Economic Modelling 136 (2024) 106749

Contents lists available at ScienceDirect

Economic Modelling

journal homepage: www.journals.elsevier.com/economic-modelling

Will voters polarize over pandemic restrictions? Theory and evidence from COVID-19[☆]

Petar Stankov

Economics Department, Royal Holloway, University of London, TW20 0EX, Egham, UK

ARTICLE INFO

Dataset link: https://data.mendeley.com/datasets/g3w7ykbz9s/2

JEL classification:
D72
I18
J22
P50

Keywords:
COVID-19
Lockdown
Fiscal transfer
Restriction
Vaccination
Political polarization

ABSTRACT

In a pandemic, voters will polarize over the choice of restrictions their society needs to impose. This is because in the eyes of voters, optimal restrictions vary with individual and aggregate productivity, vaccination status, and government fiscal transfer capacity. Specifically, less productive voters and those with a limited vaccination status prefer milder restrictions. In addition, voters in countries with superior vaccination programs prefer harsher restrictions, but the effect is weaker when their government has stronger fiscal transfer capacity. These propositions have been tested using the largest global dataset of pandemic restrictions and vaccination rates and are broadly confirmed. Consequently, the paper explains both cross-country variations in restrictions during the COVID-19 pandemic and their development as vaccination programs gained traction. It also offers policy insights into the level of restrictions for future pandemics that different voters will deem tolerable.

1. Introduction

Measures against the spread of the COVID-19 pandemic have polarized voters. The restrictions were few and far between for some and too stringent for others (Allcott et al., 2020; Cornelson and Miloucheva, 2022; Hegland et al., 2022). Using various optimization criteria, rapid progress has been achieved in identifying optimal pandemic restrictions in the literature. In early work, varieties of the susceptible–infected–recovered (SIR) model have been adopted (Kermack and McKendrick, 1927) to balance disease control and economic loss (Alvarez et al., 2021; Fajgelbaum et al., 2021). More recent contributions have incorporated group-level heterogeneity by age or health status (Acemoglu et al., 2021; Glover et al., 2023) or a natural evolution of the virus (Prieur et al., 2024) to determine optimal levels of restrictions. Empirical evidence confirming those findings emerged soon after (Blayac et al., 2022).

However, we still do not know how these restrictions are accepted by voters. Therefore, the focus of this paper is *not* on what governments should do to strike an optimal balance between disease control and economic loss. Instead, it is on what *voters think* their governments should do. We would like to know whether *politically* optimal restrictions exist and, if yes, how they depend on individual voter characteristics, such as productivity or vaccination status.

The SIR model and its timely extensions seem unable to address these questions for now. Unlike them, simple extensions of canonical political economy models like the one below hold considerable strength. Because pandemics are recurring phenomena with repeated behavioral and policy response patterns (Beach et al., 2022; Dasgupta et al., 2021), it is useful to know what level of pandemic restrictions is optimal in the eyes of voters, as already suggested by Murray (2020) and Avery et al. (2020). Knowing this level of restrictions, if it exists at all, would help governments address political polarization in future pandemics. Offering such an insight into the politics of pandemic control is the main innovation in this work.

To find the politically optimal restrictions during a pandemic, this paper uses a canonical policy preference framework (Persson and Tabellini, 2002, pp. 19–46) with heterogeneous voters, identical firms, and a government. Most choices are made by the voters. To keep the model tractable, firms are a mere mapping of labor into consumption goods. The government has both economic and political roles. Its economic role is reduced to taxing income and providing transfers,

☆ I thank Andrew Mountford, Byeongju Jeong, Jungyoon Lee, Łukasz Rachel, Michael Mandler, Michael Naef, two anonymous referees and the Editor of this journal for helpful comments and suggestions. All errors are mine.

E-mail address: petar.stankov@rhul.ac.uk.

https://doi.org/10.1016/j.econmod.2024.106749
Received 28 September 2023; Received in revised form 25 April 2024; Accepted 25 April 2024
Available online 7 May 2024

while its political role is implementing the policy choices of the median voter.

The choice of the median voter, however, depends on individual voter characteristics. Therefore, individual voter heterogeneity gains importance both in the model and in reality, for two reasons. First, as implied by Besley and Stern (2020), the aggregation of voter heterogeneity will uncover reasons behind the political polarization regarding lockdowns, with implications reaching beyond the COVID-19 pandemic. Second, if politicians commit to choosing what voters want, voter heterogeneity will deliver a variety of observed lockdown policies across and within countries over time that enable empirical tests of the model predictions.

In what follows, I derive the politically optimal level of restrictions a society will choose under fairly restrictive assumptions of pairwise choice under pure majority voting. Relaxing those assumptions then produces a tractable multiplicity of politically optimal restrictions. As each choice is feasible, individual heterogeneity and the variety of ways it can be aggregated will deliver significant cross-country differences in voter preferences concerning lockdowns. In turn, this explains the variety of lockdown and exit strategies adopted across countries over time during the COVID-19 pandemic.

Despite its simplicity, the model presented below finds the politically optimal level of restrictions in an economy with heterogeneous voters. It then proceeds with comparative statics to derive a set of empirically testable propositions naturally emerging from the model, which are then tested on COVID-19 pandemic data. The theoretical model of politically optimal lockdowns is presented in the next section, followed by its empirical counterpart before the final, concluding section.

2. The economy

The economy is populated by a continuum of rational utility maximizers. Each agent i has quasi-linear preferences over consumption c^i and leisure x^i. Consumption c^i is financed from labor income l^i and government transfers f. To finance transfers, the government taxes the labor income of each working individual at a flat rate of τ. Individuals receive labor income only if they go to work and transfers otherwise.

The government has rolled out a free and exogenously provided vaccination program. Each individual's vaccine status is represented by $v^i \in [0;1]$ and drawn from a distribution with a mean of v and a median of v^m. Vaccines are productivity-enhancing—those with a higher vaccination status can produce more than the rest. Vaccine hesitancy does not exist, but the model allows for relaxation of this assumption.

Individuals face both a budget and a time constraint. They spend time in leisure or work. We abstain from the mental health and domestic violence issues spurred by lockdowns (Altindag et al., 2022; Banks and Xu, 2020; Berniell and Facchini, 2021; Swaziek and Wozniak, 2020) and assume that leisure is a desired good, while working carries disutility. Working, furthermore, may be subject to lockdowns, $d \in [0;1]$. Lockdowns are understood as restrictions on the time when labor is allowed to go to work. Effectively, a lockdown acts like a time-tax on the working time of an individual. Similar to natural disasters, a lockdown has a directly negative effect on labor hours, as found previously by Droste and Stock (2021) and Ludvigson et al. (2021). In the absence of any restrictions, $d = 0$ and everyone can work full time. Under a full lockdown of $d = 1$, agents cannot go to work. In reality, a full lockdown is impossible, as the economy cannot shut down completely. However, a full lockdown of $d = 1$ approximates the reality where, for a brief few weeks in spring 2020, the world felt as if it was shutting down.

For simplicity, both commitment to and enforcement of lockdowns are perfect in the model, although in practice they admittedly are not. Recent works have already endogenized imperfect compliance in political economy models (Gitmez et al., 2020; Jelnov and Jelnov, 2022) and studied the origins of noncompliance (Blayac et al., 2022), the healthcare outcomes of a lack of trust (Charron et al., 2023), and effects on imperfect government commitment (Moser and Yared, 2022). Parametrizing degrees of government commitment or citizen compliance would be a straightforward extension and is therefore left out of this model to preserve its parsimony and focus on what has not been sufficiently studied so far.

In addition to their vaccination status, agents differ in their individual productivity, $\alpha^i \in [0;1]$, drawn from a distribution with a mean of α and a median of α^m. A higher α^i means an individual is more productive, which effectively adds to the individual's time endowment. Higher productivity may also correlate with better health status and younger age as in the rest of the emerging literature on the political economy of lockdowns (Boettke and Powell, 2021).

Two features of the model separate it from classic voter preference models. First, i's vaccination status v^i enhances the agent's ability to work in a way similar to a productivity parameter in a standard production function. An agent who is fully vaccinated, $v^i = 1$, can go to work and enjoy labor income. Those who do not either do not work, which means they live off transfers f, or their capacity to produce is limited because of their incomplete vaccination course. In the second feature, a lockdown affects the time that agents spend working. In a full lockdown, agents are not allowed to go to work at all, while relaxing restrictions boosts the individual time endowment.

In the following, a standard solution algorithm is used to identify a political equilibrium (Persson and Tabellini, 2002, pp. 19–46). Initially, the optimal individual-level labor supply and consumption in this economy are derived. Then, aggregation exercise is performed to derive the government budget constraint and the level of transfer, f. Based on the labor supply and transfer, the policy preferences of individual i are derived. The policy preferences then yield a single best politically feasible lockdown policy chosen by the median voter that the government implements.

To determine whether a society will polarize, politically optimal restrictions are derived for any voter i. In turn, this gives rise to multiple equilibria conditioned on the distribution of population parameters. The politically optimal lockdown policies are finally subjected to comparative statics exercises to determine their dependence on individual and aggregate heterogeneity, such as differences in productivity, vaccination rates, and income tax rates. In the model, the income tax rate is a proxy for the capacity of the government to provide fiscal transfers. This assumption is consistent with those proposed by Aizenman et al. (2019) and Besley and Persson (2014), who contend that higher-income countries have a larger fiscal space and tax their citizens at higher taxes. In turn, a higher tax rate in the model proxies a fairly large set of fiscal parameters, including a higher capacity for income transfers during the COVID crisis, that tends to absorb and offset negative shocks (Gutiérrez-Romero, 2022; Walmsley et al., 2023; Dergiades et al., 2022; Di Novi et al., 2023).

The indirect utility function of individual i is derived next. It leads to individually optimal labor and consumption in this economy. Any nontrivial derivations are left for Appendix.

3. Individual labor and consumption

Individuals choose l^i to maximize their utility $w^i = c^i + ln(x^i)$ subject to their budget and time constraints. The full problem of individual i can be formulated as follows:

$$\begin{aligned} \underset{\{l^i\}}{\text{Max.}} \quad & w^i = c^i + ln(x^i), \\ \text{s.t.} \quad & c^i \leq (1-\tau)v^i l^i + f, \\ & x^i + (1+d)l^i \leq 1 + \alpha^i. \end{aligned} \tag{1}$$

Assuming the constraints are binding and plugging c^i and x^i into the utility function produces the following indirect utility function of agent i:

$$w^i = (1-\tau)v^i l^i + f + ln[1 + \alpha^i - (1+d)l^i]. \tag{2}$$

Maximizing w^i with respect to l^i and using the individual budget constraint produces the following individually optimal levels of labor, $(l^i)^*$, and consumption, $(c^i)^*$, respectively:

$$(l^i)^* = \frac{1+\alpha^i}{1+d} - \frac{1}{(1-\tau)v^i}, \tag{3}$$

$$(c^i)^* = \frac{1+\alpha^i}{1+d}(1-\tau)v^i - 1 + f. \tag{4}$$

As both $(l^i)^*$ and $(c^i)^*$ are functions of α^i, v^i, τ and d, we could use comparative statics to find the impact of parameter changes on what citizens do in this economy. For example, it is straightforward to show that a harder lockdown reduces labor supply and consumption in line with intrapandemic evidence (Brinca et al., 2021; Chen et al., 2021; Eichenbaum et al., 2021). Although informative of individually optimal choices, such comparative statics would tell us little about individual and aggregate policy preferences. Deriving those preferences dictates aggregating individual behavior first to obtain the transfer f, which is performed next.

4. Aggregation

The optimal labor supply of voter i is aggregated in per capita terms, which helps reveal the transfer f and policy preferences for lockdown restrictions. Using $\alpha \equiv E_i\{\alpha^i\}$ and $v \equiv E_i\{v^i\}$, the per capita optimal labor supply is

$$l^* = \frac{1+\alpha}{1+d} - \frac{1}{(1-\tau)v}. \tag{5}$$

As wages are normalized to 1, l^* becomes per capita income. The government taxes income at a rate of τ to finance its per capita transfers, f. Therefore, the government's budget constraint, holding with equality, is

$$f = \tau l^*, \tag{6}$$

which, given (5), sets the level of the per capita transfer at

$$f = \tau\left[\frac{1+\alpha}{1+d} - \frac{1}{(1-\tau)v}\right]. \tag{7}$$

Plugging the individually optimal labor supply (3) and the transfer (7) into the indirect utility (2) yields the policy preferences.

5. Policy preferences

The policy preferences of individual i, conditional on their individual productivity and vaccination status, are given by $W^i(d,\tau;\alpha^i,v^i) \equiv$

$$\equiv \frac{(1+\alpha^i)(1-\tau)v^i}{1+d} + ln\left[\frac{1+d}{(1-\tau)v^i}\right] + \tau\left[\frac{1+\alpha}{1+d} - \frac{1}{(1-\tau)v}\right] - 1. \tag{8}$$

Note that in the policy preference function, both the aggregate vaccination parameter v and the aggregate productivity parameter α strictly improve the individual's welfare. This is reasonable, as the model is consistent with previously documented positive healthcare and productivity externalities (Boulier et al., 2007; Akbarpour et al., 2024).

Based on the above function, demonstrating the existence of a political equilibrium is trivial. As $W^i(\cdot) = J(d,\tau) + K(\alpha^i,v^i)L(d,\tau)$ where $J(\cdot)$ and $L(\cdot)$ are common to all voters, and $K(\alpha^i,v^i)$ is monotonic in α^i and v^i, $W^i(\cdot)$ represents intermediate preferences. Those preferences lead to a single optimal vector of policies under pure majority rule and a pairwise vote (Persson and Tabellini, 2002, pp. 19–46). This in turn coincides with the policies preferred by the median voter, $W^m(d,\tau;\alpha^m,v^m)$. As a result, we observe the following Condorcet winner, d^m:

$$d^m(\tau;\alpha^m,v^m) = (1+\alpha^m)[(1-\tau)v^m + \tau] - 1. \tag{9}$$

Four implications arise from Eq. (9). First, in an environment of zero progress with vaccination and especially low taxes—which also means a low ability of the government to finance transfers—the politically optimal level of lockdown is close to zero. Second, the politically optimal lockdown policy above depends on aggregate parameters only. This disallows politically relevant polarization, so society will unite around a single politically optimal level of restrictions, d^m. Third, a variety of such levels exist across countries and depend on cross-country differences in productivity, vaccination status, and the state of the social welfare system. Fourth, because v^m changes within a country over time, the politically optimal level of restrictions within any country varies with advancements in the vaccination program.

However, the conditions under which this political equilibrium is reached are arguably unrealistic. Pairwise votes rarely exist, and society does not choose policies under pure majority rule. The compelling evidence of a postlockdown surge in government popularity in Western Europe (Bol et al., 2021) is not equivalent to evidence of societies uniting around a single best policy. Therefore, we would like to know the politically optimal lockdown policies for voters other than the median one.

6. Optimal lockdowns away from the median

Taking the first partial derivative of the policy preference function (8) with respect to d and setting it equal to zero yields the following politically optimal level of lockdown for individual i:

$$d^*(\alpha;\tau;\alpha^i,v^i) = (1+\alpha^i)(1-\tau)v^i + (1+\alpha)\tau - 1. \tag{10}$$

Note that unlike Eq. (9), the optimal level of restrictions now varies across both aggregate and individual heterogeneity parameters. Low-productivity voters and voters with zero or incomplete vaccination will prefer different levels of restrictions to the preferences of high-productivity and high-vaccination individuals. Specifically, low-productivity voters and those expressing more vaccine hesitance prefer softer restrictions, while high-productivity voters and those with a higher vaccination status prefer harder restrictions, which offers support for the results from recent agent-based modeling work (Pangallo et al., 2024). In addition, α^i and v^i have different distributions across countries and over time, with the distribution of α^i being more stable due to the nature of the vaccination programs. As a result, their aggregation will lead to a variety of politically optimal lockdowns across countries and to a dynamic optimal level of restrictions within a country, confirming other recent results (Caulkins et al., 2023). Therefore, similar to (9), Eq. (10) can explain both cross-country and within-country differences in policy responses to the pandemic.

However, Eq. (10) is reached under far less restrictive assumptions, with neither pure majority rule nor pairwise voting in the economy leading to the heterogeneous optimum d^*. Therefore, unlike (9), Eq. (10) brings the model closer to the real world, where a polarized society unable to unite behind a single politically optimal level of restrictions is the norm—a result that has already been empirically tested (Blayac et al., 2022). Specifically, because d^* depends on individual heterogeneity, society will polarize over d^* along two dimensions: productivity and vaccination status. Higher-productivity voters and those with complete vaccination courses will generally prefer stricter lockdowns than those preferred by the rest of society. The global protests against restrictions first sparked in the United States in spring 2020 (Dyer, 2020), as well as wider evidence on the political consequences of pandemics (Lewkowicz et al., 2022), are indicative of those latent polarization trends.

The model refrains from a dynamic framework for two reasons. The first is that extensions of the canonical SIR model reach similar conclusions with added complexity but without the required emphasis on the political economy of restrictions. Second, the political economy of lockdowns is at its onset. As such, its models need parsimony to capture crucial features of the world, such as short-termism and intertemporal ignorance in the face of a crisis (Furton, 2023), and simultaneously explain considerable variations in observed data. The model presented here fits these desired characteristics. The following comparative statics exercise reveals the full range of empirical propositions following from the model, while the subsequent data section discusses those we are able to test with existing datasets.

7. Empirical propositions

We can study the impact of the model parameters on politically optimal levels of lockdown by taking the first partial derivatives of $d^*(\alpha;\tau;\alpha^i,v^i)$. The sign of the derivatives will then produce five empirical propositions, as follows:

P1: Within a given country, more productive individuals prefer stricter lockdowns, whereas less productive individuals prefer softer lockdowns. The effect is reinforced by individual vaccination status. To see this, consider $\frac{\partial d^*}{\partial \alpha^i} = (1-\tau)v^i > 0$. Differentiating further with respect to (w.r.t.) v^i yields $\frac{\partial d^*}{\partial \alpha^i}/\partial v^i = (1-\tau) > 0$.

P2: Individuals with a higher vaccination status will prefer stricter lockdowns, and those who prefer stricter lockdowns are more likely to be vaccinated. The effect is stronger for more productive individuals. To see this, consider $\frac{\partial d^*}{\partial v^i} = (1+\alpha^i)(1-\tau) > 0$. Differentiating further w.r.t. α^i produces a positive second partial derivative, with the intuition lying in the uneven impact of the COVID recession across the productivity distribution (Eichenbaum et al., 2022; Glover et al., 2023). At the same time, differentiating w.r.t τ yields a negative second partial derivative.

P3: Aggregating Proposition 2 to the country level, we find that countries with more advanced vaccination programs impose stricter lockdowns, but the effect is weaker in countries with a higher capacity for welfare transfers. To see this, it is enough to average $\frac{\partial d^*}{\partial v^i} = (1+\alpha^i)(1-\tau)$ across i to get $\frac{\partial d^*}{\partial v} = (1+\alpha)(1-\tau) > 0$ and then further differentiate w.r.t. τ.

P4: In a cross section of countries, more productive nations end up having harder restrictions. The effect is reinforced by the ability of the government to provide transfers. To see this, consider $\frac{\partial d^*}{\partial \alpha} = \tau > 0$. Differentiating further w.r.t. τ gives $1 > 0$.

P5: The impact of the government's ability to provide transfers has an ambiguous impact on the level of restrictions. Notwithstanding this observation, countries with more developed social welfare systems are more likely to prefer harder restrictions. To see this, consider $\frac{\partial d^*}{\partial \tau} = 1+\alpha-(1+\alpha^i)v^i$, which is ambiguous in the general case. Additionally, note further that $\frac{\partial d^*}{\partial \tau} < 0 \iff v^i > \frac{1+\alpha}{1+\alpha^i}$. In a cross section of countries, we typically operate with aggregate values. Therefore, if we take averages across i, we obtain $\frac{\partial d^*}{\partial \tau} < 0 \iff v > 1$. However, as $v < 1$ for all i except when everyone is fully vaccinated, $\frac{\partial d^*}{\partial \tau} \geq 0$.

The paper proceeds to test some of these propositions. The empirical methodology is presented next.

8. Empirical methodology

Testing Proposition 1 requires individual-level or regional data where the vaccine choice is observable. Initial efforts using within-country regional variations in vaccination data are under way. For the moment, we resort to testing propositions relying on cross-country time variation in the data. Propositions 2 and 3 are essentially identical except for the level at which the data are observed: Proposition 2 requires individual data, while Proposition 3 can be tested using country-level data. Proposition 4 can be tested by adding aggregate productivity measures or their proxies to the empirical model used to test Proposition 3 or introducing country-level fixed effects at a minimum. Moreover, Proposition 5 requires adding measures of the fiscal support offered by the government.

Table 1
Summary statistics.

Variable/Statistic	Mean	P50	SD	IQR	Min	Max	N
Stringency (d_{ct})	54.36	55.56	18.41	26.39	0.00	97.22	28,904
Tvac ($v_{ct}=i$)	31.34	16.35	36.45	45.07	0.00	232.72	18,421
Onedose ($v_{ct}=ii$)	19.98	11.33	21.31	30.20	0.00	116.73	17,606
Twodose ($v_{ct}=iii$)	12.81	5.71	16.50	16.92	0.00	115.99	15,213
Income support (τ_{ct})	0.88	1.00	0.76	1.00	0.00	2.00	28,836
Log-Cases ($C_{ct}=i$)	10.92	11.30	2.83	3.56	0.00	17.36	29,834
Log-Deaths ($C_{ct}=ii$)	7.03	7.18	2.71	3.75	0.00	13.32	28,713

Notes: The table presents basic summary statistics for the variables used in the estimation below. The statistics shown are (by column from left to right) variable, mean, median, standard deviation, interquartile range, minimum, maximum, and number of observations. The variables match the theory parameters and are explained in the text. *Source:* Hale et al. (2021) and Mathieu et al. (2021).

It is evident that Propositions 3, 4, and 5 can be tested by combining the COVID-19 Government Response Tracker (Hale et al., 2021) with the vaccination tracker from Our World in Data (Mathieu et al., 2021). The rest of the propositions require variation in data at individual or within-country regional levels. Therefore, testing Propositions 3, 4, and 5 is the focus of the remainder of this work.

The Hale et al. (2021) data contain daily country-level information on restrictions, the number of positive cases, and the level of income support, while the Mathieu et al. (2021) data contain daily vaccination data. The former dataset spans from January 1, 2020, to July 27, 2021, capturing the first and second waves of the COVID-19 pandemic across 186 countries and territories. The latter dataset starts when we observe the first vaccination for each country or territory, which is in winter or spring 2021 for most countries and territories. The match between the two datasets totals 30,414 country-days, as further summarized in Table 1.

The stationarity of lockdown stringency data is determined using Choi (2001) Fisher-type tests in both their augmented Dickey–Fuller and Phillips–Perron flavors. Those tests are chosen because the Levin et al. (2002) test requires strongly balanced data, while the Im et al. (2003) test cannot be performed on data with gaps. The tests include both heterogeneous intercepts to capture the variety of stringency across countries and a trend to capture the usual path of restrictions within a country. The Choi (2001) tests strongly reject the null of all panels containing unit root, prompting estimations in levels. Therefore, the model to test the above propositions is as follows:

$$d_{ct} = \lambda_c + \lambda_t + \beta_1 v_{ct} + \beta_2 v_{ct}^2 + \beta_3 (v * \tau)_{ct} + \beta_4 \tau_{ct} + \beta_5 C_{ct} + \varepsilon_{ct}, \tag{11}$$

where d_{ct} is the stringency of restrictions[1] imposed in country c on date t; λ_c and λ_t are country and time effects, respectively; v_{ct} is one of the following three vaccination indicators: (i) total number of doses administered per 100 people in the total population of the country ($Tvac$); (ii) number of people vaccinated per 100 people in the total population of the country ($Onedose$); and (iii) number of people per 100 who received all doses prescribed by the vaccination protocol ($Twodose$); τ_{ct} is income support[2] available in country c at time t;

[1] The stringency index aggregates eight containment and closure measures in response to the COVID-19 pandemic: School closing, workplace closing, cancellations of public events, restrictions on gathering, restrictions on public transportation, stay-at-home requirements, restrictions on international movement, and restrictions on international travel. For details, see Hale et al. (2021, p.530) and the accompanying Codebook.

[2] The income support tracks if the government is providing direct cash payments to people who lose their jobs or cannot work. It is measured on an ordinal scale: 0 if no income support is offered, 1 when the government is replacing less than 50% of lost salary (or if a flat sum in which case it is less than 50% median salary), and 2 when the government is replacing 50% or more of lost salary (or if a flat sum, greater than 50% median salary). For details, see Hale et al. (2021, p.530) and the accompanying Codebook.

Table 2
Lockdowns, Vaccines, and Income support.

	(1)	(2)	(3)	(4)	(5)	(6)
	C_{ct} = (i)			C_{ct} = (ii)		
	υ = (i)	υ = (ii)	υ = (iii)	υ = (i)	υ = (ii)	υ = (iii)
υ_{ct}	0.062	0.156	−0.054	0.045	0.125	−0.067
	(0.066)	(0.114)	(0.148)	(0.065)	(0.110)	(0.143)
υ^2_{ct}	0.000	−0.001	0.002	0.000	−0.000	0.002*
	(0.000)	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)
$(\upsilon * \tau)_{ct}$	−0.048*	−0.085**	−0.101	−0.068***	−0.112***	−0.152***
	(0.025)	(0.042)	(0.063)	(0.025)	(0.043)	(0.057)
τ_{ct}	2.326	2.644	2.788	2.823	2.988	3.105
	(2.062)	(2.194)	(1.975)	(2.121)	(2.313)	(1.882)
C_{ct}	11.494***	12.022***	13.934***	5.600***	6.034***	6.767***
	(2.661)	(2.566)	(2.442)	(1.316)	(1.403)	(1.748)
FE's	λ_c, λ_t	λ_c, λ_t	λ_c, λ_t	λ_c, λ_t	λ_c, λ_t	λ_c, λ_t
N	17,016	16,227	13,907	16,296	15,570	13,326
R^2	0.169	0.183	0.203	0.121	0.133	0.140

Notes: The table presents estimates from Eq. (11) using two measures of COVID-19 severity, C_{ct} = (i) and C_{ct} = (ii), respectively, and three measures of the vaccination status of the population, υ = (i), υ = (ii), and υ = (iii), respectively, as defined in the text. The dependent variable is the daily Stringency Index from Hale et al. (2021), as defined in the text. All models include both λ_c and λ_t—country- and time-fixed effects (FE's), respectively. Model (1/4) uses vaccination data υ as defined by (i), model (2/5) by (ii), and model (3/6) by (iii), respectively: (i) total vaccinations per 100 people, (ii) number of first doses per 100 people, and (iii) number of people per 100 with a complete vaccination course. Robust standard errors are given in parentheses. Symbols: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

$(\upsilon * \tau)_{ct}$ is the interaction between the respective vaccination indicator and the income support level; the severity of COVID-19 C_{ct} is measured as the natural logarithm of either (i) the number of confirmed positive COVID-19 cases ($C_{ct} = i$) or (ii) the number of confirmed deaths related to a COVID-19 illness ($C_{ct} = ii$), and ε_{ct} is an error term. The discussion of the results follows.

9. Results

The results in Table 2 provide support for the above theory. As expected, the sign of β_1 is positive. Its insignificance, however, means that advancing a country's vaccination program will rarely become a key *de facto* driver of restrictions for its voters, even though this paper and other recent works have implied otherwise (Caulkins et al., 2023). The quadratic term also appears insignificant, suggesting that any nonlinear impact of a vaccination program on the preferred level of restrictions is fairly weak.

Further, the estimates of β_3 are negative. This is in line with the expectation set out in the model; that is, countries with a more supportive social welfare system or a higher fiscal capacity for transfers will see lower level of restrictions than countries with a lower capacity for transfers. Those will typically be countries with a stronger fiscal capacity to react to crises, i.e., higher-income countries with higher income tax rates (Besley and Persson, 2014). However, it should be noted that the estimates are marginally insignificant in Model (3), which suggests that the dampening effect of social transfers on the relationship between vaccination and restrictions may not always be stable. In addition, the state of the social welfare system is never a significant stand-alone factor for lockdown decisions, consistent with recent empirical findings (Kurrild-Klitgaard, 2024). Still, in line with the expectations set out in the model, the sign of the relationship is positive.

However, in line with recent evidence (Pangallo et al., 2024; Hebert and Curry, 2022), the most potent driver of lockdown restrictions has remained the ongoing epidemiological situation. A country with positive cases numbering one standard deviation above another country will typically see its stringency index about 25–30 points higher, which is enough to move it from the 25th to the 75th percentile within the stringency distribution. The latter result is confirmed even after controlling for common daily shocks across countries.

The empirical results above broadly confirm the theoretical propositions. Novel datasets, e.g., those tracking voters or economic activity in near-real time (Vavra, 2021; Lourenço and Rua, 2021; Menezes et al., 2022), can tell us more about the heterogeneity of preferred restrictions along the productivity distribution.

Until then, the results in this paper could inform policymakers about *politically* optimal restrictions in a pandemic. The documented short-termism of politicians and technocrats in a pandemic (Furton, 2023; Karadimas, 2023) may be reinforced by the natural tendency of voters to polarize over restrictions as documented in this work. To counter such responses, policymakers could target the parameters of politically optimal restrictions, aiming to bring the politically optimal d^* closer to the one that balances the healthcare and economic costs of the pandemic. Specifically, policymakers who prefer or are advised to impose harder restrictions should engage with the more productive workers first, as they are more likely to respond positively to the policy proposals. In addition, it would be advisable to offer generous fiscal transfers at the onset of a pandemic when vaccines are not available but imposing restrictions is optimal from a healthcare perspective. This could reduce opposition to harder restrictions. Finally, as the pandemic eases, the level of restrictions that a society will tolerate sinks. The rest of the conclusions are set out below.

10. Conclusion

This paper studies the politically optimal level of restrictions in an economy with heterogeneous voters. Voters are heterogeneous along two dimensions: productivity and vaccination status. In a static voter preference framework, voters spend time at leisure and work. Work is more productive when a worker is vaccinated but is also subject to lockdown restrictions. A lockdown restricts the time individuals can spend working, effectively taxing individual time endowment.

The politically optimal levels of restrictions depend on both aggregate and individual parameters. In turn, cross-country differences in the chosen level of restrictions become easy to explain. In addition, political polarization over lockdowns becomes the norm because the aggregation of individual heterogeneity delivers various optimal lockdowns for politically relevant groups of voters.

The model predicts that more productive voters and those with a more advanced vaccination status prefer harder lockdowns. In addition, the model suggests that vaccination progress will raise the politically optimal level of restrictions, but the effect will be mitigated by the capacity of the government to offer fiscal transfers. Evidence of such mitigation was produced using global daily vaccinations and restrictions data from the second wave of the COVID-19 pandemic.

However, given the effect of the timing of restrictions and vaccinations on the success of mitigation strategies (Oraby et al., 2021), adding a dynamic angle seems like a natural place to start extending this theory. Other approaches include endogenizing the vaccination status or at least allowing parametrization of the vaccine pickup rate, possibly in response to a certain risk of infection and heightened healthcare concerns. In addition, fruitful avenues to extend the framework would be to diversify employment into a traditional and a work-from-home sector in which work-from-home workers enjoy a premium on working hours in a lockdown. This type of extension is particularly fruitful given documented cross-country differences in the capacity to work from home (Gottlieb et al., 2021), which could aggravate inequality both within and across countries (Palomino et al., 2020) and deepen political divides and polarization. A richer production side could also have firms decide on the optimal levels of labor and benefit from subsidies for furloughed labor. Finally, d could be endogenously chosen by an incumbent government and opposed by political competition. Dynamics, endogenous vaccines, richer production, and government sides, as well as a more dynamic political process, all seem promising ways to extend the current framework. In turn, this will undoubtedly add to our understanding of the politics behind restrictions in past and future pandemics.

Declaration of competing interest

None.

Data availability

I have shared the link to my data at: https://data.mendeley.com/datasets/g3w7ykbz9s/2.

Appendix. Derivations

A.1. Policy preferences $W^i(\cdot)$

Plugging the individually optimal labor supply (3) and the level of the government transfer (7) into the indirect utility function (2) yields $W^i(\cdot) = (1-\tau)v^i\left[\frac{1+\alpha^i}{1+d} - \frac{1}{(1-\tau)v^i}\right] + ln\left\{(1+\alpha^i) - (1+d)\left[\frac{1+\alpha^i}{1+d} - \frac{1}{(1-\tau)v^i}\right]\right\} + \tau\left[\frac{1+\alpha}{1+d} - \frac{1}{(1-\tau)v}\right] = (1-\tau)v^i\frac{1+\alpha^i}{1+d} - 1 + ln\left[\frac{1+d}{(1-\tau)v^i}\right] + \tau\left[\frac{1+\alpha}{1+d} - \frac{1}{(1-\tau)v}\right]$, which is (8).

A.2. Politically optimal lockdown $d^*(\cdot)$

Taking the first partial derivative of the policy preference function (8) w.r.t. d produces $\frac{\partial W^i(\cdot)}{\partial d} = -(1-\tau)v^i\frac{1+\alpha^i}{(1+d)^2} - \frac{(1+\alpha)\tau}{(1+d)^2} + \frac{1}{1+d}$. Setting this equal to 0 and solving for d produces the politically optimal level of restrictions (10).

A.3. Condorcet winner $d^m(\cdot)$

Assume that, for the median voter, $\alpha^i = \alpha = \alpha^m$ and $v^i = v = v^m$. Then, the politically optimal level of restrictions (10) simplifies to d^m.

References

Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D., 2021. Optimal targeted lockdowns in a multi-group SIR model. Am. Econ. Rev.: Insights 3 (4), 487–502.
Aizenman, J., Jinjarak, Y., Nguyen, H.T.K., Park, D., 2019. Fiscal space and government-spending and tax-rate cyclicality patterns: A cross-country comparison, 1960–2016. J. Macroecon. 60, 229–252.
Akbarpour, M., Budish, E., Dworczak, P., Kominers, S.D., 2024. An economic framework for vaccine prioritization. Q. J. Econ. 139 (1), 359–417.
Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., Yang, D., 2020. Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. J. Public Econ. 191, 104254.
Altindag, O., Erten, B., Keskin, P., 2022. Mental health costs of lockdowns: Evidence from age-specific curfews in Turkey. Am. Econ. J.: Appl. Econ. 14 (2), 320–343.
Alvarez, F., Argente, D., Lippi, F., 2021. A simple planning problem for COVID-19 lockdown, testing, and tracing. Am. Econ. Rev.: Insights 3 (3), 367–382.
Avery, C., Bossert, W., Clark, A., Ellison, G., Ellison, S.F., 2020. An economist's guide to epidemiology models of infectious disease. J. Econ. Perspect. 34 (4), 79–104.
Banks, J., Xu, X., 2020. The mental health effects of the first two months of lockdown during the COVID-19 pandemic in the UK. Fiscal Stud. 41 (3), 685–708.
Beach, B., Clay, K., Saavedra, M., 2022. The 1918 influenza pandemic and its lessons for COVID-19. J. Econ. Lit. 60 (1), 41–84.
Berniell, I., Facchini, G., 2021. COVID-19 lockdown and domestic violence: Evidence from internet-search behavior in 11 countries. Eur. Econ. Rev. 136, 103775.
Besley, T., Persson, T., 2014. Why do developing countries tax so little? J. Econ. Perspect. 28 (4), 99–120.
Besley, T., Stern, N., 2020. The economics of lockdown. Fiscal Stud. 41 (3), 493–513.
Blayac, T., Dubois, D., Duchêne, S., Nguyen-Van, P., Ventelou, B., Willinger, M., 2022. What drives the acceptability of restrictive health policies: An experimental assessment of individual preferences for anti-COVID-19 strategies. Econ. Model. 116, 106047.
Boettke, P., Powell, B., 2021. The political economy of the COVID-19 pandemic. South. Econ. J. 87 (4), 1090–1106.
Bol, D., Giani, M., Blais, A., Loewen, P.J., 2021. The effect of COVID-19 lockdowns on political support: Some good news for democracy? Eur. J. Political Res. 60 (2), 497–505.
Boulier, B.L., Datta, T.S., Goldfarb, R.S., 2007. Vaccination externalities. B.E. J. Econ. Anal. Policy 7 (1).
Brinca, P., Duarte, J.B., e Castro, M.F., 2021. Measuring labor supply and demand shocks during COVID-19. Eur. Econ. Rev. 139, 103901.
Caulkins, J., Grass, D., Feichtinger, G., Hartl, R., Kort, P., Kuhn, M., Prskawetz, A., Sanchez-Romero, M., Seidl, A., Wrzaczek, S., 2023. The hammer and the jab: Are COVID-19 lockdowns and vaccinations complements or substitutes? European J. Oper. Res. 311 (1), 233–250.
Charron, N., Lapuente, V., Rodríguez-Pose, A., 2023. Uncooperative society, uncooperative politics or both? Trust, polarization, populism and COVID-19 deaths across European regions. Eur. J. Political Res. 62 (3), 781–805.
Chen, H., Qian, W., Wen, Q., 2021. The impact of the COVID-19 pandemic on consumption: Learning from high-frequency transaction data. AEA Pap. Proc. 111, 307–311.
Choi, I., 2001. Unit root tests for panel data. J. Int. Money Finance 20 (2), 249–272.
Cornelson, K., Miloucheva, B., 2022. Political polarization and cooperation during a pandemic. Health Econ. 31 (9), 2025–2049.
Dasgupta, U., Jha, C.K., Sarangi, S., 2021. Persistent patterns of behavior: Two infectious disease outbreaks 350 years apart. Econ. Inq. 59 (2), 848–857.
Dergiades, T., Milas, C., Panagiotidis, T., 2022. Unemployment claims during COVID-19 and economic support measures in the U.S.. Econ. Model. 113, 105891.
Di Novi, C., Paruolo, P., Verzillo, S., 2023. Does labour protection influence mental-health responses to employment shocks? Evidence on older workers in Europe. Econ. Model. 126, 106406.
Droste, M., Stock, J.H., 2021. Adapting to the COVID-19 pandemic. AEA Pap. Proc. 111, 351–355.
Dyer, O., 2020. COVID-19: Trump stokes protests against social distancing measures. BMJ 369.
Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2021. The macroeconomics of epidemics. Rev. Financ. Stud. 34 (11), 5149–5187.
Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2022. Inequality in life and death. IMF Econ. Rev. 70, 68–104.
Fajgelbaum, P.D., Khandelwal, A., Kim, W., Mantovani, C., Schaal, E., 2021. Optimal lockdown in a commuting network. Am. Econ. Rev.: Insights 3 (4), 503–522.
Furton, G.L., 2023. The pox of politics: Troesken's tradeoff reexamined. Public Choice 195 (1), 169–191.
Gitmez, A., Sonin, K., Wright, A.L., 2020. Political Economy of Crisis Response. Working Papers 2020–68, Becker Friedman Institute for Research In Economics.
Glover, A., Heathcote, J., Krueger, D., Ríos-Rull, J.-V., 2023. Health versus wealth: On the distributional effects of controlling a pandemic. J. Monetary Econ. 140, 34–59.
Gottlieb, C., Grobovšek, J., Poschke, M., Saltiel, F., 2021. Working from home in developing countries. Eur. Econ. Rev. 133, 103679.
Gutiérrez-Romero, R., 2022. Conflicts increased in Africa shortly after COVID-19 lockdowns, but welfare assistance reduced fatalities. Econ. Model. 116, 105991.

Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., Tatlow, H., 2021. A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker).. Nat. Hum. Behav. 5, 529–538.
Hebert, D.J., Curry, M.D., 2022. Optimal lockdowns. Public Choice 193 (3–4), 263–274.
Hegland, A., Zhang, A.L., Zichettella, B., Pasek, J., 2022. A partisan pandemic: How COVID-19 was primed for polarization. Ann. Am. Acad. Political Soc. Sci. 700 (1), 55–72.
Im, K.S., Pesaran, M., Shin, Y., 2003. Testing for unit roots in heterogeneous panels. J. Econometrics 115 (1), 53–74.
Jelnov, A., Jelnov, P., 2022. Vaccination policy and trust. Econ. Model. 108, 105773.
Karadimas, P., 2023. Public choice theory: An explanation of the pandemic policy responses. In: The Covid-19 Pandemic: A Public Choice View. Springer Nature Switzerland, Cham, pp. 97–132.
Kermack, W.O., McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. Proc. R. Soc. Lond. 115 (772), 700–721.
Kurrild-Klitgaard, P., 2024. Size isn't everything: COVID-19 and the role of government. Public Choice.
Levin, A., Lin, C.-F., James Chu, C.-S., 2002. Unit root tests in panel data: Asymptotic and finite-sample properties. J. Econometrics 108 (1), 1–24.
Lewkowicz, J., Woźniak, M., Wrzesiński, M., 2022. COVID-19 and erosion of democracy. Econ. Model. 106, 105682.
Lourenço, N., Rua, A., 2021. The Daily Economic Indicator: tracking economic activity daily during the lockdown. Econ. Model. 100, 105500.
Ludvigson, S.C., Ma, S., Ng, S., 2021. COVID-19 and the costs of deadly disasters. AEA Pap. Proc. 111, 366–370.
Mathieu, E., Ritchie, H., Ortiz-Ospina, E., Roser, M., Hasell, J., Appel, C., Giattino, C., Rodés-Guirao, L., 2021. A global database of COVID-19 vaccinations. Nat. Hum. Behav. 5, 947–953.
Menezes, F., Figer, V., Jardim, F., Medeiros, P., 2022. A near real-time economic activity tracker for the Brazilian economy during the COVID-19 pandemic. Econ. Model. 112, 105851.
Moser, C., Yared, P., 2022. Pandemic lockdown: The role of government commitment. Rev. Econ. Dyn. 46, 27–50.
Murray, E.J., 2020. Epidemiology's time of need: COVID-19 calls for epidemic-related economics. J. Econ. Perspect. 34 (4), 105–120.
Oraby, T., Tyshenko, M., Maldonado, J., Vatcheva, K., Elsaadany, S., Alali, W., Longenecker, J., Al-Zoughool, M., 2021. Modeling the effect of lockdown timing as a COVID-19 control measure in countries with differing social contacts. Sci. Rep. 11 (1).
Palomino, J.C., Rodríguez, J.G., Sebastian, R., 2020. Wage inequality and poverty effects of lockdown and social distancing in europe. Eur. Econ. Rev. 129, 103564.
Pangallo, M., Aleta, A., del Rio-Chanona, R.M., Pichler, A., Martín-Corral, D., Chinazzi, M., Lafond, F., Ajelli, M., Moro, E., Moreno, Y., Vespignani, A., Farmer, J.D., 2024. The unequal effects of the health–economy trade-off during the COVID-19 pandemic. Nat. Hum. Behav. 8, 264–275.
Persson, T., Tabellini, G., 2002. Political Economics: Explaining Economic Policy. The MIT Press.
Prieur, F., Ruan, W., Zou, B., 2024. Optimal lockdown and vaccination policies to contain the spread of a mutating infectious disease. Econom. Theory 77, 75–126.
Swaziek, Z., Wozniak, A., 2020. Disparities old and new in US mental health during the COVID-19 pandemic. Fiscal Stud. 41 (3), 709–732.
Vavra, J., 2021. Tracking the pandemic in real time: Administrative micro data in business cycles enters the spotlight. J. Econ. Perspect. 35 (3), 47–66.
Walmsley, T., Rose, A., John, R., Wei, D., Hlávka, J.P., Machado, J., Byrd, K., 2023. Macroeconomic consequences of the COVID-19 pandemic. Econ. Model. 120, 106147.

Petar Stankov is a Senior Lecturer in Economics at Royal Holloway, University of London, holding a Ph.D. from CERGE-EI in Prague. He has published on topics such as populism, financial crises and economic reforms. Stankov has taught courses including Political Economy, Institutional Economics, Economic Growth, and Monetary Economics, among many others. He is a Senior Fellow of the HEA, a Graduate Teaching Fellow of the CERGE-EI Foundation, and an Associate of the Economics Network. He also leads the undergraduate program and is a founding co-director of the Centre for Research in Economics Education at Royal Holloway's Economics Department.