

Safe choices?

SUMMARY

In this paper, we argue that endogenous shifts in private consumption behaviour across sectors of the economy can act as a potent mitigation mechanism during an epidemic or when the economy is re-opened after a temporary lockdown. We introduce a susceptible-infected-recovered epidemiological model into a neoclassical production economy in which goods are distinguished by the degree to which they can be consumed at home rather than in a social, possibly contagious context. We demonstrate within the model, that the ‘Swedish solution’ of letting the epidemic play out without much government intervention and allowing agents to reduce their overall consumption as well as shift their consumption behaviour towards relatively safe sectors can lead to substantial mitigation of the economic and human costs of the COVID-19 crisis. We argue that significant seasonal variation in the infection risk is needed to account for the two-wave nature of the pandemic. We estimate the model on Swedish health data and show that it predicts the dynamics of weekly deaths, aggregate as well as sectoral consumption, that accord well with the empirical record and the two-waves for Sweden for 2020 and early 2021. We also characterize the allocation a social planner would choose and how it would dictate sectoral consumption patterns. In so doing, we demonstrate that the laissez-faire outcome with sectoral reallocation mitigates the economic and health crisis but possibly at the expense of unnecessary deaths and too massive a decline in economic activity.

JEL codes: E52, E30

—Dirk Krueger, Harald Uhlig and Taojun Xie

Macroeconomic dynamics and reallocation in an epidemic: evaluating the ‘Swedish solution’

Dirk Krueger, Harald Uhlig and Taojun Xie  *

University of Pennsylvania, USA, CEPR, UK and NBER, USA; University of Chicago, USA, NBER, USA and CEPR, UK and Asia Competitiveness Institute, Lee Kuan Yew School of Public Policy, National University of Singapore, Singapore

1. INTRODUCTION

The COVID-19 pandemic of 2020–21 has the world in its grip. Policymakers wrestle to find the optimal balance between economic activity to permit and the corresponding risk of death. The policy response to the disease outbreak in the early spring of 2020 was swift but varied, with many countries around the world effectively shutting down their

* D.K. and H.U. thank the National Science Foundation for support under grant SES-175708. The authors thank Chris Brunet for excellent research assistance as well as the editor, Jerome Adda, two anonymous referees and our two discussants, Moritz Kuhn and Federica Romei for very useful comments that have fundamentally changed and greatly improved the manuscript. Dynare replication codes are provided at: https://github.com/tjxie/KUX_PandemicMacro.

The Managing Editor in charge of this paper was Jerome Adda, guest editor of the special issues on the Economics of Covid-19.

economies for extended periods of time. On the other hand, other countries adopted policies that let the pandemic run its course without much government intervention. Perhaps the most pointed example in the latter group is Sweden, which has largely avoided government restrictions on economic activity and allowed people to make their own choices instead. One can reasonably argue that Sweden was not as extremely *laissez-faire* as just described: the Swedish government provided many guidelines and recommendations to its citizens. The Swedish population largely followed these to mitigate the COVID-19 spread. Nonetheless, the Swedish approach differed significantly from the lockdowns and restrictions imposed by most other countries. As such, we will interpret our analysis of Swedish data as the *laissez-faire* extreme, in which the pandemic does not move the government to intervene in the *economic* adjustments of private behaviour and in the economy as a whole. For the purposes of this work, we will refer to the *laissez-faire* model as the ‘Swedish Solution’, fully acknowledging that it describes the Swedish actual policy response only imperfectly.

In this paper, we evaluate the extent to which consumers concerned about their health, even in the absence of government intervention, will seek to mitigate economic interactions that carry the risk of infection, given the potentially disastrous consequences to their health. We then deduce the macroeconomic and public health consequences from the ‘Swedish Solution’. We pay particular attention to quantitatively matching the data on COVID-19 deaths in Sweden to the model predictions.

We conduct our analysis within a simple macroeconomic model, where agents consume and work, combined with a standard epidemiological ‘susceptible-infected-recovered’ (SIR) model modified by Eichenbaum et al. (2021) (ERT), who assumed that infections are more likely when agents are consuming or working together.¹ Individuals participating in these market activities are aware of the resulting infection- and death-risks, and thus may alter their consumption and work patterns as the epidemic unfolds, but do not take into account the externality of their behaviour on the infection risks of others. Like the agents in the model, we view the endogenous response in the behaviour of people, motivated by their own interest in preserving their life as key in understanding the spread of a pandemic and, ultimately, its economic costs. Such a model is a significant advance from the purely epidemiological models succinctly summarized in Atkeson (2020).

We depart from ERT in two crucial dimensions, by altering the economic as well as the epidemiological mechanisms. For the first dimension, that is, regarding the economic mechanism and in contrast to ERT, we assume that the economy is composed of several heterogeneous sectors that differ technologically in their infection probabilities. There are two interpretations of this assumption. One is, that very similar goods (e.g., pizza) can be consumed in the privacy of the home rather than in the marketplace (a restaurant). Alternatively, very similar work may also be performed remotely rather than in an office, for example, writing a report online. In both cases, the re-allocation of labour is

1 See Atkeson (2020) for an introduction to SIR models for economists.

fairly simple: it is the resulting goods produced and the implication for social and contagious interaction that matters. Our model's key mechanism is the endogenous, privately optimal reallocation of individuals towards less contagious economic activities and away from sectors in which consumption (or production) is subject to larger COVID-19 infection risk.

To support the empirical relevance of this mechanism, we point to [Leibovici et al. \(2020\)](#) who provide evidence for substantial heterogeneity across sectors of the United States in the degree of social interaction to facilitate the production of goods and services. [Dingel and Neimann \(2020\)](#), as well as [Mongey and Weinberg \(2020\)](#), assess what share of jobs can be performed at home, and [Toxvaerd \(2020\)](#) provides an equilibrium model in which social distancing is an endogenous outcome emerging from individually rational behaviour. Consistent with our main mechanism, [Farboodi et al. \(2020\)](#) provide evidence from US micro-data for a large reduction in social activity by private households even prior to the implementation of public stay-at-home-orders and lockdowns of economic activity. [Hodbold et al. \(2021\)](#) use data from a representative consumer survey in five European countries conducted in summer 2020, after releasing the first wave's lockdown and travel restrictions, to investigate the reasons for households' reduction in consumption in five key sectors of the economy and find that infection risk is the most often cited reason for consuming less than before, for all countries and all sectors. This lends direct support to our assumption that households internalize COVID-19 infection risk when making consumption decisions.

Following up on this argument, [Figure 1](#) displays seasonally adjusted Swedish monthly consumption expenditures on restaurants and groceries, together with total consumption. At the onset of the epidemic in March of 2020, there is clear evidence of reallocation of consumption activities away from food consumption in restaurants to food consumption at home. Aggregate consumption fell by 10% between February and April 2020, whereas consumption in restaurants collapsed by 50%. In contrast, expenditure on groceries – a proxy for food consumption at home – actually increased 5% between February and March and held steady thereafter. Similarly, when the second wave of infections and deaths commenced in late October, restaurant consumption again plummeted, whereas expenditure on groceries mildly increased and overall consumption slightly declined. It is this reallocation between sectors or modes of consumption that is at the heart of our quantitative dynamic equilibrium model.

The quantitative potency of this reallocation mechanism depends crucially on the degree to which goods in different sectors can be substituted, as measured by the elasticity of substitution across goods (or work activities), which we denote by η in our paper. If different sectors are interpreted simply as different modes (e.g., food at home and food in restaurants) in which otherwise similar goods are consumed, then this elasticity can be assumed to be high. An alternative interpretation is that these sectors represent rather distinct goods or distinct lines of work and that substitution among them may be lower. We initially followed this interpretation and chose $\eta = 3$, following [Adhmad and Riker \(2019\)](#). We also considered a higher value, $\eta = 10$ as our benchmark, following




Figure 1. Comparing two different consumption sectors in Sweden

[Fernandez-Villaverde \(2010\)](#). In both cases, we assumed that the economy comprises two equally sized sectors that simply differ in their infection risk. Our parameter implies that the infection probability in the most infectious sector (for the same consumption or work intensity) is nine times as high as in the least infectious sector.

Note that we interpret the term ‘consumption’ in this paper broadly, to include non-market social activities as well. The substitution discussion above is also applicable, for example, to talking online instead of partying with friends, or staying at home instead of congregating in parks or even sending email petitions to advocate for some cause instead of demonstrating in the street. Viewed from this perspective, infection is inexorably linked to consumption or workplace interaction and we shall assume as much in our analysis.

The second departure from ERT concerns the epidemiological mechanism. It arises from the challenge of matching both the shape and the magnitude of the first and second wave of infections and deaths in March and April of 2020 and the winter of 2020, respectively. The observed second wave in the fall and winter of 2020 is incompatible with the mechanics of a basic SIR model, which assumes constant parameters for infections, deaths and recovery, as we discuss in Section 5.1. That insight is not changed by allowing for the economic forces introduced here or in ERT or related models in the literature. Something else has to give. Closer examination of the quantitative properties of

our model leads us to assume an exogenous seasonal change in the infection parameter of the SIR portion of the model, allowing it to be high in the ‘winter’ season (October to March) and low in the ‘summer’ season (April to September).

We estimate the key model parameters on Swedish weekly data on deaths for 2020 and early 2021 by minimizing the root mean square error (RMSE) between model predictions and the data. We compare the predictions of our model to that of a basic SIR model without economic mechanisms but likewise modified to allow for the exogenous seasonal change in the infection parameter. To facilitate the comparison, we use two versions. For the first, we use our benchmark model estimates and merely ‘turn off’ the economic forces. For the second, we fit the SIR model to the first wave.² Figure 2 compares the model simulations to the weekly COVID-19 deaths rates in Sweden. It shows that our model captures the empirical record in Sweden well, both the initially small number of deaths, its explosion in late March and early April, the slower decline in May and June as well as the second wave in the late fall and winter of 2020. By contrast, the seasonal SIR epidemiological model will either capture that first wave and then massively miss the second, or miss both.

The logic is easy to understand and most clearly comes through in the top panel of Figure 2. Confronted with the rise of the death toll in both the first and the second waves, susceptible agents in our model endogenously choose consumption patterns that reduce their infection risk, thus mitigating the forces predicted by a basic or seasonal SIR model. They consume and work less (as already pointed out by ERT), and crucially, also shift their consumption in both waves towards the less infectious sectors. One can interpret the comparison of the solid line and the dashed line in the top panel as quantifying the force of our economic mechanism, which displays the power of individual choice in mitigating the number of lives lost in a pandemic. The lower panel shows that if the basic SIR model is to fit the first wave, the basic infection risk has to be tempered down relative to our model with its operative economic mitigation mechanism. But then the SIR economy enters the fall with many more individuals still susceptible to the disease, and when the infection rate increases in the fall, infections and deaths explode. In contrast, in our model, individuals facing the onset of the second wave massively reduce and reallocate consumption, thereby lowering the chance of infection and the magnitude of the second wave in the winter of 2020.

Remarkably, our model not only captures the health dynamics in Sweden well but also predicts fairly accurately the size of the initial economic crisis induced by the pandemic. In the data, aggregate consumption falls by slightly more than 10% in the second quarter of 2020. Our benchmark model predicts a decline of more than 8% in measured consumption. Especially in the second wave, the sectoral reallocation channel is

2 We also tried fitting the seasonal SIR model to both waves. Visually, it is worse than the two SIR versions presented here, as it undershoots the first wave, overshoots the second and does not get the dynamics right.




Figure 2. Comparison of KUX and SIR models

Notes: Top panel: the SIR model uses the same parameters as our benchmark ‘KUX’ model. Bottom panel: the SIR model is fit to the first wave.

quite potent, reducing the consumption recession by about 2 percentage points (relative to the model where sectoral reallocation is absent), demonstrating the importance of the sectoral reallocation mechanism for mitigating the health-induced economic crisis.

Our results are stark, partly because our analysis assumes smoothly functioning labour markets where workers can quickly reallocate to the sectors now in demand: waiters at restaurants deliver food instead, for example. It is easy to argue that the world is not as frictionless as assumed here and that the message of our paper is perhaps overly optimistic. While matters may be more complicated in practice, the key insight of our analysis is likely to remain: substitution possibilities and private incentives of agents facing infection risks are important factors when thinking about the COVID-19 epidemic, both at the onset and during its evolution.

Our analysis relates to other recent work that has emphasized the need to think about a multi-sector or heterogeneous household economy for the purpose of analysing the economic effect of the recent epidemic, such as [Acemoglu et al. \(2020\)](#), [Alvarez et al. \(2020\)](#), [Brotherhood et al. \(2020\)](#), [Glover et al. \(2020\)](#), [Guerrieri et al. \(2020\)](#) or [Kaplan et al. \(2020\)](#). However, these authors do not feature the feedback from the differential infection probabilities across sectors into the private reallocation made by agents.³ A second very active literature evaluates the impact of publicly enforced mobility restrictions and social distancing measures on the dynamics of the COVID-19 epidemic, see for example, [Correia et al. \(2020\)](#), [Fang et al. \(2020\)](#) or [Greenstone and Nigam \(2020\)](#), which in turn build on [Adda \(2016\)](#). Complementary to this work as our core argument, we emphasize that strong private incentives to redirect consumption behaviour might go a long way towards mitigating or even averting the epidemic, even in the absence of mobility restrictions or publicly enforced social distancing measures. Our paper complements the work by [Barrero et al. \(2020a\)](#) (or [2020b](#)) who document that the COVID-19 crisis is also a reallocation shock, when examining data from the US economy. For the United States, however, these sectoral reallocations may have mainly been driven by government restrictions. By contrast, our reallocation results from private concerns about infection risks only.

Ultimately, the purpose of this paper is meant to clarify the key forces rather than paint a complete, detailed and accurate quantitative picture of the Swedish experience with the epidemic. We, therefore, focus first, in the model developed in Section 2, on the infection risk in the consumption sector only. In Section 3, we provide theoretical results that demonstrate the importance of the elasticity of substitution across sectors.

3 A subset of this literature, for example, [Acemoglu et al. \(2020\)](#), [Brotherhood et al. \(2020\)](#) and [Glover et al. \(2020\)](#) emphasizes the importance of *age differences* in infection rates, case fatality rates and the potential welfare gains from age-contingent mitigation policies. This might be especially important for the first wave in Sweden where a significant share of deaths occurred among the elderly in retirement homes. We abstract from such age differences to more clearly isolate the power of sectoral reallocation induced by private incentives.

We also argue that the same mechanism is at work whether the risk of infection is located in the labour market or in the consumption goods market; in such a case, one may contend that the effect is to lower the elasticity of substitution. In Section 4, we set up the problem of a social planner who can observe which agents are infected and which are not, akin to the planning problem studied by [Alvarez et al. \(2020\)](#). One may conceptualize this as a strong government with wide testing capabilities⁴ of individuals or with powerful moral persuasion to influence infected agents to do what is good for society at large. Section 5 discusses the computation, calibration and estimation of the model. The results for the baseline economy are contained in Section 6, which shows how individually rational reallocation of economic activity across sectors is a strong mitigating force of the crisis even in the absence of explicit government intervention. Section 7 conducts a sensitivity analysis, thereby providing a deeper understanding of the forces at work as well as of the similarities and differences compared with a one-sector economy. Section 8 contrasts the outcomes of the laissez-faire economy with that of a planned economy. We show that the social planner can stop the pandemic more decisively. This should not be all that surprising: the social planner simply prevents infected agents from interacting with the susceptible part of the population (by separating the consumption of both groups across sectors), even if this imposes considerable, additional pain on the infected agents, which the social planner, of course, takes into account. More surprising, though, is that the decentralized solution with its substitution possibilities across sectors goes a long way towards the private economy achieving a fairly similar outcome relative to an economy where these substitution opportunities are absent.

2. MODEL

2.1. The macroeconomic environment

Our framework builds on ERT and shares some key model components. Time is discrete, $t = 0, 1, 2, \dots$, measuring weeks. There is a continuum $j \in [0, 1]$ of individuals, maximizing the objective function

$$U = \mathbf{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t^j, n_t^j),$$

where β denotes the discount factor, c_t^j denotes consumption of agent j and n_t^j denotes hours worked. Expectations $\mathbf{E}_0$ are taken with respect to stochastic health transitions described below in detail. Dead agents are normalized to have utility $u(c_t^j, n_t^j) \equiv 0$ as well

4 In this sense, our social planner analysis is akin in spirit to the focus on testing in [Berger et al. \(2020\)](#).

as $\dot{c}_t = \dot{n}_t = 0$. For the quantitative analysis as in ERT, we assume that preferences of living agents are given by

$$u(c, n) = \ln c - \theta \frac{n^2}{2}. \quad (1)$$

In contrast to ERT, we assume that consumption $\dot{c}_t^j$ takes the form of a bundle across a continuum of sectors $k \in [0, 1]$,

$$\dot{c}_t^j = \left(\int (\dot{c}_{tk}^j)^{1-1/\eta} dk \right)^{\eta/(\eta-1)}, \quad (2)$$

where $\eta \geq 0$ denotes the elasticity of substitution across goods and $\dot{c}_{tk}^j$ is the consumption of individuals j at date t of sector k goods. Workers can split their work across all sectors and earn a wage W_t in units of a numeraire good⁵ for a unit of labour, regardless of where they work. As the choice of the numeraire is arbitrary, we let a unit of labour denote that numeraire: thus, wages are equal to unity, $W_t = 1$.

Goods of sector k are priced at P_{tk} in terms of the numeraire, that is, in units of labour. We suppose that production of goods in sector k is linear in labour, that is, total output of goods in sector k equals the total number of hours worked there times some aggregate productivity factor A and that pricing in each sector is competitive. Thus, prices equal marginal costs and are the same across all sectors,

$$P_{tk} = P_t = 1/A.$$

The date- t budget constraint of the household is therefore⁶

$$\int \dot{c}_{tk}^j dk = A n_t^j. \quad (3)$$

2.2. The epidemic

As in ERT, we assume that the population will be divided into four groups: the ‘susceptible’ people of mass S_t who are not immune and may still contract the disease but are not currently infected, the ‘infected’ people of mass I_t , the ‘recovered’ people of mass R_t and the dead of mass D_t . We assume that the risk of becoming infected and the rate of

5 The presentation of the model is easier assuming a numeraire rather than payment in a bundle of consumption goods. We will not examine sticky prices or sticky wages in this model.

6 Different from ERT, we do not feature a tax-like general consumption discouragement and thus no government transfers. We also abstract from capital and thus from intertemporal savings decisions, as they do.

death or recovery do not depend on the sector of work but depends exclusively on consumption interactions. Our focus here is on the sectoral shift in consumption: for simplicity and in contrast to ERT, we assume that infected individuals continue to work at full productivity but that the disease can only spread due to interacting consumers. We show in Section 3.3 that this is similar to a model, where the infection can only spread via the workplace.

Different goods or, perhaps better, different ways of consuming rather similar goods vary in their degree of contagiousness.⁷ To that end, we assume that there is an increasing function $\phi : [0, 1] \rightarrow [0, 1]$, where $\phi(k)$ measures the degree of social interaction or relative contagiousness of consumption in sector k (or variety k of a consumption good). We normalize this function to integrate to unity,

$$\int \phi(k) dk = 1. \quad (4)$$

Consider an agent j , who is still ‘susceptible’: we denote this agent therefore with ‘ s ’ rather than j . This agent is consuming the bundle $(c_{jk}^s)_{k \in [0,1]}$ at date t . Symmetrically, let $(c_{jk}^i)_{k \in [0,1]}$ denote the consumption bundle of infected people. Extending ERT, we assume that the probability τ_t for an agent of type s to become infected depends on his own consumption bundle, on the total mass of infected people and their consumption choices and the degree $\phi(k)$ to which infection can be spread per unit of consumption in sector k ,

$$\tau_t = \pi_{s,t} I_t \int \phi(k) c_{jk}^s c_{jk}^i dk, \quad (5)$$

where $\pi_{s,t}$ is a parameter for the social-interaction infection risk, and which we allow to exogenously vary over time. In our quantitative exercise in Section 5, we restrict this variation to conform to a ‘winter-summer’ pattern. With Equation (5), the total number of newly infected people is given by

$$T_t = \tau_t S_t. \quad (6)$$

The dynamics of the four groups now evolves as in a standard SIR epidemiological model,

$$S_{t+1} = S_t - T_t \quad (7)$$

⁷ One could interpret one low-infectious sector as home production in the context of the model. When mapping the model to the data, however, the output of that sector would then not be included in measured National Income and Product Accounts (NIPA) market consumption expenditures.

$$I_{t+1} = I_t + T_t - (\pi_r + \pi_d)I_t \quad (8)$$

$$R_{t+1} = R_t + \pi_r I_t \quad (9)$$

$$D_{t+1} = D_t + \pi_d I_t \quad (10)$$

$$\text{Pop}_{t+1} = \text{Pop}_t - D_t, \quad (11)$$

where π_r is the recovery rate and π_d is the death rate and where Pop_t denotes the mass of the total population at date t . As in ERT, we assume that the epidemic starts from initial conditions $I_0 = \epsilon$ and $S_0 = 1 - \epsilon$, as well as $R_0 = D_0 = 0$.

2.3. Choices

We proceed to analyse the choices of the individuals.

2.3.1. Susceptible people. Denote as $U_t^s(U_t^i)$ the lifetime utility, from period t onwards, of a currently susceptible (infected) individual. As in ERT, the lifetime utility U_t^s follows the recursion

$$U_t^s = u(c_t^s, n_t^s) + \beta[(1 - \tau_t)U_{t+1}^s + \tau_t U_{t+1}^i], \quad (12)$$

where the probability τ_t is given in Equation (5) and depends on the choice of the consumption bundle $(c_{tk}^s)_{k \in [0,1]}$. An s -person maximizes the right-hand side of Equation (12) subject to the budget constraint (3) and the infection probability constraint (5), by choosing labour n_t^s , the consumption bundle $(c_{tk}^s)_{k \in [0,1]}$ and the infection probability τ_t .

The first-order condition for consumption of c_{tk}^s is

$$u_1(c_t^s, n_t^s) \cdot \left(\frac{c_t^s}{c_{tk}^s} \right)^{1/\eta} = \lambda_{bt}^s + \lambda_{\tau t} \pi_{s,t} I_t \phi(k) c_{tk}^i, \quad (13)$$

where λ_{bt}^s and $\lambda_{\tau t}$ are the Lagrange multipliers associated with the constraints (3) and (5). This equation can be rewritten as

$$u_1(c_t^s, n_t^s) \cdot \left(\frac{c_t^s}{c_{tk}^s} \right)^{1/\eta} = \lambda_{bt}^s + \nu_t \phi(k) c_{tk}^i, \quad (14)$$

where

$$\nu_t = \pi_{s,t} I_t \lambda_{\tau t}. \quad (15)$$

Equation (14) reveals that the risk of becoming infected induces an additional goods-specific component, scaled with the aggregate multiplier ν_t , compared with the usual

first-order conditions for Dixit–Stiglitz consumption aggregators (at constant prices across goods). In the absence of the impact of consumption on infection $\lambda_{\pi t} = \nu_t = 0$ and there is no consumption heterogeneity across sectors, $c_{ik}^s = c_t^s$ for all k , as in the standard model. In the presence of this effect, then susceptible households shift their consumption to sectors with a low risk of infection (i.e., those with a low $\phi(k)c_{ik}^i$).

Taking the consumption profile of infected households (c_{ik}^i) as given, by choosing her consumption portfolio a susceptible individual effectively chooses her infection probability τ_t . As in ERT, the first-order condition for τ_t reads as

$$\beta(U_{t+1}^s - U_{t+1}^i) = \lambda_{\tau t}. \quad (16)$$

The first-order condition with respect to labour is completely standard and reads as

$$u_2(c_t^s, n_t^s) + A\lambda_{bt}^s = 0. \quad (17)$$

Note that we have excluded the workplace infection, in contrast to ERT. We examine this possibility in Section 3.3. With the chosen utility function, this first-order condition simplifies to

$$\theta n_t^s = A\lambda_{bt}^s. \quad (18)$$

2.3.2. Infected people and recovered people. As in ERT, the lifetime utility of an infected person is

$$U_t^i = u(c_t^i, n_t^i) + \beta[(1 - \pi_r - \pi_d)U_{t+1}^i + \pi_r U_{t+1}^r + \pi_d \times 0]. \quad (19)$$

Taking first-order conditions with respect to the consumption choices and labour results in

$$u_1(c_t^s, n_t^s) \cdot \left(\frac{c_t^i}{c_{ik}^i} \right)^{1/\eta} = \lambda_{bt}^i, \quad (20)$$

where λ_{bt}^i is the Lagrange multiplier on Equation (3) for an infected person. This is the usual Dixit–Stiglitz CES first-order condition at constant prices, with the solution

$$c_{ik}^i \equiv c_t^i. \quad (21)$$

That is, as long as $\eta \in (0, \infty)$, infected individuals find it optimal to spread their consumption evenly across sectors, given that all sector goods have the same price, are imperfect substitutes and differential infection probabilities across sectors are irrelevant for already infected individuals. Exploiting this result and the specific form in Equation (1)

for the period utility function in Equation (20) yields $1/c_t^i = \lambda_{bt}^i$. For labour, we obtain the standard first-order condition

$$\theta n_t^i = A \lambda_{bt}^i = \frac{A}{c_t^i}. \quad (22)$$

Finally, exploiting the budget constraint (3), we arrive at the equilibrium allocations for infected people given by

$$n_t^i = \frac{1}{\sqrt{\theta}}, c_t^i = \frac{A}{\sqrt{\theta}}. \quad (23)$$

Likewise, the lifetime utility for a recovered person is

$$U_t^r = u(c_t^r, n_t^r) + \beta U_{t+1}^r. \quad (24)$$

Given our assumptions, the optimal decision for both the i group and r group is the same⁸: we will therefore use c_t^i , $c_{t,k}^i$, n_t^i and λ_{bt}^i to also denote the choices of recovered individuals.

2.4. Non-separable utility and invariance

In this section, we briefly comment on the impact of the choice of the utility function on our results. For the quantitative analysis, we will use the separable benchmark specification in Equation (1). The purpose of this section is to examine a general period utility functions $u(c, n)$, allowing for non-separability between labour and consumption, and to find conditions, when the results are not affected, up to first order, by the specific choice for $u(\cdot, \cdot)$. The technical details of this analysis are contained in Appendix B. There we show that, to a first-order approximation, two different utility functions will deliver the same values for c_t^i , $(c_{tk}^s)_{k \in [0,1]}$ and n_t^i , if the two utility functions, at the steady state, agree on $\psi = u_{11} + \frac{2}{A} u_{12} + \frac{1}{A^2} u_{22}$, on u_1 as well as on $u_{11} + u_{12}/A$. Importantly, as is clear from Equation (14), the sectoral reallocation forces determining relative consumption across the different k -sectors of the economy are operative regardless of whether the period utility function is separable in consumption and labour or not.

8 Note here that we implicitly assume that infected people will be fully at work. One might alternatively wish to assume that only a fraction of them are at work instead. Given our assumptions about excluding infections in the workplace, this does not affect the infection rate via that channel. However, lowering the amount of income of infected people lowers their consumption and thus lowers their ability to infect others in the consumption market. We do not wish to emphasize this channel, although in a somewhat richer model, people will have a buffer stock of savings and an infected person would then draw on these savings to finance consumption rather than respond to the temporary decline in labour income. Alternatively, income may fall considerably less in practice than the model would otherwise imply here, due to various social insurance policies.

2.5. Equilibrium characterization

In equilibrium, each individual solves her or his maximization problem and the labour and goods market clears in every period. Let n_{tk} be total labour employed in sector k . The market-clearing conditions then read as

$$S_t c_{tk}^s + (I_t + R_t) c_{tk}^i = A n_{tk}. \quad (25)$$

$$\int n_{tk} dk = S_t n_t^s + (I_t + R_t) n_t^i. \quad (26)$$

Given the solution to the problem of infected and recovered people, this can be simplified to

$$S_t c_{tk}^s + (I_t + R_t) \frac{A}{\sqrt{\theta}} = A n_{tk} \int n_{tk} dk = S_t n_t^s + (I_t + R_t) \frac{1}{\sqrt{\theta}}.$$

The equations can be simplified further to a set of aggregate variables as well as an equation determining the sectoral allocation; see Appendix C.

3. THEORETICAL RESULTS

3.1. The value of a statistical life

The calculations above allow us to calculate the implied value of a statistical life or VSL, using the following thought experiment. Suppose that there is no epidemic but that an agent may be exposed to some small probability $\delta > 0$ of not surviving to the next period. How much would current consumption have to be increased in order to compensate the agent for this additional risk? That is, when would a riskless scenario and a risky scenario compensated through increased current consumption be equal in the eyes of an agent?

Without the epidemic, the consumption of all agents is given by [Equation \(23\)](#) or

$$c^* = \frac{A}{\sqrt{\theta}} \text{ and } n^* = \frac{1}{\sqrt{\theta}},$$

implying current utility and lifetime utility

$$u^* = \log\left(\frac{A}{\sqrt{\theta}}\right) - \frac{1}{2} \text{ and } U^* = \frac{u^*}{1 - \beta} u^*.$$

In the riskless scenario, the agent receives lifetime utility

$$U^* = u^* + \beta U^*.$$

In the risky scenario, the agent receives the expected lifetime utility

$$\log \left(e^{\gamma} \frac{A}{\sqrt{\theta}} \right) - \frac{1}{2} + (1 - \delta)\beta U^*,$$

where γ is the compensating increase of current consumption, expressed in % (and divided by 100). Equating these two expressions, solving for γ and with the value of a statistical life expressed as the % increase in the consumption of one period required to compensate for a 1% increase in the death probability,

$$\text{VSL} := \frac{\gamma}{\delta}, \quad (27)$$

we find

$$\text{VSL} = \frac{\beta}{1 - \beta} \left(\log \left(\frac{A}{\sqrt{\theta}} \right) - \frac{1}{2} \right). \quad (28)$$

One now needs to keep in mind that the length of the period matters: for example, that percentage increase would need to be four times as large when calculated for weekly rather than monthly consumption. To calculate the actual value of a statistical life, say, in Dollars, one needs to multiply VSL with the Dollar amount of consumption of one period.

With the epidemic as described in the model, one can thus consider susceptible agents thinking of reducing the current consumption aggregate by VSL percent, if this allows them to increase their survival chances by 1% in the next period: the larger VSL, the more the agent is willing to endure a reduction in consumption. It is this trade-off that gives rise to the implied consumption dynamics in our model. Equation (28) shows that these calculations regarding the value of a statistical value of life depend on the preference and productivity parameters via β as well as the ratio $A/\sqrt{\theta}$ or steady-state consumption c^* . Furthermore, Equation (28) shows that this value can be negative in principle, if $A/\sqrt{\theta}$ is smaller than $\exp(0.5) \approx 1.65$. We will avoid that in setting our parameters, as it would imply a desire for agents to die.

3.2. Two extreme values of the elasticity of substitution η

It is instructive to consider extreme values for the elasticity of substitution η . We obtain rather tight predictions as these extremes, which turn out to be useful for understanding the quantitative results and the impact of parameter variations for the calibrated version, see Section 7.2.

The first extreme is an elasticity of substitution of zero such that the consumption aggregator is of the Leontief form.

Proposition 1. *Suppose that $\eta = 0$, that is, that the consumption aggregation in Equation (2) is Leontief. In this case, the multisector economy is equivalent to a multisector economy with a ϕ -function, which is constant and equal to 1,*

Proof. With Leontieff consumption aggregation, consumption is sector independent, $\bar{c}_{ik}^j \equiv \bar{c}_t^j$. Equations (5) and (6) now become

$$\tau_t = \pi_{s,t} I_t \int \phi(k) c_t^s c_t^i dk = \pi_{s,t} I_t c_t^s c_t^i \int \phi(k) dk = \pi_{s,t} I_t c_t^s c_t^i \quad (29)$$

and

$$T_t = \pi_{s,t} S_t I_t \int \phi(k) c_{ik}^s c_{ik}^i dk = \pi_{s,t} S_t I_t c_t^s c_t^i \quad (30)$$

Equations (29) and (30) furthermore show that the Leontief version is equivalent to the one-sector economy in ERT. The other extreme is the case where goods are perfect substitutes.

Proposition 2. *Suppose that $\eta \rightarrow \infty$, that is, that the sector-level consumption goods in Equation (2) are perfect substitutes in the limit. Let $\underline{k} = \sup_k \{k | \phi(k) = \phi(0)\}$. Assume that $\underline{k} > 0$, that is, that there is a non-zero mass of sectors with the lowest level of infection interaction. Suppose that $I_0 > 0$. Then there is a limit consumption $\bar{c}_{ik}^j$ for $j \in \{s, i, r\}$ as $\eta \rightarrow \infty$, satisfying*

$$c_{ik}^s = \begin{cases} c_t^s / \underline{k} & \text{for } k < \underline{k} \\ 0 & \text{for } k > \underline{k} \end{cases} \quad (31)$$

and

$$\bar{c}_{ik}^j \equiv \bar{c}_t^j \text{ for } j \in \{i, r\}. \quad (32)$$

Equations (5) and (6) are replaced by

$$\tau_t = \pi_{s,t} \phi(0) I_t c_t^s c_t^i \quad (33)$$

and

$$T_t = \pi_{s,t} \phi(0) S_t I_t c_t^s c_t^i. \quad (34)$$

That is, susceptible individuals only consume in the lowest infection-prone sectors with $\phi(k) = \phi(0)$, and infected (as well as recovered) individuals consume uniformly across all sectors.

Proof. Equation (32) is just Equation (21), which also holds for recovered agents: it will therefore also hold, when taking⁹ the limit $\eta \rightarrow \infty$. Equation (31) follows from

9 Note that it does not necessarily hold at the limit, as infected and recovered agents in that case are indifferent as to which goods to consume

Equation (15) together with Equation (2), taking $\eta \rightarrow \infty$. Define the consumption distribution of type $j \in \{s, i, r\}$ as $\kappa_t^j(k) = c_{tk}^j / c_t^j$ and note that

$$\int \kappa_t^j(k) dk = 1 \quad (35)$$

and that

$$\kappa_t^j(k) \geq 0, \text{ all } k \quad (36)$$

Rewrite Equations (5) and (6) as

$$\tau_t = \pi_{s,t} I_t c_t^s \int \phi(k) \kappa_t^s(k) \kappa_t^i(k) dk. \quad (37)$$

Therefore, and analogously to ERT, the total number of newly infected people is given by

$$T_t = \pi_{s,t} S_t I_t \int \phi(k) \kappa_t^s(k) \kappa_t^i(k) dk. \quad (38)$$

Equations (33) and (34) now follow from observing that $\kappa_t^i(k) \equiv 1$ and $\kappa_t^s(k) = 1/\underline{k}$ for $k \in [0, \underline{k}]$ and zero elsewhere as well as noting that $\phi(k) = \phi(0)$ for $k \in [0, \underline{k}]$. $\square$

The result of this proposition is depicted in Figure 3. Equations (33) and (34) also show that the limit is equivalent to the one-sector economy in ERT, with $\pi_{s,t}$ replaced by $\pi_{s,t}\phi(0)$. Infection only takes place in the sector with the lowest infection hazard, thus introducing the extra factor $\phi(0)$. The size of the sector, however, does not factor into the expression. The reason is that with an equal distribution of infected agents across all sectors, susceptible agents in a smaller sector meet a smaller fraction of infected agents – a mitigating force. At the same time, the consumption activity of susceptible agents in these sectors rises – an enhancing force. These two forces cancel each other out. Since the size of the sector with the lowest infection hazard does not matter at both extremes given in Propositions 1 and 2, one might conjecture that it is never relevant. However, numerical simulations indicate that larger rates of infection occur, if that sector is smaller, for substitution elasticities $0 < \eta < \infty$.

Proposition 2 above exploits the fact that infected agents wish to spread their consumption equally across all sectors for any finite value of η . However, at the limit $\eta = \infty$, infected agents are entirely indifferent. At one extreme, they might consume rather large portions of the low- k goods. At the other extreme, they stick to each other in the high-infection-risk segments and do not consume the k -goods at all. In the latter case, the infection probabilities become zero and the spread of the disease is stopped entirely. The following proposition provides the resulting range for the infection probabilities.




Figure 3. When $\eta \rightarrow \infty$

Proposition 3. Suppose that $\eta = \infty$, that is, that the sector-level consumption goods in Equation (2) are perfect substitutes. Let μ_t be any function of time satisfying

$$0 \leq \mu_t \leq \bar{\mu},$$

where $\bar{\mu}$ is defined as

$$\bar{\mu} = \frac{1}{\int \frac{1}{\phi(k)} dk} \quad (39)$$

and note that it satisfies

$$\phi(0) \leq \bar{\mu} \leq 1. \quad (40)$$

Then there is an equilibrium with Equations (5) and (6) replaced by

$$\tau_t = \pi_{s,t} \mu_t I_t c_t^s c_t^i \quad (41)$$

and

$$T_t = \pi_{s,t} \mu_t S_t I_t c_t^s c_t^i. \quad (42)$$

Proof. We first show Equation (40). For the lower bound, note that

$$\int \frac{1}{\phi(k)} \leq \int \frac{1}{\phi(0)} = \frac{1}{\phi(0)}.$$

The upper bound follows from Jensen's inequality and Equation (4). We next shall show that there is an equilibrium when μ_t equals one of the two bounds. Given the consumption distribution function κ_t^i , the problem of the susceptible agents is to choose

their own consumption distribution function κ_t^i so as to minimize Equation (37), subject to the constraints (35) and (36). The Kuhn–Tucker first-order condition implies that $\kappa_t^s(k) = 0$, unless

$$k \in \{k | \phi(k)\kappa_t^i(k) = \min \phi(k)\kappa_t^i(k)\}.$$

For $\mu_t = 0$, let infected agents consume zero, $\kappa_t^i(k) = 0$ for all k in some subset $\mathcal{K}$ of $[0, 1]$. In this case and following the argument just provided, susceptible people choose $\kappa_t^s(k) > 0$ only if $k \in \mathcal{K}$. Conversely, the worst-case scenario, in terms of infection, arises if $\phi(k)\kappa_t^i(k)$ is constant. Given Equation (35), this yields

$$\kappa_t^i(k) = \frac{\bar{\mu}}{\phi(k)}. \quad (43)$$

Given this κ_t^i function, susceptible agents are now indifferent in their consumption choice. Any κ_t^s function satisfying Equation (35) and (36) then results in

$$\int \kappa_t^s \phi(k) \kappa_t^i(k) dk = \bar{\mu}$$

and thus Equations (41) and (42) at $\mu_t = \bar{\mu}$, that is, the upper bound. Finally, let $0 < \mu_t < \bar{\mu}$ and let $\lambda = \mu_t / \bar{\mu}$. Let $\mathcal{K}$ be a measurable subset of $[0, 1]$ with mass strictly between 0 and 1. Set

$$\kappa_t^i(k) = \begin{cases} \lambda \frac{\bar{\mu}}{\phi(k)}, & \text{for } k \in \mathcal{K} \\ \tilde{\lambda} \frac{\bar{\mu}}{\phi(k)}, & \text{for } k \in [0, 1] \setminus \mathcal{K} \end{cases},$$

where $\tilde{\lambda}$ is chosen such that Equation (35) holds. Then, susceptible agents will choose $\kappa_t^s(k) = 0$ for all $k \in [0, 1] \setminus \mathcal{K}$, are indifferent between $k \in \mathcal{K}$ and Equations (41) and (42) hold true for the chosen μ_t . $\square$

The result of this proposition is depicted in Figure 4. The top graph depicts the best-case scenario when susceptible and infected agents consume entirely different bundles of goods. The bottom graph depicts the worst-case scenario when $\phi(k)c^i(k)$ is constant: in that case, a constant $c^s(k)$ is the best solution for susceptible agents. The bottom graph and the proposition show that perfect substitutability may be nearly as bad as the Leontief case, if infected people behave particularly badly and distribute their consumption according to Equation (43). Equations (41) and (42) are then the same equations as in the ERT model with $\pi_{s,t}$ replaced by $\pi_{s,t}\bar{\mu}$. On the other hand, perfect substitutability can also result in the most benign scenario of a zero spread of consumption, if infected and susceptible people simply consume different goods.




Figure 4. When $\eta = \infty$

There are fascinating policy lessons here. Given that infected people will end up seeking services and consumption, it might be best to encourage them to seek out those where the degree of interaction is high, rather than forcing all agents, including the infected agents, into the low infection transmission segments. The model here shows that this can have dramatic consequences for the spread of the disease.

3.3. Infections in the labour market

Thus far, we have assumed that infections can take place when acquiring consumption goods. We could have similarly allowed for heterogeneity in labour and assumed that it is at work in the labour market where individuals face the risk of contracting the virus.¹⁰ We explore this possibility in this section, offering two alternative approaches. We shall

10 Mulligan (2021) provides an argument and evidence that workplaces might have been safer than being at home during the pandemic.

show that the formal analysis is conceptually similar and, that the first approach is actually equivalent to the model analysed above. In economic terms and interpretation, the key distinction is arguably less in the formal differences between both versions of the model, than in the empirically plausible choice for the elasticity of substitution η : while it may be possible to easily substitute among different types of similar consumption goods ('Pizza at home' versus 'Pizza in a restaurant'), the same may not be true for work (restaurants will still have to produce the to-be-delivered pizza in the restaurant kitchen, rather than having their workers stay at home and produce in their own kitchens). Our results for the lower elasticity of substitution $\eta = 3$ may thus be more appropriate for the analysis of infection-at-the-work-place. In the extreme without substitution possibilities, we are back at the homogeneous sector case.

The formal analysis maintains the assumption that the period utility function is given by

$$u(c, n) = \log(c) - \theta \frac{n^2}{2} \quad (44)$$

but now assumes that consumption c is a homogeneous good, while household labour n is a composite of differentiated sector-specific labour n_k , $k \in [0, 1]$. For the former, production-based approach to aggregation, we assume that labour supplied by the household to the market is a Constant Elasticity of Substitution (CES) composite of sector-specific labour, that is, that $k \in [0, 1]$ as far as preferences are concerned, but that the budget constraint reads

$$c = A \left(\int n_k^{1-(1/\alpha)} dk \right)^{\alpha/(\alpha-1)}. \quad (45)$$

Assume now that infections occur in the labour market instead of through joint consumption, that is, assume that the probability of a susceptible individual to become infected is given by

$$\tau_t = \tilde{\pi}_s I_t \int \phi(k) n_{ik}^i n_{ik}^s dk. \quad (46)$$

In Appendix D, we establish the following result.

Proposition 4. *Suppose that $\tilde{\pi}_s = A^2 \pi_{s,t}$ and that $\alpha = \eta$. In this case, the production-based labour aggregation with infection in the labour market is equivalent to the consumption-infection economy described above, that is all aggregates remain the same, while $n_{ik}^s/n_t^s = c_{ik}^s/c_t^s$, when comparing the ratio of sector-specific labour to the labour aggregate in the labour-market-infection economy to the ratio of sector-specific consumption to the consumption aggregate in the consumption-infection economy.*

For this formal equivalence, it is important that the aggregation Equation (45) takes place at the household level and not at the firm level, that is, in firms hiring labour from

different households. The latter would provide an interesting alternative environment for studying the sectoral shift issues raised here but requires additional restrictions to preclude complete separation of infected and susceptible agents in equilibrium.

The household-level labour aggregation described above may be hard to envision as an environment for sector-specific contagion risk. We therefore offer a second, preference-based approach. For this, think of the household as composed of individual workers, each specialized to work in sector k , and that total household leisure, described by $\ell = f(n)$ for some strictly decreasing and differentiable function¹¹ f is a CES-aggregate of worker-specific leisure,

$$f(n) = \left(\int f(n_k)^{1-1/\alpha} dk \right)^{\alpha/(\alpha-1)} \quad (47)$$

for some elasticity of substitution $\alpha \geq 0$. The household budget constraint is $c = A \int n_k dk$. The probability of infection is given by Equation (46). This economy shares the same basic forces as the heterogeneous consumption sector economy, although its analysis is not exactly equivalent. In Appendix D, we demonstrate this more formally. The remarks here are simply meant to show that the mechanisms in both types of labour-infection-based models are quite similar to our baseline consumption-based-infection economy indeed.

A richer treatment would recognize that matters are surely more complicated than our rather stylized treatment. It is beyond the scope of this paper to investigate these issues empirically and theoretically with considerable depth. Consider our baseline example of eating Pizza at home or in a restaurant. Casual observation suggests that there has been little to no relative price shift here, which then would be indicative of the linear production and labour mobility assumption used in the paper. This also applies to, say, books purchased in a bookstore when compared with purchases by an online retailer such as Amazon, etc. However, it is possible that the linear production and labour mobility assumption is not fully justified: in theory, one would then observe relative price movements between these two types of consumption possibilities as well as corresponding movements in wages on free labour markets, but they may not have been present. This may have been true perhaps for the same reason that shortages of certain goods early on in the pandemic ('toilet paper') did not lead to price shifts either and that labour contracts tend to be long-term and that statutory wages do correspond to short-term productivities. We, therefore, skip a full quantitative analysis and do not integrate these features into the ensuing analysis.

11 Useful specifications are $f(n) = \bar{L} - n$ for some time endowment $\bar{L}$ or $f(n) = 1/n$.

4. SOCIAL PLANNING PROBLEM

The laissez-faire solution examined implies that externalities are still present, as agents do not take into account their action on the increased infection risk of others. Therefore, it is instructive to compare our results to that of a social planner ‘on the same footing’, that is, with full knowledge of who is susceptible, infected or recovered. Indeed, the increasingly prevalent opportunities for COVID-19 testing make this perhaps a more reasonable assumption rather than forcing the social planner to act behind the ‘veil of ignorance’. However, like the agents in our model, the planner cannot separate the infected from the susceptible (and recovered) when they consume. That is, the planner cannot change the consumption technology and separate infected and susceptible individuals during their consumption activities. Put differently, we do not allow for the ‘Chinese solution’, of cordoning off entire cities in which COVID-19 cases have been observed. Therefore, as in the decentralized economy, the spread of the disease while consuming can at best be mitigated, by allocating consumers to low-infectious sectors. The social planner knows all what the agents know, and can do all that agents can do, but not more and not less. As an additional point of comparison, we also consider the problem of a social planner that does not know the health status of individuals and thus cannot condition consumption allocations on this health status.¹²

For the full-knowledge situation here, the social planner maximizes date-0 aggregate social welfare W_0 , where

$$W_0 = \sum_{t=0}^{\infty} \beta^t [S_t u(c_t^s, n_t^s) + I_t u(c_t^i, n_t^i) + R_t u(c_t^r, n_t^r)]$$

subject to the following constraints, and with the respective Lagrangian multipliers, after substituting the infection risk for susceptible people, τ_b , and the number of newly infected people, T_t :

$$\mu_{f,t} : \int S_t c_{tk}^s + I_t c_{tk}^i + R_t c_{tk}^r dk = A(S_t n_t^s + I_t n_t^i + R_t n_t^r) \quad (48)$$

$$\mu_{S,t} : S_t = S_{t-1} - I_t + (1 - \pi_r - \pi_d)I_{t-1} \quad (49)$$

$$\mu_{I,t} : I_t = \pi_{s,t} S_{t-1} I_{t-1} \int \phi(k) c_{t-1,k}^s c_{t-1,k}^i dk + (1 - \pi_r - \pi_d)I_{t-1} \quad (50)$$

$$\mu_{R,t} : R_t = R_{t-1} + \pi_r I_{t-1}. \quad (51)$$

The social planner takes S_0 , I_0 and R_0 as given. It chooses the time paths of consumption for susceptible, infected and recovered people c_{kt}^x for $x \in \{s, i, r\}$, the path for labour supply n_t^x , $x \in \{s, i, r\}$ and the paths for the mass of agents in the four groups S_b , I_t

12 Appendix F spells out the details of this version of the social planner.

and R_t . The first-order conditions of the social planner's problem are presented in Appendix E.

5. COMPUTATION, CALIBRATION AND ESTIMATION

In this section, we discuss the numerical solution, the calibration and the estimation of the model.

5.1. SIR mechanics

Prior to the discussion of the calibration of our model, we now describe the basic thought process that led us to augment the basic SIR model by seasonal infection rates, that is, let $\pi_{s,t}$ vary between the summer and the winter. We seek to match our model reasonably closely to the evolution of the pandemic and macroeconomic data in Sweden. As shown in Figure 2, Sweden has been hit by two waves of infections and deaths from COVID-19, one in April and May of 2020 and then a second wave starting sometime in October 2020. For the purposes of crafting a quantitatively plausible model, these observations present two challenges, with the second following from the first. We now describe these challenges and how we have resolved them by augmenting the basic SIR model, combined with our economic mechanism of sectoral reallocation.

The first challenge emerges from the mere existence of that second wave. The standard SIR model with fixed rates of infection, recovery and death when multiplied with the appropriate population fractions, cannot by itself produce such a second wave. The additional economic choices considered here, when combined with the SIR epidemiological model will not produce a second wave either.¹³ Therefore, capturing the second wave successfully in the model requires a modification of the standard SIR mechanics: for some reason, the infection rate parameter or the death rate parameter must have increased in the fall. We resolve this challenge by introducing seasonality in the infection risk parameter $\pi_{s,t}$, which rises in late September 2020. This modelling choice can be supported by arguing that there is a seasonality to the biological infection risk, as is generally the case for influenza viruses.¹⁴

That second wave creates an additional and more subtle challenge to the standard SIR dynamics: the number of deaths in that second wave is roughly of the same order of magnitude as the number of deaths in the first wave. This implies that the susceptible fraction of the population must still have been reasonably high at the beginning of fall of 2020. Explorations of a simple SIR model with a seasonal increase in the infection risk

13 We suspect that extreme values of SIR model parameters may produce oscillating infection patterns. We have not investigated this possibility further.

14 Since in the fall and winter seasons, social interactions predominantly take place indoors and infection probabilities are higher for indoor than for outdoor interactions, we can then deduce that infection rates will rise in the colder half of the year, as we model here.

in October suggest that a rough quantitative match of the death dynamics in the second wave is impossible to obtain unless the susceptible fraction is well above 50% at the beginning of October 2020. We will now argue that without a seasonal fall in infection rates in the spring of 2020, it is hard to avoid the conclusion, based on standard SIR logic, that Sweden must have been approaching herd immunity in the fall of 2020.

To see this, consider then the dynamics of the first wave in the data. Compare, say, the week beginning 28 March with 282 COVID-19 deaths, followed by a week with 634 COVID-19 deaths (call this the ‘first phase’) to the week beginning 30 May with 267 COVID-19 deaths, followed by a week with 252 COVID-19 deaths (call this the ‘second phase’). In the standard SIR model, these numbers are proportional to the number of those infected in these weeks. These numbers therefore imply that the number of infected individuals was approximately the same at the beginning of each phase, but that the fraction of infected agents increased by a replication rate of $634/282 = 2.25$ during that phase and declined at the rate of $252/267 = 0.94$ during the second episode. The standard SIR model implies that this change in the replication rate is mostly driven by a change in the susceptible fraction of the population: in that model, the fraction of newly infected persons is proportional to the product of the fractions of infected and susceptible agents in the economy. As a back-of-the-envelope calculation, the fraction of susceptible agents must have therefore declined roughly by the factor $0.94/2.25 = 0.42$, that is, the share of the population that is still susceptible should be below 50% of the population by the end of the second phase of the first wave, given the observed dynamics of weekly deaths in Sweden. As argued before, this is impossible to reconcile with the size and shape of the second wave. Indeed, with the standard SIR mechanics, it would be hard to avoid the conclusion that Sweden was approaching herd immunity as of early fall. Clearly, that did not happen.¹⁵

There are broadly two ways to resolve this second challenge. One is to not only model a seasonal increase in infection rates in the fall to start the second wave, but also a seasonal decline in the infection risk parameter (or death parameter) in the late spring. This is the route we pursue in our benchmark calibration: we assume a step function, with $\pi_{s,t}$ at a high ‘winter’ value until 27 March, then a low ‘summer’ value until 25 September 2020, and then a reversion back to the high ‘winter’ value from 26 September onwards.¹⁶

As an alternative to this baseline scenario, we examine the power of the economic forces by turning them off until 27 March in Appendix H. Essentially, this makes our susceptible agents ignorant, myopic or irrational in the face of the initial rise of death

15 In a 24 November 2020 briefing Anders Tegnell, Sweden’s top state epidemiologist, stated ‘We see no signs of immunity in the population that are slowing down the infection right now’. Of course, the massive second wave of deaths in the winter of 2020 is the most direct evidence that herd immunity was not achieved by the fall of 2020 in Sweden.

16 It is worth noting that the observed death dynamics can be matched perfectly with a standard SIR model and an appropriately chosen time variation in $\pi_{s,t}$. Our two-value winter–summer–winter step function for $\pi_{s,t}$ puts considerable restriction on that choice.

rates during the first wave. The first phase of the first wave then evolves according to the standard SIR model (with the moderate infection rates underlying the lower panel of Figure 2) and the reversal of infections and (with a lag) weekly deaths in the first wave occurs once individuals understand the pandemic and adjust their consumption behaviour accordingly. One can debate, whether the assumption of temporary myopia (or ignorance) when confronted with an entirely novel pandemic threat is reasonable or not, in the same way one can question the fully forward-looking manner individuals behave in the benchmark. What we want to stress in both incarnations of the model is the force of the endogenous consumption choice and its power to turn a pandemic around when susceptible agents are highly aware of the infection risks associated with unchanged patterns of behaviour.

5.2. Computation of the model

The unknown sequences of variables to be computed (aside from the sector-specific consumption) are $U_t^s, c_t^s, n_t^s, \lambda_{bt}^s, \nu_t, \tau_t$. The dynamic equations determining these variables are the Bellman equation (Equation (12)), the budget constraint (Equation (3)), the infection constraint (Equation (C.8)), the share constraint (Equation (C.5)) replacing the original first-order condition with respect to consumption, the first-order condition with respect to labour (Equation (17)) and the first-order condition with respect to τ (Equation (16)) combined with Equation (15). One can easily eliminate λ_{bt}^s and n_t^s , using Equations (3) and (17), as well as eliminate ν_t with Equations (15) and (16): what remains is then a system in three unknowns U_t^s, c_t^s, τ_t and three equations, two of which are non-linear integral equations that need to be solved. The way to proceed is working backwards from a distant horizon. Knowing U_{t+1}^s allows one to compute ν_t with Equations (15) and (16). Using the two integral equations (having substituted out λ_{bt}^s and n_t^s) allows one to compute c_t^s and τ_t . From there, we can compute n_t with Equation (3) and U_t^s .

We use Dynare 4.6 to perform the calculations. We assume that the economy is initially in a steady-state and is then hit by a zero probability, MIT shock that increases the infected population from 0 to ε . There are no further aggregate shocks and we compute a perfect foresight solution with 100 periods. In order to facilitate finding a solution, a good initial guess is helpful. We first set the exogenous $\pi_{s,t}$ to a very small and constant value. With this value, the economy is close to one without a disease outbreak and using this solution as the initial guess makes it easy for Dynare to find a solution to the economy with the disease outbreak and an incrementally increased $\pi_{s,t}$, restricted to the same value $\pi_{s,\text{'winter'}}$ in the 'winter' months, October to March, and a lower value $\pi_{s,\text{'summer'}}$ in the 'summer' months, from the last week of March to September. Proceeding in steps, using the solution for the previous $\pi_{s,t}$ value as the initial state, a new solution is computed. The initial state is updated once a solution at the new $\pi_{s,t}$ value is found. This process is repeated until $\pi_{s,t}$ reaches the desired value in our calibration and estimation procedure.

5.3. Calibration and estimation

Our quantification focuses on a two-sector economy, where both sectors are of equal size denoted by v . This results in $v = 0.5$. Sector 1 has infection intensity ϕ satisfying $0 < \phi(k) = \phi < 1$ for $k \in [0, 0.5]$. Given the maintained assumption that the average $\phi(k)$ is equal to one, this implies the infection intensity for $k \in (0.5, 1]$ is calculated as $2 - \phi$. Accordingly, sector 1 is named the low-infection sector and sector 2 is the high-infection sector. The one-sector economy of ERT emerges when $\phi = 1$. For the elasticity of substitution across goods, we set $\eta = 3$ and investigate, as sensitivity analysis, $\eta = 10$, in which case the two goods are close to perfect substitutes, and thus the equilibrium allocation approaches the characterization in Proposition 2.

We choose the remaining parameters of our model in two steps. In the first step and for the ‘economics’ side, we choose the preference parameters (θ, β) and the productivity parameter A in accordance with ERT which employ a one-sector version of the model we use here to study the joint dynamics of health and economic activity during the COVID-19 crisis in the United States. We summarize these parameter choices in Table 1. To further assess the plausibility of these choices, recall from Section 3.1 that given these parameters, Equation (28) can be used to calculate the model-implied value of a statistical life VSL, measured as the percentage change in one-week consumption to compensate for a 1% change in the probability of death. Given the parameters in Table 1, we compute $VSL = 8298$, which, with weekly per-capita consumption in Sweden of about \$500, implies a value of a statistical life in Sweden of ca. \$4 million. This value strikes us as plausible and is in the ballpark of numbers often used in the economics literature.¹⁷

In a second step, we estimate the parameters governing the epidemic dynamics. Due to the probability of a significant measurement error for the number of infected individuals, especially at the beginning of the pandemic, we chose the number of weekly deaths as the most reliable target for the calibration. A time series plot of the Swedish data shows a clear seasonal pattern. As argued in Section 5.1, this led us to impose that $\pi_{s,t}$ switches from a high value $\pi_{s,\text{‘winter’}}$ to a low value $\pi_{s,\text{‘summer’}}$ in the last week of March and then back to $\pi_{s,\text{‘winter’}}$ in October.

The π_r and π_d parameters are associated with the number of days an infected individual stays infected (before recovering or dying) and the case fatality rate, by using the following equations¹⁸:

17 For example, Greenstone and Nigam (2020) set the VSL to \$11.5 million, a figure based on work by the Environmental Protection Agency (EPA). Hall and Jones (2007) discuss the empirical literature on the VSL and report numbers (in 2,000 dollars) between \$2 million and \$9 million, with substantial heterogeneity by age. In our infinite horizon economy, the VSL corresponds most closely to that of a young- to a middle-aged individual with a substantial remaining lifetime.

18 Recall that our model is solved at weekly frequency.

Table 1. Economic parameters

Parameter	Value	Description
η	3.000	Elasticity of substitution
v	0.500	Size of the low-interaction sector
A	39.835	Productivity
θ	1.275×10^{-3}	Labour supply parameter
β	$0.96^{1/52}$	Discount factor

$$\pi_r = (1 - \text{'case fatality rate'}) \times \frac{7}{\text{'days to recover or die'}}, \tag{52}$$

and

$$\pi_d = \text{'case fatality rate'} \times \frac{7}{\text{'days to recover or die'}}, \tag{53}$$

where the constant 7 transforms daily to weekly statistics. We set the number of days that one stays infected to be 14, which is consistent with prevailing quarantine periods worldwide.

We then estimate the remaining health parameters, the case fatality rate (which directly determines π_d), the initial fraction ε of infected people, the infection risk parameter $\pi_{s,t}$, the seasonal summer effect and the infection intensity ϕ of sector 1 by searching over a grid of possible values to minimize the RMSE between the simulated and actual weekly death counts between February 2020 and January 2021. The data on weekly deaths are obtained from the World Health Organization. We search within a range of the two values of 0.2% and 0.5% for the case fatality rate. For the initially infected population (per 100,000 people), the range is given by $\{1, 5, 10, 50, 100\}$. The range of possible values for the infection probability $\pi_{s,t}$ is informed by the plausible range of values for the basic reproduction number R_0 at time t . As reported by Flaxman et al. (2020), R_0 is estimated to lie between 1.4 and 4.18. In our model and absent the endogenous reduction and sectoral shift of consumption by susceptible agents, the basic reproduction number $R_{0,t}$ is related to the epidemiological parameters according to

$$R_0 = \frac{\pi_{s,t} A^2 \theta^{-1}}{\pi_r + \pi_d}. \tag{54}$$

Given the observed range of R_0 and the economic and health parameter values chosen above, the range of $\pi_{s,\text{'winter'}}$ values associated with values for R_0 in the interval (1.4, 4, 18) is $\pi_{s,\text{'winter'}} \in \left(\frac{(\pi_r + \pi_d) \times 1.4}{A^2/\theta}, \frac{(\pi_r + \pi_d) \times 4.18}{A^2/\theta} \right)$ which also depends on the value for the case fatality rate. We set the range of $\pi_{s,\text{'winter'}}$ to be $[0, 5 \times 10^{-6}]$, in steps of 2×10^{-8} , leading to a slightly wider range of possible values for R_0 than reported by Flaxman et al. (2020). The seasonal effect lowers the value of $\pi_{s,t}$ in the summer. We

search in a grid for the reduction factor $\pi_{s, \text{summer}}/\pi_{s, \text{winter}}$, in between 5% and 95%, in increments of 5 percentage points. Finally, the grid for the potential values of ϕ ranges from $\phi = 0.1$ to $\phi = 1$, with an increment of 0.1. We also conduct the exercise for an epidemiological SIR model, which is triggered by setting $\omega_t = 1$ in Appendix H for all t . The search grid for case fatality rate, I_0 , $\pi_{s, \text{winter}}$, the summer reduction factor $\pi_{s, \text{summer}}/\pi_{s, \text{winter}}$, ϕ and the epidemiological SIR setup therefore implies $2 \times 5 \times 251 \times 19 \times 11 = 524,590$ iterations in solving the model.

In each iteration, the simulated time series for weekly deaths, in the number of deaths per 100,000 people, is compared with the actual data to obtain the RMSE. The parameter values producing the best fit to the data on weekly deaths, conditional on a choice for ϕ value, are presented in Table 2. We see that a choice of $\phi = 0.6$ minimizes the RMSE between the model and the data. We, therefore, choose it as our benchmark. Note that at this value for $\phi = 0.6$, the model implies a basic reproduction number at the beginning of the pandemic of $R_0 = 3.68$.

6. RESULTS FOR THE BENCHMARK ECONOMY

We now present our results, starting with the findings for our benchmark economy with two sectors calibrated to the Swedish health data as described above. To clarify the quantitative importance of our mechanism, we also contrast the results from the benchmark economy to that of a one-sector economy inspired by the work of ERT where, by construction, the sectoral reallocation channel is absent. In this section, to facilitate the comparison between both model variants, we keep the health parameters fixed; the only difference is ϕ which takes the value $\phi = 0.6$ for our benchmark $\phi = 1$ for the representative sector economy.

Figure 5 displays the number of weekly deaths in the data and the two versions of the model. It is identical to Figure 2 in the introduction but replaces the death dynamics implied by the purely epidemiological SIR model predicted by the representative sector economy. In Figure 6, we display the health and economic dynamics in our benchmark model (blue solid line) and contrast it with the economy in which the reallocation mechanism is absent since $\phi = 1$ as in a one-sector economy (the dotted line). Finally, the left panel of Figure 7 plots the evolution of sectoral consumption across the two sectors of our benchmark economy. The right panel does the same for Swedish data.¹⁹

We observe that, as already documented in Section 1, the baseline model with sectoral allocation captures the death dynamics very well. Both the seasonality of the

19 In the data, rather than restricting attention to narrowly defined sectors (such as restaurants and groceries, as in Section 1), we now use data on all final consumption sectors of the economy, rank them by their decline in the spring of 2020 relative to December 2019 and group them into two baskets according to this decline. We then plot the time series of output for both groups relative to December 2019. The result is the right panel of Figure 7. Appendix G contains the details of how the right panel of Figure 7 is constructed.

Table 2. Results from grid search

ϕ	Case fatality (%)	$\pi_{s, \text{'winter'}} (\times 10^{-6})$	I_0 (per 100k people)	Summer	RMSE	R_0
0.1	0.2	2.44	1	−0.85	1.12	6.07
0.2	0.2	1.88	5	−0.80	0.92	4.68
0.3	0.2	1.84	5	−0.80	0.75	4.58
0.4	0.2	1.72	5	−0.75	0.61	4.28
0.5	0.2	1.50	10	−0.70	0.57	3.73
0.6	0.2	1.48	10	−0.70	0.52	3.68
0.7	0.2	1.46	10	−0.70	0.56	3.63
0.8	0.2	1.46	10	−0.70	0.62	3.63
0.9	0.2	1.46	10	−0.70	0.66	3.63
1	0.2	1.46	10	−0.70	0.68	3.63
SIR	0.2	1.00	50	−0.55	1.31	3.20

Notes: Summer is the change in the value $\pi_{s, \text{'summer'}}/\pi_{s, \text{'winter'}} - 1$ from the last week of March to September for $\pi_{s,t}$, compared with $\pi_{s, \text{'winter'}}$. R_0 is the basic reproduction number. The grid is defined as follows. Case fatality rate $\in \{0.2\%, 0.5\%\}$; $\pi_s \in [0, 5 \times 10^{-6}]$; $I_0 \in \{10^{-5}, 5 \times 10^{-5}, 10^{-4}, 5 \times 10^{-4}, 10^{-3}\}$ and Summer $\in [-0.05, -0.95]$. The RMSE is calculated from the simulated and actual deaths per 100k people. The grid search for the SIR model is done by setting $\omega_t = 1$ described in Appendix H throughout the sample period.

infection parameter $\pi_{s,t}$ as well as the endogenous adjustment of consumption decisions (relative to the purely epidemiological SIR model) is important in this regard. However, they play a different role in the first and the second wave.

When the virus breaks out in the spring of 2020, susceptible households optimally reduce consumption and substitute consumption goods from the high-infection sector with goods from the low-infection sector, see the lower middle and right panel in Figure 6 and the left panel of Figure 7. The resulting decline in aggregate consumption of ca. 10% in both the benchmark economy and the economy with two sectors accords well with the size of the economic recession in Sweden in the early part of 2020, see the right panel of Figure 7. Note that the economic consumption of a susceptible household in the model is defined as

$$c_t^s = \left(\sum_{k=1,2} 0.5(c_{tk}^s)^{1-1/\eta} \right)^{\eta/(\eta-1)}$$

(55)

where as measured consumption of that same household is given by

$$c_t^j = \sum_{k=1,2} 0.5c_{tk}^j,$$

(56)

since prices of all goods are identical. Infected and recovered households optimally spread their consumption equally across sectors, and thus for these groups, there is no difference between both consumption concepts. By the goods market-clearing condition, measured consumption (and thus measured GDP, since we abstract from investment and capital accumulation) equals aggregate labour input (and thus the last panel of




Figure 5. Weekly deaths: Swedish data and models

Figure 6 displays the change in both variables), which is the relevant model variable for comparison with the empirically observed consumption decline.

The reduction of consumption by susceptible individuals, which, early in the pandemic, accounts for most of the population, reduces infection risk and lowers the number of new infections relative to the SIR model (see again Figure 2). Besides, the first panel of Figure 6 shows that the sectoral reallocation of consumption towards less infectious sectors reduces new infections and slows down the reduction of not-yet infected (susceptible) and deceased agents from April to October 2020, relative to the representative sector economy. Quantitatively, however, this effect is small relative to the seasonal decline in the infection risk parameter $\pi_{s,t}$ in the summer of 2020. In a nutshell, this epidemiological seasonality in infection risk temporarily halts the epidemic, with economic adjustments in total and sectoral consumption playing a non-trivial but secondary role.

The situation is very different during the second wave in the winter of 2020. With a high seasonal infection parameter and associated infection risk from October 2020 on, and with more than 60% of the population still susceptible to the disease (more so in the baseline economy than in the one-sector economy, since the spring wave was somewhat smaller in that economy, see the upper panels of Figure 6), reducing consumption and changing its composition is the only way for susceptible households to avoid the strongly elevated infection risk. As the left panel of Figure 7 shows, the model-implied




Figure 6. Evolution of health distribution and economic variables: Benchmark model and homogeneous-sector economy

consumption recession, and especially the consumption reallocation across sectors, is much more potent in the second wave. Figure 8 again plots model-measured consumption in both versions of the model (as the lower middle and right panel of Figure 6 did) but zooms in on the second wave. Measured consumption in the second wave displays a significantly smaller decline in the benchmark economy than in the one-sector economy because reallocation of consumption across economic activities is an effective tool for mitigating infection risk.

We acknowledge that although the Swedish data display a significant reduction of consumption expenditures on restaurants in the second wave (see Figure 1 again in Section 1), the right panel of Figure 7 does not bear out the massive second collapse of economic activity in more broad-based sectors of the economy, as predicted by our benchmark economy, although the strong performance of other sectors predicted by the model is clearly visible in the data. This may be due to sectors already aggregating across low-infection-risk and high-infection-risk components and due to a richer structure of the economy overall. It also suggests that the frictionless reallocation of consumption across sectors of the economy envisioned by our benchmark model in light of the pandemic is perhaps too extreme an assumption. In practice, such reallocation might




Figure 7. Sectoral dynamics: Model versus data

take even more time or could be limited by fixed factors of production and by a lower willingness of individuals to substitute away from valued (yet infectious) consumption activities than we have assumed thus far.

7. INSPECTING THE ECONOMIC MECHANISM: SENSITIVITY ANALYSIS

In order to more deeply understand the economic choices and their impact, the potency of the multi-sector possibilities and the differences between the one-sector and the multi-sector version, we explore the sensitivity of our results with respect to key choices. One is the infection risk ϕ in the low-infection risk sector as well as the seasonality, or lack thereof, in the exogenous parameter $\pi_{s,t}$, see Section 7.1. The second is to vary the elasticity of substitution or the level of the exogenous parameter $\pi_{s,t}$, see Section 7.2.

7.1. Varying infection risks and seasonality

The differences in Figure 6 between the one-sector economy and the two-sector economy may be puzzlingly small: although there are considerable sectoral reallocations in




Figure 8. Dynamics of labour/measured consumption during second wave

the two-sector version, the aggregate results do not differ much, especially in the first wave.

To explore that, we first ‘crank up’ the health benefits from consuming in the low infection risk sector per lowering ϕ to $\phi = 0.2$ from $\phi = 0.6$, correspondingly raising the infection intensity of the high-infection sector to $2 - \phi = 1.8$. The impact can be seen in the top portion of Figure 9. The impact on the dynamics of the infected is now considerable: it rises more sharply and is much larger both in the first and second waves. As a result, the fraction of susceptible agents and of the deceased climb considerably more than in our benchmark economy or the one-sector economy: we only show the latter here for comparison. Nonetheless, the first wave’s impact on consumption and labour is still remarkably small: larger differences only emerge in the second wave.

To understand why this is so, it is useful to examine the first-order conditions of susceptible agents in Equations (14)–(16), that is

$$\begin{aligned}
 u_1(c_t^s, n_t^s) \cdot \left(\frac{c_t^s}{c_{tk}^i} \right)^{1/\eta} &= \lambda_{bt}^s + \nu_t \phi(k) c_{tk}^i \\
 \nu_t &= \pi_s I_t \lambda_{\tau t} \\
 \beta(U_{t+1}^s - U_{t+1}^i) &= \lambda_{\tau t}
 \end{aligned}$$




Figure 9. Two waves versus one wave. (a) Dynamics with seasonality and (b) dynamics without seasonality

Notes: The following parameters are unchanged from the baseline model: case fatality = 0.2%, $\pi_{s, \text{winter}} = 1.48 \times 10^{-6}$, $I_0 = 10$. Solid lines are for $\phi = 0.2$, while dotted lines are for $\phi = 1$. In panel (a), seasonality follows the baseline model, in which the $\pi_{s, \text{summer}}$ value drops by 70% of $\pi_{s, \text{winter}}$. In panel (b), the $\pi_{s, t}$ value is constant throughout the period.

The difference between the first-order conditions of a susceptible agent to those of an infected or recovered agent is the last term on the right in the first of these three equations, that is, the shadow value of becoming infected, characterized by ν_t . The second equation shows that this shadow value is a product of two endogenous components: the fraction of the infected I_t as well as the Lagrange multiplier $\lambda_{\tau t}$.

The last equation shows $\lambda_{\tau t}$ to be proportional to the difference in the lifetime utilities between being susceptible and infected. There are two offsetting forces determining this difference. Consider the results for the first wave as shown in the top portion of Figure 9 and compare the one-sector economy to the two-sector economy. Aggregate consumption is roughly the same, but there is an additional sectoral shift in the two-sector economy. Flow utility for susceptible agents is lower than that of infected agents: that flow utility difference alone would result in a negative value for $\lambda_{\tau t}$ and it is thus even greater for susceptible agents in the multi-sector economy than in the one-sector economy. The difference in lifetime utilities is ultimately positive due to the additional death risk that infected agents face. The infection risk is higher in the one-sector economy, as the top left panel shows. This implies that it is more valuable to be susceptible in the multi-sector economy, as infection and the death risk can be prolonged longer. If this force alone were present, $\lambda_{\tau t}$ would be higher in the multi-sector economy than the one-sector economy. The first panel of the second row in Figure 9 likely indicates that the second force is numerically stronger than the first until late in 2020.

Returning to the calculation of ν_b , we now see that the lower fraction of infected I_t in the multi-sector economy is offset by the higher life-time utility difference $\lambda_{\tau t}$ in the multi-sector economy. The resulting impact on consumption thus happens to be similar. Note that the ordering of the $\lambda_{\tau t}$ switches towards the end of 2020: it is then that larger differences in the aggregate economics appear.

When seasonality in $\pi_{s,t}$ is eliminated as in the bottom portion of Figure 9, the life-time utility difference $\lambda_{\tau t}$ becomes smaller rather than larger when moving to the two-sector economy with $\phi = 0.2$, that is, the first force of reduced flow utility becomes dominant with time. As a result, both the decrease in the fraction of infected as well as the decrease in the life-time utility difference work in the same direction: both lower the shadow value ν_t of avoiding the infection risk, thus leading to a smaller reduction in consumption than in the one-sector economy, see the middle panel of the bottom row. We concede that the lower overall consumption in the one-sector economy, in turn, generates a force towards a lower flow utility in the one-sector economy compared with the multi-sector economy: the overall net effect can be tricky to understand fully.

The point of these calculations is to argue that the various forces interact in subtle ways. Seasonality and the modest reduction in infection risk of consuming in sector 1, nearly offset each other during the first wave, as can be seen in Figure 6, but this is not necessarily so in general. They depend in subtle ways on the numerical specifications chosen.

7.2. Varying the substitution elasticity across goods

The second crucial parameter governing the sectoral substitution mechanism's quantitative importance is the elasticity of substitution across sectors, as measured by η .

The top portion of Figure 10 shows the impact of changing the elasticity of substitution η . As we approach the perfect substitution case for very high values of η , the dynamics of the first wave and 'winter phase' in March/April 2020 is reduced considerably, leading to an even faster decline of infected agents in the 'summer phase' and dramatically reducing the first wave's size. As a result, the drop in consumption is small. However, this reduction of infections in the first wave leaves a larger portion of susceptible agents for the second wave. When the next winter phase starts, the value for $\pi_{s, \text{'winter'}}$ is so high, that even a high value of η implies a positive reproduction rate of infected agents: the number of infected agents and the size of the second wave as well as the implied drop in consumption is now much larger. The details regarding the sectoral reallocation can be seen in Figure 11. When $\eta = 100$, the high-infection sectors practically cease operations, at the peak of the first wave as well as during the second wave. This may correspond to, say, movie theaters and concert venues, where indeed activity largely ceased.

The elasticity of substitution η is tightly related to the level of the infection risk parameter $\pi_{s,t}$. Proposition 1 shows that our model is equivalent to a one-sector economy when the elasticity of substitution is zero, since then, any sectoral reallocation is eliminated. When η approaches infinity, we again obtain a one-sector economy in the limit as proposition 2 shows, except that $\pi_{s,t}$ now needs to be scaled down by the infection risk $\phi(0)$ of the lowest-risk sector.

The bottom portion of Figure 10 therefore shows the results for the one-sector economy, as we scale $\pi_{s,t}$ down to $\phi_1 * \pi_{s,t} = 0.6\pi_{s,t}$, and which corresponds to $\eta \rightarrow \infty$. There are differences in the $\eta = 100$ results in the top portion. Perhaps the most interesting difference is that the onset of the second wave takes even longer in this $\pi_{s,t}$ -scaled one-sector economy, but once it starts, it appears to take off more ferociously. As an 'in-between' scenario, we also show the results for the one-sector economy, when $\pi_{s,t}$ is scaled down only half way to $0.8\pi_{s,t}$, which appears to be closer to the $\eta = 100$ case than going all way to $0.6\pi_{s,t}$. The key takeaway from this discussion is that parameter variations in the one-sector economy can be of considerable aid in understanding parameter variations in the multi-sector economy when guided by theories such as proposition 2.

8. SOCIALLY OPTIMAL ALLOCATIONS

Lastly, we explore the solution to the social planner's problem described in Section 4. Figure 12 shows the outcome of the social planner solution under our baseline




Figure 10. Top portion: heterogeneous-sector economy, variations in η and bottom portion: homogeneous-sector economy, scaling of $\pi_{s,t}$




Figure 11. Heterogeneous-sector economy: consumption dynamics

parameters.²⁰ The social planner essentially stops the pandemic dead in its tracks: the number of infected agents declines quickly and is barely noticeable within a few weeks after the start of the outbreak; note the extremely small scale of the y -axes in Figure 12. The social planner achieves this outcome by restricting consumption of infected agents in a Draconian manner, thereby hugely mitigating the infection risk and stopping infections at the onset. Compared with the competitive equilibrium outcome, the solution of a planner with the power to distinguish between the health status of infected and susceptible agents is even more successful in averting the epidemic.

Figure 13 further illustrates how the consumption of infected people is restricted and how the allocation depends on the degree of substitutability between the two sectors. In the baseline scenario, per capita consumption of an infected household is restricted to 5.8% of its pre-pandemic steady-state competitive equilibrium value during the winter

20 The social planner solution for the homogeneous sector case (equivalently, the $\eta = 0$ case) and the $\eta = 10$ case are practically identical, in terms of aggregates. Since the planner effectively stops the epidemic and thus infected individuals are virtually absent from the economy, the consumption dynamics of susceptible individuals making up close to 100% of the economy throughout the epidemic is close to independent of η . As far as aggregates are concerned, sector heterogeneity plays practically no role. Therefore, we only display the time series of these aggregates for the benchmark economy.




Figure 12. Heterogeneous-sector economy: social planning solution

Note: The tiny numbers on the vertical axes.

of 2020/21. Consumption of infected individuals in the highly infectious sector is effectively driven to zero, especially if the two sectors can be substituted easily (the $\eta = 10$ case). Essentially, the planner insulates susceptible individuals facing infection risk from infected individuals in the dangerous sector. In the benchmark case with a lower elasticity of substitution, $\eta = 3$, the social planner does not impose quite as drastic a difference across the sectors (since this would be very costly in terms of lifetime utility of the infected individuals, which the planner values as well). In the homogeneous sector case (equivalent to the case for $\eta = 0$ in the two-sector economy), consumption in both sectors is the same, as the dashed black lines show. Now, consumption for infected individuals is reduced to only 4.6% of the non-infectious steady-state competitive equilibrium. It is in this treatment of the infected agents that the sectoral substitution possibilities matter considerably since without any other means to insulate the infected from the susceptible population, the planner finds it optimal to reduce consumption activity of the infected population across the board to avoid the outbreak of the epidemic.

It is particularly interesting that the social planner's restriction on infected agents' consumption is also contingent on the seasonal effect of the infection risk. During the summer, when the value of $\pi_{s, \text{'winter'}}$ is 70% lower than the winter value $\pi_{s, \text{'winter'}}$, infected




Figure 13. Per capita consumption of infected people

agents are allowed to consume about two times more in both sectors. The consumption in the low-infection sector increases to 37% of the pre-pandemic competitive equilibrium and consumption in the high-infection sector rises to 4.5% of the former steady state.

One should take the social planner solution with a pinch of salt, of course, rather than as a normative prescription of what a government should actually do in practice. One objection may be that such a social planner solution is hard to achieve without widespread testing that allows the planner to identify the health status of individuals. A more reasonable specification may be to assume that the social planner does not know who is infected and who is not, see Appendix F. In this case, consumption is constrained to be the same across health types.²¹ It turns out that even such a weaker social planner can stop the pandemic early on and keep it from resurfacing, see Figure 14. The switch from winter to summer during that first wave helps: the social planner allows the

21 In this version of the planner problem, we assume that households also do not know their health type. Otherwise the planner might try to devise a revelation mechanism eliciting that information from individuals by offering consumption-labour allocations that depend on the health reports of individuals.

fraction of infected to rise a bit at first, then squashes the pandemic during the summer phase. What is particularly remarkable here is that the social planner imposes a reduction in economic activity, in particular in the high-infection sector, not only early on in the pandemic when there are still more than, say, 10 in a 100,000 infected agents in population, but also during that second wave when infected agents seem to have all but disappeared. The social planner nonetheless implements the reduction in activity in order to not let the pandemic rise again at all, despite imposing the costs of reduced economic activity on all members of the population. This rather large cost is avoided in the problem of the omniscient social planner analysed previously. One may wish to read our findings as pointing to a large social value of making testing widely available.

Conversely, a really powerful social planner that can change how individuals consume would entirely separate the infected and recovered people from those who are susceptible (always under the maintained assumption that infected individuals cannot pretend to be healthy in order to be eligible for a more favourable consumption allocation). If this is technologically feasible, the disease cannot spread any further and no consumption decline for the infected is needed. The formulation of our benchmark social planner problem precludes this possibility. In summary, our calculations and this remark show that the possibility of containing the pandemic depends crucially on the tools available to the planner. They may involve imposing considerable hardship on a few (the initially infected) to rescue the many.

9. CONCLUSION

In this paper, we have analysed the economic *laissez-faire* approach or ‘Swedish’ solution to the COVID-19 pandemic. We have built on the macroeconomic-cum-SIR model of ERT, but departed from their analysis in two key dimensions. Regarding the economic mechanism, we permit substitution of consumption across sectors with different degrees of infection probabilities. Regarding the epidemiological mechanism, we allow for exogenous seasonality in the key infection risk parameter. After presenting theoretical results for limiting values of the elasticity of substitution between sectors, we quantified the model by calibrating our epidemiological parameters to the Swedish death dynamics. We found that our model fits that death dynamics substantially better than a plain-vanilla SIR model, while also providing a reasonably good fit to the overall consumption decline in Sweden during the first wave of the epidemic. With our sensitivity analysis, we show that the sectoral substitution possibility is a powerful force in mitigating the speed of the epidemic and the depth of the economic decline.

As a result and while the COVID-19 death toll in Sweden was substantial and tragic, it was perhaps not as dramatic as some observers have feared. As of 8 March 2021 and according to the data provided by John Hopkins University, the Swedish death toll stands at about 1,264 per million, roughly the same as similarly open Brazil or subject-to-lockdown France, but below the United States with 1,592 per million, Italy with 1,656 per million or, at the sad top, the Czech Republic with 2,035 per million. The death toll is higher than in, say, Germany with 866 fatalities per million or neighbouring




Figure 14. Social planner solutions without knowing the agent types




Figure 15. Model simulation, when agents ignore the infection risks until late March 2020

Note: $\phi_1 = 0.1$, case fatality rate 0.5%, $\pi_s = 0.78 \times 10^{-6}$, $I_0 = 100$, winter scale 2.

Norway with just 118 per million. Given these numbers and given the substantial role of the change in the exogenous (and thus difficult-to-predict) infection risk parameter π_s , ‘winter’, we are hesitant to judge whether Sweden could have done better with the kinds of policies adopted in, say, Norway, Germany, France or the United States and what the final death toll numbers will be when the epidemic has run its course around the world. As a warning sign, it is worth remembering that the Czech Republic, the current COVID-19 death leader as of 8 March 2021, had a nearly negligible death toll as late as September 2020.

Neither do we wish to argue that a laissez-faire approach is the best approach available. In the competitive equilibrium, agents do not take the externality – that their infectious state and infection risk impose on others – into account. We, therefore, juxtapose our laissez-faire solution with the solution of a social planner, who knows the health status of agents and can dictate their consumption and labour choices but cannot otherwise separate individuals. The social planner would drastically lower the consumption of infected agents, thereby avoiding the economic decline and the pandemic dynamics almost entirely. The planner would therefore massively reduce the number of deaths linked with COVID-19, relative to the competitive equilibrium.

APPENDIX A: TWO-SECTOR SIMULATIONS

The consumer interaction indicator $\phi(k)$ is defined piece-wisely as

$$\phi(k) = \begin{cases} \phi & k \in [0, v) \\ 2 - \phi & k \in [v, 1] \end{cases},$$

where v is the size of the sector with lower consumer interactions. For each sector $j \in \{1, 2\}$, there is a first-order condition with respect to c_{jt}^x , where $x \in \{s, i, r\}$. The equations for infected and recovered people are substituted out, because their consumption and labour are constant. The following equations consist of the system delivering the paths of key variables

$$\frac{1}{c_t^s} \left(\frac{c_t^s}{c_{1t}^s} \right)^{1/\eta} = \frac{\theta}{A} n_t^s + \pi_{s,t} I_t \lambda_{\tau,t} \phi_1 \frac{A}{\sqrt{\theta}} \quad (A.1)$$

$$\frac{1}{c_t^s} \left(\frac{c_t^s}{c_{2t}^s} \right)^{1/\eta} = \frac{\theta}{A} n_t^s + \pi_{s,t} I_t \lambda_{\tau,t} \phi_2 \frac{A}{\sqrt{\theta}} \quad (A.2)$$

$$v_1 c_{1t}^s + v_2 c_{2t}^s = A n_t^s \quad (A.3)$$

$$c_t^s = [v_1 c_{1t}^{s^{1-1/\eta}} + v_2 c_{2t}^{s^{1-1/\eta}}]^{\frac{\eta}{\eta-1}} \quad (A.4)$$

$$\lambda_{\tau,t} = -\beta(U_{t+1}^i - U_{t+1}^s) \quad (A.5)$$

$$U_t^s = u(c_t^s, n_t^s) + \beta[(1 - \tau_t)U_{t+1}^s + \tau_t U_{t+1}^i] \quad (A.6)$$

$$U_t^i = u(c_t^i, n_t^i) + \beta[(1 - \pi_d)U_{t+1}^i + \pi_r U_{t+1}^r] \quad (A.7)$$

$$U_t^r = u(c_t^r, n_t^r) + \beta U_{t+1}^r \quad (A.8)$$

$$\tau_t = \frac{A}{\sqrt{\theta}} \pi_{s,t} I_t (\phi_1 v_1 c_{1t}^s + \phi_2 v_2 c_{2t}^s) \quad (A.9)$$

$$T_t = \tau_t S_t \quad (A.10)$$

$$S_t = 1 - I_t - R_t - D_t \quad (A.11)$$

$$R_t = R_{t-1} + \pi_r I_{t-1} \quad (A.12)$$

$$D_t = D_{t-1} + \pi_d I_{t-1} \quad (A.13)$$

$$I_t = T_{t-1} + (1 - \pi_d - \pi_r) I_{t-1} + 1_{t-1} \varepsilon. \quad (A.14)$$

Note that the time convention of disease dynamics is modified for implementation in Dynare. An MIT shock of size ε is added to Equation (A.14) in period 1. The paths of aggregate consumption and labour are given by

$$C_t = S_t c_t^s + (I_t + R_t) \frac{A}{\sqrt{\theta}} \quad (A.15)$$

$$N_t = S_t n_t^s + (I_t + R_t) \frac{1}{\sqrt{\theta}}. \quad (A.16)$$

APPENDIX B: DETAILS OF THE ANALYSIS OF NONSEPARABLE UTILITY

In this appendix, we provide the details of the analysis for a general period utility functions $u(c, n)$ in Section 2.4, allowing for non-separability between labour and consumption. For examining the dynamics, we first treat ν_t defined in Equation (15) as given. With this as well as $c_{tk}^i \equiv c^*$, the maximization problem over labour and the consumption varieties is a static problem, where the two first-order conditions (14) and (17) can be combined into

$$u_1(c_t^s, n_t^s) \cdot \left(\frac{c_t^s}{c_{tk}^s} \right)^{1/\eta} + \frac{1}{A} u_2(c_t^s, n_t^s) = \nu_t \phi(k) c^*. \quad (B.1)$$

In steady state or without infection risk, $\nu_t = 0$. We shall examine a local analysis of Equation (B.1) for small deviations of ν_t from zero. For any power κ , we approximate²² $\int (c_{tk}^s)^\kappa dk \approx (c_t^s)^\kappa$. Replace the derivatives of the utility functions with their first-order approximations

$$u_1(c_t^s, n_t^s) \approx u_1 + (u_{11} + \frac{1}{A} u_{12})(c_t^s - c^*) \quad (B.2)$$

$$u_2(c_t^s, n_t^s) \approx u_2 + (u_{12} + \frac{1}{A} u_{22})(c_t^s - c^*), \quad (B.3)$$

where u_1, u_{11}, u_{12} and u_{22} denote the first and second derivatives at the steady state and where we have exploited that $c_t^s \approx A n_t^s$ per $\int (c_{tk}^s)^\kappa dk \approx (c_t^s)^\kappa$ for $\kappa = 1 - (1/\eta)$. Given ν_t and up to this first-order approximation, Equation (B.1) now shows that the cross-sectoral allocation of consumption across the different consumption varieties driven by the relative infection risks $\phi(k)$ is the same for all utility functions, for which the right-hand sides of Equations (B.2) and (B.3) is the same. These right-hand sides are not properties of the utility function and the steady state alone, however: they still involve the deviation $c_t^s - c^*$ of total consumption from the steady state. To complete the analysis, we need to calculate these deviations. Replacing Equations (B.2) and (B.3) in Equation (B.1), noting $u_1 = -u_2/A$, integrating over k and exploiting $\int (c_{tk}^s)^\kappa dk \approx (c_t^s)^\kappa$ for $\kappa = -(1/\eta)$, delivers

$$\psi(c_t^s - c^*) \approx \nu_t \phi(k) c^*, \quad (B.4)$$

where [Database]

22 This is a first-order approximation, since $f'(0) = 0$, when $f(x) = \int (c_t^s + x \delta_k)^\kappa dk$ for some δ_k with $\int \delta_k dk = 0$.

$$\psi = u_{11} + \frac{2}{A}u_{12} + \frac{1}{A^2}u_{22}. \quad (B.5)$$

Hence, if ψ takes the same value for two utility functions, then $c_t^s - c^*$ is the same up to a first-order approximation, given ν_t . Moreover, if two utility functions also take the same value for u_1 as well as $u_{11} + u_{12}/A$, then the right-hand sides of [Equations \(B.2\)](#) and [\(B.3\)](#) are the same for all values of $c_t^s - c^*$, exploiting $u_1 = -u_2/A$ in the derivation.²³ We note that our benchmark utility specification (1) and $c^* = A/\sqrt{\theta}$, $n^* = 1/\sqrt{\theta}$ implies $\psi = -2\theta/A^2$ as well as

$$u_1(c_t^s, n_t^s) \approx \frac{\sqrt{\theta}}{A} + \frac{\theta}{A^2}(c_t^s - c^*) \approx \frac{-1}{A}u_2(c_t^s, n_t^s) \quad (B.6)$$

as first-order approximations to the derivatives. Any utility function with $\psi = -2\theta/A^2$, $u_1 = \sqrt{\theta}/A$ and $u_{11} + u_{12}/A = \theta/A^2$ delivers the same dynamics, given ν_t and up to the approximations imposed.

It remains to check that the paths for ν_t are approximately identical. For that, one has to examine the differences in the utilities $U_t^s - U_t^i$, see [Equation \(15\)](#) and [\(16\)](#), and thus the differences in the flow utilities $u(c_t^s, n_t^s) - u(c^*, n^*)$. A second-order approximation and exploiting $u_1 = -u_2/A$ as well as $c_t^s \approx An_t^s$ delivers

$$u(c_t^s, n_t^s) - u(c^*, n^*) \approx \psi(c_t^s - c^*)^2. \quad (B.7)$$

Thus, the utility differences $U_t^s - U_t^i$ and therefore the path for ν_t also remain unchanged up to our approximations, provided two utility functions agree on the value for ψ .

APPENDIX C: ELIMINATING c_{tk}^s

Note that $c_{tk}^i = A/\sqrt{\theta}$. Let us reexamine [Equation \(14\)](#) and write it as

$$\left(\frac{c_t^s}{c_{tk}^s}\right)^{1/\eta} = x_{tk}, \quad (C.1)$$

where we define

23 Note that equating u_1 is easily achieved by scaling the utility function: we are only comparing steady-state values, not the functional forms.

$$x_{tk} = c_t^s \left(\lambda_{bt}^s + \nu_t \phi(k) A / \sqrt{\theta} \right). \quad (C.2)$$

Rewrite Equation (C.1) as

$$c_{tk}^s = x_{tk}^{-\eta} c_t^s. \quad (C.3)$$

Thus,

$$(c_{tk}^s)^{1-1/\eta} = x_{tk}^{1-\eta} (c_t^s)^{1-1/\eta} \quad (C.4)$$

and integrate

$$\int (c_{tk}^s)^{1-1/\eta} dk = \int x_{tk}^{1-\eta} dk \times (c_t^s)^{1-1/\eta}.$$

Taking this to the power $\eta/(\eta-1)$ finally yields

$$c_t^s = \left(\int x_{tk}^{1-\eta} dk \right)^{\eta/(\eta-1)} c_t^s$$

or the constraint

$$1 = \left(\int x_{tk}^{1-\eta} dk \right)^{\eta/(\eta-1)}.$$

This can be simplified to

$$1 = \int x_{tk}^{1-\eta} dk \quad (C.5)$$

or

$$c_t^s = \left(\int \left(\lambda_{bt}^s + \nu_t \phi(k) A / \sqrt{\theta} \right)^{1-\eta} dk \right)^{1/(1-\eta)}. \quad (C.6)$$

Thus, Equations (5) and (6) can be rewritten as

$$\tau_t = \pi_{s,t} I_t \int \phi(k) x_{tk}^\eta c_t^s A / \sqrt{\theta} dk \quad (C.7)$$

$$= \pi_{s,t} I_t \int \phi(k) \left(\lambda_{bt}^s + \nu_t \phi(k) A / \sqrt{\theta} \right)^{-\eta} (c_t^s)^{1-\eta} A / \sqrt{\theta} dk \quad (C.8)$$

and

$$T_t = \pi_{s,t} S_t I_t \int \phi(k) x_{tk}^\eta c_t^s A / \sqrt{\theta} dk \quad (C.9)$$

$$= \pi_{s,t} S_t I_t \int \phi(k) \left(\lambda_{bt}^s + \nu_t \phi(k) A / \sqrt{\theta} \right)^{-\eta} (c_t^s)^{1-\eta} A / \sqrt{\theta} dk. \quad (C.10)$$

APPENDIX D: DETAILS FOR THE HETEROGENEOUS LABOUR ECONOMY

Proof of Proposition 4. To see the similarities and differences between the heterogeneous consumption- and heterogeneous labour economy more formally, observe that the first-order conditions for infected and recovered agents are unchanged. In particular, we obtain $c_t^i = c_{tk}^i \equiv A / \sqrt{\theta}$ and $n_t^i = n_{tk}^i = 1 / \sqrt{\theta}$, regardless as to whether consumption or labour is heterogeneous and regardless of the particular approach taken for labour heterogeneity. It, therefore, suffices to examine the first-order conditions for susceptible agents.

For the consumption-infection baseline model, the first-order conditions can be summarized by

$$\frac{1}{c_t^s} \left(\frac{c_t^s}{c_{tk}^s} \right)^{1/\eta} - \frac{\theta}{A} n_t^s = \pi_{s,t} \phi_k \lambda_{\tau t} \frac{A}{\sqrt{\theta}} I_t, \quad (D.1)$$

while the aggregation constraint and budget constraint are

$$\begin{aligned} c_t^s &= \left(\int (c_{tk}^s)^{1-1/\eta} dk \right)^{\eta/(\eta-1)} \\ A n_t^s &= \int c_{tk}^s dk \end{aligned}$$

Define $n_{tk}^s = c_{tk}^s / A$ and rewrite these two equations equivalently as

$$c_t^s = A \left(\int (n_{tk}^s)^{1-1/\eta} dk \right)^{\eta/(\eta-1)} \quad (D.2)$$

$$n_t^s = \int n_{tk}^s dk. \quad (D.3)$$

Recall that the infection probability is given by Equation (5). Substituting out the solution for c_t^i as well as c_{tk}^i with $A n_{tk}^s = c_{tk}^s$, it can be restated as

$$\tau_t = \pi_{s,t} I_t \int \phi(k) A n_{tk}^s dk \frac{A}{\sqrt{\theta}}. \quad (D.4)$$

For the production-based heterogeneous-labour model, the first-order conditions can be summarized by

$$\frac{1}{c_t^s} \left(\frac{n_t^s}{n_{tk}^s} \right)^{1/\alpha} - \frac{\theta}{A} n_t^s = \frac{\tilde{\pi}_s}{A} \phi_k \lambda_{\tau t} \frac{1}{\sqrt{\theta}} I_t, \quad (D.5)$$

while the aggregation constraint and budget constraint are

$$c_t^s = A \left(\int (n_{tk}^s)^{1-1/\alpha} dk \right)^{\alpha/(\alpha-1)} \quad (D.6)$$

$$n_t^s = \int n_{tk}^s dk. \quad (D.7)$$

Recall that the infection probability is given by Equation (46), which can be restated as

$$\tau_t = \tilde{\pi}_s I_t \int \phi(k) n_{tk}^s dk \frac{1}{\sqrt{\theta}}. \quad (D.8)$$

For $\alpha = \eta$, $\tilde{\pi}_s = A^2 \pi_{s,t}$ and $n_{tk}^s/n_t^s = c_{tk}^s/c_t^s$, Equation (D.5) becomes Equation (D.1), Equation (D.6) becomes Equation (D.2) and Equation (D.8) becomes Equation (D.4). $\square$

For the preference-based heterogeneous-labour formulation, the first-order conditions can be summarized by

$$\frac{1}{c_t^s} - \frac{\theta}{A} n_t^s \frac{f'(n_{tk}^s)}{f'(n_t^s)} \left(\frac{f(n_t^s)}{f(n_{tk}^s)} \right)^{1/\alpha} = \frac{\tilde{\pi}_s}{A} \phi_k \lambda_{\tau t} \frac{1}{\sqrt{\theta}} I_t, \quad (D.9)$$

while the aggregation constraint and budget constraint are

$$c_t^s = A \int n_{tk}^s dk \quad (D.10)$$

$$n_t^s = \left(\int (n_{tk}^s)^{1-1/\alpha} dk \right)^{\alpha/(\alpha-1)}. \quad (D.11)$$

While there is considerable formal similarity, there no longer is a formal equivalence as in Proposition 4.

APPENDIX E: FIRST-ORDER CONDITIONS OF THE SOCIAL PLANNER PROBLEM

The social planner's problem in the main text yields the following first-order conditions:

$$\begin{aligned}
\left(\frac{\partial}{\partial c_t^s} : \right) u_{1,t}^s \left(\frac{c_t^s}{c_{tk}^s}\right)^{1/\eta} + \mu_{f,t} &= \beta \pi_{s,t} \phi \mu_{I,t+1} I_t c_{tk}^i \\
\left(\frac{\partial}{\partial c_t^i} : \right) u_{1,t}^i \left(\frac{c_t^i}{c_{tk}^i}\right)^{1/\eta} + \mu_{f,t} &= \beta \pi_{s,t} \phi \mu_{I,t+1} S_t c_{tk}^s \\
\left(\frac{\partial}{\partial c_t^r} : \right) u_{1,t}^r \left(\frac{c_t^r}{c_{tk}^r}\right)^{1/\eta} + \mu_{f,t} &= 0 \\
\left(\frac{\partial}{\partial n_t^s} : \right) u_{2,t}^s &= \mu_{f,t} A \\
\left(\frac{\partial}{\partial n_t^i} : \right) u_{2,t}^i &= \mu_{f,t} A \\
\left(\frac{\partial}{\partial n_t^r} : \right) u_{2,t}^r &= \mu_{f,t} A \\
\left(\frac{\partial}{\partial S_t} : \right) u(c_t^s, n_t^s) + \mu_{f,t} \int c_{tk}^s dk + \mu_{S,t} &= \mu_{f,t} A n_t^s \\
&\quad + \beta [\mu_{S,t+1} + \mu_{I,t+1} \pi_{s,t} I_t \int \phi(k) c_{tk}^s c_{tk}^i dk] \\
\left(\frac{\partial}{\partial I_t} : \right) u(c_t^i, n_t^i) + \mu_{f,t} \int c_{tk}^i dk + \mu_{I,t} &= \mu_{f,t} A n_t^i - \mu_{S,t} \\
&\quad + \beta [(\mu_{S,t+1} + \mu_{I,t+1})(1 - \pi_r - \pi_d) \\
&\quad + \pi_r \mu_{R,t+1} + \mu_{I,t+1} \pi_{s,t} S_t \int \phi(k) c_{tk}^s c_{tk}^i dk] \\
\left(\frac{\partial}{\partial R_t} : \right) u(c_t^r, n_t^r) + \mu_{f,t} \int c_{tk}^r dk + \mu_{R,t} &= \mu_{f,t} A n_t^r + \beta \mu_{R,t+1}.
\end{aligned}$$

APPENDIX F: THE SOCIAL PLANNER'S PROBLEM WITHOUT KNOWING THE AGENTS TYPES

In this section, we provide a version of the social planner's problem, when the social planner does not know the agent type. We then also assume that agents do not know their type either: thus, there is no revelation-mechanism, say, that would enable the social planner to elicit the type information from the agent. Compared with the formulation in Section 4, the social planning problem amounts to imposing the restriction that the consumption and labour choices do not depend on the type of the agent.

Formally, the social planner maximizes date-0 aggregate social welfare W_0 , where

$$W_0 = \sum_{t=0}^{\infty} \beta^t (S_t + I_t + R_t) u(c_t, n_t)$$

subject to

$$\mu_{f,t} : \int c_{ik} dk = A n_t \quad (F.1)$$

$$\mu_{S,t} : S_t = S_{t-1} - I_t + (1 - \pi_r - \pi_d) I_{t-1} \quad (F.2)$$

$$\mu_{I,t} : I_t = \pi_s S_{t-1} I_{t-1} \int \phi(k) (c_{t-1,k})^2 dk + (1 - \pi_r - \pi_d) I_{t-1} \quad (F.3)$$

$$\mu_{R,t} : R_t = R_{t-1} + \pi_r I_{t-1}. \quad (F.4)$$

The social planner takes S_0 , I_0 and R_0 as given. It chooses the time paths of consumptions c_{kt} and labour supply n_t as well as the paths for the mass of agents in the four groups S_t , I_t and R_t . The first-order conditions of the social planner's problem are

$$\begin{aligned} \left(\frac{\partial}{\partial c_t} : \right) u_{1,t} \left(\frac{c_t}{c_{ik}} \right)^{1/\eta} + \mu_{f,t} &= \beta \pi_s \phi \mu_{I,t+1} I_t c_{ik} \\ \left(\frac{\partial}{\partial n_t} : \right) u_{2,t} &= \mu_{f,t} A \\ \left(\frac{\partial}{\partial S_t} : \right) u(c_t, n_t) + \mu_{S,t} &= \beta [\mu_{S,t+1} + \mu_{I,t+1} \pi_s I_t \int \phi(k) (c_{ik})^2 dk] \\ \left(\frac{\partial}{\partial I_t} : \right) u(c_t, n_t) + \mu_{I,t} &= -\mu_{S,t} \\ &\quad + \beta [(\mu_{S,t+1} + \mu_{I,t+1})(1 - \pi_r - \pi_d) \\ &\quad + \pi_r \mu_{R,t+1} + \mu_{I,t+1} \pi_s S_t \int \phi(k) (c_{ik})^2 dk] \\ \left(\frac{\partial}{\partial R_t} : \right) u(c_t, n_t) &= \beta \mu_{R,t+1} \end{aligned}$$

APPENDIX G: REPLICATING DATA

G.1. Replicating Figure 1

Navigate to the *Statistical Database* module of the *Statistics Sweden* website and use their GUI to extract all *Fixed Price* denominated rows from the Monthly Indicator for Household Consumption.²⁴ This index details trends in retail trade, wholesale trade and

other service industries, and spans 20 years from January 2000 to July 2020. Transpose and subset this data frame until only three columns remain: *Total Consumption*, *Restaurants* and *Groceries*.

Control for seasonality (e.g., Christmas spending spikes) in these three columns by following along with [Sax and Eddelbuettel \(2018\)](#) in their aptly named paper, *Seasonal Adjustment by X-13ARIMA-SEATS in R*, which presents a model selection framework for linear regression models with ARIMA errors. This paper builds upon a SEATS procedure originally developed by [Gómez and Maravall \(1997\)](#) at the Bank of Spain.

Once deseasonalized, rebase each of the three columns so that December 2019 is always equal to 100, by dividing all the points in each series by the value of said series at 2019M12 and multiplying by 100. The resulting deseasonalized and rebased dataset gives us a clearer picture of what is happening. Plotting the last 12 rows of these three columns should now flawlessly replicate [Figure 1](#).

G.2. Replicating [Figure 7](#)

Extract two new data frames from the *Statistics Sweden* repository: GDP (Production Approach) by Industrial Classification²⁵ and Turnover Index for the Service Sector,²⁶ denoted, respectively, in *Constant Prices (SEK Millions)* and *Volume (Fixed Prices)*. While GDP (Production Approach) measures GDP as the sum of the values added by all sectors of the economy, the Turnover Index is a survey tracking monthly changes in invoice sales. The turnover Index is used as an input for national accounts statistics. Go ahead and deseasonalize/rebase each column in Turnover Index just as in [Figure 1](#).

A ‘mean response to COVID-19’ for each industry is then constructed, defined as the mean value of the index for each sector during the first 6 months of 2020. This metric is used to ordinally rank each sector in the Turnover Index, from those most impacted by COVID-19 (e.g., airlines/hotels) to those least impacted (e.g., veterinarians and IT services).

Next, filter the GDP (Production Approach) data frame so that it contains only values for 2019Q1, 2019Q2, 2019Q3 and 2019Q4, then sum these four quarters to obtain total GDP for 2019. This nominal amount for each sector is then divided by the entire sum for all sectors so that the weights sum to one. We have now ascertained the relative size of each sector.

We now wish to link/merge these two data frames (GDP and Turnover); at first glance, this might seem straightforward as they both rely on NACE (Nomenclature des Activités Économiques dans la Communauté Européenne) classification codes. However, upon closer inspection, one realizes that the variables in the two datasets do

25 www.statistikdatabasen.scb.se/pxweb/en/ssd/START__NR__NR0103__NR0103A/NR0103ENS2010T06Kv/

26 www.statistikdatabasen.scb.se/pxweb/en/ssd/START__HA__HA0101__HA0101B/DivtjansterM07X/

not match up perfectly; some miscellaneous data cleaning practices and subjective choices are required to build our own bespoke consumption bundle, a bundle seeking to proxy for *Total Consumption* in Figure 1. Many columns are dropped outright because they are contained in one data frame but not the other; other columns are dropped because they are superfluous or do not proxy for consumption. Other columns must be aggregated, for example, sectors 94, 95 and 96 in the Turnover Index are averaged into one new variable dubbed 94–96, because while GDP (Production Approach) does not contain any of 94, 95 or 96, it does contain a variable called 94–96. This aggregation is a workaround that lets us compare the two data frames. The entire cleaning and linking process is described in detail in our accompanying R code, and whittles our initial data frames down from 60+ miscellaneous sectors down to a core of 14 sectors. Depending on the choices you (the replicator) make, your particular dataset may have roughly more or less than 14 sectors, depending on how you aggregate/drop columns and how you choose to define total consumption.

Finally, classify all sectors in this newly linked data frame into those that boomed the most versus those that declined the most due to COVID-19. This is done simply by sorting the data frame on ‘mean response to COVID-19’, as calculated two paragraphs ago. Add together all the sectors with the highest mean response until their cumulative sum of GDP contributions gets just over 0.5, and plot the time series. Congratulations, you have now successfully replicated Figure 2!

APPENDIX H: MYOPIC INDIVIDUALS AND THE FIRST WAVE

In this appendix, we discuss in more detail how the assumption of initially myopic (or unaware) individuals in March 2020 can generate a steep rise in infections and deaths, and the recognition that there is a pandemic that is linked to consumption behaviour in early April can lead to an end of the first wave. Formally, we turn off the economic forces until 27 March. This temporary myopia version of the model introduces a switch to susceptible agents’ utility function as

$$U_t^s = u(c_t^s, n_t^s) + \beta \{ (1 - \omega_t) [(1 - \tau_t) U_{t+1}^i + \tau_t U_{t+1}^r] + \omega_t U_{t+1}^s \},$$

where $\omega_t = 1$ when susceptible agents are unaware of the pandemic. Hence, the first-order condition with respect to τ_t yields

$$\lambda_{\tau,t} = \begin{cases} \beta (U_{t+1}^s - U_{t+1}^i) & \text{when } \omega_t = 0 \\ 0 & \text{when } \omega_t = 1 \end{cases}.$$

It follows that in periods when $\omega_t = 1$, the consumption demand for the susceptible agents in Equation (13) resembles that for the infected agents in Equation (20). No decline in the infection parameter π_s in April is needed in this version of the model to qualitatively (but not necessarily quantitatively) capture both waves in the data, as Figure 15

shows. While the overall fit is clearly worse than in our benchmark analysis shown in Figure 2 of the main text, the power of ‘turning on’ the economic forces in April is nonetheless clearly visible in the decline of deaths at the end of the first wave.

REFERENCES

- Acemoglu, D., V. Chernozhukov, I. Werning and M. Whinston (2020). Optimal targeted lockdowns in a multi-group sir model. Draft.
- Adda, J. (2016). ‘Economic activity and the spread of viral diseases: evidence from high frequency data’, *Quarterly Journal of Economics*, 131, 891–941.
- Adhmad, S. and D. Riker (2019). ‘A method for estimating the elasticity of substitution and import sensitivity by industry’, Economics Working Paper Series 2019-05-B, U.S. International Trade Commission.
- Alvarez, F., D. Argente and F. Lippi (2020). ‘A simple planning problem for covid-19 lockdown’, Draft, University of Chicago.
- Atkeson, A. (2020). ‘What will be the economic impact of covid-19 in the US? Rough estimates of disease scenarios’, NBER WP No. 26867.
- Barrero, J., N. Bloom and S.J. Davis (2020a). ‘Covid-19 is also a reallocation shock’. Draft, University of Chicago, Brookings Papers on Economic Activity.
- Barrero, J., N. Bloom and S. J. Davis (2020b). ‘Covid-19 is also a reallocation shock’, Research Briefs in Economic Policy, CATO Institute (232).
- Berger, D., K. Herkenhoff and S. Mongey (2020). ‘An SEIR infectious disease model with testing and conditional quarantine’, BFI Working Paper No. 2020–25.
- Brotherhood, L., P. Kircher, C. Santos and M. Tertilt (2020). ‘An economic model of the covid-19 epidemic: The importance of testing and age-specific policies’, Draft.
- Correia, S., S. Luck and E. Verner (2020). ‘Pandemics depress the economy, public health interventions do not: Evidence from the 1918 flu’, Draft.
- Dingel, J. and B. Neimann (2020). ‘How many jobs can be done at home?’, Covid Economics: Vetted and Real-Time Papers.
- Eichenbaum, M.S., S. Rebelo and M. Trabandt (2021). ‘The macroeconomics of epidemics’, *The Review of Financial Studies*, hhab040.
- Fang, H., L. Wang and Y. Yang (2020). ‘Human mobility restrictions and the spread of the novel coronavirus (2019-ncov) in China’, Working Paper 26906, National Bureau of Economic Research.
- Farboodi, M., G. Jarosch and R. Shimer (2020). ‘Internal and external effects of social distancing in a pandemic’, Working Paper, University of Chicago.
- Fernandez-Villaverde, J. (2010). ‘The econometrics of DSGE models’, *SERIEs*, 1, 3–49.
- Flaxman, S., S. Mishra, A. Gandy, H.J.T. Unwin, H. Coupland, T.A. Mellan, H. Zhu, T. Berah, J.W. Eaton, P.N.P. Guzman, N. Schmit, L. Callizo, K.E.C. Ainslie, M. Baguelin, I. Blake, A. Boonyasiri, O. Boyd, L. Cattarino, C. Ciavarella, L. Cooper, Z. Cucunubá, G. Cuomo-Dannenburg, A. Dighe, B. Djaafara, I. Dorigatti, S. van Elsland, R. FitzJohn, H. Fu, K. Gaythorpe, L. Geidelberg, N. Grassly, W. Green, T. Hallett, A. Hamlet, W. Hinsley, B. Jeffrey, D. Jorgensen, E. Knock, D. Laydon, G. Nedjati-Gilani, P. Nouvellet, K. Parag, I. Siveroni, H. Thompson, R. Verity, E. Volz, P.G.T. Walker, C. Walters, H. Wang, Y. Wang, O. Watson, C. Whittaker, P. Winskill, X. Xi, A. Ghani, C.A. Donnelly, S. Riley, L.C. Okell, M.A.C. Vollmer, N.M. Ferguson and S. Bhatt (2020). ‘Report 13: Estimating the number of infections and the impact of non-pharmaceutical interventions on covid-19 in 11 European countries’, Imperial College Covid-19 Response Team.
- Glover, A., J. Heathcote, D. Krüger and J.-V. Rios-Rull (2020). ‘Health versus wealth: On the distributional effects of controlling a pandemic’, Draft.
- Gómez, V. and A. Maravall (1997). ‘Programs tramo and seats: instructions for the user’, Mimeo, Banco de España.
- Greenstone, M. and V. Nigam (2020). ‘Does social distancing matter? (March 30, 2020)’. University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2020-26. Available at SSRN:<https://ssrn.com/abstract=3561244> or <http://dx.doi.org/10.2139/ssrn.3561244>.

- Guerrieri, V., G. Lorenzoni, L. Straub and I. Werning (2020). 'Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages?', Draft, MIT.
- Hall, R. and C. Jones (2007). 'The value of life and the rise in health spending', *Quarterly Journal of Economics*, 122, 39–72.
- Hodbod, A., C. Hommes, S. Huber and I. Salle (2021). 'Zombies ahead: The covid-19 consumption game-changer evidence from a large-scale multi-country survey', Economics Working Paper, University of Amsterdam.
- Kaplan, G., B. Moll and G. Violante (2020). 'Pandemics according to HANK', Draft, University of Chicago.
- Leibovici, F., A.M. Santacreu and M. Famiglietti (2020). 'Social distancing and contact-intensive occupations', St. Louis FED.
- Mongey, S. and A. Weinberg (2020). 'Characteristics of workers in low work-from-home and high personal-proximity occupations', Draft, University of Chicago.
- Mulligan, C. (2021). 'The backward art of slowing the spread? congregation efficiencies during covid-19', Working paper, University of Chicago.
- Sax, C. and D. Eddelbuettel (2018). 'Seasonal adjustment by x-13ARIMA-SEATS in R', *Journal of Statistical Software*, 87, 1–17.
- Toxvaerd, F. (2020). 'Equilibrium social distancing', Cambridge-INET Working Paper 2020/08.