al., 2018) Given $\varepsilon_P = 11$ this then implies that $\chi_p = 110$.

The parameters governing the technology shock and the monetary policy shock are taken as standard values from the literature. In the simulations autocorrelation parameters $\rho = 0.95$ and $\iota = 0.65$ are used together with standard errors of $\sigma_z = 0.007$ and $\sigma_{R^n} = 0.0028$.

2.4 Generating Epidemic Shocks

The simulation of the model requires one to generate the epidemic shocks Γ_t. Epidemic shocks are drawn in three broad steps,

1. Simulate the epidemic model assuming $\mathcal{X} = 1$, i.e. $\beta_t^- = \beta_0 \; \forall \; t$. Set T to be the twice as long as the duration of the modelled epidemic.

   Setting T to be twice the length of the epidemic ensures that the model is not always in an epidemic state and alternates randomly between periods of epidemic and periods without disease.

2. Generating random epidemics.

   (a) Generate an initial epidemic start date τ_1 by drawing a random number from a Uniform$(1, T)$ distribution

   (b) Sequentially draw $\tau_t = \tau_{t-1} + \epsilon_\tau$ where $\epsilon_\tau \sim U\,(1, 20)$.
   This step captures the fact that epidemics have a natural ordering in time. It ensures that time moves in the right direction, i.e. it rules out the possibility of jumping to a point to the left of the distribution as this not possible conditional on your starting point unless one epidemic has ended and another randomly begun.

   (c) If $\tau_t > T$ then redraw τ_t from a Uniform$(1, T)$ distribution

   (d) Repeat these steps until there is a τ_t for each simulation period

3. Generating the Epidemic Shock for simulation period t

   (a) Draw the values for $S\,(\tau_t)\,, E\,(\tau_t)\,, I\,(\tau_t)\,, R\,(\tau_t)\,, D\,(\tau_t)\,, N\,(\tau_t)$ and set these as the start of period t values for each epidemic state

   (b) Using the realised values for $\mathcal{X}_t$ from the period t simulation of the model, forecast the end of period values using the epidemic model

   (c) The end of period values are used to define an epidemic model consistent Γ_{t+1} using the relationship,
   $$\Gamma_{t+1} = 1 + \beta_0 \mathcal{X}_{t+1}^- \mathcal{S}_{t+1}^- \frac{\mathcal{I}_{t+1}^-}{\mathcal{N}_{t+1}^-} \tag{2.4}$$
   where start of period values for $t + 1$ are equivalent to end of period values simulated

   (d) Repeat for each simulated time period

Recall that the belief is formed on information available at the beginning of the period so that Γ_t pre-determined in any given period, i.e. it is a state variable in period t. Since Γ_t is a state variable in period t economic decisions are constrained by it so one can simulate economic decisions based on this state and use this information to update the state in the subsequent

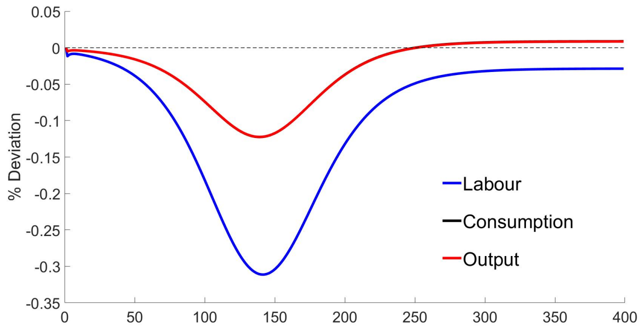

Figure 1: Aggregate Responses to Epidemic

period. The values that are drawn from the epidemic model at random constitute the epidemic shock since agents forecasts of the start of period values differ from what they observe. Agents then use the epidemic model to generate a consistent forecast of the end of period values in order to update the state variable for the subsequent period, i.e. Γ_{t+1}.

3 Epidemic Dynamics

Macroeconomic dynamics for in response to the epidemic are driven by the response of labour to the evolution of the epidemic. In order to study the dynamics I assume that initially 0.001% of the population are infected, this allows one to see the dynamics clearly. The dynamics of variables at the highest level of aggregation, i.e. a single aggregate series for labour, output and consumption, can be seen in Figure 1.

At the aggregate level it is seen that there is a gradual fall driven by the growing disutility from working in market labour as the epidemic spreads. As the epidemic eases, labour recovers as the fear of contracting the disease falls causing labour supply to increase. Aggregate labour never fully recovers to its pre-epidemic levels due to deaths, while the marginally higher level of output and consumption is the result of higher income as the marginal product of labour increases due to death.

The aggregate variables mask the underlying group level dynamics. Figure 2 highlights that dynamics of labour and consumption to the epidemic are very different when one considers the group level and labour type.[6] Figure 2a shows that it is Type-2 households, i.e. those that can only engage in market labour, that bear the brunt of the epidemic shock. The Type-2 households see their labour fall by about 4 times that of the Type-1 household, and a post-epidemic steady-

[6]The disaggregated impulse responses have been corrected for the different population sizes as the epidemic progresses.

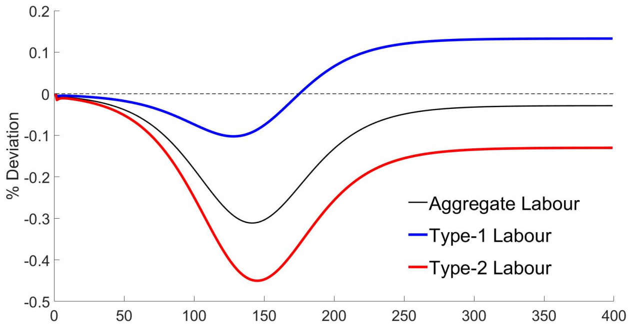

(a) Labour - Group Level

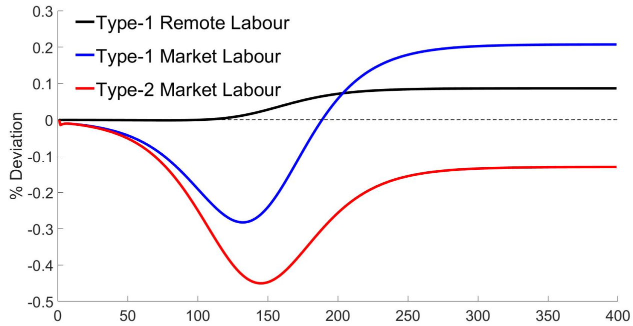

(b) Labour - Type and Group Level

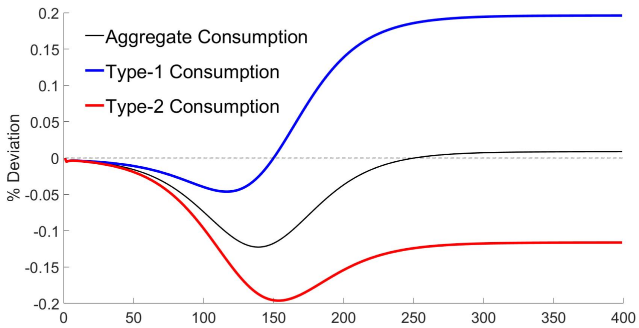

(c) Consumption - Group Level

Figure 2: Group Level Responses to Epidemic

state below the pre-epidemic steady state. Figure 2b shows that Type-1 households initially do not change their remote labour supply but only substitute away from market labour to remote labour to offset their increasing disutility of working in market labour as the epidemic evolves. Type-1 households, regardless of the wage rates, are also more wealthy post-epidemic than their Type-2 counterparts given that they supply more of both types of labour than their pre-epidemic steady states. As will be shown below this is driven by the higher death rates among the Type-2 agents. This leads to the observed consumption responses where it is again evident the Type-1 households are uniformly better off than their Type-2 counterparts. Thus a more unequal society materialises post-epidemic.

So far it has been seen that the epidemic leads to worsening economic outcomes for the Type-2 household. But what about health outcomes? Figure 3 highlights that, when compared to a standard SEIR model of the epidemic, the dual feedback between economic activity and the epidemic leads to a significant 'flattening of the curve' and most importantly an overall reduction in deaths. Specifically, the epidemic has longer duration under the current model but that

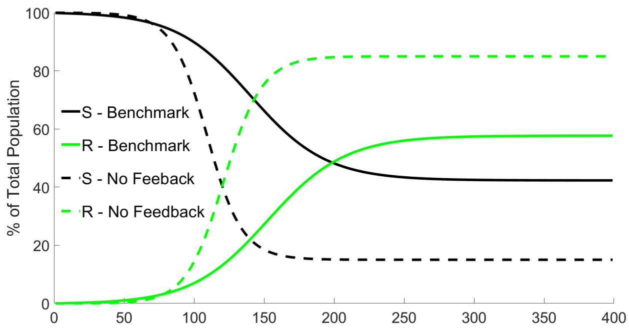

(a) Susceptible and Recovered

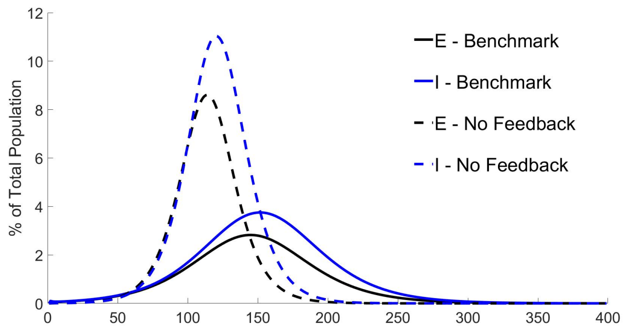

(b) Exposed and Infected

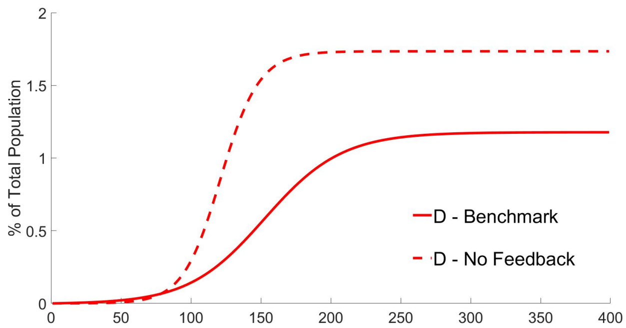

(c) Dead

Figure 3: Evolution of Epidemic for Total Population
Note: 'Benchmark' refers to the epidemiological model outlined in the model section. 'NoFeedback' refers to an epidemiological model where $\mathcal{X}_t = 1; \forall\, t$.

a significantly lower proportion of the the population is ultimately exposed and infected with the disease. This lower level of exposed and infected populations drives the significantly lower number of deaths. While outside the scope of this paper, such a flattening effect places a lower burden on the health sector as not only are less people ultimately cared for, but it also evolves over a longer period of time placing less strain on capacity in the health sector. This shows that in the absence of any policy intervention the epidemic will be flatter than that predicted by the standard SEIR model with no feedback. This is does not mean that government intervention is not necessary to combat an epidemic like Covid-19, but rather that health systems are being put under extreme strain even in a world where the curve is much flatter than the pure epidemiological model predicts.

Turning to health outcomes from the epidemic, the next thing to consider is whether the epidemic has a differential impact on the two households. Figure 4 shows the group level evolu-

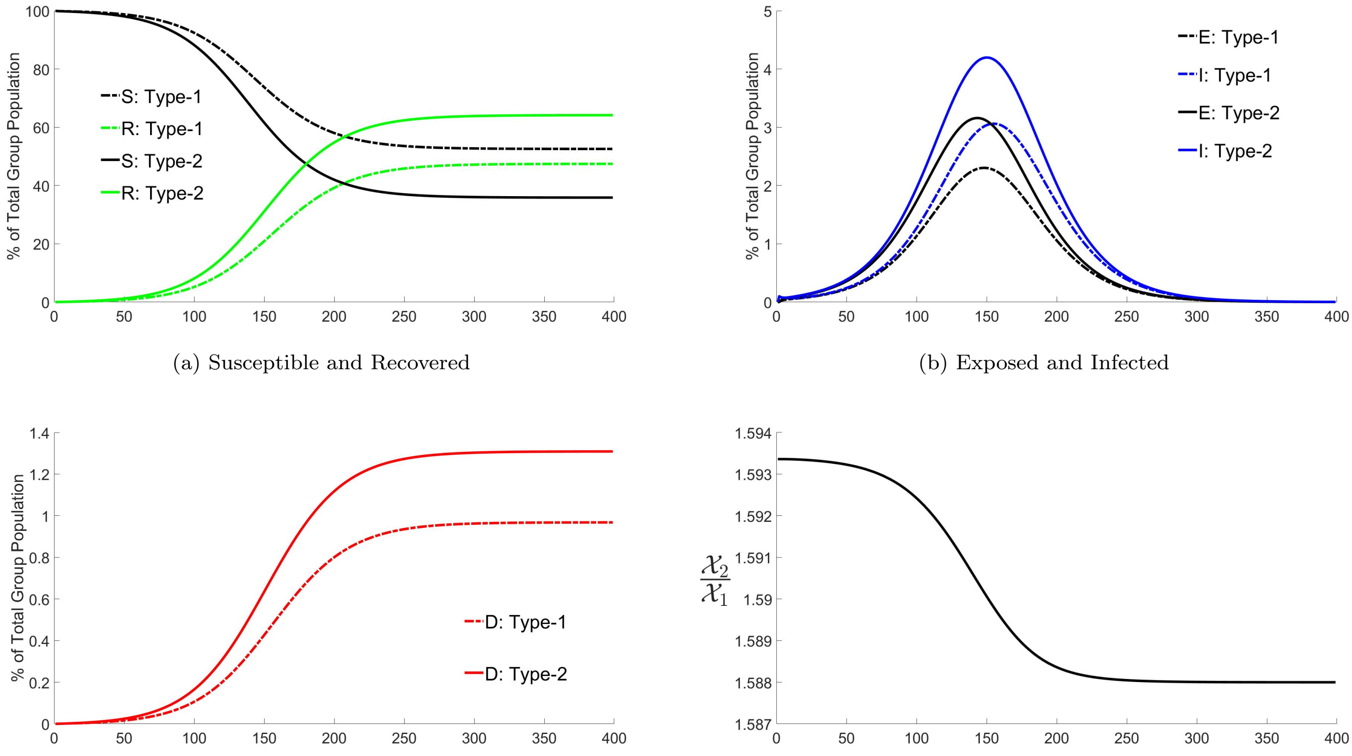

(a) Susceptible and Recovered

(b) Exposed and Infected

(c) Dead

(d) Exposure

Figure 4: Evolution of Epidemic at Group Level

tion of the epidemic correcting for different group sizes. It is immediately apparent that there are significant differences in the impact of the epidemic on the two groups, were the impact symmetric then the figures in Figure 4 would have identical shape and location. Type-2 households have higher exposure, infection and deaths compared to Type-1 agents even when correcting for different population sizes. This is wholly due to Type-2 households having a greater exposure to the epidemic due to their inability to work from home. Figure 4d shows that not only do the Type-2 households have a higher exposure rate pre-epidemic but they continue to have a higher exposure during the epidemic and also post-epidemic despite a higher death rate. Hence the experience of the epidemic is worse for Type-2 households.

In order to get some sense of the quantitative implications impulse response analysis is conducted where a value of initial infections is chosen that leads to a $\sim 13.5\%$ fall in labour over 180 days. This value is chosen to mimic the sharp rise in the unemployment rate in the US since the Covid-19 pandemic and associated containment measures. Any quantitative calculations for the current Covid-19 pandemic must be treated very cautiously, the results presented below are

Table 2: Quantitative Impact of Epidemic (180 Days From Initial Infection)

Output Loss	Total Death (% of Total Pop.)			Consumption Decline (%)		Labour Decline		
	Aggregate	Type-1	Type-2	Type-1	Type-2	Aggregate	Type-1	Type-2
5.473	0.350	0.102	0.248	2.723	7.351	13.499	5.296	18.674

no exception. The results of this quantitative calculation are presented in Table 2, and further highlight the different experience of the epidemic by the two types of household in the economy. There is clearly a large increase in inequality as Type-2 households see their consumption levels fall by more than double that of the Type-1 agents. Critically, note that while Type-2 households constitute 61.44% of the population, they account for 71.21% of the total deaths 180 days after the start of the epidemic. In the next section I consider how containment policies might change these values, and how these changes affect the observed group differences.

4 Containment Policies

The results presented so far are in stark contrast with the argument that an epidemic is a great leveller because disease does not discriminate between different people. While this may be true from a purely scientific point of view, it fails to take into account the fact that some groups are more at risk due to their economic circumstances and the opportunities available to them in the labour market. The results of this model highlight that epidemic shocks are unambiguously bad for Type-2 households - they have both worse economic and health outcomes.

In the wake of the Covid-19 pandemic many governments imposed stringent containment policies with the aim of reducing the spread of the disease. From an epidemiological perspective such containment measures are aimed at reducing the Basic Reproduction Rate, R_0. Containment policy is defined as any policy that seeks to reduce R_0 by imposing certain social restrictions in the interest of public health. Such policies include, but are not limited to, social distancing, curfews, quarantine, restriction of non-essential services, restriction on local travel, closure of borders (both domestic and international), wearing of face masks, compulsory sanitising of hands, etc., all of which were implemented to varying degrees in most countries.

Containment measures have the effect of reducing the amount of market labour available in the economy while it leaves amount of remote labour unaffected. Containment measures, $\mu_{t,j}$, are introduced into the model by modifying the budget constraint of the household and the production function of the firm to capture the exogenous reduction in market labour availability. The containment policy may be symmetric, ($\mu_{t,j} = \mu_t \ \forall \ j$) or non-symmetric ($\mu_{t,j}$ differs for each group). The modified budget constraints and production function are given by,

$$
\begin{aligned}
c_{t,j} + \frac{B_{t+1,j}}{R_t^n P_t} &= (1-\mu_{t,j}) \frac{W_t^M}{P_t} n_{t,j}^M + \kappa \frac{W_t^R}{P_t} n_{t,j}^R + \frac{B_{t,j}}{P_t} + D_{t,j} \\
y_t &= z_t \left[\left(1 - \sum_{j \in J} w_{t,j} \mu_{t,j} \right) n_t^M \right]^{\alpha_M} \left(n_t^R \right)^{\alpha_R}
\end{aligned}
$$

The introduction of containment measures modifies the household equilibrium condition for market labour, (1.4), and the aggregate resources constraint, 1.28. Containment measures leave the equilibrium conditions of the firm unchanged as the production function is Cobb-Douglas.[7] The modified household condition for market labour and aggregate resources constraint are,

$$
(1-\mu_{t,j}) \frac{W_t^M}{P_t} = \chi_j^M \Gamma_t \left(n_{t,j}^M \right)^{\psi} \tag{4.1}
$$

$$
y_t = c_t + (1-\kappa) \frac{W_t^R}{P_t} n_t^R + \mu_{t,1} (1-\theta_t) \frac{W_t^M}{P_t} n_{t,1}^M + \mu_{t,2} \theta_t \frac{W_t^M}{P_t} n_{t,2}^M \tag{4.2}
$$

So containment measures reduce market labour and consumption in equilibrium via the introduction of a containment wedge. Eichenbaum et al. (2020) view containment measures as akin to a tax. Using this idea containment measures are introduced as a tax on market labour. This tax revenue is used to provide transfers to those agents prevented from participating in the labour market through welfare schemes outside of the model.

Let us consider containment policy measures that reduce R_0 to 50% of its baseline value at full implementation. Full implementation of the policy occurs with a lag to more accurately mimic the actual behaviour of policy makers where decisions on lock-down policies take time and lag the start of an epidemic. The following containment scenarios[8] are considered in this section:

1. Symmetric Strict Containment: Full implementation 90 days from start of epidemic and lasting for 1 year, $\mu_t = \frac{\mathcal{I}_{t,1}^- + \mathcal{I}_{t,2}^-}{\mathcal{N}_t^-}$
2. Symmetric Early Exit: Full implementation 90 days from start of epidemic and lasting for 60 days, $\mu_t = \frac{\mathcal{I}_{t,1}^- + \mathcal{I}_{t,2}^-}{\mathcal{N}_t^-}$

[7] This is a special case in for Cobb-Douglas production functions. If one considered more general CES production functions then the containment measures would affect the equilibrium conditions of the firm and introduce another wedge on the production side. The study of CES production functions is beyond the scope of the present paper and left for future work.

[8] Early Warning and Phased Early Exit scenarios. The Early Warning scenario was identical to Early Exit with added light containment implementation for 30 days before strict lock-down. The Phased Early Exit scenario was identical to Early Exit except it was followed by stepwise increase of 10% every 15 days until R_0 at 90% of baseline. Both of these additional scenarios led to results very similar to the Early Exit scenario and have been omitted for clarity.

3. Non-Symmetric Early Exit: Full implementation 90 days from start of epidemic and lasting for 60 days, $\mu_{t,1} = \frac{\mathcal{I}^{-}_{t,1}}{\mathcal{N}^{-}_{t}}$, $\mu_{t,2} = \frac{\mathcal{I}^{-}_{t,2}}{\mathcal{N}^{-}_{t}}$

The containment policies considered differ in the sophistication of the information available to the government. Non-symmetric policy is more information intensive since it requires the policy maker to have information about agent types, while symmetric policy only requires information at the aggregate level. It is assumed that the government has access to perfect testing each period so that it implements policy by first removing all infected agents from the labour force, and then randomly removing other groups to meet the policy rate. Finally, the simulations assume that containment, under any scenario, will have long lasting behavioural impact leading to a reduction of R_0 to 90% of its baseline value once the containment policy ends. The Symmetric Early Exit containment scenario best mimics the types of policy that have been implemented by governments thus far in the Covid-19 pandemic.

The macroeconomic and epidemic responses to these lock-down policy scenarios are presented in Figure 5 and Figure 6. Containment policy changes the evolution of the epidemic by reducing the R_0 to 50% of its original value and restricting labour supply plays a significant role in allowing this to occur. These figures highlight a key temporal trade-off between macroeconomic variables and epidemic variables, i.e. large short-term macroeconomic losses in order to contain the spread of the epidemic. Considering the specific containment policies in this section, all of the containment responses in Figure 5 have large declines early but then intersect their respective benchmark curves before these have reached their minimum. So that after a period of large economic decline, macroeconomic variables perform better than if no containment policy were implemented. This occurs because the containment policy arrests the rise in Γ_t and moves it closer to unity much faster (see Figure 7) thereby reducing the disutility agents experience from working in market labour and hence increasing their supply of market labour. This macroeconomic loss is traded-off against the significant improvement in death rates as seen in Figure 6.

Strict Containment succeeds in eliminating the epidemic, but there is an endogenous second-wave of the epidemic in the Early Exit scenarios. This is an interesting response of the model is a consequence of the Early Exit policy not being long enough to remove all infected agents from the population. It does reduce the number sufficiently that the 'fear' of catching the disease falls, moving Γ_t closer to unity and increasing market labour supply. However, unlike Strict Containment, the presence of infected agents post containment means that as market labour increases the transmission of the disease increases as well. The second wave of the disease, while inevitable under Early Exit scenarios, is not as large as would have eventuated had no lock-down policy been implemented, and the economic impact is consequently smaller.

Turning to the quantitative impact, output decline, consumption decline at agent level and death rates 180 days from the start of the epidemic are computed; this compares all containment

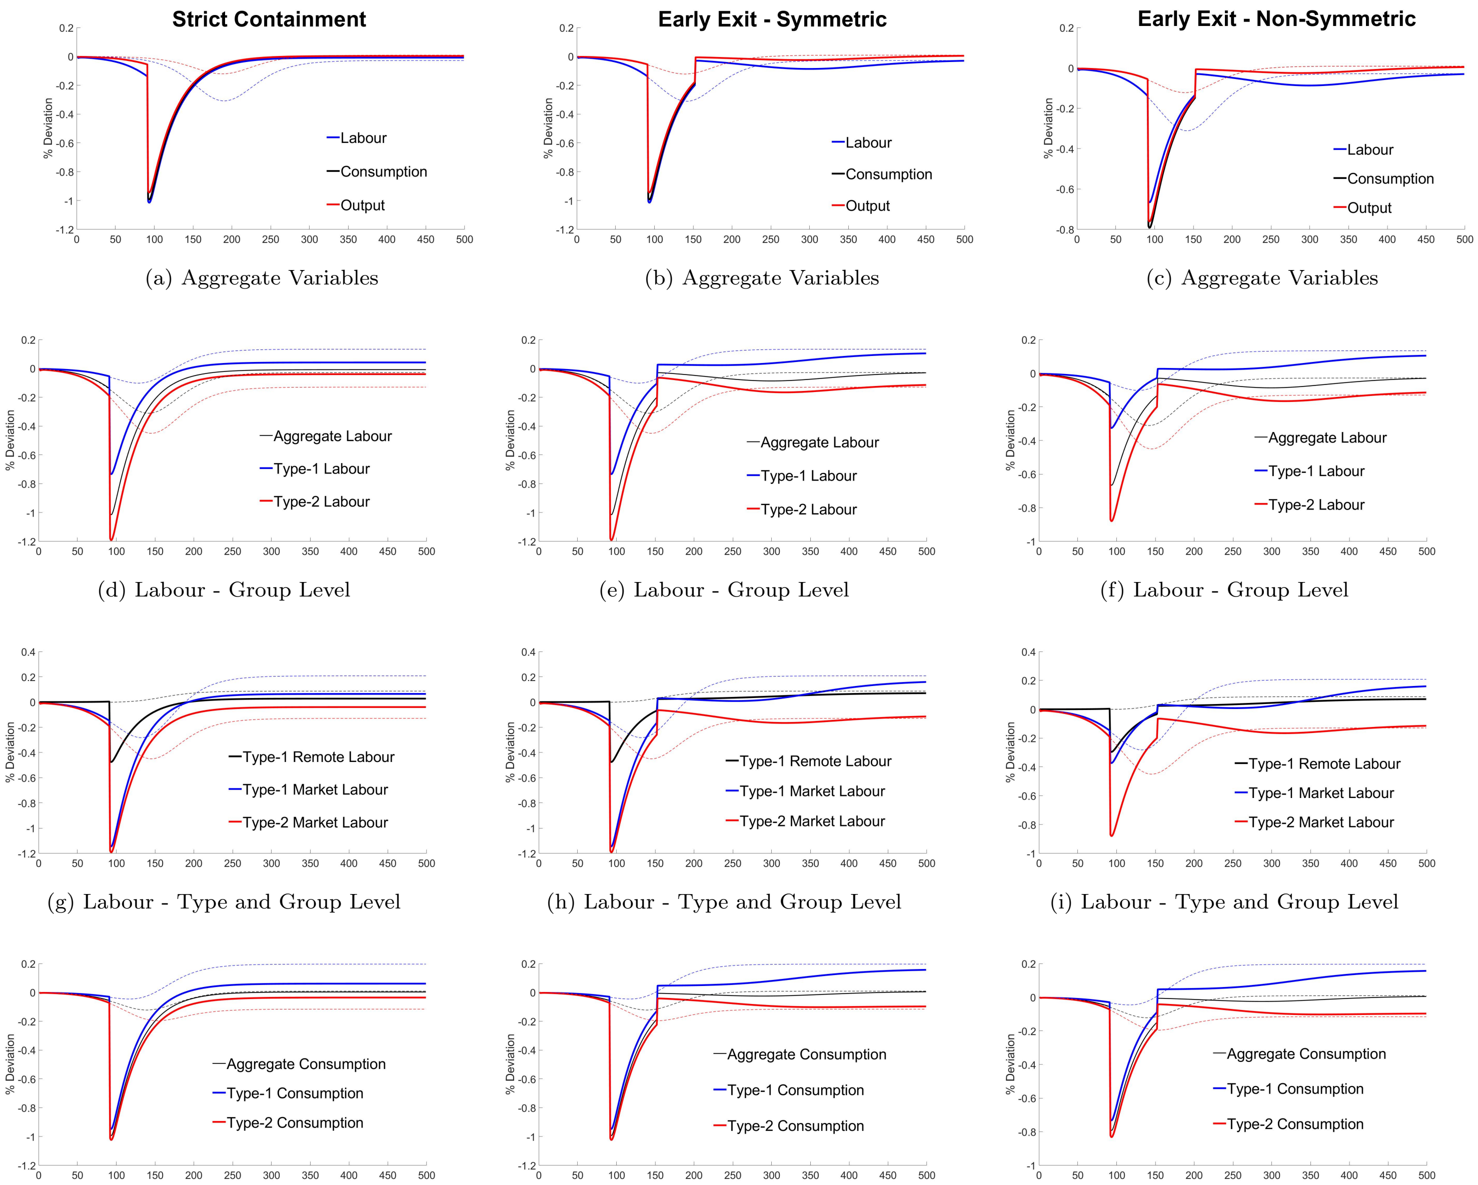

(a) Aggregate Variables (b) Aggregate Variables (c) Aggregate Variables

(d) Labour - Group Level (e) Labour - Group Level (f) Labour - Group Level

(g) Labour - Type and Group Level (h) Labour - Type and Group Level (i) Labour - Type and Group Level

(j) Consumption - Group Level (k) Consumption - Group Level (l) Consumption - Group Level

Figure 5: Macroeconomic Response to Lock-down Policy
Note: The responses to the 'Strict Lock-Down' and 'Early Exit' policies are provided by the thick coloured lines, while the dashed lines are reproductions of the 'Benchmark', i.e. no containment policy, case for ease of comparison.

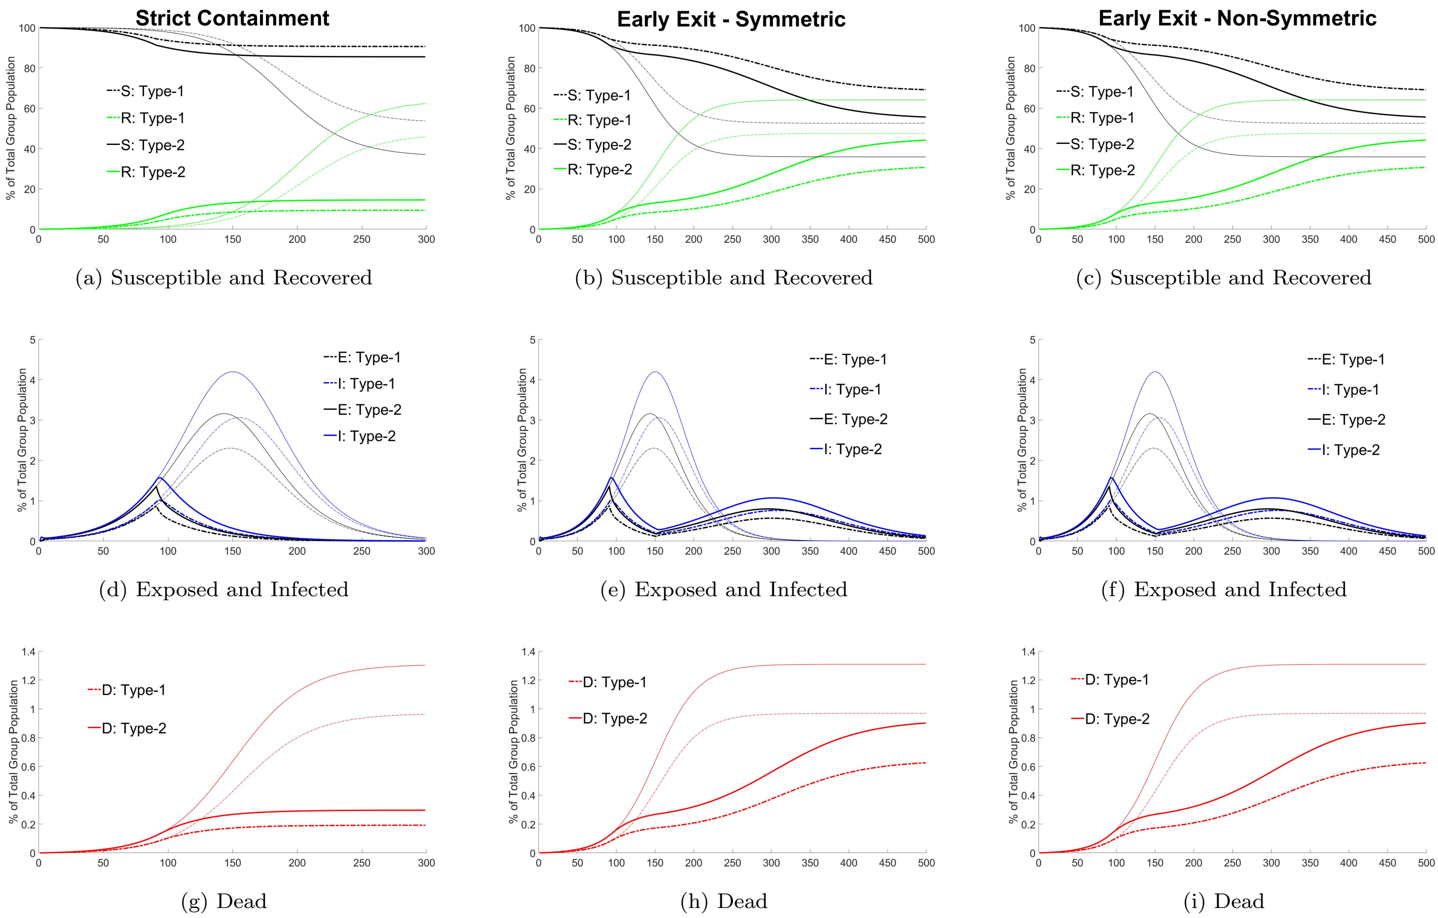

(a) Susceptible and Recovered (b) Susceptible and Recovered (c) Susceptible and Recovered

(d) Exposed and Infected (e) Exposed and Infected (f) Exposed and Infected

(g) Dead (h) Dead (i) Dead

Figure 6: Epidemic Response to Lock-down Policy

Note: The responses to the 'Strict Lock-Down' and 'Early Exit' policies are provided by the thick coloured lines, while the dashed lines are reproductions of the 'Benchmark', i.e. no containment policy, case for ease of comparison.

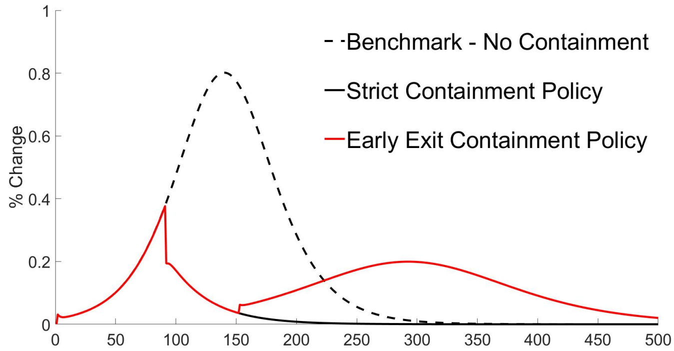

Figure 7: Containment Response of Γ_t

policies 30 days after the Early Exit scenario ends. The initial infection rate is once again set so that there is a $\sim 13.5\%$ fall in labour over the first 180 days of the epidemic in the Benchmark Case. The quantitative measures for each policy 180 days from the onset of the epidemic presented in Table 3. The quantitative impact further highlights the trade-off that policy makers make between economic prosperity and saving lives in policy responses to an epidemic. One could adopt a laissez-faire policy that minimises the output loss from the epidemic, but this also results in an unacceptably high death rate and cases on the rise. The containment policies accept a lower level of output in order to significantly reduce the death rate. All of the containment scenarios studied reduce the aggregate death rate by around 85% of the benchmark value at a cost of increased output loss between $0.04 - 2.18\%$. While death rates do fall for both groups it must be stressed that Type-2 agents still have death rates double that of their Type-1 counterparts as they can only supply risky market labour. Additionally, the implementation of containment policies leads to more equal falls in consumption thereby minimising the increased inequality resulting from the epidemic. In particular note that the Early Exit scenario with non-symmetric intervention significantly reduces the death rate while only reducing output by slightly more than the no policy benchmark. It outperforms the Early Exit scenario with symmetric intervention because it restricts less people from participating in the labour force. This is due to the fact that the two household types are differentially impacted by the epidemic due to the labour opportunities available to them. Remote labour opportunities mean that less Type-1 households get infected and so a symmetric policy restricts some Type-1 households who are not in the $\mathcal{I}$ group. Thus if policy makers do not take into account the asymmetry in the experience of the epidemic at the group level there is ultimately a larger loss in output.

The implementation of non-symmetric policy is fraught with other issues key among them being timely availability of accurate of information at the group level and fairness considerations. Implementation of the non-symmetric policy requires access to accurate daily data of

Table 3: Quantitative Impact of Containment (180 Days From Initial Infection)

	Output Loss %	Total Death (% of Total Pop.)			Consumption Decline (% of Group Pop.)	
		Aggregate	Type-1	Type-2	Type-1	Type-2
Benchmark Model	5.473	0.350	0.102	0.248	2.723	7.351
Strict Containment	7.650	0.048	0.014	0.034	6.970	8.747
Early Exit: Symmetric	6.736	0.051	0.015	0.036	6.016	7.772
Early Exit: Non-Symmetric	5.514	0.051	0.015	0.037	4.582	6.495

infections at the group level. This information ideal is clearly impractical and unlikely to materialise in the real world. The study of the containment policies taking into account of information availability and the presence of lags is an extension left for future work. Even if one had access to perfect data, issues surrounding about treating different groups differently would likely render non-symmetric policies difficult to implement on fairness grounds and would likely draw significant political backlash. The non-symmetric policy does not materially change the death rate from the symmetric case.

The Early Exit policies appear to tread the fine line between draconian measures that have a single minded focus on saving lives at the expense of liberty and maximising economic prosperity during the epidemic. Despite the inevitable second wave, such policies provides policy makers the breathing space needed to increase capacity in the health system, research the disease to improve care provided to those infected, search for a vaccine, etc. All of these efforts may allow policy makers to fundamentally alter the evolution of the epidemic post-containment. For example, research into improved care could reduce the time spent in the infectious state, or the development of a viable vaccine could remove whole swathes of the population from the susceptible group. The moral of the Early Exit story is that if policy makers aim to exit from their containment measures early they should use the time afforded to them to invest in programmes that will help to fundamentally alter the evolution of the epidemic post exit from containment.

5 Conclusion

The main conclusion of this paper is that entrenched differences between groups in the lead up to an epidemic can have a significant bearing on how individual groups experience the epidemic both in terms of health outcomes and economic prosperity. The main results of the paper show that such differences can be a powerful driving force behind post-epidemic inequality, while the study of containment scenarios highlight that government intervention can help to minimise any post-epidemic inequality.

During the outbreak of disease engaging in labour that requires social proximity to others is a risky activity as it increases the chances of contracting the disease. This paper has highlighted that the labour market opportunities available to an agent has a significant bearing on how that agent experiences the outbreak of an epidemic. It shows that when agents cannot engage in remote labour and must supply labour in the market during an epidemic that these agents end up experiencing higher death rates vis-à-vis their make up of the population as they must engage in labour that requires social contact. These agent also suffer larger declines in labour and consumption over the course of the epidemic. Post-epidemic, these agents are significantly worse off than their counterparts who have the opportunity to work from home and a more unequal society emerges.

The paper then went on to show that simple containment policies, while leading to larger losses in economic prosperity as measured by output loss, can significantly reduce death rates across the population and bring the death rates of the two groups closer together. Containment policies continue to see workers who can only supply market labour with death rates double that of those who have the opportunity to work from home. Again this is due to these agents having to continue engaging in risky social interactions by supplying market labour. Containment policies also appear to reduce the inequality that emerges with falls in consumption more similar across agent types.

There are many possible extensions to the model and analysis presented in this paper. Below I briefly discuss three interesting extensions of the current paper inspired both by the analysis presented in this paper and research efforts by others in the quickly evolving economic literature on the Covid-19 pandemic. These extensions are the presence of a social consumption channel, the inclusion of capital, and information constraints due to lags and noisy data.

This paper has only considered one dimension along which agents may experience an epidemic - the labour market. This paper has abstracted away from the consumption of social goods (e.g. haircuts, dining at restaurants and shopping physically in shops) which are another form of risky social economic interaction during a pandemic. Eichenbaum et al. (2020) use this channel to great effect and it plays an important role in the story presented in that paper. The absence of an income effect on labour supply due to the assumption of GHH preferences to get a clearer view of the mechanism at work prevents this channel from playing any role. An income effect would lead to a muting of the dynamics presented in this paper. The interaction of the consumption channel with the labour channel studied in this paper could result in some interesting dynamics once an income effect channel is opened up.

This paper does not have a role for capital accumulation which can be an important channel during a pandemic. Bodenstein et al. (2020) allow for capital accumulation to capture the idea that capital allows both firms and households to minimise engaging in risky labour supply by substituting capital for labour in production. The presence of capital in the present model could possibly amplify the fall in labour during a pandemic which could improve the health outcomes in terms of lower death rates as agents minimise their supply of risky labour supply. However, it might also lead to a muting of the labour response if capital rental rates increase during a pandemic leading to a demand for more labour which would then lead to a deterioration of health outcomes as agents must engage in more risky labour.

The final extensions to be discussed relate to the information constraints faced by both households and governments. The current paper assumes that both households and agents have access to accurate data at the beginning of the period. However, as the current pandemic has shown, during a crisis there are likely to be significant delays in access to accurate data by both households and the government. Lags in access to data and noisy data could have a significant impact on the dynamics presented here as agents make decisions based on information about the epidemic that is not current and/or accurate. In the present model this could lead to agents oversupplying market labour thereby exacerbating the spread of the epidemic, or failing to return to market labour quickly enough at the end of the epidemic thereby making the recovery slower.

References

Adhikari, S. P., S. Meng, Y. Wu, Y. Mao, R. Ye, Q. Wang, C. Sun, S. Sylvia, S. Rozelle, H. Raat, et al. (2020). Epidemiology, causes, clinical manifestation and diagnosis, prevention and control of coronavirus disease (covid-19) during the early outbreak period: a scoping review. *Infectious diseases of poverty 9*(1), 1–12.

Bodenstein, M., G. Corsetti, and L. Guerrieri (2020, April). Social Distancing and Supply Disruptions in a Pandemic. Finance and Economics Discussion Series 2020-031, Board of Governors of the Federal Reserve System (U.S.).

den Haan, W. J. and A. Marcet (1990, January). Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics 8*(1), 31–34.

Dingel, J. I. and B. Neiman (2020, April). How Many Jobs Can be Done at Home? NBER Working Papers 26948, National Bureau of Economic Research, Inc.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020, March). The macroeconomics of epidemics. Working Paper 26882, National Bureau of Economic Research.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Ríos-Rull (2020, April). Controlling a Pandemic. NBER Working Papers 27046, National Bureau of Economic Research, Inc.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988, June). Investment, capacity utilization, and the real business cycle. *American Economic Review 78*(3), 402–17.

Guan, W., Z. Ni, Y. Hu, W. Liang, C. Ou, J. He, L. Liu, H. Shan, C. Lei, D. S. Hui, et al. (2020). Clinical characteristics of coronavirus disease 2019 in china. *New England journal of medicine 382*(18), 1708–1720.

He, X., E. H. Lau, P. Wu, X. Deng, J. Wang, X. Hao, Y. C. Lau, J. Y. Wong, Y. Guan, X. Tan, et al. (2020). Temporal dynamics in viral shedding and transmissibility of covid-19. *Nature medicine 26*(5), 672–675.

Judd, K. L., L. Maliar, and S. Maliar (2011, July). Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics 2*(2), 173–210.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary Policy According to HANK. *American Economic Review 108*(3), 697–743.

Kaplan, G., B. Moll, and G. L. Violanter (31 March 2020). Pandemics according to hank. *Powerpoint Presentation, LSE,*.

Kermack, W. O. and A. G. McKendrick (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character 115*(772), 700–721.

Kremer, M. (1996). Integrating behavioral choice into epidemiological models of aids. *The Quarterly Journal of Economics 111*(2), 549–573.

Maliar, L. and S. Maliar (2014). Numerical methods for large scale dynamic economic models. *Handbook of computational economics 3*, 325–477.

Rotemberg, J. J. (1982). Monopolistic price adjustment and aggregate output. *The Review of Economic Studies 49*(4), pp. 517–531.

Schorfheide, F. (2008). DSGE model-based estimation of the New Keynesian Phillips curve. *Economic Quarterly 94*(Fall), 397–433.

Wang, H., Z. Wang, Y. Dong, R. Chang, C. Xu, X. Yu, S. Zhang, L. Tsamlag, M. Shang, J. Huang, et al. (2020). Phase-adjusted estimation of the number of coronavirus disease 2019 cases in wuhan, china. *Cell discovery 6*(1), 1–8.