
iRisk WORKING PAPER SERIES

2020-iRisk-01

Rational policymaking during a pandemic

Loïc Berger

CNRS, IÉSEG School of Management, University of Lille, UMR 9221 -LEM, 59000 Lille, France; RFF-CMCC European Institute on Economics and the Environment (EIEE), Centro Euro-Mediterraneo sui Cambiamenti Climatici, 20123 Milan, Italy

Nicolas Berger

Faculty of Public Health and Policy, London School of Hygiene & Tropical Medicine, WC1H 9SH, London, UK; Sciensano (Belgian Scientific Institute of Public Health), 1050 Brussels, Belgium

Valentina Bosetti

RFF-CMCC European Institute on Economics and the Environment (EIEE), Centro Euro-Mediterraneo sui Cambiamenti Climatici, 20123 Milan, Italy; Department of Economics and IGIER, Bocconi University, 20136 Milan, Italy

Itzhak Gilboa

HEC, Paris-Saclay, 78350 Jouy-en-Josas, France; Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 69978, Israel

Lars Peter Hansen

Department of Economics, University of Chicago, Chicago, IL 60637; Department of Statistics, University of Chicago, Chicago, IL 60637; Booth School of Business, University of Chicago, Chicago, IL 60637

Christopher Jarvis

Department of Infectious Disease Epidemiology, London School of Hygiene & Tropical Medicine, WC1E 7HT, London, UK

Massimo Marinacci

Department of Decision Sciences and IGIER, Università Bocconi, 20136 Milan, Italy

Richard D. Smith

Faculty of Public Health and Policy, London School of Hygiene & Tropical Medicine, WC1H 9SH, London, UK; College of Medicine and Health, University of Exeter, Exeter, EX1 2LU, UK

IÉSEG School of Management Lille Catholic University 3, rue de la Digue F-59000 Lille Tel: 33(0)3 20 54 58 92
www.ieseeg.fr

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of IÉSEG School of Management or its partner institutions.

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorization of the author(s).

For all questions related to author rights and copyrights, please contact directly the author(s).

Rational policymaking during a pandemic

Loïc Berger^{a,b,*}

Nicolas Berger^{c,d}

Valentina Bosetti^{b,e}

Itzhak Gilboa^{f,g}

Lars Peter Hansen^{h,i,j,*}

Christopher Jarvis^k

Massimo Marinacci^l Richard D. Smith^{c,m}

Abstract

Policymaking during a pandemic can be extremely challenging. As COVID-19 is a new disease and its global impacts are unprecedented, decisions are taken in a highly uncertain, complex, and rapidly changing environment. In such a context, in which human lives and the economy are at stake, we argue that using ideas and constructs from modern decision theory, even informally, will make policymaking a more responsible and transparent process.

Keywords: model uncertainty, ambiguity, robustness, decision rules

JEL Classification: D81, I18

^a CNRS, IESEG School of Management, University of Lille, UMR 9221 - LEM, 59000 Lille, France; (l.berger@ieseg.fr) ^b RFF-CMCC European Institute on Economics and the Environment (EIEE), Centro Euro-Mediterraneo sui Cambiamenti Climatici, 20123 Milan, Italy; ^c Faculty of Public Health and Policy, London School of Hygiene & Tropical Medicine, WC1H 9SH, London, UK (Nicolas.Berger@lshtm.ac.uk); ^d Sciensano (Belgian Scientific Institute of Public Health), 1050 Brussels, Belgium; ^e Department of Economics and IGIER, Bocconi University, 20136 Milan, Italy (valentina.bosetti@unibocconi.it); ^f HEC, Paris-Saclay, 78350 Jouy-en-Josas, France; ^g Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 69978, Israel; ^h Department of Economics, University of Chicago, Chicago, IL 60637; ⁱ Department of Statistics, University of Chicago, Chicago, IL 60637 (lhansen@uchicago.edu); ^j Booth School of Business, University of Chicago, Chicago, IL 60637; ^k Department of Infectious Disease Epidemiology, London School of Hygiene & Tropical Medicine, WC1E 7HT, London, UK (Christopher.Jarvis@lshtm.ac.uk); ^l Department of Decision Sciences and IGIER, Università Bocconi, 20136 Milan, Italy (massimo.marinacci@unibocconi.it); ^m College of Medicine and Health, University of Exeter, Exeter, EX1 2LU, UK (rich.smith@exeter.ac.uk). * Corresponding authors: Loïc Berger : (l.berger@ieseg.fr), and Lars Peter Hansen (lhansen@uchicago.edu). This work was supported by the Alfred P. Sloan Foundation Grant [G-2018- 11113], the AXA Chairs for Decision Sciences at HEC and in Risk at Bocconi University, the European Research Council (ERC) under the European Union's [Seventh Framework Programme (FP7-2007-2013)] (Grant agreement No. 336703) and [Horizon 2020 research and innovation programme] (Grant agreement No. 670337), the French Agence Nationale de la Recherche (ANR), under grants ANR-17-CE03-0008-01, ANR-11-IDEX-0003, ANR-11-LABX-0047, the Foerder Institute at Tel-Aviv University, and the Israel Science Foundation (ISF) under grant 1077/17.

1. Introduction

The coronavirus disease 2019 (COVID-19) pandemic exposes decision problems faced by governments and international organizations. Policymakers are charged with taking actions to protect their population from the disease while lacking reliable information on the virus and its transmission mechanisms, on the effectiveness of possible measures, and their (direct and indirect) health and socio-economic consequences. The rational policy decision would combine the best available scientific evidence—typically provided by expert opinions and modeling studies. But in an uncertain and rapidly changing environment, the pertinent evidence is highly fluid, making it challenging to produce scientifically-grounded predictions of the outcomes of alternative courses of action.

A great deal of attention has been paid to how policymakers have handled uncertainty in the COVID-19 response (Lazzerini and Putoto 2020; Chater 2020; Anderson et al. 2020; Emanuel et al. 2020; Hansen 2020; Manski 2020). Policymakers have been confronted with very different views on the potential outbreak scenarios stemming from divergent experts' assessments or differing modeling predictions. In the face of such uncertainty, policymakers may respond by attempting to balance the alternative perspectives, or they may fully embrace one without a concern that this can vastly misrepresent our underlying knowledge base (Johnson-Laird 2010). This tendency to lock on to a single narrative—or more generally, this inability to handle uncertainty—may result in overlooking valuable insights from alternative sources, and thus in misinterpreting the state of the COVID-19 outbreak, potentially leading to suboptimal decisions with possibly disastrous consequences (Chater 2020; Cancryn 2020; Rucker et al. 2020; The Editors 2020).

This paper argues that insights from decision theory provide a valuable way to frame policy challenges and ambitions. Even if the decision theory constructs are ultimately used only informally in practice, they offer a useful guide for transparent policymaking that copes with the severe uncertainty in sensible ways. First, we outline a framework to understand and guide decision-making under uncertainty in the COVID-19 pandemic context. Second, we show how formal decision rules could be used to guide policymaking and illustrate their use with the example of school closures. These decision rules allow policymakers to recognize that they do not know which of the many potential scenarios is 'correct' and to act accordingly by taking precautionary and robust decisions, i.e., that remain valid for a wide range of futures and keep options open (Lempert and Collins 2007). Third, we discuss new directions to define a more

transparent approach for communicating the degree of certainty in scientific findings and knowledge, particularly relevant to decision-makers managing pandemics.

2. Decision under uncertainty

2.1. The policymaker's problem(s)

The decision-making problem faced by a high-level government policymaker during a crisis like the COVID-19 pandemic is not trivial. In the first stage, when a new infectious disease appears, the policymaker may attempt to contain the outbreak by taking early actions to control onwards transmission (e.g., isolation of confirmed and suspected cases, contact tracing). If this phase is unsuccessful, policymakers face a second-stage decision problem that consists of determining the appropriate level, timing, and duration of interventions to mitigate the course of clinical infection. These interventions may include banning mass gatherings, closing schools, and more extreme 'lockdown' restrictions.

While these measures are expected to reduce the pandemic's health burden by lowering the peak incidence, they also impose costs on society. For instance, they may have adverse impacts on mental health, domestic abuse, and job loss at a more personal level. Moreover, there are societal losses due to the immediate reduced economic activity coupled with a potentially prolonged recession and adverse impacts on longer-term health and social gradients. Policymakers must thus promptly cope with a complex and multi-faceted picture of direct and indirect, proximal and distal, health, and socio-economic trade-offs. In the acute phase of the pandemic, the trade-off between reducing mortality and morbidity and its associated socio-economic consequences may seem relatively straightforward. Still, once out of this critical phase, most trade-offs are difficult and costly. How should the policymaker decide when and how to introduce or relax measures in a justifiable way, not just from a health and economic perspective, but politically? The answer critically depends on the prioritization and balance of potentially conflicting objectives (Hollingsworth et al. 2011).

2.2. Scientific evidence and the role of modeling

Scientific knowledge is foundational to the prevention, management, and treatment of global outbreaks. Some of this evidence can be summarized in pandemic preparedness and response plans (at both international and national levels) or might be directly obtained from panels of scientists with expertise in relevant areas of research, such as epidemiologists, infectious disease modelers, and social scientists. An essential part of the scientific evidence comes from quantitative models (Morgan 2019). Quantitative models are abstract representations of reality that provide a logically consistent way to organize thinking about the

relationships among variables of interest. They combine what is known in general with what is known about the current outbreak to produce predictions to help guide policy decisions (Den Boon et al. 2019).

Epidemiological models, e.g., (Ferguson et al. 2020; Davies et al. 2020), have been used to guide decision-making by assessing what is likely to happen to the transmission of the virus if policy interventions –either independently or in combination– were put in place. Such public health-oriented models are particularly useful in the short term to project the direct consequences of policy interventions on the epidemic trajectory and to guide decisions on resource allocations (Holmdahl and Buckee 2020). As the measures put in place also largely affect the economic environment, decision-makers must, at least implicitly, confront trade-offs in the health and non-health-related economic consequences. To weigh these trade-offs necessarily requires more than epidemiological models. For example, health policy analysis models, such as computable general equilibrium models, are used to simultaneously estimate the direct and indirect impacts of the outbreak on various aspects of the economy, such as labor supply, government budgets, or household consumption (M. Keogh-Brown et al. 2020). More recent integrated assessment models combine economics and epidemiology by incorporating simplified epidemiological models of contagion within stylized dynamic economic frameworks. Such models address critical policy challenges by explicitly modeling dynamic adjustment paths and endogenous responses to changing incentives. They have been used to investigate the optimal policy response or alternative macroeconomic policies' effectiveness to the economic shocks due to the COVID-19 pandemic (Eichenbaum, Rebelo, and Trabandt 2020; Guerrieri et al. 2020; Alvarez, Argente, and Lippi 2020; Thunström et al. 2020). However, these different modeling approaches do not formally incorporate uncertainty; instead, they treat it ex-post, for example using sensitivity analyses.

2.3. Uncertainty

Decisions within a pandemic context have to be made under overwhelming time pressure and amid high scientific uncertainty, with minimal quality evidence, and potential disagreements among experts and models. In the COVID-19 outbreak, there was uncertainty about the virus's essential characteristics, such as its transmissibility, severity, and natural history (Anderson et al. 2020; Hellewell et al. 2020; R. Li et al. 2020). This state of knowledge translates into uncertainty about the system dynamics, which renders uncertain the consequences of alternative policy interventions, such as closing down schools or wearing masks in public. At a later stage of the pandemic, information overload becomes an issue, making it more difficult for the decision-maker to identify useful and good-quality evidence. The consequence is that, given the many uncertainties they are built on, no single model can be genuinely predictive in

the context of an outbreak management strategy. Yet, if their results are used as insights providing potential quantitative stories among alternative ones, models can offer policymakers guidance by helping them understand the fragments of information available, uncover what might be going on, and eventually determine the appropriate policy response. The distinction between three layers of uncertainty—uncertainty *within* models, *across* models, and *about* models—can help the policymaker understand the extent of the problem (Hansen 2014; Marinacci 2015; Hansen and Marinacci 2016; Aydogan et al. 2018).

Uncertainty *within* models reflects the standard notion of risk: uncertain outcomes with known probabilities. Models may include random shocks or impulses with prespecified distributions. It is the modeling counterpart to flipping coins or rolling dice in which we have full confidence in the probability assessment.

Uncertainty *across* models encompasses both the unknown parameters for a family of models or more discrete modeling differences in specification. Thus, it relates to unknown inputs needed to construct fully specified probability models. In the COVID-19 context, this corresponds for example to the uncertainty of some model parameters, such as how much transmission occurs in different age groups or how infectious people can be before they have symptoms. Existing data, if available and reliable, can help calibrate these model inputs. An additional challenge for the policymaker is the proliferation of modeling groups, researchers, and experts in various disciplines (epidemiology, economics, and other social sciences). Each of these provides forecasts and projections about the disease's evolution and/or its socio-economic consequences. This uncertainty *across* models and their consequent predictions may be difficult to handle by policymakers, especially as one approach is not necessarily superior to another but simply adds another perspective (S.-L. Li et al. 2017). There is no single 'view.' Analysis of this form of uncertainty is typically the focal point of statistical approaches. Bayesian analyses, for instance, confront this via the use of subjective probabilities, whereas robust Bayesians explore sensitivity to prior inputs. Decision theory explores the ramifications of subjective uncertainty, as there might be substantial variation in the recommendations *across* different models and experts, reflecting other specific choices and assumptions regarding modeling type and structure.

Finally, as models are, by design, simplifications of more complex phenomena, they are necessarily misspecified, at least along some dimensions. For instance, they might not mention certain variables that matter, which modelers are or are not aware of, or they may be limited in the scope of functional relationships considered, unknown forms of specification and measurement errors, and so forth. Consequently, there is also uncertainty *about* the models' assumptions and structures. It might sometimes be challenging, even for experts, to assess

the merits and limits of alternative models and predictions^{*}. This is what we mean in our reference to uncertainty *about* models.

3. How to make rational decisions under uncertainty?

Now that we have characterized the elements of the decision problem under uncertainty (see Figure 1), the question remains on how to make the best possible decision? In other words, how should the policymaker proceed to aggregate the different (and usually conflicting) scientific findings, model results, and expert opinions—which are all uncertain by construction and by lack of reliable data—and ultimately determine policy? Insights from modern decision theory are of the most significant value at this stage. They propose normative guidelines and “rules,” to help policymakers make the best, i.e., the most rational decision under uncertainty.




Figure 1: Overview of the decision problem under uncertainty

^{*} Note that another way to see this additional layer of uncertainty is as uncertainty over predictions of alternative models that have not been developed yet.

3.1. How can formal decision rules be useful?

The formal decision rules proposed by decision theorists are powerful, mathematically-founded[†] tools that relate theoretical constructs and choice procedures to presumably observable data. Making a decision based on such rules is equivalent to complying implicitly with a set of general consistency conditions or principles governing human behavior. During a crisis such as the COVID-19 pandemic, using decision theory as a formal guide will lend credibility to policymaking by ensuring that the resulting actions are coherent and defensible. To illustrate how decision theory can serve as a *coherence test* (Itzhak Gilboa and Samuelson 2020), imagine the case of a policymaker trying to determine what the best response to the current pandemic is. The decision-makers can make up their minds by whatever mix of intuition, expert advice, imitation, and quantitative model results they have available, and then check their judgment by asking whether they can justify the decision using a formal decision rule. Conceptually, it can be seen as a form of dialogue between the policymakers and decision theory, in which an attempt to justify a tentative decision helps to clarify the problem and, perhaps, leads to a different conclusion (Itzhak Gilboa, Rouziou, and Sibony 2018). Used this way, formal decision rules may help policymakers clarify the problem they are dealing with, test their intuition, eliminate strictly dominated options, and avoid reasoning mistakes and pitfalls that have been documented in psychological studies (e.g., confirmation bias, optimism bias, representativeness heuristic, prospect theory, etc.) (Cairney and Kwiatkowski 2017).

Finally, because committees might investigate how decisions were taken during the crisis, for example, about how lockdown measures were implemented and lifted, policymakers are held to account for the actions they took. A formal decision model can play an essential role in defending one's choice and generating ex-post justifiability. For example, it could help a policymaker, who had to decide which neighborhoods to keep under lockdown and which not, to explain the process that led to such decisions to citizens who might think they have not been treated fairly.

3.2. Which decision rule to follow?

As decision theory proposes a variety of different rules for decision making under uncertainty, the call for using decision theory begs the question, which rules to follow? The answer depends, in our opinion, on the society or organization in question. Decision theory should offer a gamut of models, and the people for whom decisions are made should find acceptable the model that is considered to “provide a justification” for a given decision. Thus, the answer

[†] Typically, each of these rules results from an axiomatization (i.e., an equivalence result taking the form of a theorem that relates a theoretical description of decision-making to conditions on observable data (Itzhak Gilboa et al. 2019)).

ultimately depends on the policymakers' characteristics, e.g., which conditions or behavioral principles they want to comply with, how prudent they want the policy to be, or what answer their constituency expects to receive. In Figure 2, we present a simple example of school closures' decision problem during the COVID-19 pandemic. We use this to demonstrate how distinct quantitative model outputs (some of which represent "best guesses" while others represent "reasonable worst-case"[‡] possibilities) can be combined and used in formal decision rules, and what the resulting recommendations in terms of policy responses are.

The policymaker's problem consists in finding the right balance between protecting the health and preventing economic and social disruptions by choosing whether and for how long to keep schools closed, given the scarce scientific evidence and the disagreement that may exist across model projections; possibly leading to significantly different policies.

The decision rules that we present differ primarily in how they handle probabilities. According to the Bayesian view, which holds that any source of uncertainty can be quantified probabilistically, the policymaker should always have well-defined probabilities about the impacts of the measures taken. If they rely on quantitative model outputs or expert advice to obtain different estimates, then they should attach a well-defined probability weight to each of these and compute an average. Thus, in the absence of objective probabilities, the decision-makers have their own subjective probabilities to guide decisions.

However, it may not always be rational to follow this approach (Itzhak Gilboa, Postlewaite, and Schmeidler 2008; 2009; 2012; Itzhak Gilboa and Marinacci 2013). Its limitation stems from its inability to distinguish between uncertainty *across* models (which has an epistemic nature, and is due to limited knowledge or ignorance) and uncertainty *within* models (which as an aleatory nature, and is due to the intrinsic randomness in the world). In the response to the COVID-19 outbreak, the Bayesian approach requires the policymaker to express probabilistic beliefs (about the impact of a policy, about the correctness of a given model, etc.), without being told which probability it makes sense to adopt, nor being allowed to say "I don't know". Because of the disagreements that may exist across different model outputs, or expert opinions, another path may be to acknowledge one's ignorance and relax the assumption that we can associate precise probabilities to any event. Modern decision theory proposes decision rules in line with this non-Bayesian approach. The axiomatic approach on which it is founded serves as an essential guide in understanding the merits and limitations of alternative ways to confront uncertainty formally. While we do not see this theory as settling on a single recipe for all decision problems, it adds important clarity to the rationale behind alternative decision rules.

[‡] As these projections are typically premised on "reasonable" bounds in terms of their model inputs.

Case study: Decisions about school closures and their length during COVID-19 pandemic

In this case study, we explore the problem of a policymaker having to make a decision about school closures and their length during the COVID-19 pandemic and illustrate the difference of policy prescribed by different decision rules.

1. Context: By end of April 2020, 191 countries had implemented national school closures in response to the COVID-19 pandemic (United Nations Educational Scientific and Cultural Organization 2020). Yet the effectiveness of such a measure is highly uncertain, due to the lack of data on the relative contribution of school closures to transmission control, and conflicting modeling results (Viner et al. 2020).


2. The policymaker's problem: Decisions about closures and their length involve a series of trade-offs. The policy assessment thus involves weighting the benefits and costs of alternative courses of action. On the one hand, school closures can slow the pandemic and its impact by reducing child-child transmission, thus delaying the pandemic peak that overwhelms health care services, and therefore ultimately reducing morbidity and associated mortality. If this is the case, such interventions bring clear health benefits for society and avoid unsustainable demands on the health system. On the other hand, school closure can have high direct and indirect health and socio-economic costs. For example, they may increase child-adult transmission, reduce the ability of healthcare and key workers to work, and thus reduce the capacity of healthcare (Bayham and Fenichel 2020; Brooks et al. 2020). The economic costs of lengthy school closures are also high (Sadique, Adams, and Edmunds 2008; M. R. Keogh-Brown et al. 2010; Lempel, Epstein, and Hammond 2009), generated for example through absenteeism by working parents, loss of education, etc.

3. Uncertainty: The evidence supporting national closure of schools in the COVID-19 pandemic context is very weak. In particular, evidence of COVID-19 transmission through child-child contact or through schools is not available at the time of decision (Viner et al. 2020). As a consequence, it is unclear whether school closures are effective in the COVID-19 pandemic (Bayham and Fenichel 2020).

4. The formal decision problem

4.1 Setup

In this



example, the policy action is to choose whether to and how long to close schools. The consequence includes both the benefits of the action (lives saved, reduction in future cases and beds needed, etc.) and its costs (reduction in education for children, health care workers, and other key workers not working, lives cost due to changes in disease dynamics, etc.). The consequence also depends on the realization of a state of the environment. For example, the way the number of deaths, beds, and cases is affected by school closures depends on the biology of the virus and baseline transmission dynamics, which are outside the decision-maker's control. The consequence function then relates actions and states to consequences. It can for

4.2 Model uncertainty

As there is no evidence supporting school closures at the early stage of the COVID-19 pandemic, policymakers rely on different epidemiological model projections and/or the advice of experts to assess the effectiveness of the measure. For example, imagine three different projections.

- The first projection (**Model 1**) is based on the only evidence we have, which is the one coming from influenza outbreaks for which the majority of transmission is between children (Jackson C, Mangtani P, and Vynnycky E 2014). Closing schools is thus the biggest contributor to reducing R^* to below 1 and it may be the only intervention that could do so. In this case, the benefit is proportional to the duration of school closure.
- Alternatively, the second projection (**Model 2**) relies on some previous coronavirus outbreaks, for which evidence suggests minimal transmission between children (Wong et al. 2003). Here, R^* cannot be reduced below 1, school closures do not affect the size of the epidemic, and therefore do not bring any benefits.
- The third scenario (**Model 3**) projects that some child to child transmission happens so that closing schools contributes to reducing R^* to below 1 and reduces the size of the epidemic. However, this only works in combination with other measures (Prem et al. 2020). Without it, R^* would be above 1, but as an isolated measure school closure does not have such a big effect (Ferguson et al. 2020). Under this scenario, the effectiveness of school closures is important at

4.3 Policy objective

In view of the scarce evidence concerning the use and effectiveness of school closures during the COVID-19 pandemic, as well as the disagreement that may exist across model projections and/or expert opinions, policymakers have to find the right balance between protecting the health and preventing economic and social disruptions. Choosing the appropriate length of this common-sense measure may be exceptionally challenging as lengthy school closures bear very high costs and can therefore substantially reduce any benefit to health systems and populations, whereas earlier relaxation of the measure increases the risk that transmission surged again, leading to a second peak.

Preferences The choice made by the policymaker depends ultimately on her preferences, such as the degree to which she likes/dislikes uncertainty (Berger and Bosetti 2020). These preferences are represented numerically via a decision rule V . The value $V(a)$ attained by selecting an action a may be interpreted in welfare terms.

Optimum Before making a decision, the policymaker knows all the elements of the decision problem. After the decision, she only observes the consequence of the

	Children are infectious	Children are not infectious
No school closures	<ul style="list-style-type: none"> • Very high death toll (health care system overwhelmed) • Very high socio-economic costs 	<ul style="list-style-type: none"> • Status quo
Short-term closures	<ul style="list-style-type: none"> • High death toll (health care system under stress due to a second wave after schools reopen) • High socio-economic costs (due to the second wave) 	<ul style="list-style-type: none"> • Moderate socio-economic costs
Long-term closures	<ul style="list-style-type: none"> • Moderate death toll (health care system operational) • High socio-economic costs (workforce affected) 	<ul style="list-style-type: none"> • High socio-economic costs

example be the net benefit (benefits-costs) of school closures, expressed in monetary terms.

the beginning but declines as time goes on.


Assuming there is no uncertainty regarding the costs of school closures, the collection of potential models characterizing a combination of health/economy environment, M consists of three elements.

intervention chosen. In formal terms, the objective of the policymaker is to select the action a (in our case, the duration of school closures) that is optimal according to her preferences, in the sense that it is preferred to any other action available.

4.4 Illustration

The marginal benefit (MB) of school closure is constant and positive (model 1), null (model 2), or positive and decreasing with the duration of the measure (model 3). In contrast, the marginal cost (MC) is increasing.

It is desirable to maintain the school closed as long as the MB outweighs the MC. So, if the 'true' MB was known, it would be easy to find the optimal duration of school closures: 0 if model 2 is the correct one, 10 if it is model 3, and 20 in the case of model 1. Yet, in reality, there is a lot of uncertainty.



5. Decision rules and optimal choices

To address the epistemic uncertainty across models, the policymaker may follow different decision rules. These rules differ on whether the policymaker may quantify her belief about which is the correct model. If it exists, we let $\mu(m)$ be the policymaker's subjective belief that m is the true model.

5.1 Subjective Expected utility rule

The expected utility has long been the standard way to consider rational decision making under uncertainty (L. J. Savage 1954). We assume that a utility function u translates economic monetary consequences into utility levels. This function captures attitudes towards uncertainty within models. For each action a and each model m , we can compute the expected reward associated with a given action: $R(a, m) = \sum_s u_0(a, s) p(m, s)$. Given that different models exist, an expected reward is considered for each possible model. They are averaged according to the policymaker's beliefs. The subjective expected utility criterion (Cerreia-Vioglio et al. 2013) is

$V_{SEU}(a) = \sum_m R(a, m) \mu(m)$. For example, if the prior distribution is uniform (i.e. $\mu(1) = \mu(2) = \mu(3) = \frac{1}{3}$), and u is linear, the optimal policy is a 10-week closure.

5.2 Smooth ambiguity rule

The smooth ambiguity criterion (Klibanoff, Marinacci, and Mukerji 2005) proposes another way to distinguish attitudes toward uncertainty within and across models. It takes the form

$V_{SAB}(a) = \sum_m \phi(R(a, m)) \mu(m)$, where the concavity of ϕ reflects uncertainty aversion (i.e. being more averse to uncertainty across models than u within models). In our example, if the prior distribution is uniform, and ϕ is logarithmic, the optimal policy is $a=12$ weeks.

5.3 Maxmin rules

The maxmin rule (Wald 1950) is an extremely cautious rule that makes the decision-maker consider only the model providing the lowest expected reward. This is the case when the "worst" health/economy model is considered (i.e. model 1). Here, prior probabilities do not play any role when choosing the optimal policy. The maxmin criterion is written $V_{MM}(a) = \min_m R(a, m)$. In our example, the optimal policy is to keep the school closed for 20 weeks.

A more general version of this rule, which is due to (Hurwicz 1951), consists in considering both the worst and the best possible models. In this case, the criterion is written $V_{CDMM}(a) = \alpha \min_m R(a, m) + (1 - \alpha) \max_m R(a, m)$, where α is the coefficient of pessimism.

5.4 Multiple priors rules

An alternative approach (I. Gilboa and Schmeidler 1989) relaxes the assumption that the policymaker can quantify uncertainty across models through a single probability distribution μ . Instead, because she does not have sufficient information, the policymaker may have multiple priors over the different models. The multiple priors decision rule is $V_{MP}(a) = \min_{\mu \in \mathcal{L}} \sum_m R(a, m) \mu(m)$, where \mathcal{L}

is the set of priors. In contrast with the maxmin rule, this criterion considers the least favorable among all the subjective expected utilities determined by each prior μ . In our example, a particular prior distribution may be the uniform that gives equal weights, $\mu(m) = 1/3$, to all the possible models, while another prior may not consider model 2 to be plausible (in which case, some $\mu(2)=0$). In this case, the multiple priors rule leads to an optimal policy of 13 weeks. Note, finally, that a

more general version of this rule, where both the "max" and the "min" appear with weights α and $1 - \alpha$, is known as the α -maxmin multiple priors rule (Ghirardato, Maccheroni, and Marinacci 2004).

Figure 2: Case study. School closures and their length during the COVID-19 pandemic. Details are provided in the SI.

3.3. Discussion

The decision rules presented in Figure 2 are fully compatible with normative interpretations and could be particularly useful to design robust policies in this COVID-19 pandemic context. They assume that policymakers cope with uncertainty without reducing everything to risk, a pretension that tacitly presumes better information than they typically have. When exploring alternative courses of action, policymakers are necessarily unsure of the consequences. In such a context, sticking to the Bayesian expected utility paradigm not only requires substantive expertise (in weighting the pros and cons of alternative models) but also overshadows the policymaker's reaction to the variability that may exist across models. While we focus, in Figure 2, on a subset of decision rules, which can be checked for logical consistency, it should be clear, however, that other criteria, such as minmax regret (Leonard J Savage 1951), also exist and have been used in some applied contexts (Smith 1996; Manski 2019b). As mentioned above, we believe that a decision criterion is also a matter of personal preferences, which should somehow be aggregated over the different individuals for whom the decisions are made. Thus, the examples used in this paper are inevitably subjective, too.

We recognize the challenges in using decision theory when the decision-making process itself is complicated, and many participants are involved with potentially different incentives. Nevertheless, we also see value to its use in less formal ways as guideposts to prudent decision-making and as a sensible way of framing the uncertainties in the trade-offs that policymakers are presented with.

In this example, the decision problem setup has been deliberately kept to minimal complexity to focus on the decision theory aspects. In particular, the set of actions is here limited to a single intervention (the duration of the school closure). In reality, the decision problem would, of course, require a much higher dimensional space (e.g., selective local closures, school dismissal, etc.), the interaction with other social distancing measures, or the ability to integrate start and stop times. Along the same lines, time constraints, learning, and dynamic considerations have been assumed away for the sake of tractability. In reality, it should be clear that the existence of deadlines could restrict which actions are feasible so that different sets of actions may correspond to different timings.

Similarly, as time passes, experts learn more about virus transmission and disease dynamics, which ultimately leads them to update their projections. Different “updating rules” allow incorporating such new information into the decision-making process. Our general message is the same as the one concerning the decision rules: the decision-maker should make her choice of the updating rule and be able to justify her decision based on this rule and the conditions that it does or doesn't satisfy. Finally, for expositional simplicity, we also abstracted from

concerns about model misspecification, while recognizing this to be an integral part of how decision-makers should view the alternative models or perspectives that they confronted.

4. Concluding remarks

During a period of crisis, policymakers, who make decisions on behalf of others, may be required to provide a protocol that suggests a decision-theoretic model supporting their decisions. Decision theory can contribute to a pandemic response by providing a way to organize a large amount of potentially conflicting scientific knowledge and providing rules for evaluating response options and turning them into concrete decision-making.

In this *perspective*, we have highlighted the importance of quantitative modeling to support policy decisions (the same recommendation has also been made in other public health contexts (Manski 2019b)). This use of models is common in different macroeconomic settings, including the assessments of monetary and fiscal policies. Some may see quantitative modeling as problematic because it requires seemingly arbitrary subjective judgments about the correctness of the different model specifications, leading them to prefer qualitative approaches. Even qualitative methods cannot escape the need for subjective inputs, however. Restricting scientific inputs to be only qualitative limits severely potentially valuable inputs into prudent policymaking. Instead, we argue in favor of using quantitative models and data, including explicit information about our underlying knowledge's limits. We propose decision rules that incorporate the decisionmaker's confidence in her subjective probabilities, thus rendering the decision-making process based on formal quantitative rules, both robust and prudent.

In practical terms, ensuring that policy options are in line with formal decision rules could be achieved by having a decision analyst in the group of advisors to nurture a dialogue between policymakers and decision theory. This dialogue could clarify the trade-offs and encourage a more sanguine response to the uncertainties present when assessing the alternative courses of action and result in an improved policy outcome (Itzhak Gilboa and Samuelson 2020; Itzhak Gilboa et al. 2018).

To make the decision-making process under uncertainty more efficient, we also suggest acknowledging and communicating the various uncertainties transparently (Manski 2019a). For example, illustrating, quantifying, and discussing the multiple sources of uncertainty may help policymakers better understand their choices' potential impact. To this aim, modelers should provide all information needed to reconstruct the analysis, including information about model structures, assumptions, and parameter values. Moreover, the way uncertainty around

these choices affects model results needs to be accurately communicated, such as systematically reporting uncertainty boundaries around the estimates provided (World Health Organization and other 2019). Scientific and policy advisors would then need to synthesize all this information (32) –coming from diverse sources across different disciplines, possibly of different quality—to help policymakers turning it into actionable information for decisions, while making sure the complete range of uncertainty (including *within* and *across* models) is clearly reported and understood properly (Spiegelhalter, Pearson, and Short 2011; Bosetti et al. 2017).

One possible way to go is to enhance standardization by developing and adopting standard metrics to evaluate and communicate the degree of certainty in key findings. While several approaches have been proposed (van der Bles et al. 2019), insights could, for example, be gained from the virtues and the shortcomings of the reports of the Intergovernmental Panel on Climate Change (IPCC) (Mastrandrea et al. 2010). Another way is to develop further communication and collaboration between model developers and decision-makers to improve the quality and utility of models and the decisions they support (Rivers et al. 2019).

Finally, while policymakers are responsible for making decisions, they are also responsible for communicating to professionals and the public. The way individuals respond to advice and measures selected is as vital as government actions, if not more (3). Communication should thus be an essential part of the policy response to uncertainty. In particular, government communication strategies to keep the public informed of what we (do not (World Health Organization and 2017)) know should balance the costs and benefits of revealing information (how much, and in what form) (Aikman et al. 2011).

As government strategies have been extensively debated in the media and models have become more scrutinized, one lesson learned from the COVID-19 management experience may be that policymakers and experts must increase their approaches' transparency. Using the constructs from decision theory in policymaking, even in an informal way, will help ensure prudent navigation through the uncertainty that pervades this and possibly future pandemics. Being open about the degree of uncertainty surrounding the scientific evidence used to guide policy choices and allowing for the assumptions of the models used or for the decision-making process itself to be challenged is a valuable way of retaining public trust (Fiske and Dupree 2014). At the same time, it is essential to counteract what is too often displayed by self-described experts who seek to influence policymakers and the public by projecting a pretense of knowledge that is likely to be false.

References

- Aikman, David, Philip Barrett, Sujit Kapadia, Mervyn King, James Proudman, Tim Taylor, Iain De Weymarn, and Tony Yates. 2011. 'Uncertainty in Macroeconomic Policy-Making: Art or Science?' *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369 (1956): 4798–4817.
- Alvarez, Fernando E, David Argente, and Francesco Lippi. 2020. 'A Simple Planning Problem for Covid-19 Lockdown'. *American Economic Review: Insights* forthcoming.
- Anderson, Roy M, Hans Heesterbeek, Don Klinkenberg, and T Déirdre Hollingsworth. 2020. 'How Will Country-Based Mitigation Measures Influence the Course of the COVID-19 Epidemic?' *The Lancet* 395 (10228): 931–934.
- Aydogan, Ilke, Loïc Berger, Valentina Bosetti, and Ning Liu. 2018. 'Three Layers of Uncertainty and the Role of Model Misspecification: An Experiment'. Bocconi University.
- Bayham, Jude, and Eli P Fenichel. 2020. 'Impact of School Closures for COVID-19 on the US Health-Care Workforce and Net Mortality: A Modelling Study'. *The Lancet Public Health* 5 (5): E271–78.
- Berger, Loïc, and Valentina Bosetti. 2020. 'Are Policymakers Ambiguity Averse?' *The Economic Journal* 130: 331–355.
- Bles, Anne Marthe van der, Sander van der Linden, Alexandra LJ Freeman, James Mitchell, Ana B Galvao, Lisa Zaval, and David J Spiegelhalter. 2019. 'Communicating Uncertainty about Facts, Numbers and Science'. *Royal Society Open Science* 6 (5): 181870.
- Bosetti, Valentina, Elke Weber, Loïc Berger, David V Budescu, Ning Liu, and Massimo Tavoni. 2017. 'COP21 Climate Negotiators' responses to Climate Model Forecasts'. *Nature Climate Change* 7 (3): 185–190.
- Brooks, Samantha K, Louise E Smith, Rebecca K Webster, Dale Weston, Lisa Woodland, Ian Hall, and G James Rubin. 2020. 'The Impact of Unplanned School Closure on Children's Social Contact: Rapid Evidence Review'. *Eurosurveillance* 25 (13): 2000188.
- Cairney, Paul, and Richard Kwiatkowski. 2017. 'How to Communicate Effectively with Policymakers: Combine Insights from Psychology and Policy Studies'. *Palgrave Communications* 3 (1): 37.
- Cancryn, Adam. 2020. 'How Overly Optimistic Modeling Distorted Trump Team's Coronavirus Response'. Politico. 2020. <https://www.politico.com/news/2020/04/24/trump-coronavirus-model-207582>.
- Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci, and Luigi Montrucchio. 2013. 'Classical Subjective Expected Utility'. *Proceedings of the National Academy of Sciences* 110 (17): 6754–6759.
- Chater, Nick. 2020. 'Facing up to the Uncertainties of COVID-19'. *Nature Human Behaviour*, 1–1.
- Davies, Nicholas G, Adam J Kucharski, Rosalind M Eggo, Amy Gimma, W John Edmunds, CMMID COVID-19 Working Group, and others. 2020. 'Effect of Non-Pharmaceutical Interventions on COVID-19 Cases, Deaths, and Demand for Hospital Services in the UK: A Modelling Study'. *The Lancet Public Health* 5 (7): e375–85.
- Den Boon, Saskia, Mark Jit, Marc Brisson, Graham Medley, Philippe Beutels, Richard White, Stefan Flasche, et al. 2019. 'Guidelines for Multi-Model Comparisons of the Impact of Infectious Disease Interventions'. *BMC Medicine* 17 (1): 163.
- Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt. 2020. 'The Macroeconomics of Epidemics'. National Bureau of Economic Research.

- Emanuel, Ezekiel J., Govind Persad, Ross Upshur, Beatriz Thome, Michael Parker, Aaron Glickman, Cathy Zhang, Connor Boyle, Maxwell Smith, and James P. Phillips. 2020. 'Fair Allocation of Scarce Medical Resources in the Time of Covid-19'. *New England Journal of Medicine* 382 (21): 2049–55. <https://doi.org/10.1056/NEJMs2005114>.
- Ferguson, Neil, Daniel Laydon, Gemma Nedjati Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, et al. 2020. 'Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID19 Mortality and Healthcare Demand'.
- Fiske, Susan T, and Cydney Dupree. 2014. 'Gaining Trust as Well as Respect in Communicating to Motivated Audiences about Science Topics'. *Proceedings of the National Academy of Sciences* 111 (Supplement 4): 13593–13597.
- Ghirardato, P., F. Maccheroni, and M. Marinacci. 2004. 'Differentiating Ambiguity and Ambiguity Attitude'. *Journal of Economic Theory* 118 (2): 133–173.
- Gilboa, I., and D. Schmeidler. 1989. 'Maxmin Expected Utility with a Non-Unique Prior'. *Journal of Mathematical Economics* 18 (2): 141–154.
- Gilboa, Itzhak, and Massimo Marinacci. 2013. 'Ambiguity and the Bayesian Paradigm'. In *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society*. D. Acemoglu, M. Arellano, and E. Dekel (Eds.). New York: Cambridge University Press.
- Gilboa, Itzhak, Andrew Postlewaite, Larry Samuelson, and David Schmeidler. 2018. 'Economics: Between Prediction and Criticism'. *International Economic Review* 59 (2): 367–390.
- . 2019. 'What Are Axiomatizations Good For?' *Theory and Decision* 86 (3–4): 339–359.
- Gilboa, Itzhak, Andrew Postlewaite, and David Schmeidler. 2009. 'Is It Always Rational to Satisfy Savage's Axioms?' *Economics and Philosophy* 25 (03): 285–296.
- . 2012. 'Rationality of Belief or: Why Savage's Axioms Are Neither Necessary nor Sufficient for Rationality'. *Synthese* 187 (1): 11–31.
- Gilboa, Itzhak, Andrew W Postlewaite, and David Schmeidler. 2008. 'Probability and Uncertainty in Economic Modeling'. *The Journal of Economic Perspectives* 22 (3): 173–188.
- Gilboa, Itzhak, Maria Rouziou, and Olivier Sibony. 2018. 'Decision Theory Made Relevant: Between the Software and the Shrink'. *Research in Economics* 72 (2): 240–250.
- Gilboa, Itzhak, and Larry Samuelson. 2020. 'What Were You Thinking? Decision Theory as Coherence Test'.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2020. 'Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?' National Bureau of Economic Research.
- Hansen, Lars Peter. 2014. 'Nobel Lecture: Uncertainty Outside and Inside Economic Models'. *Journal of Political Economy* 122 (5): 945–987.
- . 2020. 'Using Quantitative Models to Guide Policy Amid COVID-19 Uncertainty'.
- Hansen, Lars Peter, and Massimo Marinacci. 2016. 'Ambiguity Aversion and Model Misspecification: An Economic Perspective'. *Statistical Science* 31: 511–515.
- Hellewell, Joel, Sam Abbott, Amy Gimma, Nikos I Bosse, Christopher I Jarvis, Timothy W Russell, James D Munday, et al. 2020. 'Feasibility of Controlling COVID-19 Outbreaks by Isolation of Cases and Contacts'. *The Lancet Global Health* 8 (4): e488–e496.
- Hollingsworth, T Déirdre, Don Klinkenberg, Hans Heesterbeek, and Roy M Anderson. 2011. 'Mitigation Strategies for Pandemic Influenza A: Balancing Conflicting Policy Objectives'. *PLoS Computational Biology* 7 (2).

- Holmdahl, Inga, and Caroline Buckee. 2020. 'Wrong but Useful – What Covid-19 Epidemiologic Models Can and Cannot Tell Us'. *New England Journal of Medicine* 384 (4): 303–5. <https://doi.org/10.1056/NEJMp2016822>.
- Hurwicz, Leonid. 1951. 'Some Specification Problems and Applications to Econometric Models'. *Econometrica* 19 (3): 343–344.
- Jackson C, P, Mangtani P, and Vynnycky E. 2014. *Impact of School Closures on an Influenza Pandemic: Scientific Evidence Base Review*. London: Public Health England.
- Johnson-Laird, Philip N. 2010. 'Mental Models and Human Reasoning'. *Proceedings of the National Academy of Sciences* 107 (43): 18243–18250.
- Keogh-Brown, Marcus, Henning Jensen, William Edmunds, and Richard; Smith. 2020. 'The Impact of Covid-19, Associated Behaviours and Policies on the UK Economy: A Computable General Equilibrium Model'. *Social Science & Medicine*.
- Keogh-Brown, Marcus Richard, Richard D Smith, John W Edmunds, and Philippe Beutels. 2010. 'The Macroeconomic Impact of Pandemic Influenza: Estimates from Models of the United Kingdom, France, Belgium and The Netherlands'. *The European Journal of Health Economics* 11 (6): 543–554.
- Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. 'A Smooth Model of Decision Making under Ambiguity'. *Econometrica* 73: 1849–1892.
- Lazzerini, Marzia, and Giovanni Putoto. 2020. 'COVID-19 in Italy: Momentous Decisions and Many Uncertainties'. *The Lancet Global Health* 8 (5): E641–E642.
- Lempel, Howard, Joshua M Epstein, and Ross A Hammond. 2009. 'Economic Cost and Health Care Workforce Effects of School Closures in the US'. *PLoS Currents* 1.
- Lempert, Robert J, and Myles T Collins. 2007. 'Managing the Risk of Uncertain Threshold Responses: Comparison of Robust, Optimum, and Precautionary Approaches'. *Risk Analysis* 27 (4): 1009–1026.
- Li, Ruiyun, Sen Pei, Bin Chen, Yimeng Song, Tao Zhang, Wan Yang, and Jeffrey Shaman. 2020. 'Substantial Undocumented Infection Facilitates the Rapid Dissemination of Novel Coronavirus (SARS-CoV-2)'. *Science* 368 (6490): 489–493.
- Li, Shou-Li, Ottar N. Bjørnstad, Matthew J. Ferrari, Riley Mummah, Michael C. Runge, Christopher J. Fonnesbeck, Michael J. Tildesley, William J. M. Probert, and Katriona Shea. 2017. 'Essential Information: Uncertainty and Optimal Control of Ebola Outbreaks'. *Proceedings of the National Academy of Sciences* 114 (22): 5659–5664. <https://doi.org/10.1073/pnas.1617482114>.
- Manski, Charles F. 2019a. 'Communicating Uncertainty in Policy Analysis'. *Proceedings of the National Academy of Sciences* 116 (16): 7634–7641.
- . 2019b. *Patient Care under Uncertainty*. Princeton University Press.
- Manski, Charles F. 2020. 'Forming COVID-19 Policy Under Uncertainty'. *Journal of Benefit-Cost Analysis*, 1–16. <https://doi.org/10.1017/bca.2020.20>.
- Marinacci, Massimo. 2015. 'Model Uncertainty'. *Journal of the European Economic Association* 13 (6): 1022–1100. <https://doi.org/10.1111/jeea.12164>.
- Mastrandrea, Michael D, Christopher B Field, Thomas F Stocker, Ottmar Edenhofer, Kristie L Ebi, David J Frame, Hermann Held, et al. 2010. 'Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties'.
- Morgan, Oliver. 2019. 'How Decision Makers Can Use Quantitative Approaches to Guide Outbreak Responses'. *Philosophical Transactions of the Royal Society B* 374 (1776): 20180365.
- Prem, Kiesha, Yang Liu, Timothy W Russell, Adam J Kucharski, Rosalind M Eggo, Nicholas Davies, Stefan Flasche, et al. 2020. 'The Effect of Control Strategies to Reduce Social

- Mixing on Outcomes of the COVID-19 Epidemic in Wuhan, China: A Modelling Study'. *The Lancet Public Health* 5 (5): E261–70.
- Rivers, Caitlin, Jean-Paul Chretien, Steven Riley, Julie A Pavlin, Alexandra Woodward, David Brett-Major, Irina Maljkovic Berry, et al. 2019. 'Using "Outbreak Science" to Strengthen the Use of Models during Epidemics'. *Nature Communications* 10 (1): 1–3.
- Rucker, Philip, Josh Dawsey, Yasmeen Abutaleb, Robert Costa, and Lena H. Sun. 2020. '34 Days of Pandemic: Inside Trump's Desperate Attempts to Reopen America'. *The Washington Post*. 2020. https://www.washingtonpost.com/politics/34-days-of-pandemic-inside-trumps-desperate-attempts-to-reopen-america/2020/05/02/e99911f4-8b54-11ea-9dfd-990f9dcc71fc_story.html.
- Sadique, Md Z, Elisabeth J Adams, and William J Edmunds. 2008. 'Estimating the Costs of School Closure for Mitigating an Influenza Pandemic'. *BMC Public Health* 8 (1): 135.
- Savage, L. J. 1954. *The Foundations of Statistics*. New York: J. Wiley.
- Savage, Leonard J. 1951. 'The Theory of Statistical Decision'. *Journal of the American Statistical Association* 46 (253): 55–67.
- Smith, Richard David. 1996. 'Is Regret Theory an Alternative Basis for Estimating the Value of Healthcare Interventions?' *Health Policy* 37 (2): 105–15. [https://doi.org/10.1016/S0168-8510\(96\)90055-X](https://doi.org/10.1016/S0168-8510(96)90055-X).
- Spiegelhalter, David, Mike Pearson, and Ian Short. 2011. 'Visualizing Uncertainty about the Future'. *Science* 333 (6048): 1393–1400.
- The Editors. 2020. 'Dying in a Leadership Vacuum'. *New England Journal of Medicine* 383 (15): 1479–80. <https://doi.org/10.1056/NEJMe2029812>.
- Thunström, Linda, Stephen C Newbold, David Finnoff, Madison Ashworth, and Jason F Shogren. 2020. 'The Benefits and Costs of Using Social Distancing to Flatten the Curve for COVID-19'. *Journal of Benefit-Cost Analysis*, 1–27.
- United Nations Educational Scientific and Cultural Organization. 2020. *COVID-19 Educational Disruption and Response*. <https://en.unesco.org/themes/education-emergencies/coronavirus-school-closures>.
- Viner, Russell M, Simon J Russell, Helen Croker, Jessica Packer, Joseph Ward, Claire Stansfield, Oliver Mytton, Chris Bonell, and Robert Booy. 2020. 'School Closure and Management Practices during Coronavirus Outbreaks Including COVID-19: A Rapid Systematic Review'. *The Lancet Child & Adolescent Health* 4 (5): P397-404.
- Wald, A. 1950. *Statistical Decision Functions*. New York: John Wiley & Sons.
- Wong, Gary WK, Albert M Li, PC Ng, and Tai F Fok. 2003. 'Severe Acute Respiratory Syndrome in Children'. *Pediatric Pulmonology* 36 (4): 261–266.
- World Health Organization and. 2017. 'Communicating Risk in Public Health Emergencies: A WHO Guideline for Emergency Risk Communication (ERC) Policy and Practice'. World Health Organization.
- World Health Organization and other. 2019. 'WHO Guide for Standardization of Economic Evaluations of Immunization Programmes'. World Health Organization.

SI – Methods supplement to Figure 2: deciding about school closures and their length during COVID-19 pandemic

Overview

This Appendix supports the case study presented in Figure 2 of the manuscript. Its purpose is to show how decision theory could be used in a context where distinct model projections exist. Using a simple example of a decision problem of school closures during the COVID-19 pandemic, we highlight what are the resulting recommendations from different formal decision rules in terms of policy responses.

It is important to note that the decision problem presented in this case study is necessarily simplistic and should be used for demonstrative purposes only.

1 Background

Context By end of April 2020, 191 countries had implemented national school closures in response to the COVID-19 pandemic (United Nations Educational Scientific and Cultural Organization, 2020). Yet the effectiveness of such a measure is highly uncertain, due to the lack of data on the relative contribution of school closures to transmission control, and conflicting modelling results (Viner et al., 2020).

The policymaker's problem Decisions about closures and their length involve a series of trade-offs. The policy assessment thus involves weighting benefits and costs of alternative courses of action. On the one hand, school closures can slow the pandemic and its impact by reducing child-child transmission, thus delaying the pandemic peak that overwhelms health care services, and therefore ultimately reducing morbidity and associated mortality. If this is the case, such interventions bring clear health benefits for the society and avoid unsustainable demands on the health system. On the other hand, school closures can have high direct and indirect health and socio-economic costs. For example, they may increase child-adult transmission, reduce the ability of healthcare and key workers to work and thus reduce the capacity of healthcare (Brooks et al., 2020; Bayham and Fenichel, 2020). Economic costs of lengthy school closures are also high (Sadique et al., 2008; Lempel et al., 2009; Smith et al., 2011), generated for example through absenteeism by working parents, loss of education, etc.

Uncertainty The evidence supporting national closure of schools in the COVID-19 pandemic context was very weak. In particular, evidence of COVID-19 transmission through child-child contact or through schools was not available at the time of decision (Viner

et al., 2020). As a consequence, it was unclear whether school closures would be effective in the COVID-19 pandemic (Bayham and Fenichel, 2020).

Framework We use a framework that decomposes uncertainty into distinct layers of analysis: (i) uncertainty *within* models (also called *risk*, *aleatory* uncertainty, or *physical* uncertainty), (ii) uncertainty *across* models (also called *model ambiguity*, or *model uncertainty*), and (iii) uncertainty *about* models (also called *model misspecification*).¹

We consider a general decision problem in which consequences depend on the states of the environment that are viewed as realizations of an underlying economic or physical generative mechanism (Marinacci, 2015). A *model* is a probability distribution induced by such a mechanism. It describes states' variability by combining a structural component based on theoretical knowledge (e.g. economic or physical) and a random component coming from, for example, measurement errors or minor omitted explanatory variables (Koopmans, 1947; Marschak, 1953). We assume that decision makers (DMs) posit a collection of such models. Uncertainty *across* model therefore results from the uncertainty about the true underlying mechanism: within the posited collection, there is uncertainty about which model actually governs states' realizations. However, even after a model is specified, there is still uncertainty *within* model, i.e. about which specific state will actually obtain; this is the notion of *risk* typically considered in economics. Finally, the third layer of uncertainty (*about* models), arises as the true model might not belong to the posited collection of models, reflecting the idea that all posited models have an inherent approximate nature.

2 Decision making under uncertainty

2.1 The structure of a decision problem

The general problem that a DM, in particular a policymaker, faces is to choose an *action* a within a set A of possible alternative actions, whose *consequences* $c \in C$ depend on the realization of a *state of the environment* $s \in S$ which is outside the DM's control. The relationship among consequences, actions and states is described by a consequence function $\rho : A \times S \rightarrow C$, where $c = \rho(a, s)$ is the consequence of action a when state s obtains. DMs have a (complete and transitive) *preference relation* \mathbf{t} over actions that describes how they rank the different alternative actions.² The quintet $(A, S, C, \rho, \mathbf{t})$ characterizes the decision problem under uncertainty. The aim of the DM is to select the action \hat{a} that is *optimal* according to her preference, that is, such that $\hat{a} \mathbf{t} a$ for all actions

¹See Arrow (1951); Hansen (2014); Marinacci (2015); Hansen and Marinacci (2016) for a discussion, and Aydogan et al. (2018) for empirical evidence on the distinction between these layers.

²As is usual, we write $\mathbf{a} \mathbf{t} \mathbf{b}$ if the DM prefers action \mathbf{a} to action \mathbf{b} (i.e., either strictly prefers action \mathbf{a} to action \mathbf{b} , $\mathbf{a} \succ \mathbf{b}$, or is indifferent between the two, $\mathbf{a} \sim \mathbf{b}$).

$a \in A$. The preference \mathbf{t} is assumed to admit a numerical representation via a *decision criterion* $V: A \rightarrow \mathbb{R}$, with

$$a \mathbf{t} b \Leftrightarrow V(a) \geq V(b)$$

for all actions $a, b \in A$. This numerical representation permits to formulate the decision problem as an optimization problem

$$\max_a V(a) \quad \text{sub } a \in A. \quad (1)$$

Optimal actions \hat{a} are the solutions of this problem. To find an optimal action thus amounts to solve this optimization problem.

The DM may address, especially in policy problems, state uncertainty through the guise of models. Based on ex ante scientific and socio-economic information, the DM might be able to posit a set of probability models $M \subseteq \Delta(S)$ describing the likelihoods of the different states. This set of models is taken as a datum of the decision problem, which is now characterized by a sextet $(A, S, C, \rho, \mathbf{t}, M)$. It is often assumed, following Wald (1950), that the correct model belongs to the set of models that the DM posits, thus abstracting from model misspecification issues.

2.2 Decision criteria

The form of the decision criterion V determines the nature of decision problem [1]. Different possible criteria have been proposed in the literature. The ones we consider are: the subjective expected utility (SEU) criterion, which dates back to the seminal works of von Neumann and Morgenstern (1947); Wald (1950); Savage (1954); Marschak and Radner (1972), and has recently been revisited by Cerreia-Vioglio et al. (2013) to accommodate explicitly the presence of uncertainty across models; the maxmin criterion of Wald (1950); the smooth ambiguity criterion, developed by Klibanoff et al. (2005); and the multiple priors criterion proposed by Gilboa and Schmeidler (1989). For the sake of brevity, the more general α -versions of the maxmin criterion (Hurwicz, 1951) and of the multiple priors criterion (Ghirardato et al., 2004) are only discussed in the main text.

2.3 Making decisions in a pandemic

2.3.1 States and consequences

With the letter R we denote a rate of contagion within a given population, i.e., the average number of individuals infected per single case. The baseline rate of contagion, denoted by R_0 , is called *basic reproduction number*. It applies to a population never exposed to the virus, where everyone is susceptible,³ and depends on the biology of the

³An individual is *susceptible* if has no immune protection against the virus.

virus as well as on the natural (pre-pandemic) socio-economic structure that characterizes the population (Ng and Wen, 2019). The biology of the virus determines its ability to infect (i.e., the probability of infection per interaction) and the duration of infectiousness.⁴ The natural socio-economic structure determines the natural social distancing and, through it, the average number of interactions per individual (Dietz, 1993; Delamater et al., 2019). For instance, natural social distancing might be higher in Northern than in Southern European countries.

These natural factors, biological and socio-economic, determine R_0 . It is the natural, *ex ante*, rate at which the pandemic progresses, without private and public decisions that respond to it. *Ex post*, after these decisions are put in place and affect the biological and socio-economic factors that determine R_0 , the relevant rate of contagion becomes the *effective reproduction* (or *reproductive*) number R_e (Wallinga and Lipsitch, 2007). For example, school closure is a public decision that may increase social distancing (a socio-economic factor), while a diligent use of protective gear is a private decision that may decrease the virus ability to infect (a biological factor).

Here we focus on public decisions, policies, and assume that private ones are subsumed by them.⁵ A policy translates a basic reproduction number R_0 into an effective one R_e . Yet, how this translation occurs often remains uncertain. For instance, evidence on the effectiveness of school closure policies for COVID-19 comes from influenza outbreaks, but the ability of children to transmit the disease greatly varies across coronaviruses Wong et al. (2003). For this reason, we represent how a policy α maps R_0 into R_e via the relation $R_e = f(\alpha, R_0, \theta_r, \varepsilon_r)$, where θ_r is a structural parameter and ε_r is a shock.⁶ We assume that $\partial f / \partial \alpha \leq 0$ and $\partial f / \partial R_0 > 0$. For example, if the relation is linear we have

$$R_e = \theta_{r,1}\alpha + \theta_{r,2}R_0 + \varepsilon_r, \quad (2)$$

with $\theta_{r,1} \leq 0$ and $\theta_{r,2} > 0$.⁷ In the baseline scenario without policy intervention – i.e., when $\alpha = 0$ – the effective reproduction number R_e is determined by: (i) the natural proportion $\theta_{r,2}$ of the population that is susceptible, (ii) the basic reproduction number R_0 that summarizes the biological and socio-economic factors previously discussed, (iii) a shock ε_r that accounts for minor omitted variables. The economic damage D , in monetary terms (e.g., loss of GDP)⁸, associated with the pandemic is determined by the rate of

⁴By *interaction* we mean a contact amenable to virus transmission (in terms of closeness and duration).

⁵A highly non-trivial assumption that, for instance, requires people to use diligently protective gear if asked by local or national authorities.

⁶Throughout, shocks have zero mean and unit variance.

⁷For simplicity, we allow R_e to be negative (otherwise, we should add constraints that preserve the positivity of R_e , something that we prefer to abstract from).

⁸The scalar D lumps together economic consequences and health effects (through a monetary measure, e.g. GDP loss). Yet, in principle D may be a multidimensional vector with different components, say health and economic ones, that the decision criterion then trades off.

contagion R_e via a function $D = g(R_e, \theta_d, \varepsilon_d)$, where θ_d is a structural parameter and ε_d is a shock. This function represents the ability of health and economic systems to cope with the pandemic. We assume that $\partial g / \partial R_e > 0$. For example, assuming a quadratic damage function we have:

$$D = \theta_{d,1} R_e^2 + \theta_{d,2} R_e + \varepsilon_d, \quad (3)$$

with $2\theta_{d,1} R_e + \theta_{d,2} > 0$. The economic damage D associated with a policy a is then $D = g(R_e, \theta_d, \varepsilon_d) = g(f(a, R_0, \theta_r, \varepsilon_r), \theta_d, \varepsilon_d)$. A policy affects, according to relation f , the effective reproduction number and, through it, determines an economic damage according to relation g .

In the linear-quadratic example, we have

$$D = \kappa_1 a^2 + \kappa_2 a + \kappa_3, \quad (4)$$

where

$$\kappa_1 = \theta_{d,1} \theta_r^2, \quad (5)$$

$$\kappa_2 = \theta_{r,1} (2\theta_{d,1} (\theta_{r,2} R_0 + \varepsilon_r) + \theta_{d,2}) \quad (6)$$

$$\kappa_3 = \theta_{d,1} (\theta_{r,2} R_0 + \varepsilon_r)^2 + \theta_{d,2} (\theta_{r,2} R_0 + \varepsilon_r) + \varepsilon_d. \quad (7)$$

Policy a has an uncertain implementation cost C that, for instance, for school closures includes, as previously mentioned, absenteeism by working parents and loss of education. This cost is represented by a function $C = h(a, \theta_c, \varepsilon_c)$. We assume that costs grow more than proportionally, so that $\partial h / \partial a > 0$ and $\partial^2 h / \partial a^2 > 0$ (e.g., the cost of school closures grows more than proportionally with its duration). For example, a quadratic cost function is

$$C = \theta_{c,1} a^2 + \theta_{c,2} a + \varepsilon_c, \quad (8)$$

with $2\theta_{c,1} a + \theta_{c,2} > 0$ and $\theta_{c,1} > 0$. We also assume that the policymaker knows the functional forms of the relations f , g and h (e.g., whether they are linear or quadratic) but not their structural parameters. This lack of knowledge, along with that of the basic reproduction number R_0 and of the shocks' value ε , prevents the DM to know the actions' consequences. States thus have the form $s = (R_0, \varepsilon, \theta) \in S$ with both random and structural components. In particular, the vector $\varepsilon = (\varepsilon_r, \varepsilon_d, \varepsilon_c) \in E$ represents the shocks affecting the health and economic systems, while the vector $\theta = (\theta_r, \theta_d, \theta_c) \in \Theta$ specifies the structural coefficients parametrizing a model population. If we denote by $B = -D$ the benefit of policy a as its ability to reduce the economic damages due to the pandemic, its consequence is the difference $B - C$ between its benefits and costs. The consequence function is then $\rho(a, \varepsilon, \theta) = -g(f(a, R_0, \theta_r, \varepsilon_r), \theta_d, \varepsilon_d) - h(a, \theta_c, \varepsilon_c)$. In the linear-quadratic

example, it becomes

$$\rho(a, \varepsilon, \theta) = -(\kappa_1 + \theta_{c,1}) a^2 - (\kappa_2 + \theta_{c,2}) a - \kappa_3 + \varepsilon_c$$

2.3.2 Models and beliefs

Shocks have the form $\varepsilon_r = \sigma_r w_r$; $\varepsilon_d = \sigma_d w_d$; $\varepsilon_c = \sigma_c w_c$, where w_r , w_d and w_c are uncorrelated “white noises” with zero mean and unit variance. The vector parameter $\sigma = (\sigma_r, \sigma_d, \sigma_c) \in \Sigma$ then specifies the standard deviations of shocks. To ease the analysis, we assume that their distribution q_σ is known, up to their standard deviations σ . We also assume that the distribution p_ξ of the rate R_0 is indexed by a parameter $\xi \in \Xi$ that accounts for different epidemiological views on the quantification of the basic reproduction number. With this, the positive scalar $m(\varepsilon, \theta, R_0)$ gives the joint probability of shock ε , parameter θ and rate R_0 under a posited model $m \in M$. We adopt the model factorization $m = q_\sigma \times \delta_\theta \times p_\xi$, that is,⁹

$$m(\varepsilon, \theta, R_0) = \begin{cases} q_\sigma(\varepsilon) p_\xi(R_0) & \text{if } \theta = \theta \\ 0 & \text{else} \end{cases} \quad (9)$$

where $q_\sigma(\varepsilon)$ is the probability of shock ε under the standard deviation specification σ , while $p_\xi(R_0)$ is the probability that R_0 is the basic reproduction number according to epidemiological view ξ . We can thus index models as $m_{\theta, \sigma, \xi} = q_\sigma \times \delta_\theta \times p_\xi$ and denote by $M = \{m_{\theta, \sigma, \xi}\}$ the set of models that the policymaker posits. Because of the factorization, the policymaker’s subjective belief $\mu(m)$ that m is the correct model is actually over the values of θ , σ and ξ and so has the form $\mu(\theta, \sigma, \xi)$. A convenient separable form is, with an abuse of notation, $\mu(\theta, \xi) = \mu(\theta, \sigma) \mu(\xi)$.

For example, consider the pandemic decision problem $(A, S, C, \rho, \mathbf{t}, M)$ and the set of models $M = \{m_{\theta, \sigma, \xi}\}$ that the policymaker posits. Assume that a von Neumann-Morgenstern utility function $u : C \rightarrow \mathbb{R}$ translates economic consequences, measured in monetary terms, into utility levels. This function captures attitudes toward risk (i.e. uncertainty within models). The *expected reward* of action a under model $m \in M$ is

$$\begin{aligned} R(a, \theta, \sigma, \xi) &= \int_{\varepsilon, R_0} u(\rho(a, \theta, \varepsilon, R_0)) m_{\theta, \sigma, \xi}(\varepsilon, R_0) \\ &= \int_{\varepsilon, R_0} u(\rho(a, \theta, \varepsilon, R_0)) q_\sigma(\varepsilon) p_\xi(R_0). \end{aligned}$$

⁹Here δ_θ is the probability distribution concentrated on θ , i.e., $\delta_\theta(\theta) = 1$ and $\delta_\theta(\theta) = 0$ if $\theta \neq \theta$.

2.4 Numerical example

In the case study presented in Figure 2 of the manuscript, the policymaker must decide whether to and how long to close school for. Closing schools is costly (e.g. it increases child-adult transmission, reduces the ability of healthcare and key workers to work and the capacity of healthcare, generates economic costs through absenteeism by working parents, loss of education, etc.), but it helps slow the pandemic and its impact by reducing child-child transmission, thus delaying the pandemic peak that overwhelms health care services, and therefore ultimately reducing morbidity and associated mortality.

Here, we illustrate how different decision rules may be used in this specific example, in which there is only structural uncertainty about the benefits of school closures. The cost function is assumed to be known, so that there are only three different models in the set M .

- **Model 1** is based on the evidence coming from influenza outbreaks, for which the majority of transmission is between children (Mangtani et al., 2014). According to this model, closing schools would be the biggest contributor to reducing R_e to below 1 and it may be the only intervention that could do so. In this case the benefit would, for example, be proportional to the duration of school closure. In the linear-quadratic example, this would imply that $\kappa_1 = 0$ and $\kappa_2 < 0$, so that benefits positively depend on the action α : stronger measures reduce the effective reproductive number R_e , and thus the economic damage D of the pandemic.
- **Model 2**, instead, relies on some previous coronavirus outbreaks, for which evidence suggests minimal transmission between children (Wong et al., 2003). In this case, R_e cannot be reduced below 1, school closures do not affect the size of the epidemic, and therefore do not bring any benefits. This is for example the case if $\kappa_1 = \kappa_2 = 0$ in the linear-quadratic setup, so that the economic damages, and thus the benefits, are unaffected by the policy action α .
- **Model 3** projects that some child to child transmission happens so that closing schools contributes to reducing R_e to below 1 and reduces the size of the epidemic. However, this only works in combination with other measures (Prem et al., 2020). Without it, R_e would be above 1 but as an isolated measure school closures do not have such a big effect (Ferguson et al., 2020). Under this scenario, the effectiveness of school closures is important at the beginning, but declines as time goes on. In the linear-quadratic example, this would imply that $\kappa_1 > 0$ and $\kappa_2 < 0$.

For simplicity, we assume that there is no uncertainty within models ($\varepsilon_r = \varepsilon_d = \varepsilon_c = 0$).

The illustrative benefit and cost functions we used are the following:

- $B(a) = 4a + 20$ in the case of model 1,
- $B(a) = 100$ in the case of model 2,
- $B(a) = -0.1a^2 + 4a + 70$ in the case of model 3,
- $C(a) = 0.1a^2 + 10$.

Consider the decision problem $(A, S, C, \rho, \dots, M)$. In our case, we restrict the action space so that $A = [0, 20]$. For each of these 3 models $m_{\theta, \sigma, \xi}$, it is possible to compute the expected reward $R(a, \theta, \sigma, \xi)$ associated with a school closure policy. The policymaker, however, does not know which is the correct one. The expected reward is, in that sense, itself uncertain because it depends on the values of the different structural parameters used. For each particular model representing the net overall monetary benefits of the school closure policy, it is possible to determine the optimal action to put in place. Table 1 presents the expected rewards, together with their associated optimal actions \hat{a} in the case of linear utility u .

Table 1: Example of expected rewards and their associated optimal actions with linear utility u

	$R(a, \theta, \sigma, \xi)$	\hat{a}
Model 1	$-0.1a^2 + 4a + 10$	20
Model 2	$-0.1a^2 + 90$	0
Model 3	$-0.2a^2 + 4a + 60$	10

If the policymaker considers uncertainty within and across models in the same way, she aggregates the expected rewards by taking a weighted average over them, where the weights represent the degrees of belief in each specific model. The decision criterion in this case is the classical SEU criterion of Savage (1954). For example, under a uniform prior over the possible models, i.e. if $\mu(m) = 1/3$ for all m , the optimal decision is a school closure policy $\hat{a}_{seu} = 10$. It therefore means that, putting the same weight on the three different models given by three different sources, or a single model (such as model 3) on which all experts would agree, would lead exactly to the same optimal school closure policy. Instead, if the policymaker decides to behave extremely precautionary by taking into account only the model providing the lowest expected reward, she only considers Model 1 and decides to close schools for the maximum length $\hat{a}_{mxm} = 20$. This policymaker is extremely averse to uncertainty across models, and in consequence, uses the maxmin decision rule of Wald (1950). Alternatively, if the policymaker is averse to uncertainty in the sense of disliking more uncertainty across than within models but is not as precautionary as a maxmin policymaker, she may follow the smooth ambiguity criterion

of Klibanoff et al. (2005). In such a case, the optimal length of school closures is longer than under expected utility. It approximately corresponds to 12, when the ambiguity function φ is logarithmic. Finally, if the policymaker has multiple prior probability measures over the models, she can compute the expected utility for each of them, and considers only the one providing the lowest level of subjective expected utility. For example, imagine two distinct priors: the uniform prior, in which the 3 models are weighted equally, and the prior that considers model 2 as implausible, but models 1 and 3 as equally likely (i.e., this prior puts a weight 0 on model 2 and a weight 0.5 over the two other models). The optimal length of school closures under the multiple priors model of Gilboa and Schmeidler (1989) in this situation is higher than under subjective expected utility. It corresponds to closing schools for approximately 13 weeks. Table 2 summarizes the optimal decisions for each of these decision rules.

Table 2: Example of optimal policies depending on the decision rules followed

Decision rules (criterion)	Optimal policy
V_{seu}	$\hat{a}_{seu} = 10$
V_{mxm}	$\hat{a}_{mxm} = 20$
V_{smt}	$\hat{a}_{smt} = 11.65$
V_{mp}	$\hat{a}_{mp} = 13.33$

References

- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica* 19, 404–437.
- Aydogan, I., L. Berger, V. Bosetti, and N. Liu (2018). Three layers of uncertainty and the role of model misspecification: an experiment. IGIER Working Paper 623, Bocconi University.
- Bayham, J. and E. P. Fenichel (2020). Impact of school closures for covid-19 on the us health-care workforce and net mortality: a modelling study. *The Lancet Public Health*.
- Brooks, S. K., L. E. Smith, R. K. Webster, D. Weston, L. Woodland, I. Hall, and G. J. Rubin (2020). The impact of unplanned school closure on children’s social contact: rapid evidence review. *Eurosurveillance* 25 (13), 2000188.
- Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2013). Classical subjective expected utility. *Proceedings of the National Academy of Sciences* 110(17), 6754–6759.
- Delamater, P. L., E. J. Street, T. F. Leslie, Y. T. Yang, and K. H. Jacobsen (2019). Complexity of the basic reproduction number (r_0). *Emerging infectious diseases* 25 (1), 1.
- Dietz, K. (1993). The estimation of the basic reproduction number for infectious diseases. *Statistical methods in medical research* 2 (1), 23–41.
- Ferguson, N., D. Laydon, G. Nedjati Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunuba Perez, G. Cuomo-Dannenburg, et al. (2020). Report 9: Impact of non-pharmaceutical interventions (npis) to reduce covid19 mortality and healthcare demand.
- Ghirardato, P., F. Maccheroni, and M. Marinacci (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118 (2), 133–173.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18 (2), 141–154.
- Hansen, L. P. (2014). Nobel lecture: Uncertainty outside and inside economic models. *Journal of Political Economy* 122 (5), 945–987.
- Hansen, L. P. and M. Marinacci (2016). Ambiguity aversion and model misspecification: An economic perspective. *Statistical Science* 31, 511–515.
- Hurwicz, L. (1951). Some specification problems and applications to econometric models. *Econometrica* 19 (3), 343–344.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica* 73, 1849–1892.
- Koopmans, T. C. (1947). Measurement without theory. *The Review of Economics and Statistics* 29 (3), 161–172.
- Lempel, H., J. M. Epstein, and R. A. Hammond (2009). Economic cost and health care workforce effects of school closures in the us. *PLoS currents* 1.
- Mangtani, P. et al. (2014). Impact of school closures on an influenza pandemic: Scientific evidence base review.
- Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association* 13 (6), 1022–1100.

- Marschak, J. (1953). Economic measurements for policy and prediction. In *Studies in Econometric Method* (W. Hood and T. J. Koopmans, eds.), pp. 1–26. Wiley, New York.
- Marschak, J. and R. Radner (1972). *Economic theory of teams*. Yale University Press, New Haven.
- Ng, T.-C. and T.-H. Wen (2019). Spatially adjusted time-varying reproductive numbers: Understanding the geographical expansion of urban dengue outbreaks. *Scientific Reports* 9 (1), 1–12.
- Prem, K., Y. Liu, T. W. Russell, A. J. Kucharski, R. M. Eggo, N. Davies, S. Flasche, S. Clifford, C. A. Pearson, J. D. Munday, et al. (2020). The effect of control strategies to reduce social mixing on outcomes of the covid-19 epidemic in wuhan, china: a modelling study. *The Lancet Public Health*.
- Sadique, M. Z., E. J. Adams, and W. J. Edmunds (2008). Estimating the costs of school closure for mitigating an influenza pandemic. *BMC public health* 8 (1), 135.
- Savage, L. (1954). *The Foundations of Statistics*. New York: J. Wiley. second revised edition, 1972.
- Smith, R. D., M. R. Keogh-Brown, and T. Barnett (2011). Estimating the economic impact of pandemic influenza: An application of the computable general equilibrium model to the uk. *Social Science & Medicine* 73 (2), 235–244.
- United Nations Educational Scientific and Cultural Organization (2020). Covid-19 educational disruption and response. [Online; accessed April 24, 2020].
- Viner, R. M., S. J. Russell, H. Croker, J. Packer, J. Ward, C. Stansfield, O. Mytton, C. Bonell, and R. Booy (2020). School closure and management practices during coronavirus outbreaks including covid-19: a rapid systematic review. *The Lancet Child & Adolescent Health*.
- von Neumann, J. and O. Morgenstern (1947). *Theory of Games and Economic Behavior*. 2nd ed. Princeton University Press.
- Wald, A. (1950). *Statistical decision functions*. New York: John Wiley & Sons.
- Wallinga, J. and M. Lipsitch (2007). How generation intervals shape the relationship between growth rates and reproductive numbers. *Proceedings of the Royal Society B: Biological Sciences* 274 (1609), 599–604.
- Wong, G. W., A. M. Li, P. Ng, and T. F. Fok (2003). Severe acute respiratory syndrome in children. *Pediatric pulmonology* 36 (4), 261–266.