Journal of Econometrics 220 (2021) 86–105

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Consumer panic in the COVID-19 pandemic

Michael Keane, Timothy Neal *

University of New South Wales & CEPAR, Sydney, NSW 2000, Australia

ARTICLE INFO

Article history:
Received 9 May 2020
Received in revised form 22 July 2020
Accepted 27 July 2020
Available online 25 August 2020

JEL classification:
C23
C51
D12
M31

Keywords:
Coronavirus
Hoarding
Consumption
Panel data
Panic buying

ABSTRACT

We develop an econometric model of consumer panic (or panic buying) during the COVID-19 pandemic. Using Google search data on relevant keywords, we construct a daily index of consumer panic for 54 countries from January 1st to April 30th 2020. We also assemble data on government policy announcements and daily COVID-19 cases for all countries. Our panic index reveals widespread consumer panic in most countries, primarily during March, but with significant variation in the timing and severity of panic between countries. Our model implies that both domestic and world virus transmission contribute significantly to consumer panic. But government policy is also important: Internal movement restrictions – whether announced by domestic or foreign governments – generate substantial short run panic that largely vanishes in a week to ten days. Internal movement restrictions announced early in the pandemic generated more panic than those announced later. Stimulus announcements had smaller impacts, and travel restrictions do not appear to generate consumer panic.

1. Introduction

We present a model of consumer panic during the COVID-19 pandemic. Panic buying of storable consumer goods is a common phenomenon during natural disasters and man-made crises. Examples include both World Wars (Hughes 1988), the Great East Japan Earthquake in 2011,[1] and the hyperinflation in Zimbabwe (Musvanhiri 2017). It may be provoked by much less: In the United States, Johnny Carson joked about a shortage of toilet paper on *The Tonight Show* in 1973, leading to panic buying and actual shortages (Malcolm 1974). Thus, it is not surprising that the COVID-19 pandemic has caused consumers in many countries to engage in panic buying of storable consumer goods like toilet paper, rice and pasta — see e.g. Knoll (2020) and Rieder (2020). Here we develop a predictive model of how government policies such as social distancing, lockdowns and travel restrictions, as well as growth in COVID-19 cases, generate such behaviour.

To clarify the discussion, it is necessary to define "panic buying". First, we need to understand why consumers hold inventories of storable consumer goods in normal times. Erdem et al. (2003) – henceforth EIK – estimated a structural model of optimal consumer demand for storable goods in a stationary environment. Consumers have two motives for building up inventories in excess of current consumption needs: (i) as a buffer to protect against stock outs given uncertainty about future usage needs, and (ii) it is optimal to stock up on storable goods when confronted with a "deal" (i.e., an instance when the good is offered by retailers at a relatively low price).

In natural disasters or crises, however, consumers are commonly observed to stock up on consumer goods to an extent that greatly exceeds levels observed in normal times. During the COVID-19 pandemic, IRI (2020) document a sharp spike

* Corresponding author.
E-mail addresses: m.keane@unsw.edu.au (M. Keane), timothy.neal@unsw.edu.au (T. Neal).

[1] See Hori and Iwamoto (2014), Ishida et al. (2013), and Kurihara et al. (2012) .

https://doi.org/10.1016/j.jeconom.2020.07.045
0304-4076/

in grocery spending, and increases in stockpiling, in several European countries as well as the US during mid-March as the crisis intensified. For the week ending March 15, spending on paper products (including toilet paper) was up 50% (year-on-year) in Italy, 108% in France, 109% in Germany, 134% in the UK, 210% in Spain and 217% in the US. These countries also saw sharp increases in sales of packaged necessities, such as pasta, rice and flour, as well as fresh fruit and vegetables. By April these sales spikes had greatly subsided.

There are psychological and economic explanations for such stockpiling behaviour in a crisis. A common psychological explanation is that stockpiling storable goods gives consumers a sense of control over the risky situation created by a crisis (Grohol 2020). But as noted by Hansman et al. (2020), there is nothing intrinsically irrational or "panicky" about stocking up on storable consumer goods in a crisis. There are two main economic explanations for a jump in optimal inventory holdings in a crisis situation. First, in the EIK inventory model, any potential supply disruption that increases stock out risk, or any restriction on movement that increases the cost of store visits, will have the effect of increasing optimal inventory.[2] Second, as emphasized by Hansman et al. (2020), a crisis often leads to higher expected future prices, making the current price look like a "deal" that calls for stocking up as optimal behaviour.[3]

Our primary goal in this paper is to develop a predictive model of how government policies impact on panic buying, so we need not take a stand on whether it is driven by psychological or economic factors. Regardless of its cause, the phenomenon of panic buying is socially costly for several reasons. A crisis-induced jump in consumer demand above expected historical levels may often lead to retail store stocks-outs in the short run. This is especially true in the case of just-in-time supply chains where little inventory is available to handle jumps in demand. During the COVID-19 pandemic, serious stock-out situations have arisen in many countries for consumer staples like toilet paper, rice and pasta. Stock-outs are costly for consumers in general, and particularly costly for vulnerable groups like the elderly and the disabled for whom shopping can be challenging.

Shortages created by panic buying also force consumers to devote extra time and effort to shopping, diverting time from welfare-improving activities like work, leisure, and sleep, and generating psychic costs of anxiety and stress. Shortages may heighten anxiety about the pandemic – and the government's response – among the general population. Furthermore, as stressed by Hansman et al. (2020), reputable retailers avoid price increases during crises, both due to legal constraints and long run reputational concerns. This can lead to speculative buying and the emergence of black markets. Moreover, consumer panic may directly promote virus transmission, by causing people to flock to the supermarket before the onset of a lockdown.

We develop a measure of consumer panic for 54 countries over the first four months of 2020 using Google search data. Our measure shows consumers in most large economies experienced panic in response to the COVID-19 pandemic. Much of the panic occurs in March, consistent with IRI (2020) data showing sharp increases in March grocery sales. We find heterogeneity in the timing and severity of consumer panic across countries. Some, such as Australia and the US, experienced more panic than others and some panicked earlier while others much later. In general, panic appeared earlier in Asia than the rest of the world, and richer countries tended to panic more than poorer ones.

It is useful to consider how government policy may contribute to consumer panic. The policy response to COVID-19 has focused on containment of viral transmission and efforts to prop up the economy. We construct three daily measures of government policy during the pandemic: 'internal restrictions' which curtail freedom of movement and association within a city or greater region, 'travel restrictions' which limit or prevent people from entering the country, and 'stimulus announcements' (i.e., fiscal or monetary measures).

Our measures show that most policy changes occurred between the 13th and 24th of March, but there is substantial heterogeneity across countries. Some, like Brazil and South Korea, did not impose severe restrictions on movement, while others like Spain and Peru imposed strict lockdowns. Countries such as Italy and Norway imposed internal restrictions 'early' (relative to other countries), while others like Singapore, Mexico and India imposed restrictions quite late. Some like the US and Canada allowed states/provinces to gradually implement restrictions (or not), while many others such as France and Argentina announced lockdowns at the federal level.

Using this data, we develop a dynamic model of consumer panic as a function of policy change and the spread of COVID-19. Our model features domestic policy changes, the average change of policy in other countries, a nonlinear specification of domestic and overseas COVID-19 case increases, and interaction terms that allow for the effect of domestic policy changes to vary by the average state of policy overseas.

We highlight some key findings: First, the announcement of internal movement restrictions clearly generates increased consumer panic, and the magnitude of the effect is large. Stimulus announcements have smaller effects, and we find no evidence that travel restrictions generate panic. Second, consumers are also very sensitive to the announcement of internal restrictions overseas. Third, the context of domestic policy announcements matters: if internal restrictions are announced

[2] In the EIK model, cost of a store visit and stock-out risk are two key parameters driving optimal inventory: Optimal inventory is increasing in the cost of a store visit (i.e., if store visits were costless, one could buy consumer goods on a just-in-time basis, keeping inventories near zero). The COVID-19 pandemic increased the cost of a store visit in three ways: (i) shopping may lead to contact with infected customers, (ii) infected consumers must quarantine at home, making store visits impossible for a time, (iii) government lockdown policies may make going to the store more difficult. Optimal inventory is also increasing in stock-out risk. If a crisis raises stock-out risk – either due to supply disruptions or a crisis-induced jump in demand – it also leads to higher optimal inventory.

[3] The expectation of future price increases may even lead to speculative buying, where some people attempt to buy up inventory for subsequent resale at a higher price.

early (relative to other countries) it has more of an effect on panic than if they are announced later. Lastly, both domestic and overseas COVID-19 cases affect domestic consumer panic (in a nonlinear manner).

Thus, the announcement of internal movement restrictions can be expected to induce panic buying in the short-run, particularly if the policy is announced before similar measures in other countries. We find the panic response is sudden, strong, and dissipates in a week to ten days. Of course, as the goal of internal movement restrictions is to contain virus transmission, early implementation is still likely to be optimal despite negative short-run consequences. Nevertheless, our results suggest that measures to prevent shortages in the face of consumer panic, such as rationing or priority access to essential goods for vulnerable groups, ought to be pursued concurrently.

As travel restrictions have little effect on consumer panic, implementing them early may have been a good strategy, assuming they are effective in suppressing virus transmission.[4] Further research is needed on the effects of alternative policies on virus transmission and unemployment, to determine the optimal policy mix that achieves suppression of virus transmission while minimizing economic disruption and consumer panic.

The outline of the paper is as follows. Section 2 describes how we construct measures of panic and government policy during the COVID-19 pandemic, and describes key features of the data. Section 3 describes the econometric methods we use to model consumer panic. Section 4 presents our results, and Section 5 concludes.

2. Measuring panic and policy during the COVID-19 pandemic

2.1. Measuring consumer panic

Consumer panic by its very nature is subject to sudden daily changes. And the policy and virus transmission variables whose impact on panic we seek to analyse varied at high frequency during the pandemic. Thus, to model the dynamics of the panic process, we need a daily measure of consumer panic. We also require a measure that is available for many countries, in order to exploit heterogeneity in virus transmission and policy response. While supermarket scanner data on grocery sales is available on a daily basis for some countries – see IRI (2020) – it is not available for most of the countries we wish to include in the analysis. Accordingly, in this article we construct a high frequency measure of consumer panic using Google search data.

Choi and Varian (2012) show Google search data is useful in "now-casting" economic variables such as consumer confidence and unemployment claims. Da et al. (2011) use it to obtain a measure of investor attention that predicts demand for stocks, while Goel et al. (2010) use it to predict near-future consumer behaviour.[5] Our use of Google search data to develop a measure consumer panic is a simple extension of such ideas.

To construct an index of consumer panic during the COVID-19 pandemic, we seek a set of keywords or phrases that can now-cast sharp increases in demand for supermarket goods indicative of 'panic buying' and stockpiling. The first set of keywords we collect are search terms consumers are more likely to use if they are concerned about or aware of panic buying behaviour at supermarkets (either first-hand or through media reports). These are 'panic buying', 'panic', 'hoarding', and 'supermarket'.

We also consider keywords for specific product categories in which the (IRI, 2020) data show a surge in demand in mid-March, indicative of stockpiling or panic buying, for the several countries where that data is available. These are 'toilet paper' 'pasta', 'rice', 'flour', 'vegetables', and 'fruit'.[6] Consumers are more likely to search these terms online when there is increased demand for the product, and also when it is impossible to source the product from their local supermarket due to empty shelves.

The daily search data is collected through the Google Health Trends API, which provides the proportion of searches within a country/day matching any specified word or phrase. The data give the probability that a 'short search-session,' defined as a few consecutive searches by a single user, includes a search for that word/phrase. By reporting probabilities, the data adjusts for differences in population and Google usage between countries. The data are based on a random sample of Google web searches, and are available for most countries, with days defined by the UTC timezone.

We assemble the search data for 54 countries, including all major economies save for Mainland China,[7] covering the period from January 1st to April 31st, 2020 ($T = 122$).[8] For countries that are not majority English-speaking, we translated

---

4 Chinazzi et al. (2020) study effects of travel restrictions on virus transmission, while Fang et al. (2020) study effects of internal movement restrictions. The former argue that travel restrictions have modest effects unless combined with transmission-reduction interventions (e.g., internal movement restrictions).

5 Google search data has also been used in other sciences: Ginsberg et al. (2009) use it to predict influenza outbreaks, and Lampos et al. (2020) use it to estimate COVID-19 cases.

6 The demand for health care products such as face masks and hand sanitizer also surged with the onset of the pandemic. We choose not to include these products as their *usage* also increases dramatically during a pandemic, so a surge in demand is not necessarily indicative of stockpiling/hoarding behaviour.

7 We cannot include Mainland China as the government restricts the use of the Google search engine.

8 The countries are: Argentina, Australia, Austria, Belgium, Bolivia, Brazil, Canada, Chile, Colombia, Denmark, the Dominican Republic, Ecuador, Egypt, Finland, France, Germany, Greece, Guatemala, Hong Kong SAR, India, Iran, Ireland, Italy, Japan, Kenya, Malaysia, Mexico, Morocco, the Netherlands, New Zealand, Norway, Panama, Peru, the Philippines, Poland, Portugal, Qatar, Romania, Russia, Saudi Arabia, Singapore, South Africa, South Korea, Spain, Sweden, Switzerland, Taiwan, Thailand, Turkey, the United Arab Emirates, the United Kingdom, the United States, Uruguay, and Vietnam.

Table 1
Descriptive statistics of Google trends keywords.

Search term	Mean	Median	Max	Std. Dev.	CV
Toilet Paper	0.009	0.003	0.985	0.028	3.293
Panic	0.010	0.007	0.098	0.011	1.143
Panic Buying	0.002	0.000	0.140	0.005	2.506
Hoarding	0.002	0.001	0.065	0.003	1.520
Supermarket	0.029	0.019	0.239	0.028	0.981
Pasta	0.047	0.033	0.616	0.050	1.066
Rice	0.070	0.056	0.376	0.057	0.814
Flour	0.031	0.021	0.263	0.033	1.039
Vegetables	0.025	0.020	0.210	0.021	0.851
Fruit	0.039	0.037	0.220	0.023	0.589

Note: CV refers to the Coefficient of Variation.

Table 2
Factor analysis statistics.

	Pooled		Pooled	Country-Specific	
	Search factor	Panic factor	Panic factor	Mean	Std. Dev.
Eigenvalues	2.504	0.897	0.904	1.621	
Factor Loadings:					
Toilet Paper	0.031	0.395	0.395	0.600	0.248
Panic	0.067	0.442	0.449	0.519	0.263
Panic Buying	0.024	0.535	0.531	0.596	0.313
Hoarding	0.041	0.450	0.452	0.383	0.318
Supermarket	0.447	0.205	0.247	0.358	0.237
Pasta	0.602	−0.046			
Rice	0.749	−0.070			
Flour	0.850	−0.025			
Vegetables	0.532	−0.021			
Fruit	0.605	−0.085			
Factor Scores:					
Toilet Paper	0.006	0.211	0.211	0.269	0.143
Panic	0.018	0.244	0.250	0.223	0.138
Panic Buying	0.012	0.319	0.317	0.289	0.176
Hoarding	0.012	0.250	0.251	0.137	0.121
Supermarket	0.090	0.130	0.126	0.131	0.134
Pasta	0.132	−0.026			
Rice	0.243	−0.060			
Flour	0.440	−0.010			
Vegetables	0.119	−0.007			
Fruit	0.145	−0.055			
Mean	1.354	0.702	0.703	1.305	
Std. Dev.	0.838	0.660	0.662	0.990	

Note: The eigenvalue for Factor 3 when all 10 terms are included is 0.048, and the eigenvalue for Factor 2 when 5 terms are included is 0.001.

our ten search terms into 23 languages in consultation with native speaking colleagues. Example translations can be found in Table 8 in the Online Appendix. If a country has multiple languages we add up the search probabilities for all major languages. For instance, for Switzerland, we add up the probabilities of the German, French, Italian and English search terms. Table 1 reports descriptive statistics on search frequencies for each search term.

We use factor analysis to determine the best way to combine the ten search terms into a coherent index of panic buying. Using the pooled data for all countries, we construct the correlation matrix of daily movements in the ten search terms.[9] The results from factor analysis of the correlation matrix are reported in Table 2. The data is very well described by a two-factor model, as the eigenvalue of the 3rd factor is essentially zero. Interestingly, toilet paper and the panic related search terms (panic, panic buying, hoarding) load on one factor, which we label the 'panic' factor, while the other consumer goods (pasta, rice, etc.) load on a different factor, which we label the 'search' factor (for reasons we make clear below). The term 'supermarket' has a non-negligible loading on both factors.

9 Before calculating the correlation matrix we remove country specific means of the probabilities of each search term. This adjusts for baseline differences in search frequencies across countries.

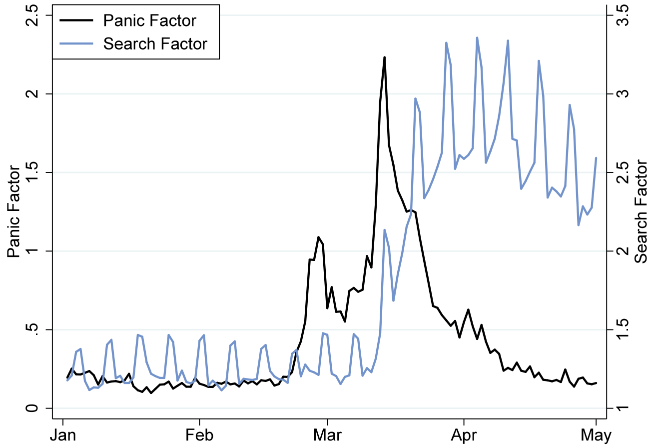

Fig. 1. Factor Analysis on Google Search Terms. Note: This figure presents the average of the 'Panic Factor' and 'Search Factor' for the United States, the United Kingdom, Germany, France, Italy, and Spain.

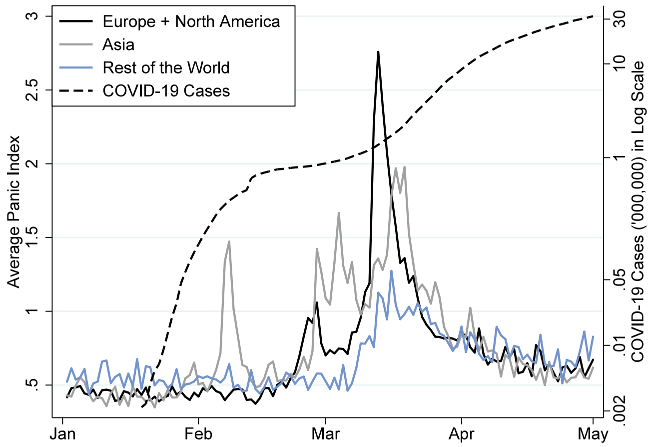

Fig. 2. Average Panic and Worldwide Cases of COVID-19. Note: This figure plots the average panic factor across North America and European Countries, Asian countries (including Oceania), and the Rest of the World. On the secondary *y*-axis is the global confirmed cases of COVID-19 in log scale.

We use the factor scores to optimally combine our ten variables into estimates of the two factors for each country and each day. In Fig. 1 we plot the two factors averaged across the subset of countries analysed in the (IRI, 2020) data. The 'panic' factor closely tracks the timing of stockpiling behaviour described in the IRI data. Stockpiling begins in late February, and spikes tremendously in mid-March. It then subsides rather quickly, and sales return to almost normal levels by early April.

In contrast, the 'search' factor behaves very differently. It rises gradually in mid-March, and peaks in early April, at a high level that is sustained through April. Interestingly, the IRI data shows a *sustained* increase in eCommerce (online shopping activity) starting in mid-March and sustained through April (and beyond). The most plausible explanation of the behaviour of the 'search' factor is that it captures increased online search activity for consumer goods, generated by the sustained increase in eCommerce induced by lockdowns.[10] Thus, we conclude that the 'panic' factor is more representative of the consumer panic/stockpiling phenomenon that we are trying to measure.

To construct the panic index for our analysis, we run a factor analysis on the correlation matrix of the five variables we include (i.e., toilet paper and panic related search terms). The third column of Table 2 reports the results. The data is well described by a 1-factor model. This suggests it is sensible to combine the five variables into an index of the single underlying construct that causes them to co-move, which we label the 'panic index.' The scoring coefficients give the optimal weights for combining the variables.

Fig. 2 plots the panic index against log global COVID-19 cases. We group countries into three regions: Europe and North America, Asia and Oceania, and Rest of the World, and show the average index for each region. The first panic occurred

[10] The 'search' factor also exhibits weekend spikes, indicative of more eCommerce activity on weekends.

in Singapore and Taiwan from February 6th to 10th, causing the panic index for Asia to triple its normal level. During this time COVID-19 cases in China climbed from 20,000 to over 40,000, and cases were starting to spread in other East Asian countries. Panic quickly returned to normal levels in Asia until February 28th, when it jumped again and remained high during the first week of March. By this point COVID-19 cases in China had plateaued above 80,000, but cases outside China were growing quickly, particularly in Italy.

After a brief spike in late-February, many countries in North America and Europe began to experience high levels of panic in mid-March, with the average index increasing from its normal level of about 0.40 to a peak of about 2.80 on March 11th. Consistent with this, the (IRI, 2020) data show a massive spike supermarket sales in the week ending March 15. In the second and third weeks of March, reported COVID-19 cases were growing exponentially in many countries of the world, and unprecedented policy responses were being announced. The remaining countries in the sample (in Central and South America, Africa, and the Middle East) experienced a smaller surge in panic in mid-March. By mid-April, the panic index returns to near normal levels.

There was also significant heterogeneity in the panic index across countries within regions.[11] This leads to the question of why panic was more severe in some countries than others, and why some countries panicked earlier and others later. The differential spread of COVID-19 across countries, and their diverse range of policy responses, provide an excellent opportunity to study how different factors drive the spread of consumer panic.

2.2. Measuring government policy

The pandemic has induced strong responses from governments: Drastic changes to the functioning of society have often been announced only one or a few days before implementation. Because of this we measure policy changes at their announcement. As our outcome of interest is consumer panic, which as we have seen is subject to sudden daily changes, it is important to track policy at a daily frequency as well.

We categorize policy announcements into three broad types: (1) 'internal restrictions' which curtail freedom of movement and association within a city or greater region, (2) 'travel restrictions' which limit or prevent people from entering the country, and (3) 'stimulus announcements' which are fiscal and monetary policies. We now explain how we code the policy announcement variables, noting it must involve some subjectivity:

We define the index for 'internal restrictions' as $Internal_{ct} = Schools_{ct} + Gatherings_{ct} + Movement_{ct}$, where $Schools_{ct} = 1$ if there is a federal closure of primary and secondary schools in country c on day t, and 0 otherwise, $Gatherings_{ct} = 1$ if there is a ban on large gatherings (more than 500 people), 2 if there is a ban on smaller gatherings (50 to 500 people), and 0 otherwise. $Movement_{ct} = 1$ if the government strongly encourages work from home where feasible, there are heavy restrictions on the use of public spaces, and most retail and entertainment businesses are closed. $Movement_{ct} = 2$ if, in addition, many non-essential industries are shut down and a majority of individuals are prevented from working, and 0 otherwise. We occasionally add or deduct a half point if a policy position is difficult to categorize (e.g., if some State governments within a country adopt a policy but not others). $Internal_{ct}$ can range from 0 to 5.

Our index of 'travel restrictions' is $Travel_{ct} = (ChinaBan_{ct} + IranBan_{ct} + ItalyBan_{ct} + SouthKoreaBan_{ct})/2 + Noncitizens_{ct} + Citizens_{ct}$ where $ChinaBan_{ct} = 1$ if travellers from China must self-isolate for 14 days, 2 if they are banned from entering entirely, and 0 otherwise. Likewise for $IranBan_{ct}$, $ItalyBan_{ct}$, and $SouthKoreaBan_{ct}$. These are the four countries most affected early in the pandemic. $Noncitizens_{ct} = 1$ if non-citizen arrivals must self-isolate for 14 days, 2 if non-citizens are banned from entering entirely, and 0 otherwise. Likewise, $Citizens_{ct} = 1$ if all citizen entrants must self-isolate for 14 days, 2 if citizens are effectively prevented from re-entering, and 0 otherwise. Lastly, our index for stimulus announcements $Stim_{ct}$ is the sum of all announcements of significant changes in fiscal or monetary policy, collected from an array of media websites.

Fig. 3 describes the evolution of our policy measures over the sample period, averaging values for the 54 countries across the same three regions as in Fig. 2. Looking at internal restrictions in the top-left panel, it is clear that most policy change occurs between the 13th and 24th of March, when the average index increased from one to over three. The Asian countries imposed some internal restrictions much earlier, but settled at a lower average level. Travel Restrictions behave similarly, with many countries completely closing their borders around mid-March, while Asian countries tended to implement restrictions sooner. The Asian countries adopted an incremental approach to stimulus, with frequent announcements of small policy changes, while the other regions tended to announce large stimulus packages in Mid-March.

Fig. 4 plots the number of policy announcements over time. The peak frequency is clearly in mid-March. Nevertheless, many internal restriction announcements occurred in early or late March (with some even in early April and February). For travel restrictions, several countries such as Australia and New Zealand gradually ramped up entry restrictions from late January to early March.

An interesting question is whether countries that announced strong containment policies earlier than most others, which we will call 'early countries,' had different consumer panic patterns than those who announced their policies late, which we will call 'late countries.' Fig. 5 reports average internal restrictions and average panic across seven notably early countries and seven notably late countries. The early countries adopted tight internal restrictions by early March,

[11] An Appendix with graphs of the panic index for every country is available on request.

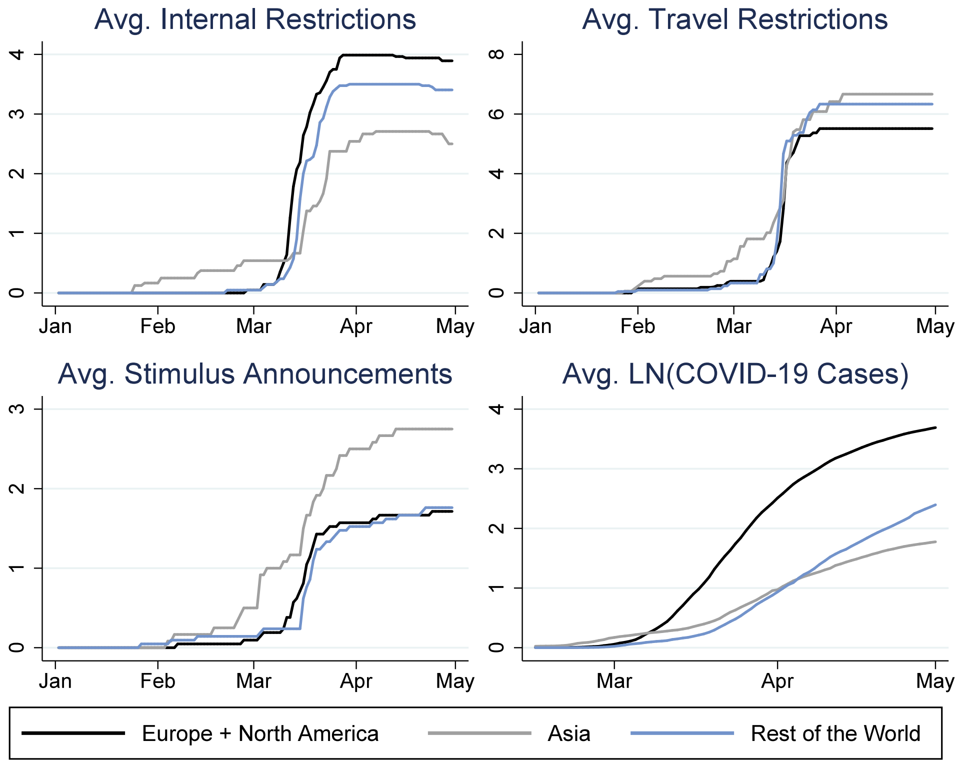

Fig. 3. Average Policy and COVID19 Cases by Region. Note: This figure plots the average value of the three policy variables constructed in this article and log of COVID-19 cases across three regions in the sample of 54 countries.

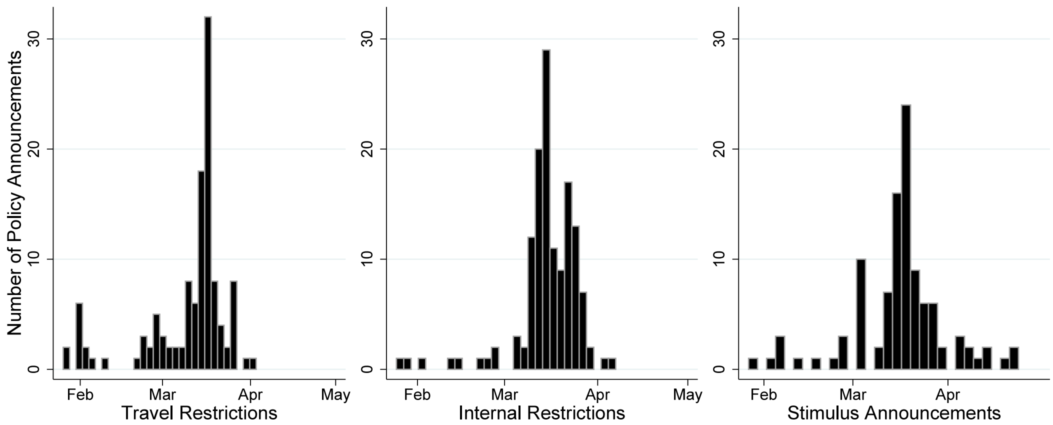

Fig. 4. The Timing of Policy Announcements. Note: This figure shows the number of new policy announcements by day..

and consumer panic tracks the policy index very closely for these countries until the panic decays in late March and April. The late countries did not adopt tight restrictions on movement until late-March. In the late countries panic peaks well prior to the implementation of tight policy measures.

This suggests two key hypotheses about the relationship between internal restrictions and panic: First, that consumer panic reacts not only to domestic policy change but also overseas policy change. Second, that the effect of internal restriction announcements on panic may depend on the international context in which it is announced, with late implementation leading to less consumer panic in the short-run.

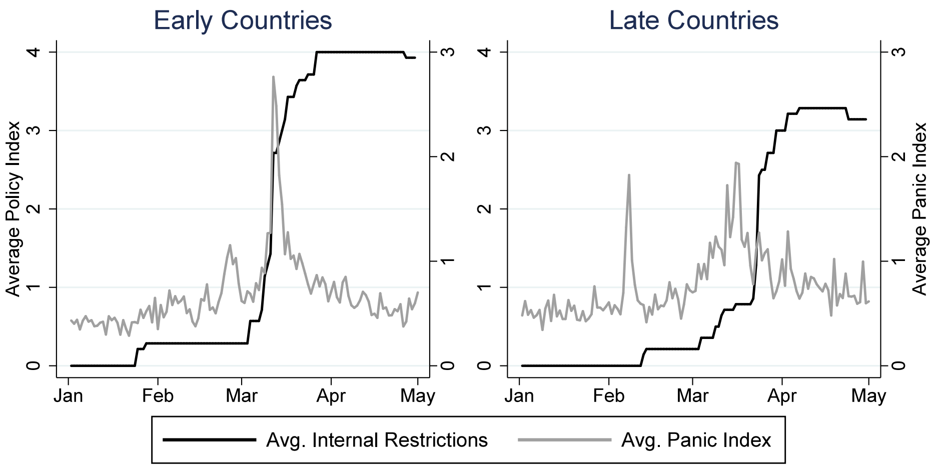

Fig. 5. Average Panic and Internal Restrictions by Group. Note: 'Early Countries' are Austria, Belgium, Ireland, Italy, Hong Kong, Norway and Romania. 'Late Countries' are India, Mexico, Qatar, Singapore, Sweden, the UAE and Vietnam..

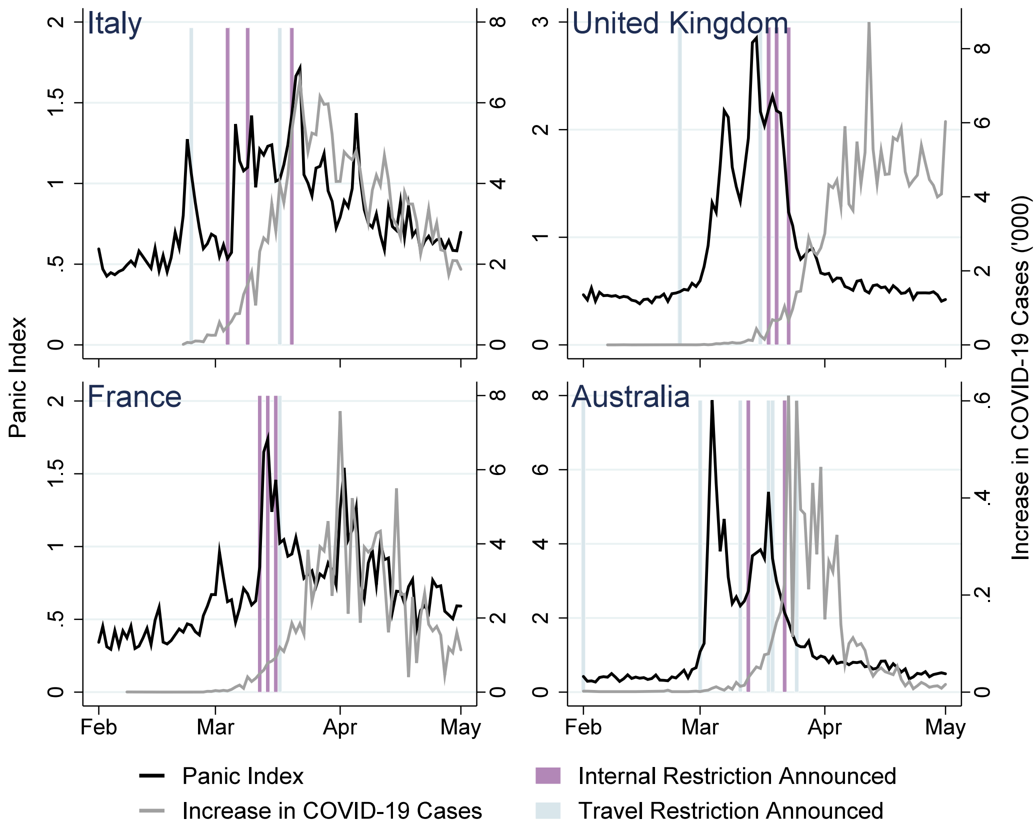

Fig. 6. Panic, COVID-19, and Government Policy in Select Countries.

Finally, Fig. 6 presents the experience of four countries in detail: Italy, the UK, France, and Australia. For each country we show the panic index, the daily increase in COVID-19 cases (thousands), and announcements of internal and travel restrictions.

In Italy there was a brief spike in panic on Feb. 23–24, when outbreaks were reported in northern Italy and several governors imposed interstate travel restrictions. The real surge in panic began in the second week of March, and it climbed until March 22, when our panic index peaked at three times its normal level. In early March the government began implementing restrictions on gatherings, and on movement in the worst affected parts of the country. On March 20 a national lockdown was announced, and soon after that panic reached its peak. Since then, panic has trended down in a way entirely consistent with declines in the number of new cases per day from about 6000 to about 2000.

In France panic surges in close proximity to announcements of internal restrictions. A nationwide lockdown was announced on March 16th, which is the same day that panic reaches its peak of 1.75. The increase also coincides with increases in domestic cases. In contrast, in the United Kingdom panic peaks roughly at same time as in Italy and France even though the virus started spreading much later, and the peak of panic clearly precedes the announcement of domestic restrictions.

Australia is notable for the incredible speed and scale with which panic took hold in early March. The panic index increased from 1.3 on March 2nd to 7.9 on the 4th, which is an order of magnitude higher than the peak in almost any other country. This does not correspond to any significant increase in domestic COVID-19 cases, nor to any important policy announcement.[12] While a travel ban from Iran was announced on the 1st of March, it seems unlikely that this was important enough to be the direct cause. Restrictions on gatherings were announced on March 13th, along with a series of escalating travel restrictions, which coincides with a second peak in the middle of March.

In these four examples, we see a great deal of heterogeneity in the timing and severity of consumer panic, COVID-19 transmission rates, and containment policy. It would be exceedingly difficult to make conclusions on the likely causes of panic surges without a multivariate statistical model of panic, to which we now turn.

3. Methodology

3.1. Conceptual framework

It is useful to present a simple conceptual framework to motivate our empirical specification. To set ideas, let P_{ct} denote the level of consumer panic for a typical consumer in country c on day t of the pandemic. Let S_{ct} denote the severity of the pandemic, and let R_{ct} denote the policy regime in place in country c on day t. Then we may write:

$$P_{ct} = F\left(E(S_{ct}|I_{ct}), E(S_{ct+1}, \ldots |I_{ct}), R_{ct}, E(R_{ct+1}, \ldots |I_{ct}), \mu_c\right) \tag{3.1}$$

Here $E(S_{ct}|I_{ct})$ denotes the *perceived* severity of the pandemic on day t, conditional on the available information set I_{ct}. Similarly, $E(S_{ct+1}, \ldots |I_{ct})$ denotes the expected severity of the pandemic in future periods. μ_c is a range of country-specific factors that affect the propensity to panic (e.g., cultural norms).

Thus, expectations of severity and policy both affect panic. The information set is:

$$I_{ct} = \left(C_{ct}, C_{c,t-1}, \ldots, C_{ft}, C_{f,t-1}, \ldots, R_{ct}, R_{c,t-1}, \ldots, R_{ft}, R_{f,t-1}\right) \tag{3.2}$$

where C_{ct} denotes the COVID-19 caseload in country c on day t, and C_{ft} denotes a vector of caseloads in foreign countries. Similarly, R_{ct} denotes the policy regime in country c on day t, and R_{ft} denotes a vector of policy regimes in foreign countries. Thus, the information set includes current and lagged caseloads both domestic and foreign, as well as current and lagged policies, both domestic and foreign.

In writing (3.1)–(3.2) we assume consumers have incomplete information about the severity of the crisis, which seems obvious given that even epidemiologists had great uncertainty about the virulence and transmissibility of COVID-19. We assume consumers use COVID-19 caseloads, both domestic and foreign, to infer the severity of the pandemic. Consumers also use government policy as a signal of the severity of the pandemic, under the reasonable assumption that governments and the experts they consult have information about the pandemic exceeding that of the typical consumer.[13]

Eq. (3.1) is the "structural" equation for consumer panic. Estimation would require us to make assumptions on (i) how consumers form expectations and (ii) how perceived and expected future severity and policy map into consumer panic. Specifying the structure is complicated by the fact that we do not necessarily believe panic buying is driven by purely economic factors. Thus, we choose to instead work with the "reduced form" equation for consumer panic, obtained by substituting out for the unobserved expectation terms in (3.1) using the information set in (3.2), giving:

$$P_{ct} = f\left(C_{ct}, C_{c,t-1}, \ldots, C_{ft}, C_{f,t-1}, \ldots, R_{ct}, R_{c,t-1}, \ldots, R_{ft}, R_{f,t-1}, \ldots, \mu_c\right) \tag{3.3}$$

This is a valid reduced form if the caseload and policy variables are exogenous to the panic process (an issue we address in Section 4). Of course, without specifying the structure we cannot know the functional form of $f(.)$. So in practice, we will seek an "approximate reduced form" model that approximates (3.3) using a flexible function of the caseload and policy variables, and that passes a variety of stringent specification tests.

In this framework the policy regime R_{ct} affects consumer panic through four channels. First, there is the direct effect of the policy. Second, there is the effect of the policy on expected future policy. Third, there is the effect operating through

[12] Indeed, it is important to note that axes in Fig. 6 are scaled differently for Australia: It has an order of magnitude higher peak panic index and lower peak caseload than the other three countries!

[13] See Avery et al. (2020) for a discussion of alternative epidemiological models, as well the uncertainty about the inputs to those models. Their paper discusses how the caseload projections of the Imperial College model in Ferguson et al. (2020) significantly influenced several leading countries, such as the US and UK, to impose strict social distancing rules after an initial response that was more moderate.

the impact of policy on the perceived current severity of the pandemic $E(S_{ct}|I_{ct})$. Fourth, there is the impact of policy on the expected severity of the pandemic in the future $E(S_{ct+1},\ldots|I_{ct})$.

As a simple example, consider an internal movement restriction that makes it more difficult for consumers to leave home and visit the store. First, as discussed in the introduction, anything that raises the cost of store visits will cause consumers to desire higher inventories. Thus, a new internal restriction may generate "panic" buying (i.e., heavy stockpiling) in the short run. Second, if the policy leads consumers to expect even tighter restrictions in the future, it will encourage even more demand today. Third, if consumers have incomplete information about the severity of the crisis and use government policy as a signal, it is plausible that a new internal movement restriction may cause panic by causing consumers to infer that the crisis is worse than previously thought. Fourth, stronger government action may increase confidence about the future course of the pandemic, reducing fears of future shortages. Through this channel, a new internal movement restriction may reduce consumer panic.[14]

The overall impact of any policy move on consumer panic will depend on the balance of these four forces. And this depends on consumers' information set. Early in the pandemic, when there is great uncertainty about its severity, the third channel is likely to be very strong, so we would not be surprised if strong policy action induces consumer panic. Later in the crisis, when consumers have better information, we hypothesize that this third channel is weaker, so that policy actions will be less likely to induce panic.

The number of COVID-19 cases reported in foreign countries, as well as policy actions of foreign governments, provide signals about the current and likely future severity of the pandemic in one's own country. Policy actions of foreign governments may also provide signals about the likely future policies one's own government will adopt. Thus, foreign government policies R_{ft} and foreign caseloads C_{ft} may also affect domestic panic through channels two through four described above.

A key issue is whether panic induced by domestic policy action depends on timing relative to other countries. We hypothesize that if a country acts early, the second channel (policy action signals a more severe crisis) is likely to dominate, so early action will induce short-run panic. We also hypothesize that if a country acts relatively late, the severity of the pandemic will already be fairly well understood, so this signalling mechanism will be less strong. So we expect later action to induce less consumer panic. To capture such effects, it is important that when we specify the model in (3.3) we allow the policy regime in foreign countries to moderate the impact of domestic policy.

3.2. Reduced form model of consumer panic

Guided by this conceptual framework, we proceed to specify a reduced form model of consumer panic. A priori, we decided to model consumer panic as a dynamic process in which the *level* of panic in country c on day t depends on its own lagged level, as well as *changes* in the forcing variables (i.e., domestic and foreign policy regimes, domestic and foreign cases) from day $t-1$ to day t. This is because, based on the data description in Section 2, we expect that changes in policy and COVID-19 cases have strong short-run effects on panic, but that these effects die off quickly. After some experimentation with functional form, we arrived at the following model as our main specification:

$$\begin{aligned} lnP_{ct} &= \mu_c + \sum_{\ell=0}^{1}\left(\boldsymbol{\beta}_\ell \Delta\mathbf{R}_{c,t-\ell} + \boldsymbol{\gamma}_\ell \Delta\bar{\mathbf{R}}_{f,t-\ell} + \boldsymbol{\lambda}_\ell(\Delta\mathbf{R}_{c,t-\ell} * \bar{\mathbf{R}}_{f,t-\ell})\right) \\ &+ \theta_1 \Delta ln[C_{ct}+\alpha] + \theta_2 B(\Delta C_{ct}|\lambda) + \theta_3 \Delta ln[C_{ct}+\alpha]\cdot B(\Delta C_{ct}|\lambda) \\ &+ \theta_4 \Delta ln[C_{ft}+\alpha] + \theta_5 B(\Delta C_{ft}|\lambda) + \theta_6 \Delta C_{CN,t} + \theta_7 C_{CN,t}\cdot I_{c\in Asia} \\ &+ \sum_{\ell=1}^{2}\rho_\ell lnP_{c,t-\ell} + \boldsymbol{\psi}\mathbf{day}_t + e_{ct} \end{aligned} \tag{3.4}$$

Here lnP_{ct} is the log of the panic index for country c on day t, and the model includes two lags of the dependent variable. $\mathbf{R}_{ct} = [Internal_{ct}, Travel_{ct}, Stimulus_{ct}]$ is a vector of our three policy variables. These enter the model in first difference form ($\Delta\mathbf{R}_{ct}$). We include both the current and first-lag daily difference. $\mathbf{R}_{ft}$ is a vector of the three policy variables for foreign countries, measured as the simple average across countries outside of country c.[15] We also enter this variable in first-differenced form, including one lag.

The term $\boldsymbol{\lambda}_\ell(\Delta\mathbf{R}_{c,t-\ell} * \mathbf{R}_{f,t-\ell})$ allows the effect of domestic policy changes to vary depending on the international context, captured by the international average $\mathbf{R}_{ft}$. If $\lambda_{p\ell} < 0$ then the effect of announced changes in policy $\Delta\mathbf{R}_{c,t-\ell}$ on

14 The severity of the pandemic at time t will be some function of past severity and past containment policy $S_{ct} = g(S_{ct-1}, R_{ct-1})$. In Section 2.2 we defined the internal restriction policy measure R_{ct} so larger values imply stricter constraints on movement. So it is reasonable to assume that $\partial S_{ct}/\partial R_{ct-1} < 0$. A similar assumption makes sense for travel restrictions. But in the case of stimulus announcements it is not clear why they would affect virus transmission, a point we return to when interpreting the results.

15 We experimented with weighted averages of international policy variables, based on criteria such as distance, GDP, population, and the incremental R^2 when panic in country c is regressed on measures for each foreign country separately. But these refinements had little impact on the results.

panic is reduced the more governments overseas have already announced that change in policy. We hypothesize that $\lambda_{p\ell} < 0$ based on the discussion in Section 3.1.

We let changes in the number of confirmed domestic COVID-19 cases (C_{ct}) enter the model very flexibly. Early in the pandemic, when caseloads are low, large daily percentage increases are sometimes observed. Later in the pandemic, when caseloads are high, daily percentage changes are typically small, but the absolute changes can be large. We have no strong prior on whether consumer panic responds more to the former or latter.

Thus, we include both the terms $\Delta ln[C_{ct} + \alpha]$ and $B(\Delta C_{ct}|\lambda)$. The former is the daily percentage change in domestic cases plus parameter α, which moderates the impact of large percentage changes from low levels, and the latter is the Box–Cox transform of the daily absolute change in domestic cases, with Box–Cox parameter λ. We run a grid search to determine the values of α and λ that maximize R^2, and settle on 50 and 0.30. We also include the interaction term, $\Delta ln[C_{ct} + \alpha] * B(\Delta C_{ct}|\lambda)$, which allows the impact of a given percentage change to be greater if it coincides with a larger absolute change.

We also allow foreign reported cases (C_{ft}) to influence domestic panic. We enter changes in foreign cases in the same flexible way as domestic cases.

We cannot include mainland China in our analysis as it restricts access to Google. However, as the pandemic originated in China, we include the number of confirmed changes in mainland China ($C_{CN,t}$) as a driver of consumer panic in other countries. We also included an interaction between the number of Chinese cases and an indicator for whether country c is in East or Southeast Asia, to allow for the possibility that Chinese cases have a larger effect on consumers in nearby countries. Finally, $\mathbf{day}_t$ is a vector of day-of-the week dummies to capture daily differences in search activity.

The term μ_c is a country specific fixed effect meant to capture differences across countries in the baseline level of Google search activity for panic related terms. We estimate the model by fixed effects, relying on the fact that we have a long panel ($T = 122$) so that the so-called (Nickell, 1981) bias that arises from applying a fixed effects estimator to a lagged dependent variable model in a panel with small T is rendered negligible.[16]

We assume the error e_{ct} satisfies the usual assumptions: serially uncorrelated, homoskedastic, cross-sectionally independent, and independent of the regressors. We found evidence of conditional heteroskedasticity, and applied a weighted least squares procedure described below. Given this, we show in Section 4 that the model passes a stringent set of specification tests, so the testable assumptions on the errors are not rejected.

Finally, to avoid letting the log panic index be too sensitive to days with exceptionally low search activity, we add .01 to the index before taking the log, so we actually work with $ln(P_{ct} + .01)$. This value is below the 1st percentile of the distribution of P_{ct} across all countries and days in the sample. Recall from Table 2 that P_{ct} has a mean of .703.

4. Results

Here we present results from estimating the consumer panic model of Eq. (3.4) on the panel data set described in Section 2. First we present results from a simplified model that includes only domestic policy and COVID-19 case data. Then we present results from our full model that includes international variables. Finally we present results from alternative specifications that (i) allow for panic across some countries to be "related" (in a sense defined later), and (ii) allow for country-specific weights in the panic index.

4.1. Domestic model

The first two columns of Table 3, labelled "Model 1", present a restricted model that includes only domestic cases and policy variables, omitting the foreign variables from Eq. (3.4). Both current and lagged changes in internal restrictions are highly significant. The point estimates imply that if a government increases internal restrictions by 1 on a 0–5 scale (e.g. closing schools or restricting gatherings) it causes the panic index to increase by roughly 13% on the same day, followed by a greater increase on the second day and then a gradual decline. Stimulus announcements and travel restrictions are insignificant. Both the percentage and (Box–Cox transformed) absolute changes in COVID-19 cases are significant determinants of panic, as is their interaction.

The results from this simplified model are suspect however, as the specification tests point to some problems. On the plus side, we find only weak serial correlation in the residuals, suggesting our lag structure is adequate. Additional lags of the dependent and independent variables were not significant here or in the specifications discussed below. However, the (Pesaran, 2015) test overwhelmingly rejects the null hypothesis of "weak" cross-sectional dependence, indicating the residuals contain important shocks that are correlated across countries. And a (Sims, 1972) non-causality test finds that leads of the policy variables are highly significant in the panic equation, thus rejecting the hypothesis that the policy variables are strictly exogenous with respect to consumer panic.

[16] Given large T, consistency of the fixed effects estimator requires that the policy and caseload variables be predetermined — strict exogeneity is not essentially.

Table 3
Models of Panic during the COVID-19 Pandemic (Pooled Factors).

Models of $ln(panic_{c,t})$:	Model 1		Model 2		Model 3	
	β	S.E.	β	S.E.	β	S.E.
Internal Restrictions:						
$\Delta Internal_{c,t}$	0.130	0.025	0.191	0.072	0.178	0.064
$\Delta Internal_{c,t-1}$	0.139	0.030	0.261	0.075	0.210	0.063
$\Delta Internal_{f,t}$			0.457	0.182	0.445	0.176
$\Delta Internal_{f,t-1}$			0.106	0.224	0.170	0.214
$\Delta Internal_{c,t} * Internal_{f,t}$			−0.055	0.029	−0.047	0.026
$\Delta Internal_{c,t-1} * Internal_{f,t-1}$			−0.086	0.030	−0.063	0.027
Stimulus Announcements:						
$\Delta Stim_{c,t}$	0.074	0.045	0.237	0.099	0.177	0.095
$\Delta Stim_{c,t-1}$	0.073	0.044	0.092	0.094	0.075	0.097
$\Delta Stim_{f,t}$			0.111	0.283	0.179	0.276
$\Delta Stim_{f,t-1}$			0.015	0.268	0.000	0.263
$\Delta Stim_{c,t} * Stim_{f,t}$			−0.207	0.080	−0.165	0.076
$\Delta Stim_{c,t-1} * Stim_{f,t-1}$			−0.068	0.075	−0.060	0.072
Travel Restrictions:						
$\Delta Travel_{c,t}$	0.021	0.013	0.034	0.034	0.034	0.035
$\Delta Travel_{c,t-1}$	0.014	0.015	−0.001	0.037	0.009	0.036
$\Delta Travel_{f,t}$			0.025	0.100	0.041	0.098
$\Delta Travel_{f,t-1}$			−0.052	0.081	−0.060	0.081
$\Delta Travel_{c,t} * Travel_{f,t}$			−0.008	0.008	−0.007	0.008
$\Delta Travel_{c,t-1} * Travel_{f,t-1}$			−0.001	0.008	−0.003	0.008
Domestic COVID-19 Cases:						
$\Delta ln(C_{c,t})$	1.381	0.140	0.513	0.147	0.394	0.144
$\Delta C_{c,t}$	0.004	0.001	0.002	0.001	0.002	0.001
$\Delta C_{c,t} * \Delta ln(C_{c,t})$	−0.027	0.009	−0.010	0.008	−0.006	0.009
International COVID-19 Cases:						
$\Delta ln(C_{f,t})$			1.952	0.293	1.891	0.290
$\Delta C_{f,t}$			0.001	0.001	0.003	0.001
$\Delta C_{China,t}$			−0.008	0.007	−0.007	0.006
$\Delta C_{China,t} * Asia_c$			0.049	0.017	0.067	0.017
Panic Terms:						
$ln(panic_{c,t-1})$	0.285	0.012	0.237	0.012	0.205	0.012
$ln(panic_{c,t-2})$	0.242	0.012	0.199	0.012	0.174	0.012
Related Country Panic					0.712	0.058
Diagnostics:						
AR1 Term of $\hat{e}_{c,t}$	−0.038	0.012	−0.011	0.012	−0.013	0.012
CSD Test (w/ p-value)	13.687	0.000	2.923	0.003	0.809	0.418
Sims Test (w/ p-value)	6.907	0.000	0.515	0.865	0.448	0.909
R^2	0.381		0.404		0.423	

Note: WLS is used for estimation, where $N = 54$ and $T = 119$. The results for additional terms are available in the Online Appendix. The CSD Test refers to the (Pesaran, 2015) test with a null of weak cross-sectional dependence. The Sims Non-Causality Test refers to a F-test on the significance of three leads of the policy variables as suggested by Sims (1972).

4.2. Main specification

We present our full model of (3.4) in the columns of Table 3, labelled "Model 2". This model adds international COVID-19 cases and policy variables. If we estimate the model by applying a within transformation and running OLS (obtaining the fixed effects estimator via the Frisch–Waugh theorem) there is (weak) evidence of conditional heteroskedasticity. We model this heteroskedasticity as shown in Table 6 of the Online Appendix and then apply weighted least squares (WLS).[17] We use WLS to estimate all our models.

In the full model both the current and lagged change in internal restrictions are again highly significant. As we see in Table 3, the average level of internal restrictions in *foreign* countries is also a significant positive determinant of domestic consumer panic, suggesting that domestic consumers use foreign government policy as a signal of the severity of the pandemic and/or to predict their own government's future policy.

[17] The variance estimates indicate the variability of consumer panic increases in the short-run if internal restrictions are tightened. This effect is less if more countries have already adopted internal restrictions. But the R^2 of the variance equation is only 0.012.

Furthermore, the interactions between domestic internal restrictions and the average level of restrictions in foreign countries are negative and significant. This supports our key hypothesis that internal restrictions tend to cause less domestic consumer panic if they are introduced relatively late in the pandemic compared to other countries.

We also find that domestic stimulus announcements have a positive short-run effect on consumer panic, diminishing if more countries have already announced such policies. In addition, both domestic and international percentage changes in COVID-19 cases have significant positive effects on consumer panic. The absolute change in domestic cases also has a positive effect, as does the change in Chinese cases in Asian countries only.

It is difficult to interpret the magnitudes of the policy and caseload effects on consumer panic, due to the interaction terms and the lagged dependent variables. So we defer that discussion until Section 4.4 when we look at impulse response functions.

4.3. Specification tests

Our full model exhibits no statistically or quantitatively significant correlation in the errors. It easily passes a Sims non-causality test, so we cannot reject that the policy and caseload variables are strictly exogenous with respect to consumer panic. The fact this is only true after we include the international caseload and policy variables means these variables can account for the error component that predicts future policy. This is intuitive in the pandemic context, as both foreign caseload and policy changes clearly have strong influences on future domestic policy decisions. So controlling for the international variables purges the residuals of the component that predicts future domestic policy.

The Pesaran cross-sectional dependence test is still significant. But the test statistic, distributed N(0,1), drops from 13.7 to 2.9, so evidence of strong cross-sectional dependence is far weaker than in Model 1. As Pesaran (2015) notes, the test over-rejects when T is large relative to N (as we have here), so it is actually quite impressive that the CSD test statistic for Model 2 is only 2.9. This result is intuitive in the pandemic context, as it seems clear that foreign caseload and policy changes influence domestic consumer panic. Indeed, many countries experience spikes in panic prior to having substantial numbers of confirmed domestic cases. Hence, controlling for the international variables purges the residuals of the component that induces most of the cross-sectional dependence.

Rejecting the null of weak cross-sectional dependence does not necessarily imply the estimates are inconsistent. Of the 1431 pairs of residuals between countries only 14 have a correlation above 0.35, while 80% have a correlation below 0.10. So we doubt the cross-section dependence is strong enough to cause significant bias. Most pairs with non-trivial correlations make intuitive sense: Many are within the Anglophone countries (Australia, Canada, US, UK), or countries with close cultural ties (e.g., Germany/Austria). This suggests we may purge Model 2 of cross-section dependence by allowing the panic of related countries to respond to each other's panic, for the 14 country pairs with high correlations. The results are shown in the columns of Table 3 labelled 'Model 3.' The coefficient on 'related country' panic is highly significant, and the CSD test statistic drops to .81 with a p-value of .42. But the parameter estimates change little from those of Model 2, suggesting cross-sectional dependence is not biasing the Model 2 results.

Finally, we examine the issue of stationarity. Using the (Pesaran, 2007) panel unit-root test in the presence of cross-section dependence, we reject the null of homogeneous non-stationarity (i.e., all countries are non-stationary) for all variables. If we run separate Augmented Dickey–Fuller tests by country, the only variable for which non-stationarity is a concern is the (Box–Cox transformed) absolute change of cases, which shows evidence of a unit root in 76% of countries.[18] Given that all covariates are in difference form, it seems logically impossible that any would have a unit root given large enough T.

4.4. Impulse response functions

Fig. 7 shows impulse response functions for domestic internal movement restrictions, evaluated at the average international level of internal restrictions on three dates: February 21 (early), March 13 (middle) and April 1 (late). If a government increases internal restrictions by one unit (e.g. closing schools or restricting gatherings) on Feb 21 it causes the panic index to increase by approximately 19% on the same day. The effect peaks at 30% on day 2, and then gradually vanishes after about a week to ten days. But if a government delays until March 13 the peak is only 20% on day 2, and if the government delays until April 1st the impulse response function shows no significant effects (although the short-run impact is imprecisely estimated).

Our point estimates imply that stimulus announcements also have a positive short-run effect on consumer panic. Again, the interaction with the international average of stimulus announcements is also significant, implying the effect is lessened if the announcement comes "late" relative to other countries. But the impulse response functions on the right side of Fig. 7 show the effects of stimulus announcements are both weaker and less precisely estimated than those of internal movement restrictions. As was the case with internal restrictions, we find the effect of stimulus announcements on panic diminishes if they come later in the pandemic. We find no statistically or quantitatively significant effects of travel restrictions (impulse response not shown).

[18] The log(panic) variable only shows evidence of non-stationarity in 15% of countries, the change in internal restrictions, 5.5%, the change in travel restrictions, 5.5%, the change in stimulus, 5.5%, and the log change of cases, 11%.

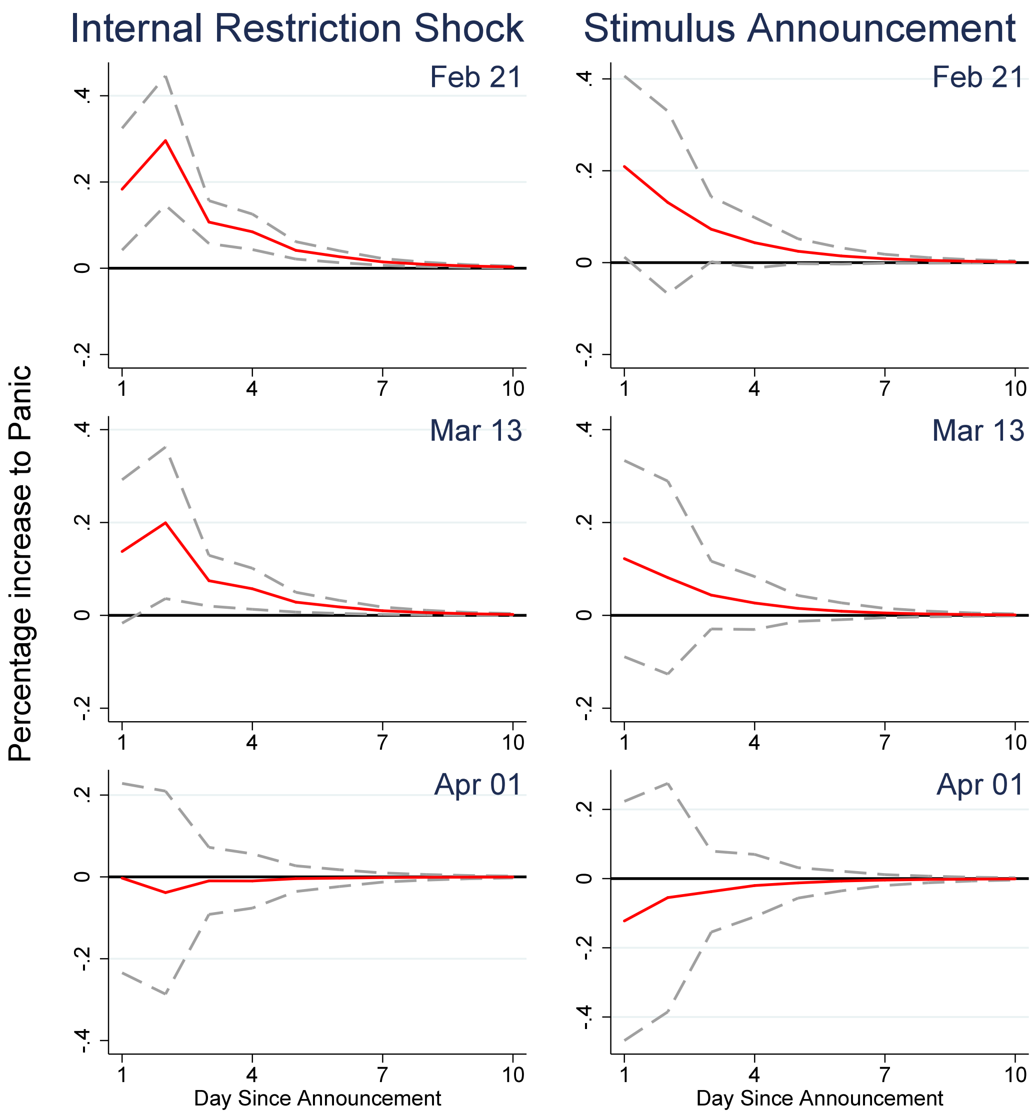

Fig. 7. Impulse Response Functions For Policy Change.

Finally, we consider the role of COVID-19 caseloads. Both percentage and absolute changes in domestic COVID-19 cases are significant and positive drivers of panic, as is the interaction between the two. Given our flexible specification, the overall effect of a given increase in cases depends on both the percentage and absolute change. The impulse response functions in the left panel of Fig. 8 show the effects of average increases of COVID-19 cases in February, March and April. The average (across all countries) of the daily increases were 1% on a base of 15 cases in February, 11% on 2,900 cases in March, and 7.3% on 18,600 cases in April. Thus, both percentage and absolute changes tend to be small in February. The largest percentage increases tend to be in March, and the largest absolute increases tend to be in April.

Our model implies that the typical March daily increase in caseloads would have increased the panic index by about 8% immediately, with an effect that dies off over time, becoming negligible after about a week. Effects of a typical change of caseloads in April are very similar.

International COVID-19 cases are very significant determinants of domestic panic, both statistically and quantitatively. Impulse response functions in the right panel of Fig. 8 show effects of average increases in worldwide COVID-19 cases in February, March and April.[19] Our model implies the typical March daily increase in caseloads increased the panic index by

[19] As we noted earlier, the average (across all countries) of the daily increases are 1% on a base of 15 cases in February, 11% on 2,900 cases in March, and 7.3% on 18,600 cases in April. The international average caseload changes used to construct the right panel of Fig. 8 are the same as the average domestic changes use to construct the left panel.

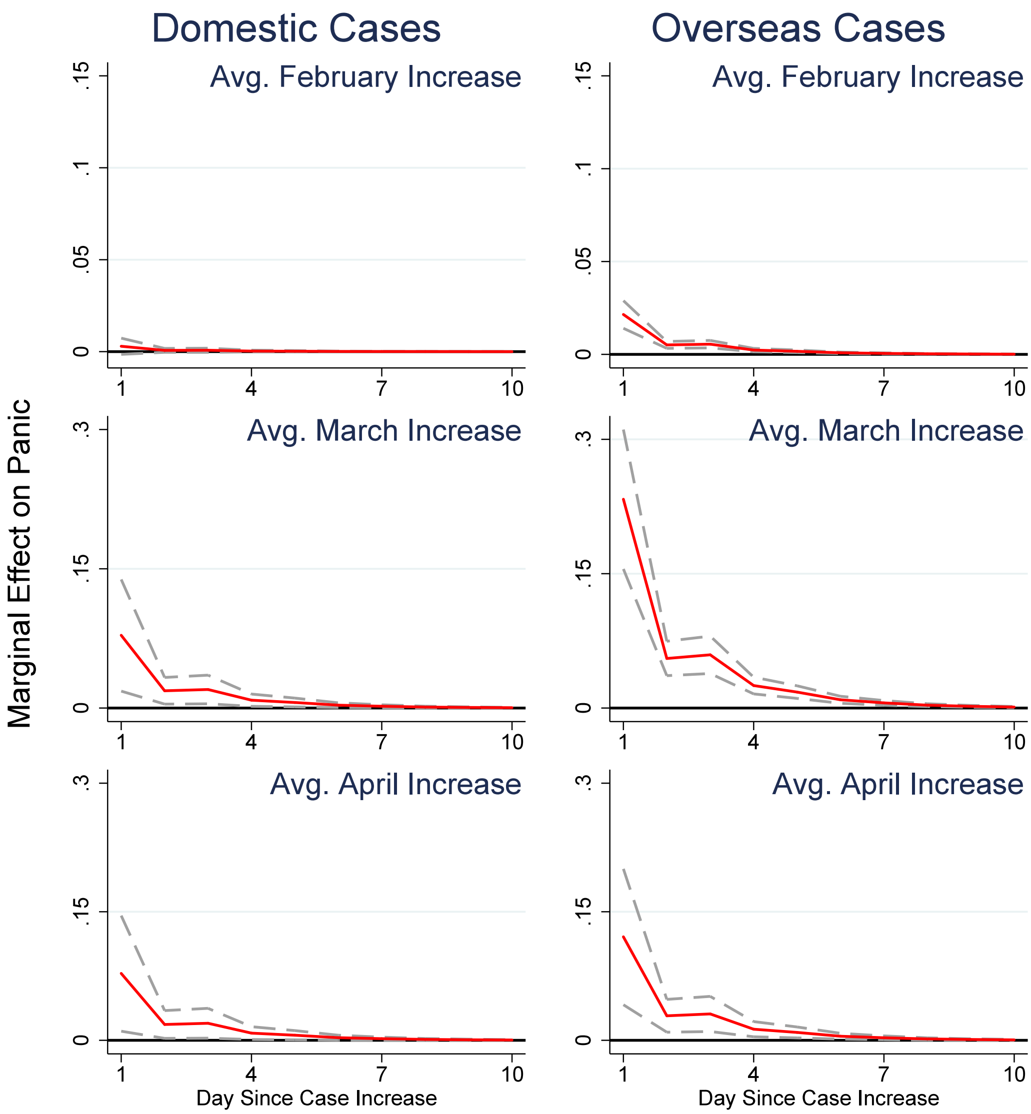

Fig. 8. Impulse Response Functions For COVID-19 Case Increases. Note: The average daily increase in February is a percentage increase of 1% and absolute increase of 2, in March it is a percentage increase of 11% and absolute increase of 314, and in April it is a percentage increase of 5.5% and absolute increase of 1,564.

about 23% immediately, with an effect that dies off over time, becoming negligible after about a week. Effects of typical April changes in caseloads are about 50% smaller. This combined with the fact that most policy announcements were concentrated in March (see Fig. 4) is consistent with the fact that most countries saw very low levels of the panic index in February, following by peaks of the panic index during March, with the index declining through April (see Figs. 2, 5, 6).

Comparing the impulse response functions for domestic vs. international cases, we see that international cases have substantially larger effects on domestic panic than domestic cases. This makes sense, given that large caseloads were concentrated in relatively few countries, while panic was widespread across many countries. Overall, we have found that both international policy regimes and international cases are important drivers of domestic panic, suggesting that consumers do pay attention to international conditions.

4.5. Model fit

Fig. 9 illustrates how the model fits the time series of the panic index data for selected countries. The figure presents fitted values of the daily panic index based on current and lagged values of the forcing variables (i.e., policy variables and

COVID-19 cases). We present both conditional predictions, where the two lags of the panic index in Eq. (3.4) are always set at their true values, and unconditional predictions, where we plug in the two previous day's predicted panic indices from the model. Italy, France and Brazil are examples of countries where the model provides very good predictions of the path of the panic index based on the forcing variables.

Our model fits the UK and US somewhat less well, as it can only generate about half to two-thirds of the peak increase in panic. This is largely due to the fact that in the UK and US our panic index rose to substantially higher levels than in most other countries, and this is not easy to explain based on the forcing variables.[20]

Australia is an example of a country where the panic index had a massive spike on a particular day (March 4) and the model is not able to explain why this occurred. As we noted in Section 2, the increase in the panic index in Australia was an order of magnitude greater than we observe in almost any other country, so it is not surprising that our model cannot explain it. Other countries with hard to explain massive spikes are Japan, Taiwan, Singapore and Germany.[21] It is of course not surprising that a variable like "panic" is sometimes hard to predict based only on observed forcing variables, so we would argue that – with the exception of this set of countries – the model generally fits quite well. (An Appendix showing the fit for all 54 countries is available on request).

We also report results from holding out the April data in estimation, and using the model to forecast the path of panic during April. The red line in Fig. 9 shows the forecasts for the selected countries. The forecasts seem reasonably accurate, except that for the UK our model understates the rate at which panic dropping in April, while for Brazil we predict a decline in early April that did not materialize.

4.6. Country specific panic index

Given differences in language and culture, it seems likely that the optimal weighting of the five Google search terms that enter the panic index may differ by country. Thus, we also tried using country-specific factor analyses to form country specific panic indices. The columns of Table 2 labelled 'Country-Specific' describe the results. Clearly there is a great deal of cross-country heterogeneity in the weightings. We re-ran all our analysis using these country-specific panic indices and report the results in Table 4. The results are remarkably similar to those using homogeneous weights in Table 4. Thus, our results are quite robust to this alternative method of constructing the panic index. The only significant difference is that the estimated effect of the lagged change in international restrictions is greater here, and it becomes significant.

We also tried a simplified version of the panic index that weighs the five search terms equally (rather than using factor scores as weights). That also gave very similar results.

4.7. Discussion

As we discussed in Section 3.1, there are four channels through which internal restrictions may plausibly affect consumer panic, and the overall effect depends on the balance of these four. Recall that internal restrictions: (i) increase the cost of store visits, which increases desired inventories, leading to a short-run spike in demand, (ii) increase expected future costs of store visits, with a similar effect, (iii) increase perceived severity of the pandemic, which may increase consumer panic, and (iv) reduce expected future severity of the pandemic, which may reduce consumer panic. Estimates from a reduced form model cannot disentangle these different channels, but our model implies that the three panic increasing effects dominate until very late March and early April.

Our model also implies that stimulus announcements have a significant positive effect on consumer panic in the short-run, but the point estimates imply this is much weaker than the effect of internal restrictions. A weaker effect seems plausible, as stimulus announcements do not have any plausible effects on the cost of store visits, so channels (i)-(ii) are not operative. Their only positive effect on consumer panic would seem to operate through channel (iii), as they may signal greater severity of the pandemic.

Our model implies that travel restrictions have no significant effect on consumer panic. Again, travel restrictions have no plausible effects on costs of store visits, so channels (i)-(ii) are not operative. Our results suggest that the signalling effect of travel restrictions, operating through channel (iii), is fully counter-balanced by channel (iv), whereby consumers expect that travel restrictions will help mitigate the severity of the pandemic in the future, thus reducing panic.

Another key finding is that both foreign cases and the policy decisions of foreign governments contribute to domestic panic. A model that fails to account for this was found to be seriously misspecified, as it suffered from meaningful cross-sectional dependence and a failure of the strict exogeneity assumption for the domestic policy variables. This is because the foreign caseload and policy variables contribute not only to domestic consumer panic in the short run, but also to future domestic policy decisions.

20 Also, referring back to Fig. 6, one can see that in Italy and France domestic cases began to rise and internal restrictions were announced prior to the large spike in panic in early to mid-March. But in the UK panic increased sharply prior to the emergence of substantial domestic cases or internal restrictions. So the increase in panic in the UK in early March is driven largely by international cases and policies.

21 In addition, the model under-predicts the magnitude of early panic spikes in the Netherlands, US, UK, Canada, Mexico. But, unlike Australia, Japan, Taiwan and Singapore, it underestimates the magnitude of the spikes rather than missing them altogether.

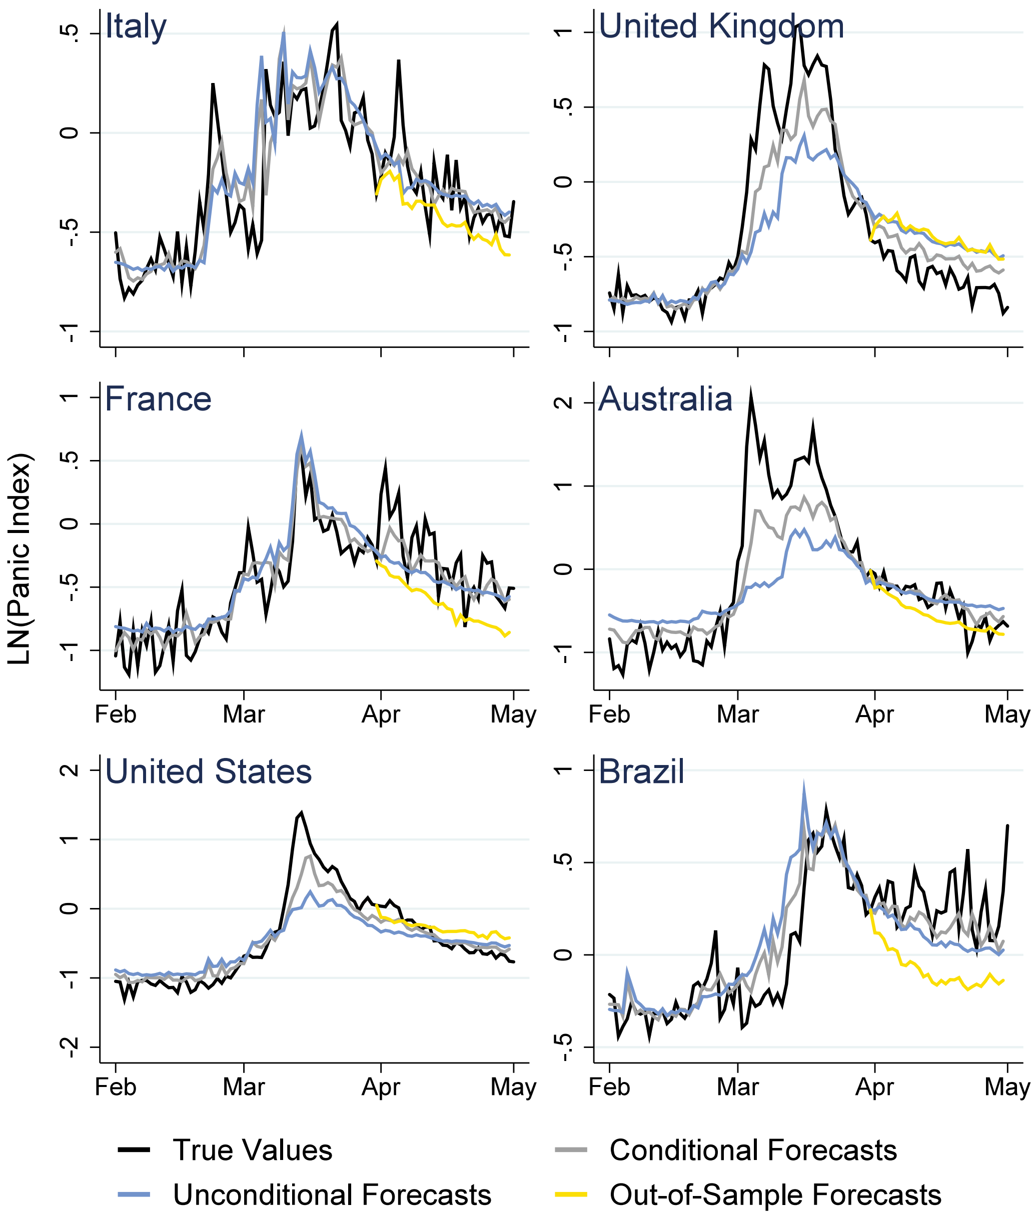

Fig. 9. Predictions of Log Panic from Model 2.

5. Conclusion

Panic buying of storable consumer goods is a common phenomenon during natural disasters and man-made crises. But consumer panic has been little studied from an econometric point of view (A notable exception being Hansman et al. 2020). This is because of the difficulty of obtaining suitable data, and the lack of variation in the determinants of panic. Typical panics are one-off events where consumers in a particular location react to news of a crisis (e.g., an approaching hurricane) by stockpiling consumer goods. The nature of such a panic makes it difficult to study using conventional econometric methods: First, relevant high frequency data on consumer response is difficult to obtain because of the very suddenness of the event. Second, panic events (like a hurricane) typically affect a particular group pf consumers in a particular location at about the same time. Thus it is hard to find variation in the forcing variables that drive the panic.

The COVID-19 pandemic provides a unique opportunity to study consumer panic for two reasons: First, the panic has affected nearly every country on earth, but at different times and to different degrees. Thus, there is a great deal of variation across locations in the timing and severity of the panic-inducing event. Furthermore, governments have

Table 4
Models of Panic during the COVID-19 Pandemic (Country-Specific Factors).

Models of $ln(panic_{c,t})$:	Model 1		Model 2		Model 3	
	β	S.E.	β	S.E.	β	S.E.
Internal Restrictions:						
$\Delta Internal_{c,t}$	0.142	0.025	0.188	0.064	0.197	0.064
$\Delta Internal_{c,t-1}$	0.152	0.030	0.260	0.065	0.253	0.061
$\Delta Internal_{f,t}$			0.496	0.183	0.489	0.176
$\Delta Internal_{f,t-1}$			0.438	0.217	0.419	0.207
$\Delta Internal_{c,t} * Internal_{f,t}$			−0.056	0.026	−0.055	0.026
$\Delta Internal_{c,t-1} * Internal_{f,t-1}$			−0.086	0.026	−0.081	0.024
Stimulus Announcements:						
$\Delta Stim_{c,t}$	0.063	0.044	0.178	0.095	0.149	0.084
$\Delta Stim_{c,t-1}$	0.042	0.048	0.064	0.097	0.049	0.089
$\Delta Stim_{f,t}$			−0.052	0.274	−0.055	0.266
$\Delta Stim_{f,t-1}$			0.302	0.264	0.307	0.257
$\Delta Stim_{c,t} * Stim_{f,t}$			−0.166	0.078	−0.146	0.072
$\Delta Stim_{c,t-1} * Stim_{f,t-1}$			−0.081	0.076	−0.066	0.072
Travel Restrictions:						
$\Delta Travel_{c,t}$	0.020	0.013	0.051	0.030	0.050	0.030
$\Delta Travel_{c,t-1}$	0.023	0.015	0.038	0.036	0.040	0.035
$\Delta Travel_{f,t}$			0.025	0.101	0.050	0.098
$\Delta Travel_{f,t-1}$			−0.125	0.085	−0.139	0.083
$\Delta Travel_{c,t} * Travel_{f,t}$			−0.014	0.008	−0.014	0.008
$\Delta Travel_{c,t-1} * Travel_{f,t-1}$			−0.011	0.008	−0.012	0.008
Domestic COVID-19 Cases:						
$\Delta ln(C_{c,t})$	1.791	0.141	0.847	0.144	0.774	0.136
$\Delta C_{c,t}$	0.004	0.001	0.002	0.001	0.002	0.001
$\Delta C_{c,t} * \Delta ln(C_{c,t})$	−0.047	0.009	−0.027	0.008	−0.026	0.008
International COVID-19 Cases:						
$\Delta ln(C_{f,t})$			2.171	0.288	2.273	0.280
$\Delta C_{f,t}$			0.002	0.001	0.003	0.001
$\Delta C_{China,t}$			−0.007	0.006	−0.006	0.006
$\Delta C_{China,t} * Asia_c$			0.066	0.017	0.071	0.016
Panic Terms:						
$ln(panic_{c,t-1})$	0.312	0.012	0.241	0.012	0.232	0.012
$ln(panic_{c,t-2})$	0.229	0.012	0.167	0.012	0.159	0.012
Related Country Panic					0.492	0.042
Diagnostics:						
AR1 Term of $\hat{e}_{c,t}$	−0.048	0.012	−0.012	0.012	−0.027	0.012
CSD Test (w/ p-value)	17.639	0.000	3.531	0.000	0.620	0.535
Sims Test (w/ p-value)	9.430	0.000	0.638	0.765	1.030	0.413
R^2	0.432		0.463		0.476	

Note: WLS is used for estimation, where $N = 54$ and $T = 119$. The results for additional terms are available in the Online Appendix. The CSD Test refers to the (Pesaran, 2015) test with a null of weak cross-sectional dependence. The Sims Non-Causality Test refers to a F-test on the significance of three leads of the policy variables as suggested by Sims (1972).

responded to the pandemic with a variety of different policies implemented at different times and with differing strictness, generating variation in the policy drivers of panic.

The first contribution of this article is to construct a daily index of consumer panic for 54 countries, covering the key period from January 1st to April 30th 2020. Spikes in the index align well with the timing of stockpiling as measured in the (IRI, 2020) data for the subset of countries where IRI data is available. We also construct daily data on government policy announcements in response to the pandemic, including internal movement restrictions, travel restrictions and economic stimulus. And we collect daily data on COVID-19 cases for all 54 countries.

The result is a high-frequency panel dataset on consumer panic, government policy and COVID-19 cases that covers many countries over a key four month period of the pandemic. By exploiting the heterogeneity in the timing and severity of panic, along with the heterogeneity in governmental policy response and virus transmission, a panel data model is better placed to overcome issues of spurious regression that would inevitably arise from a time series or cross-sectional study of the same problem.

The second contribution of this article is to build a dynamic model that allows for panic to respond to domestic and international policy change, and domestic and international virus transmission. Importantly, our model allows the effect of domestic policy to depend on the international context. We show that a model that attempts to predict panic based only on domestic events is seriously misspecified, but that a model that incorporates the international context passes a stringent set of specification tests.

Our results show that the announcement of internal movement restrictions generates considerable consumer panic in the short-run, but the effect largely vanishes after a week to ten days. Consumer panic also responds to announcements of internal movement restrictions in foreign countries, suggesting that consumers take these as a signal of the severity of the pandemic and/or likely future policy in their home country. Furthermore, internal movement restrictions generate less panic if they are implemented relatively late compared to other countries.

Of the other two policies we consider, stimulus announcements appear to have a positive short-run effect on panic, but it is very short lived. And there is no evidence that travel restrictions generate significant consumer panic.

Consumer panic is also sensitive to both domestic and worldwide COVID-19 cases. The significance of cases in foreign countries suggests that consumers take these as a signal of the severity of the pandemic and/or likely future policy in their home country.

It is important for governments and retailers to better understand how to minimize consumer panic in a pandemic environment. For example, in the early stages of a pandemic, when internal restrictions are first being contemplated, large retailers would be well advised to carry extra stock of key consumer goods and/or to prepare to implement purchase quantity constraints. Given that the panic induced by internal restrictions is sudden but short lived, having extra stock to handle an initial surge in demand, or containing the initial surge via purchase quantity constraints, could prevent shortages and hence prevent consumer panic from taking hold in the first place.

It is particularly important to learn lessons from the COVID-19 pandemic given that future pandemics are essentially inevitable, whether in the near term from a second wave or new strain of COVID-19, or longer term as new virus strains emerge. This article is the first attempt to model the determinants of consumer panic in a pandemic environment, but there is significant scope for this to be an area of active future research.

Acknowledgements

This research was supported by Australian Research Council grants FL110100247 and CE170100005. The authors acknowledge Yuanyuan Deng, Inka Eberhardt, Katja Hanewald, Heloise Labit Hardy, Shahin Kiaei, George Kurdna, Seojeong Lee, Shiko Maruyama, Peyman Firouzi Naeim, Miguel Olivo-Villabrille, Thanh Phan, Olena Stavrunova, and Nada Wasi for assistance with translations.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jeconom.2020.07.045.

References

Avery, C., Bossert, W., Clark, A., Ellison, G., Ellison, S.F., 2020. Policy implications of models of the spread of coronavirus: perspectives and opportunities for economists. In: NBER Working Paper No. 27007. National Bureau of Economic Research.

Chinazzi, M., Davis, J.T., Ajelli, M., Gioannini, C., Litvinova, M., Merler, S., Piontti, A.P.y., Mu, K., Rossi, L., Sun, K., et al., 2020. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. Science 368 (6489), 395–400.

Choi, H., Varian, H., 2012. Predicting the present with google trends. Econ. Rec. 88, 2–9.

Da, Z., Engelberg, J., Gao, P., 2011. In search of attention. J. Finance 66 (5), 1461–1499.

Erdem, T., Imai, S., Keane, M.P., 2003. Brand and quantity choice dynamics under price uncertainty. Quant. Market. Econ. 1 (1), 5–64.

Fang, H., Wang, L., Yang, Y., 2020. Human mobility restrictions and the spread of the novel coronavirus (2019-ncov) in China. In: NBER Working Paper No. 26906. National Bureau of Economic Research.

Ferguson, N., Laydon, D., Nedjati Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunuba Perez, Z., Cuomo-Dannenburg, G., et al., 2020. Report 9: Impact of Non-Pharmaceutical Interventions (Npis) to Reduce COVID19 Mortality and Healthcare Demand. Imperial College London.

Ginsberg, J., Mohebbi, M.H., Patel, R.S., Brammer, L., Smolinski, M.S., Brilliant, L., 2009. Detecting influenza epidemics using search engine query data. Nature 457 (7232), 1012–1014.

Goel, S., Hofman, J.M., Lahaie, S., Pennock, D.M., Watts, D.J., 2010. Predicting consumer behavior with web search. Proc. Natl. Acad. Sci. 107 (41), 17486–17490.

Grohol, J., 2020. Panic buying: The psychology of hoarding toilet paper, beans and soup. PsychCentral. Available at: https://psychcentral.com/blog/panic-buying-the-psychology-of-hoarding-toilet-paper-beans-soup/.

Hansman, C., Hong, H., de Paula, A., Singh, V., 2020. A sticky-price view of hoarding. In: NBER Working Paper No. 27051. National Bureau of Economic Research.

Hori, M., Iwamoto, K., 2014. The run on daily foods and goods after the 2011 Tohoku earthquake: a fact finding analysis based on homescan data. Japan. Political Econ. 40 (1), 69–113.

Hughes, W.T., 1988. A tribute to toilet paper. Rev. Infect. Dis. 10 (1), 218–222.

IRI, 2020. COVID-19 Impact: Consumer Spending Tracker for Measured Channels. Information Resources Inc., Available at: https://www.iriworldwide.com/IRI/media/Library/2020-04-30-IRI-BCG-COVID-Global-Consumer-Spend-Tracker.pdf.

Ishida, T., Maruyama, A., Kurihara, S., 2013. Consumer reaction to the great east Japan earthquake. J. Life Sci. 7 (8), 883.

Knoll, C., 2020. Panicked shoppers empty shelves as coronavirus anxiety rises. N.Y. Times. Available at: https://www.nytimes.com/2020/03/13/nyregion/coronavirus-panic-buying.html.

Kurihara, S., Maruyama, A., Luloff, A., 2012. Analysis of consumer behavior in the Tokyo metropolitan area after the Great East Japan Earthquake. J. Food Syst. Res. 18 (4), 415–426.

Lampos, V., Moura, S., Yom-Tov, E., Cox, I.J., McKendry, R., Edelstein, M., 2020. Tracking COVID-19 using online search. ArXiv preprint arXiv:2003.08086.

Malcolm, A., 1974. The ‘shortage’ of bathroom tissue: Classic study in rumor. N.Y. Times. Available at: https://www.nytimes.com/1974/02/03/archives/the-shortage-of-bathroom-tissue-a-classic-study-in-rumor-shortage.html.
Musvanhiri, P., 2017. Zimbabweans stock up on essentials amid fears of hyperinflation. Deutsche Welles. Available at: https://www.dw.com/en/zimbabweans-stock-up-on-essentials-amid-fears-of-hyperinflation/a-40724529.
Nickell, S., 1981. Biases in dynamic models with fixed effects. Econometrica 1417–1426.
Pesaran, M.H., 2007. A simple panel unit root test in the presence of cross-section dependence. J. Appl. Econometrics 22 (2), 265–312.
Pesaran, M., 2015. Testing weak cross-sectional dependence in large panels. Econometric Rev. 34 (6–10), 1089–1117.
Rieder, K., 2020. What we may learn from historical financial crises to understand and mitigate COVID-19 panic buying. VOXeu. Available at: https://voxeu.org/article/mitigating-covid-19-panic-buying-lessons-historical-financial-crises.
Sims, C.A., 1972. Money, income, and causality. Am. Econ. Rev. 62 (4), 540–552.