



Exponential-growth prediction bias and compliance with safety measures related to COVID-19

Ritwik Banerjee^{a,*}, Joydeep Bhattacharya^b, Priyama Majumdar^c

^a Indian Institute of Management Bangalore, India

^b Iowa State University, USA

^c University of Warwick, UK

ARTICLE INFO

Keywords:

COVID
Exponential growth bias
WHO safety Measures
Health communication
Graphical communication
Nudges

ABSTRACT

Objective: We define prediction bias as the systematic error arising from an incorrect prediction of the number of positive COVID cases x -weeks hence when presented with y -weeks of prior, actual data on the same. Our objective is to investigate the importance of an exponential-growth prediction bias (EGPB) in understanding why the COVID-19 outbreak has exploded. To that end, our goal is to document EGPB in the comprehension of disease data, study how it evolves as the epidemic progresses, and connect it with compliance of personal safety guidelines such as the use of face coverings and social distancing. We also investigate whether a behavioral nudge, cost less to implement, can significantly reduce EGPB.

Rationale: The scientific basis for our inquiry is the received wisdom that infectious disease spread, especially in the initial stages, follows an exponential function meaning few positive cases can explode into a widespread pandemic if the disease is sufficiently transmittable. If people suffer from EGPB, they will likely make incorrect judgments about their infection risk, which in turn, may lead to reduced compliance of safety protocols.

Method: To collect data on prediction bias, we ran an incentivized, experiment on a global, online platform with participation from people in forty-three countries, each at different stages of progression of COVID-19. We also constructed several indices of compliance by surveying participants about their frequency of hand-washing and use of sanitizers and masks; their willingness to pay for masks; their view about the social appropriateness of others' behavior; and their like/dislike of government responses. The prediction data was used to construct several measures of EGPB. Our experimental design permits us to identify the root of under-prediction as EGPB arising from the general tendency to underestimate the speed at which exponential processes unfold.

Results: Respondents make predictions about the path of the disease using a model that is substantially less convex than the actual data generating process. This creates significant EGPB, which, in turn, is significantly and negatively associated with non-compliance with safety measures. The bias is significantly higher for respondents from countries at a later stage relative to those at an early stage of disease progression. A simple behavioral nudge that shows prior data in terms of raw numbers, as opposed to a graph, causally reduces EGPB.

Conclusion: Behavioral biases concerning the comprehension of disease data are quantitatively important, and act as severe impediments to effective policy action against the spread of COVID-19. Clear communication of future infection risk via raw numbers could increase the accuracy of risk perception, in turn, facilitating compliance with suggested protective behaviors.

1. 1 Introduction

The COVID-19 outbreak is a global pandemic, adversely affecting the lives of millions of people around the world. Not a lot is yet known about the virus and how it operates. With considerable uncertainty about the arrival of a vaccine, there is complete agreement on the need for people

to strictly follow WHO guidelines regarding frequent washing of hands, use of hand sanitizers and face masks, social distancing, and if needed, self-quarantine. Of course, unsurprisingly, not everyone complies with these guidelines, at least not with the seriousness with which they need to be followed (Cummins, 2020; Lunn et al., 2020; Pinsker, 2020; van Bavel et al., 2020). This sort of (non)compliance is an active decision

* Corresponding author.

E-mail address: ritwikbanerjee@iimb.ac.in (R. Banerjee).

<https://doi.org/10.1016/j.socscimed.2020.113473>

Received in revised form 16 September 2020; Accepted 21 October 2020

Available online 28 October 2020

0277-9536/© 2020 Elsevier Ltd. All rights reserved.

predicated upon, at least, one influential variable, the accuracy with which an individual perceives her likelihood of getting infected (loosely, her risk perception). All else the same, if her risk perception is high, she is more likely to show compliance. This point was noted in a classic article about the 1918 influenza pandemic in *Science* “People do not appreciate the risks they run” (Soper, 1919).

A large body of work in behavioral psychology and economics has documented how the accuracy of risk perception may be compromised by a whole host of behavioral biases (such as optimism, overconfidence, and so on). This study focuses attention on one such bias: exponential-growth prediction bias (EGPB), the “pervasive tendency to linearize exponential functions when assessing them intuitively,” which leads to “a systematic tendency to underestimate a future value given a present value” (Goda et al., 2019; Stango and Zinman, 2009). Existing analysis of the exponential growth bias relies on quantifying the difficulty most people have with compound interest rates (Levy and Tasoff, 2016). General difficulty with discriminating linear from non-linear processes is documented in Cordes et al. (2019) and shows up even as early as in pre-kindergarten students (Ebersbach et al., 2010). Wagenaar and Timmers (1979) make an early and influential attempt at understanding the relative importance of numerical versus non-numerical (e.g., graphical) means of data delivery and its effect on the perception of exponentiality. Why is a study of this bias critical to our understanding of the current pandemic? Prior epidemiological studies (see Keeling and Rohani (2011); Thomas (1996), and more recently, Li et al. (2020) for Wuhan, China) document how the spread of infectious diseases, especially in the initial stages, often follows an exponential function. This information means a few positive cases, initially can explode into a pandemic if the disease is sufficiently contagious. To illustrate this point in the current context, focus attention on Figure A1, which plots the time trajectory of reported cases for four countries, Germany, the United States, France, and Spain. In each case, it is apparent that the growth trajectory is exponential. But do humans see it that way? Pinsker (2020) argues, no: “[t]he human brain can have trouble keeping pace with such rapid growth” and that “people tend to underestimate the speed at which exponential processes—such as a disease outbreak—unfold.” This, in other words, is EGPB, a behavioral failure to “read the tea leaves” correctly, which may lead to inaccurate infection risk perception. Ours is a first attempt at documenting EGPB in the comprehension of disease data, to study how it evolves as the epidemic progresses, and to connect it with compliance with personal safety guidelines.

We define prediction bias as the systematic error arising from under or over -prediction of the number of COVID-19 positive detections x -weeks hence when presented with y -weeks of prior, actual data on the same. We call it EGPB if the actual data follows an exponential function, and the predictions fail to appreciate the extent of the true convexity. Our analysis works off the premise that those who suffer from EGPB will significantly underestimate how quickly a disease spreads, fail to perceive the onrushing infection risk, and hence, show low compliance with safety measures.

We use data from an online experiment to investigate three pressing questions of significant policy relevance:

1. How much of individual-level compliance with WHO guidelines can be explained by the bias associated with predicting the number of COVID-19 cases, after controlling for demographic and cultural variables?
2. We identify three distinct stages a country can be in: Stage 1 with less than 100 positive detections, Stage 2 with between 100 and 999, and finally, Stage 3 with 1000 or more as of March 21, 2020. Does EGPB diminish as a country moves through different stages of the disease?
3. Does a simple nudge in terms of how disease data is presented help mitigate EGPB?

We ran an incentivized, data-collection survey on Amazon’s Mechanical Turk, an online platform, with participation from people in 43

countries. The survey is not nationally representative, but since the samples were collected on the same platform, they permit a relatively clean comparison. More importantly, MTurk facilitates access to a global pool of participants who reside in countries at different stages of the disease. This approach provides a unique opportunity to study how EGPB may vary with the stages of the disease progression. Examples of Stage 1 countries in our sample include Netherlands, South Africa, and Bangladesh; Stage 2 includes India, Mexico, and Turkey; and Stage 3 includes U.S., Italy, and Germany. The list of countries and the corresponding stages are reported in Table A1.

Besides collecting data on prediction bias, we also asked participants about their frequency of handwashing and the use of sanitizers and masks, their willingness to pay for masks, their view about the social appropriateness of others’ behavior; and their like/dislike of government responses. Other demographic information was collected as well. Using this information, we generated composite indices measuring individual attitudes regarding their i) own compliance, ii) appropriateness of violation of WHO measures, and iii) satisfaction with the government’s performance. Taken together, these indices give us a broad sense of “compliance”. The prediction data was used to construct several measures of EGPB and featured alongside the compliance measure as regressors in a multivariate regression model.

Our main results are as follows. First, we document the presence of EGPB as it pertains to forecasting the x weeks-ahead path of the disease. Second, the “degree of convexity” reflected in the predicted path of the disease is significantly and substantially lower than the actual path. (We use the term “degree of convexity” to mean the rate of change of the gradient of the data function.) This finding connects with the first result: the source of the prediction bias is the “lower convexity” of the mental model used. Pennycook et al. (2020) found that overall cognitive sophistication (the composite of four measures) was a strong negative predictor of COVID-19 misperceptions. To the extent discriminating between linear and exponential processes is correlated with cognitive sophistication, our results are in sync with their findings. Strikingly though, their measure of cognitive sophistication was not a reliable or consistent predictor of COVID-19 risk perceptions or behavior change intentions. Similarly, Stanley et al. (2020) found that “individuals less willing to engage in effortful, deliberative, and reflective cognitive processes were [...] less likely to have recently engaged in social-distancing and handwashing.” Third, EGPB is significantly lower for participants from countries at an early stage relative to those at a later stage of the disease. Fourth, we find our measures of EGPB are significant predictors of compliance: higher bias is associated with lower own safety compliance, higher approval of a violation of safety measures, and greater satisfaction with the government’s policy response to the pandemic.

Fifth, we find that providing disease trajectory-information using y -weeks prior data in the form of raw numbers *causally* reduces EGPB more than delivering the same via a graph. Like us, Wagenaar and Timmers (1979), and more recently, Levy and Tasoff (2017), find “exponential-growth bias is unlikely to be eliminated by simple “nudges” such as a graphical intervention.”

How do our results connect to the literature? Our finding that EGPB is a significant predictor of compliance is an important, practical contribution to discussions of health policy (Bischoff et al., 2000; Callaghan T et al., 2019; Lyons et al., 2020). That EGPB in our sample is significantly lower for participants from countries at an early (as opposed to a later stage) stage of the disease suggests the bias is stubborn; it does not go away once the raw numbers become bigger. If the bias persists, and people comply less and less, the disease spreads faster, and a vicious cycle is created. In this sense, our work showcases a personal, psychological dimension of epidemiology, one that can link itself on to the forces of social epidemiology (Donovan and Blake, 1992; Kawachi and Subramanian, 2018) and precipitate terrible outcomes. Finally, our demonstration of the fact that data presented with raw numbers (as opposed to graphs) causally reduces EGPB contributes to

the field of health communication.

There is a small yet burgeoning literature studying the role of biases in the context of the COVID-19 outbreak. Our study was designed concurrently and independently of this literature and is closest in spirit to an important contribution by Wise et al. (2020) who, like us, use an online sample (but only of U.S. households) between March 11, the day when the WHO declared COVID-19 a pandemic, and March 16. Their primary focus is on deciphering to what extent individuals are aware of their risk of contracting the disease, their chance of passing it on, and the extent to which their perception of risk predicts protective behaviors. We do not directly elicit peoples' risk perception as it may be contaminated by the simultaneous presence of multiple, cognitive biases (such as optimism or overconfidence; see Wise et al. (2020) and Stanley et al. (2020)). Our measure of EGPB is, arguably, better. First, we use actual data from a real, unnamed country, one different from the one where the subject resides to avoid confounds such as local knowledge, perceived efficiency of administration, etc. in the estimate of EGPB. Our subjects, instead, make predictions based on data from a country *other* than that of their residence. Also, our measures of EGPB are composite in nature. While composite, summary measures are useful to policymakers, the benefit of the granular measures in Wise et al. (2020), capturing own risk of getting infected versus the risk of infecting others, is also high. Second, our measure of perception bias, EGPB, is rooted in the linear vs. exponential heuristic literature with a clear theoretical underpinning. And third, our measures of prediction bias are less likely to be "pure noise" since they are the result of an incentivized elicitation process.

A related paper by Fetzer et al. (2020) is devoted to unearthing prediction bias in the context of a fictitious disease over several days under several formulaic scenarios. We depart from their work in several ways. First, our design is not aimed at understanding whether participants can "do the math" and possess enough cognitive ability to figure out where a hypothetical series is headed. Instead, it seeks to detect bias in realistic environments with actual data on the disease growth, where general cognitive sophistication may play a role (Pennycook et al., 2020; Stanley et al., 2020). Second, we go beyond detecting exponential growth prediction bias and identify how the bias can affect an important health outcome, namely, compliance behavior. Third, our results suggest participants do not seem to have a linear model in mind, but a model whose curvature is less than that of the growth path of the actual data.

A final note on our contribution to the literature is in order. Extant macro-finance research on EGPB focuses exclusively on the inability of laypeople to comprehend the power of compound interest rates and its implication on lower savings, lower net worth, and so on. Our contribution, focused entirely in the health/epidemiology domain, is to show that the inability to foresee the future path of the disease correctly can have negative implications for compliance and that, in and of itself, may raise the future growth rate of the disease.

2. Method

We show our participants data on the actual number of COVID-19 (same as in Figure A1 from four countries majorly affected by the virus as of March 21, 2020, namely, Germany, the U.S.A., France, and Spain. In our experiment, participants perform four prediction tasks using data from two randomly chosen countries out of these four; country names are not revealed. More precisely, if a participant belongs to any of the four countries, say X, she is not be shown numbers from X; instead, she is shown numbers from two countries randomly chosen from the set W, Y, and Z. This approach ensures any prior information she has about disease progression in her country, X, does not contaminate her prediction. This strategy also prevents other confounds (such as perceived efficiency of the government and quality of general health care infrastructure) in one's country to enter the prediction calculus. In each task, they are shown three, actual weekly data points of COVID-19 cases. Subsequently, they are asked to predict the number of cases for

Weeks 4 and 5, the actual numbers of which are known to us, the researchers. This means, we, the researchers, have full knowledge of the true, underlying data-generating process and participants know that we know. Participants are paid \$0.55 as a participation fee, and the prediction task is accuracy-rewarding: two of the four prediction tasks are randomly picked, and if the participant's prediction is within 5% of the actual number, she is paid an additional \$0.25 for that task. Participants can earn a maximum of \$1 for the entire experiment, which lasted for about 7 min. The research was conducted through oTree, a web-based experimental platform (Chen et al., 2016).

Formally, denote the actual and the predicted number of COVID-19 cases in Week i by N_i and P_i , respectively, for $i = 1, 2, 3, 4, 5$. As noted earlier, the participants observe three data points on the number of COVID-19 cases in three consecutive weeks (N_1, N_2, N_3) and are asked to make their predictions for Week 4 and Week 5. Since the actual number of infected individuals at any point is unknown, we go by the official statistics on reported cases. Interestingly, respondents in our sample report that their belief about the true infection rate is, on average, 10% higher than the official statistics. To make sure the biases are comparable, we represent them relative to the maximum possible error a participant can make. For example, the bias for Week 4 is defined as the difference between the actual number (N_4) and the predicted number in Week 4 (P_4), relative to the difference between the actual number in Week 4 (N_4) and the actual number in Week 3 (N_3). In other words, the actual prediction error relative to the maximum possible error in Week 4 may be interpreted as the Bias for Week 4 with respect to Week 3 i.e., $Bias_{43} \equiv \frac{N_4 - P_4}{N_4 - N_3}$. Similarly, we define Bias for Week 5 with respect to Week 3 as $Bias_{53} \equiv \frac{N_5 - P_5}{N_5 - N_3}$. Finally, we analyze the results in terms of $Bias_{avg}$, which is the average of $Bias_{43}$ and $Bias_{53}$. In addition to the afore discussed measures of bias arising from the incentivized prediction task, we also obtain a measure of bias relating to one's ability to predict the actual number of cases in one's country a week hence – OwnBias. For a summary of the different definitions, refer to Table 1.

We survey three parts, restricting each piece to only those participants who are registered in countries belonging to a particular stage. MTurk facilitates participation restriction based on specific geographical criteria, the I.P. address, and the initial registration information. Since our inferences are at a stage-level, we have no country-specific restrictions within each stage. Consequently, while the number of participants in each stage is balanced, the number of participants from countries within a stage is not.

Notice, from Figure A1, the early phases of the spread of COVID-19, is described by notably less convexity than the later phases. To understand whether prediction accuracy varies with phases of the disease, we implement a within-subject variation in the four prediction tasks. In two consecutive tasks, the participants are shown data either from the early or late phases from two randomly selected countries. Asking them to respond to both early and later phases ensures that their home country experience is mimicked in at least one of the tasks. If, in the first two tasks, a participant is shown data from two countries at an early phase, in the next two tasks, they are shown numbers from the later phases of the same two countries.

We implement this design so that we can explore whether the nature of representation (graphical or numeric) of the actual data helps mitigate EGPB. To that end, some participants are randomly shown the exact number of COVID-19 cases graphically (as is the dominant form of representation of the data in print or online media). In contrast, others are shown the same in terms of raw numbers. We implement this method through a between-subject design, which allows us to estimate the causal effect of the nudge in mitigating the EGPB. Screen 5 of the experimental instruction given in online Appendix 2 presents an example of the two forms of data representation.

Then, we administered a short survey to capture participants' protection behavior and compliance with the WHO guidelines. The survey details, along with the experimental instructions, are included in

Table 1
Summary statistics.

Variable	Definition	Median	Mean Absolute Deviation from Median
Bias43	Difference between the log of actual and predicted number in Week 4, relative to the change in log of actual number of COVID-19 cases between Week 3 and 4	0.42	0.35
Bias53	$Bias_{43} = \frac{N_4 - P_4}{N_4 - N_3}$ Difference between the log of actual and predicted number in Week 5, relative to the change in log of actual number of COVID-19 cases between Week 4 and Week 5	0.46	0.29
AverageBias OwnBias	$Bias_{53} = \frac{N_5 - P_5}{N_5 - P_3}$ Average of Bias43 and Bias53 Difference between the log of actual and predicted number of COVID-19 cases one week later in one's own country	0.43	0.31
Own Country Information Bias	Difference between the log of actual and perceived number of COVID-19 cases on the day of response in one's own country	0.001	0.1
Variable	Definition	Median	Standard Deviation
Actual Realized Compliance	First Principal Component of Questions 1a-1g in Screen 13	0.12	1.68
Appropriateness of violation of safety norms	First Principal Component of Questions 2a-2d in Screen 14	-0.14	1.53
Agreeableness with government performance	First Principal Component of Questions 3a-3b in Screen 14	-0.05	1.35
Female	= 1, if gender is Female	0.27	0.44
Age	age in years	34.19	9.15
Education	= 0, if highest educational level is up to class 12 = 1, if highest education level is bachelor's degree = 2, if highest education level is master's degree or above	1.12	0.65
Income	Log of monthly family income (PPP USD)	8.39	1.79
Health	Health condition on a scale of 0-5 [0 if very poor health, 5 if very good health]	4.12	0.75
Health Insurance	= 1, if the person has health insurance	0.71	0.45
Perceived effectiveness	Perceived effectiveness of the safety measures being proposed to counter the spread of COVID-19	4.3	0.92
N – Stage 1	Sample Size	334	
	Sample size from countries with COVID-19 cases less than 100, as on 21st March	121	
N – Stage 2	Sample size from countries with COVID-19 cases between 100 and 999, both numbers included, as on March 21	108	
N – Stage 3	Sample size from countries with COVID-19 cases more than 999, as on 21st March	105	
Countries	Number of countries represented in our sample	43	

Appendix 2. Participants are asked to report their frequency of hand-washing, use of sanitizers and masks, their opinion about the appropriateness of public gatherings, their perception about the desirability and efficacy of their own government's responses, and so on. Necessary demographic information is also collected. We categorize the questions into three broad headings, namely, Actual Realized Compliance, Appropriateness of Violation of Safety Norms, and Agreeableness with Government Performance; taken together, they offer a complete sense of participants' compliance with safety guidelines. The definitions are given in Table 1. Besides, we randomize the order of whether a participant is shown the prediction part first or the survey part first to mitigate experimenter demand effects. The overview of the experimental design appears in Table A6.

In sum, the experimental design allows us to answer the three main research questions in this study: 1) Do people exhibit EGPB in the context of the COVID-19 spread? 2) Is this bias a significant predictor of lax attitude towards COVID-19 related safety norms? 3) Can a simple nudge in the way the data is presented help mitigate the bias?

3. Results

Our first set of results documents the existence of the afore discussed prediction biases. Fig. 1A presents the number of COVID-19 cases in early and later phases, for all the five weeks, along with the mean, median, and interquartile range of the predicted number of cases in Weeks 4 and 5, averaged over all the four countries. We present the country-specific predictions on the log-transformed data for each of the four countries separately in Fig. 1B. The mean prediction in the early phase panel of Fig. 1 exceeds the actual number because of outliers. Except for the one case, the medians and the means of the predictions lie well below the actual numbers of COVID-19 cases. To avoid outlier-driven distortions, we rely on the median measure in all subsequent analyses. The actual numbers, averaged across all four countries, are 1540 in Week 4 and 9189 in Week 5 for the early phase; the respective medians are 521 and 1081. The corresponding numbers in Week 4 and Week 5 are 17450 and 55934, while the median predictions are 9500 and 18000, respectively.

As discussed earlier and shown in Figure A1, the underlying data generating process for the actual spread of COVID-19 is convex. An interesting question is, is the prediction model used by the median individual also convex, or is it linear? To test this, we compute the ratio of the slopes of the line segment connecting N_4 and N_5 and that connecting N_3 and N_4 (i.e., $\eta \equiv \frac{\text{slope}(N_4, N_5)}{\text{slope}(N_3, N_4)}$). We compare η with the ratio of slopes of the line segment connecting P_4 and P_5 and that connecting N_3 and P_4 (i.e., $\rho \equiv \frac{\text{slope}(P_4, P_5)}{\text{slope}(N_3, P_4)}$). Column (3) in Table A2 reports $\eta - \rho$. We test $H_0: \eta - \rho = 0$ and find it is significantly different from zero for all the countries, for both phases (except in one case). This finding indicates that the underlying prediction model used by the median individual is significantly “less convex” than the underlying data generating process. We further test if the participants' prediction model is linear by comparing slopes of the individual linear pieces connecting N_3, P_4, P_5 with the slope of the best linear fit of N_1, N_2, N_3 ($\hat{\beta}$). The null is that the piece-wise slopes, N_3, P_4 and P_4, P_5 are indeed equal to the slope of the linear fit. The non-parametric equality-of-median-test rejects the null. Further, we statistically compare the Euclidian distance between the median prediction and the linear fit, and the median prediction and the actual data. We find the prediction model to be significantly closer to the linear fit. These results are not reported but are available upon request.

How does the “degree of convexity” of the predictions vary with stages of the disease a participant witnesses? We find that participants from Stage 3 countries (relative to those from Stages 1 and 2), for all our bias measures, make predictions that are closer to the best-fit linear model, while predictions of those from Stage 1 countries are closer to the




Fig. 1. Actual and predicted number of COVID-19 cases pooled across four countries. *Note.* Panel A above plots the actual number and the predictions of COVID-19 cases in the linear scale, while Panel B plots the same data on a logarithmic scale. Both the panels use data pooled across the four countries. The predictions for early (later) phase cases are given on the left (right). Each graph presents participants' 25th percentile, mean, median, and 75th percentile prediction of the number of COVID-19 cases on Week 4 and Week 5.

actual, exponential data. This has the *prima facie* implication that people in advanced stages of the disease outbreak may perceive the growth path as less, not more, “convex”.

We carry out the rest of the analysis on log-transformed data with median as the primary statistic. In Fig. 2, we analyze whether differences between predictions and actual numbers, when transformed as biases, are significantly different from zero. Median $Bias_{43}$ is positive and significantly different from zero at 0.42 (Wilcoxon signed-rank test, $p < 0.01$), meaning the median participant exhibits 42% under-prediction. Similarly, $Bias_{53}$ and the *AverageBias* are positive and significant at 0.46 and 0.44 (Wilcoxon signed-rank tests, $ps < 0.01$). The participants exhibit substantial prediction bias. Interestingly, $Bias_{43}$ is significantly smaller than $Bias_{53}$ (non-parametric median test, $p < 0.01$),

suggesting that the size of the bias increases with time. This is a consequence of the fact that people use a prediction model that has a smaller “degree of convexity” than the actual growth path of the disease. How does EGPB vary with gender and education levels? While we do not find any evidence of gender differences in EGPB, we find, on average, those with education levels bachelors or above are significantly more biased than those with lower education levels (t -test, $p = 0.03$). The main focus of the study is on EGPB arising from the incentivized prediction, but we also elicit participants' beliefs about the number of COVID-19 cases on the day of the experiment in their own country and their prediction about the same seven days hence. As Fig. 2 indicates, participants' (non-incentivized) prediction about the number of cases in their own country, seven days hence, reveals a 49% under-prediction.




Fig. 2. Prediction bias. *Note.* The figure reports the median values of different measures of biases. The error bars represent the 95% confidence intervals from Kendall's τ test for the hypothesis that the median is zero.

Our Amazon-mTurk participant pool comprises people from 43 countries from different stages of COVID-19 spread at the time of data collection, which gives us a unique opportunity to examine how EGPB in our sample compare across people from countries at different stages of the disease. Fig. 3 compares the biases across Stage 1, Stage 2, and Stage 3, and separately plots the regression coefficients of the stage dummies for each definition of bias. While the incentivized-bias measures are not

different between Stage 2 and Stage 1, the OwnBias is significantly higher in Stage 2 than in Stage 1 ($p = 0.02$). Fig. 3 shows that $Bias_{43}$, $Bias_{53}$, $AverageBias$ and $OwnBias$ are significantly higher in Stage 3 than in Stage 1 ($p < 0.05$ for all the specifications). The pattern is less clear when Stage 3 is compared with Stage 2, but $AverageBias$ is significantly higher in Stage 3 compared to Stage 2 in the most stringent specification. These results suggest that biases, as per our different measures, are considerably higher for countries at Stage 3 than those in Stage 1 or 2. Relatedly, participants from Stage 3 countries use a prediction model that is closer to the best-fit linear model than the actual, exponential data relative to those from Stages 1 and 2. (The results, not reported, are available upon request.) Table A3 reports estimates and standard errors. The regression results in Model (2), (4), and (6) control for age, gender, health condition, health insurance, education level, income, and log number of reported COVID-19 cases as of March 21.

Having established the presence of a significant EGPB, we ask if it is a significant predictor of near-contemporaneous (one-week past) compliance with safety measures. We create compliance indices based on responses to a menu of questions asked of the participants. Appendix 2 presents the entire experimental protocol, along with the survey questions used. We categorize the questions into three indices: Actual Realized Compliance, Appropriateness of Violation of Safety Norms, and Agreeableness with Government Performance. The indices are constructed by taking the first principal component of the relevant set of variables. Table 1 lays out the summary statistics of the compliance indices. Next, we regress these indices on our EGPB measures. Each column of Fig. 4 corresponds to each of the compliance indices and plots the estimated regression coefficient for each definition of prediction bias. Models 1 and 2 run the regression without and with control




Fig. 3. Variation of prediction bias between countries at different stages of COVID-19 spread. *Note.* This figure plots differences in median EGPB across countries in Stages 1, 2, and 3 of COVID-19 spread for each of the four measures of bias. Model 1 (Model 2) shows the pairwise differences in EGPB between the three stages estimated from a median regression without (with) controls. The control variables include age, gender, health, health insurance, education level, income, treatment and log of reported COVID-19 cases as on March 21 (fixed for each country). The specification in Model 2 for OwnBias in (iv), additionally controls for an individual's information bias. The error bars show 95% confidence interval. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Fig. 4. EGPB and COVID-19 compliance. *Note.* This figure plots the coefficients estimated from OLS regressions between EGPB and compliance, measured via three indices: Actual realized compliance, Appropriateness of violation of safety norms, and Agreeableness with government performance. Model 1 (Model 2) shows the estimates from specification without (with) controls. The control variables include age, gender, health, health insurance, education level, income, perceived effectiveness of the safety measures, treatment, and log of reported COVID-19 cases as on March 21 (fixed for each country). The specification in Model 2 for OwnBias in (iv), additionally controls for an individual's information bias. # $p = 0.11$, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

variables, respectively. As the figure illustrates and Table A4 confirms, EGPB is a negative predictor of Actual Realized Compliance, indicating that the higher the bias, the lower is the self-reported measure of compliance with safety norms. Note, Actual Realized Compliance is an index constructed from the following seven components: frequent washing of hands, use of sanitizers, staying at home, avoiding social gathering, maintaining a meter distance, minimizing contact, and wearing masks. To delve deeper, we further regress EGPB on each of these components and plot the coefficients in Figure A3 in the Appendix. The results show that EGPB is negatively correlated with Staying at home, Minimizing contact, and Avoid social gathering; however, surprisingly, it is positively correlated with Wearing masks. The latter may be a result of the combination of our finding that staying at home increases with EGPB and the near-universal policy mandate on mask usage. In other words, those who are disinclined to stay at home may be more likely to require a mask to access public transport or public places.

A higher EGPB also predicts a higher Appropriateness of Violation of Safety Norms as the plotted coefficients in the second column of Fig. 4 reveals. This means an individual who shows EGPB is also likely to view violations of safety norms, such as the avoidance of public gatherings, working from home, and so on, as not very alarming. Finally, EGPB is also a significant predictor of Agreeableness with Government Performance, implying, higher bias is associated with greater satisfaction with the performance of the government concerning the measures taken vis-a-vis COVID-19. The regression results appear in Table A4, where col (2), (4), and (6) controls for age, gender, health, health insurance, education level, income, treatment, and log of reported COVID-19 cases as of March 21. In addition to these variables, the results for OwnBias also controls for own-country information bias, defined as the difference

between the log of the actual and perceived number of COVID-19 cases on the day of response in one's own country.

Our analysis, thus far, suggests that people make predictions about the disease using a model that is substantially less convex than the true data generating process. This creates significant prediction biases, which in turn, are significantly associated with non-compliance with safety measures. Given this link, we ask: could a simple, perceptual nudge help reduce the prediction bias? To that end, we use a randomized experimental design in which some randomly chosen participants are shown N_1 , N_2 , N_3 in terms of raw numbers, while the rest are shown the same in graphical form.

The median biases across graphical and numerical treatments are presented in Fig. 5. Presentation of the past data in a numerical form significantly decreases the bias, however defined, relative to the graphical representation. The treatment effects from ordinary least square regressions are reported in Table A5. This result holds with and without a set of controls. The treatment effect is negative and significant at the 1% level when the regression controls for age, gender, health, health insurance, education level, and log number of COVID-19 cases as of March 21 (fixed for each country). The regressions, when separately run for the early and later phases, reveal the same pattern. We do not find any heterogeneity in the additional robustness checks we run to see if the treatment effect varies with age, education levels, and phases. These are not reported in the paper but are available on request.

4. Discussion

The critical question, of course, is why do we see early-stage participants show more EGPB than late-stage ones (see Fig. 3)? We




Fig. 5. Prediction bias across treatments. *Note.* Participants are randomized into graphical and numerical treatments, where they are shown the data from Weeks 1, 2, and 3 in the form of raw numbers and graphs, respectively. The figure plots the median prediction biases for the graphical and numerical treatments. The error bars represent the 95% confidence intervals from a Kendall's Tau test for the hypothesis that the median is zero. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

speculate there are two opposing, informational forces at work. One, people in late-stage countries may know more about the underlying data generating process, and that may reduce their EGPB. Two, people in early-stage countries may be overly (possibly, irrationally) scared about the disease (presumably due to limited and/or incorrect information), and that makes the disease hyper salient in their minds. One possible explanation on the flip side is behavioral or caution fatigue. In our setting, this possibility could afflict people who diligently followed safety protocols early on, got tired or stressed after a while, and subsequently resumed their pre-COVID course of life (Brooks et al., 2020). This phenomenon can raise EGPB in late-stage countries. Which of these two forces dominate is ultimately an empirical question. Surprisingly, it turns out that, in our sample, the second force dominates, meaning people in later-stage countries demonstrate larger EGPB relative to those at early stages.

Do we see evidence of such behavioral fatigue? While we do not have a direct measure of behavioral fatigue, we have a proxy relating to awareness: information bias, the difference between the log of actual and log of the perceived number of cases. It turns out participants from countries in stage 1 and 2 have significantly lower information bias relative to those from countries in stage 3 (t -test $ps < 0.01$). This evidence is merely suggestive that behavioral fatigue or caution fatigue may have affected the later stage countries, which may, in turn, have led to the larger EGPB observed.

Of course, the aforementioned second force may not always dominate the first one. For instance, if a sustained public health campaign manages to maintain the salience of the disease in late-stage countries, EGPB in those countries will likely be lower. Relatedly, it is well known that, often in a pandemic, the infection spread comes in waves. We conjecture that EGPB will be lower in subsequent waves if there is a sudden spike in infections, and the disease returns to becoming salient.

A key strength of our research, one that differentiates it from existing work, is our use of incentivized prediction elicitation. Another is that unlike existing computations of prediction bias using false information, we show actual data from COVID infections and detect prediction bias, not just a mathematical inability to compute exponential progression. In Fetzter et al. (2020), for example, “participants were instructed to assume that on day 1, one person has the fictitious disease. Furthermore, they were told to assume that each day a newly infected person infects two healthy people and then stops being contagious. Participants were

further told that on day 2, 3 people would be infected by the disease as the person who had the disease on day 1 spread it to two other people on day 2. Participants were then asked to predict the count of total people infected with the fictitious disease on days 5, 10, and 20.” In our tasks, participants did not have to speculate. The numbers on which prediction tasks were based were publicly available, and participants were informed that the researchers knew the actual numbers. Also, the timing of our survey was crucial: we wanted to know if the prediction biases changed in real-time as a country moved from under a hundred infections to over a thousand infections. By now, most countries have moved on to Stage 3, but our data on biases from earlier stages may be useful for a general understanding of the path of future epidemics. The fact we detect EGPB, not only in early-stage countries but also in countries in the thick of the pandemic, suggests such biases can help explain the dramatic spike in cases in some countries.

5. Limitations

A few caveats are in order. First, there are well-understood problems (representativeness and external validity) of conducting global surveys on Amazon's MTurk that are pertinent to our study. In our defense, though, these problems are present in most online surveys that researchers are forced to rely on while the pandemic rages on. Not to mention, arguably, the quality of self-report data collected through online platforms is superior to those obtained through face-to-face interviews, as there is no social desirability confounds in the former. Of course, the MTurk environment is not perfect; for example, it cannot capture the importance of powerful emotions such as shame or guilt (from, say, not wearing a mask) on compliance behavior. We contend that such measurement errors in our compliance measure are unlikely to affect the high and the low EGPB participants differentially.

It is interesting to note that EGPB has a flavor similar to the familiar Dunning-Kruger effect (Kruger and Dunning, 1999). The Dunning-Kruger effect captures the tendency for some people, mostly less competent, to overrate their skill, expertise, and performance and has been found to have significant public health implications, including attitude towards vaccination (Motta et al., 2018). By contrast, EGPB is concerned with a bias in judging a data generating process exogenous to the person being surveyed, one they had no hand in generating. It does not concern itself with the participants' judgment about themselves, their ability, or their power to judge their ability. It is, of course, possible that the two biases may be present in the same person – a connection our study was not designed to explore.

Our work was focused and designed to detect prediction biases and to see if they were significant predictors of compliance behavior (Fig. 4). It was not designed to make definitive causal statements connecting prediction bias or the nudge with compliance. In particular, one association is a bit perplexing: why do participants from countries at a later stage of the disease show more substantial prediction biases and less “convexity”. We provide suggestive evidence this finding may be due to “behavioral fatigue”. It is reassuring that such prediction biases are causally reduced by health communication via raw as opposed to graphical data. This point suggests that data shown via raw numbers make quite an impact on people's risk perception and should be presented alongside familiar “flatten-the-curve” style graphics. This idea is likely true of any COVID-related data, say death or recovery rates, if they follow paths similar to that of infections. A related limitation of our study is the inability to answer the question, are people when presented with numeric (as opposed to graphical) data more likely to comply with safety protocols in the future? The issue again is, even though we observe EGPB decreasing in the raw-numbers treatment, that reduction does not *causally* imply higher compliance, since our compliance data capture past behavior.

6. Conclusions

Our study documents evidence of exponential-growth prediction bias in the context of the spread of COVID-19. The results show that such biases are greater in the late-stage countries than in the early stage countries. We further show that such biases are negatively associated with WHO recommended safety compliance measures. A simple nudge related to presentation of data on COVID-19 positive cases in the form of numbers decrease EGPB relative to graphical presentation.

Future work should examine whether simple nudges can reduce EGB and improve compliance. Illustrations of such nudges may be found in Banerjee and Majumdar (2020) and Lammers et al. (2020). If successful, they can generate enormous welfare gains and produce transformative implications for social science-based, medical research, and health communication. These gains can come in the form of infections and fatalities avoided. In turn, it may spur business investment, boost aggregate demand and hiring, healing lives, and preserving livelihoods in the process. To get a firmer sense of the magnitude of these gains would require a structural model with optimizing agents who make compliance decisions given their budgets and the infection risk they perceive (see for instance, Bhattacharya et al., 2020). While such a task is outside our current scope, it is worthy of future attention.

Acknowledgement

This research has received human subjects' approval from the IRB at Iowa State University. We thank Sher Afghan Asad, Helle Bunzel, Matthew Levy, Kavitha Sandeep, and especially two anonymous referees for their comments and suggestions. Partial funding from the Indian Institute of Management, Bangalore and the NSF (grant # 2029313) is gratefully acknowledged.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.socscimed.2020.113473>.

References

- Banerjee, Ritwik, Majumdar, Priyama, 2020. Exponential Growth Bias in the Prediction of COVID-19 Spread and Economic Expectation, vol. 13664. IZA DP No.
- Bhattacharya, Joydeep, Chakraborty, Shankha, Yu, Xiumei, 2020. A Rational-Choice Model of Covid-19 Transmission with Endogenous Quarantining and Two-Sided Prevention. Economics Working Papers: Department of Economics, Iowa State University, p. 20016.
- Bischoff, Werner E., Reynolds, Tammy M., Sessler, Curtis N., Edmond, Michael B., Wenzel, Richard P., 2000. Handwashing Compliance by Health Care Workers: the Impact of Introducing an Accessible, Alcohol-Based Hand Antiseptic. *Archives of Internal Medicine*.
- Brooks, Samantha K., Webster, Rebecca K., Smith, Louise E., Woodland, Lisa, Wessely, Simon, Greenberg, Neil, James Rubin, Gideon, 2020. The Psychological Impact of Quarantine and How to Reduce it: Rapid Review of the Evidence.
- Callaghan, T Motta M., Sylvester, S., Lunz Trujillo, K., Blackburn, C.C., 2019. Parent psychology and the decision to delay childhood vaccination. *Soc. Sci. Med.* 238.
- Chen, Daniel L., Martin, Schonger, Chris, Wickens, 2016. oTree-An open-source platform for laboratory, online, and field experiments. *J. Behav. Exp. Finance* 9, 88–97.
- Cordes, Henning, Bryan, Foltice, Langer, Thomas, 2019. Misperception of exponential growth: are people aware of their errors? *Decis. Anal.* 16, 261–280.
- Cummins, Eleanor, 2020. I'll Do what I Want - Why the People Ignoring Social Distancing Orders Just Won't Listen. *Vox*.
- Donovan, Jenny L., Blake, David R., 1992. Patient non-compliance: deviance or reasoned decision-making? *Soc. Sci. Med.* 34 (5), 507–513.
- Ebersbach, Mirjam, Van Dooren, Wim, Goudriaan, Margje N., Verschaffel, Lieven, 2010. Discriminating Nonlinearity from Linearity: its Cognitive Foundations in Five-Year-Olds. " *Mathematical Thinking and Learning*.
- Fetzer, Thiemo, Hensel, Lukas, Johannes, Hermle, Roth, Christopher, 2020. Coronavirus Perceptions and Economic Anxiety.
- Goda, Gopi Shah, Levy, Matthew, Flaherty Manchester, Colleen, Sojourner, Aaron, Tasoff, Joshua, July 2019. Predicting retirement savings using survey measures of exponential growth bias and present bias. *Econ. Inq.* 57 (3), 1636–1658.
- Kawachi, Ichiro, Subramanian, S.V., 2018. Social epidemiology for the 21st century. *Soc. Sci. Med.* 196, 240–245.
- Keeling, Matt J., Rohani, Pejman, 2011. Modeling infectious Diseases in Humans and Animals.
- Kruger, J., Dunning, David, 1999. Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Psychol. Sci.* 77 (6), 1121–1134.
- Lammers, Joris, Jan, Crusius, Gast, Anne, 2020. Correcting misperceptions of exponential coronavirus growth increases support for social distancing. In: *Proceedings of the National Academy of Sciences of the United States of America*.
- Levy, Matthew, Tasoff, Joshua, 2016. Exponential - growth bias and lifecycle consumption. *J. Eur. Econ. Assoc.* 14 (3), 545–583.
- Levy, Matthew R., Tasoff, Joshua, 2017. Exponential-growth bias and overconfidence. *J. Econ. Psychol.* 58, 1–14.
- Li, Qun, et al., 2020. Early transmission dynamics in wuhan, China, of novel coronavirus-infected pneumonia. *N. Engl. J. Med.* 382 (13), 1199–1207. PMID: 31995857.
- Lunn, Pete, Timmons, Shane, Cameron, Belton, Martina, Barjakova, Hannah, Julienne, Ciaran, Lavin, 2020. Motivating Social Distancing during the COVID-19 Pandemic: an Online Experiment. *ESRI Working Paper*, p. 658.
- Lyons, Benjamin A., Vittorio, Merola, Jason, Reifler, 2020. Shifting medical guidelines: compliance and spillover effects for revised antibiotic recommendations. *Soc. Sci. Med.* 255.
- Motta, M., Callaghan, T., Sylvester, S., 2018. Knowing less but presuming more: Dunning-Kruger effects and the endorsement of anti-vaccine policy attitudes. *Soc. Sci. Med.* 211, 274–281.
- Pennycook, Gordon, Mcphetres, Jonathon, Bago, Bence, Rand, David, 04 2020. Predictors of Attitudes and Misperceptions about COVID-19 in Canada, the U.K., and the U.S.A.
- Pinsker, Joe, 2020. The People Ignoring Social Distancing. *The Atlantic*.
- Soper, George A., 1919. The lessons of the pandemic. *Science* 49, 501–506.
- Stango, Victor, Zinman, Jonathan, 2009. Exponential growth bias and household finance. *J. Finance* 64 (6), 2807–2849.
- Stanley, Matthew, Barr, Nathaniel, Peters, Kelly, Paul, Seli, 2020. Changes in Risk Perception and Protective Behavior during the First Week of the COVID-19 Pandemic in the United States.
- Thomas, Richard, 1996. Modelling space-time HIV/AIDS dynamics: applications to disease control. *Soc. Sci. Med.* 43 (3).
- van Bavel, Jay, J., Baicker, Katherine, Boggio, Paulo, Capraro, Valerio, Cichocka, Aleksandra, Crockett, Molly, Cikara, Mina, Crum, Alia, Douglas, Karen, James, Druckman, et al., Mar 2020. Using Social and Behavioral Science to Support COVID-19 Pandemic Response.
- Wagenaar, Willem A., Timmers, Han, 1979. The pond-and-duckweed problem; Three experiments on the misperception of exponential growth. *Acta Psychol.* 43, 239–251.
- Wise, Toby, Zbozinek, Tomislav D., Michelini, Giorgia, Hagan, Cindy C., Dean, Mobbs, 2020. Changes in Risk Perception and Protective Behavior during the First Week of the COVID-19 Pandemic in the United States.