Turking in the time of COVID

Antonio A. Arechar^{1,2,3} & David G. Rand^{1,4,5}

On March 16, 2020, the US Government introduced strict social distancing protocols for the United States in an effort to stem the spread of the COVID-19 pandemic. This had an immediate major effect on the job market, with millions of Americans forced to find alternative ways to make a living from home. As online labor markets like Amazon Mechanical Turk (MTurk) play a major role in social science research, concerns have been raised that the pandemic may be reducing the diversity of subjects participating in experiments. Here, we investigate this possibility empirically. Specifically, we look at 15,539 responses gathered in 23 studies run on MTurk between February and July 2020, examining the distribution of gender, age, ethnicity, political preferences, and analytic cognitive style. We find notable changes on some of the measures following the imposition of nationwide social distancing: participants are more likely to be less reflective (as measured by the Cognitive Reflection Test), and somewhat less likely to be white, Democrats (traditionally over-represented on MTurk), and experienced with MTurk. Most of these differences are explained by an influx of new participants into the MTurk subject pool who are more diverse and representative – but also less attentive – than previous MTurkers.

First version: June 3, 2020

This version: March 23, 2021

Keywords: online research; COVID-19; demographics; diversity; Amazon Mechanical Turk

✉ antonio@arechar.com, drand@mit.edu

¹Sloan School of Management, Massachusetts Institute of Technology

²Center for Research and Teaching in Economics, CIDE

³Center for Decision Research and Experimental Economics, University of Nottingham

⁴Institute for Data, Systems, and Society, Massachusetts Institute of Technology

⁵Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

Introduction

In March 2020, the COVID-19 pandemic began to spread across the United States. On March 16, the US Government announced new measures to curtail the virus, such as limiting travel, restricting public gatherings, and closing schools. This nation-wide social distancing policy – hereafter “quarantine” for simplicity – had an immediate impact on the job market, with millions of Americans suddenly unemployed.

Concerns have been raised that this change in the labor market would also affect the subject pool available on crowdsourcing marketplaces, such as Amazon Mechanical Turk (MTurk), which are cornerstones for subject recruitment across the experimental social sciences (1). For instance, subject pools may become less representative, if people experiencing economic hardship due to the pandemic no longer have access to sufficiently reliable internet to participate in online studies (2); remain unaltered (3); or in fact become more diverse as the need for other forms of employment might drive new workers to such marketplaces.

Indeed, the quarantine may be changing the pool of workers available to take part in academic studies in two ways: either by influencing which existing workers participate, or by causing an influx of new ones. Depending on their nature and prevalence, samples drawn by researchers might appear more or less representative than or as representative as before the pandemic started. Hence, in order to assess any potential change in sample compositions, it appears imperative to first disentangle both effects. However, studying the influx of new participants is hard because researcher groups traditionally only recruit restricted samples with previous experience and acquired reputation on platforms (4).

Here we consider the impact of the current pandemic on subject pool composition by looking at 15,539 responses given to studies conducted by our group on MTurk between February and July 2020. Moreover, by exploiting a recent change in our recruitment approach, we are better suited to attract new participants, to track them over time (5) and, relative to usual recruitment restrictions, to anticipate eventual surges – as many of the new accounts considered would meet their criteria within a few weeks (6).

We look at changes in participant demographics and the overall quality of the data produced, which can have important implications for social science researchers assessing representativeness and generalizability of results, as well as for comparing results from studies run pre-quarantine versus during (or after) quarantine. Specifically, we examine the distribution of gender, age, ethnicity, political preferences, and analytic cognitive style (as measured by the Cognitive Reflection Test, CRT (7)) before and after the social distancing measures were introduced on March 16. In addition, we distinguish specific trends by type of participants: either first recruited using traditional qualifications (“baseline”), or first recruited with no restrictions pre- or post-quarantine introduction (“pre” and “post”, respectively).

We find important differences on all measures during quarantine when comparing *between* types. In particular, baseline participants are somewhat less diverse than unrestricted ones, as they tend to be more reflective, female, white, older, and in favor of the Democratic party; pre-quarantine participants are also more reflective and whiter than post-quarantine ones. When comparing *within* participants before and during the pandemic, we note no evidence of structural changes for the unrestricted sample, whereas for the returning baseline sample, we find it to be younger, more ethnically-diverse, and less reflective.

At the same time, we note a potential trade-off between diversity and overall data quality. Responses by unrestricted participants are noisier than the ones provided by the baseline, whereas responses from post-quarantine participants are somewhat noisier than their pre-quarantine counterparts.

We conclude by suggesting that researchers take these changes in the MTurk subject pool – as well as the potential for changes in other variables that we did not measure – into account. This pandemic has changed the way we interact with each other, as it has for the people we often employ as primary input in research.

Methods

We first identify all workers taking part in our MTurk studies in 2020. Based on the number of tasks taken, their date, and the criteria used for initial recruitment, we then classify participants into three categories: (i) “baseline” workers who had at least 100 studies taken on the platform and at least 95% of them approved at the time of initial recruitment – which occurred before February 24, the date when we changed our recruitment policy and allowed everyone located in the US to participate; (ii) workers recruited with no restrictions “pre”-quarantine (i.e. between February 24 and March 15); and (iii) workers first recruited “post”-quarantine. We finally selected the 23 studies conducted since February 24 with relevant demographic data, unrestricted recruitment, and at least 100 participants in them to identify our target sample of 15,925 observations (see Table S1 for study details).

Each of the individual difference variables we analyze appears in at least 12 of the 23 studies surveyed (see Fig. S1 for distributions over time). Gender and ethnicity had different scales across studies; for comparability, we contrast male vs female and white vs non-white, respectively. Political preference appeared in some studies on a 6-point Likert scale and in others on a 7-point one, so we set both as percentages with a minimum of 0 (“Strongly Democrat”) and a maximum of 1 (“Strongly Republican”).

Since we match records from two different datasets (one from Qualtrics to identify participants’ choices, and one from MTurk to identify their category), we lose 386 (2%) observations where participants completed part of the study but did not submit it back to MTurk – leaving a final sample of 15,539 observations. Most of those entries have missing values in the individual difference variables but, among those with data, we find that non-complete respondents differ from complete ones in that they are represented by a less white and less reflective population ($p < 0.003$).

Most, but not all, participants appear only once in the collated dataset (69% once, 22% twice, 6% three times, and the remaining 3% four times or more). Hence, all significance values reported are based on linear regressions with standard errors clustered on subject. Unless otherwise stated, dependent variables are the individual difference variables, whereas the independent variables are dummies for baseline and (unrestricted) pre-quarantine samples. In addition, we include a linear time trend to account for potential temporal patterns (see Fig. S1) but results remain qualitatively similar excluding it. Data and code are available at: <https://osf.io/trfmv/>

Results

We find a surge of brand-new participants since the COVID-19-related quarantine was announced. The share of new accounts per study was previously 17%, and it increased to 34% during quarantine. Relatedly, over 60% of participants in the last few weeks of the data gathered were recent and recruited unrestrictedly (Fig. 1a). It is thus apparent that the sample composition of our studies drastically changed during quarantine, mainly driven by new workers who joined after its announcement.

Are post-quarantine workers different? A comparison between the three types of workers outlined above reveals that, during quarantine, post-quarantine workers are more likely to be of non-white ethnicity and to perform more poorly on the CRT relative to baseline and pre-quarantine samples, and more likely to be

younger, male and favorable towards the Republican party relative to baseline; whereas baseline workers are more likely to be reflective, favorable towards the Democratic party, female, and older than the unrestricted samples recruited pre-quarantine ($p < .010$; Fig. 1b and Table S2). Unrestricted pre- and post-quarantine samples are particularly closer to each other with regards to age and gender.

Figure 1. a) Share of workers per category over time. Bubbles represent the fraction and relative size of new accounts appearing in a given day. The red dotted line represents the day when quarantine was introduced. The areas in gray, blue and red represent the fraction of baseline, pre-, and post-quarantine workers per day, excluding four days where sample size was less than 43 (i.e. 1, 1, 4, and 10). **b)** Standardized mean differences between baseline, pre-, and post-quarantine workers. CRT: number of correct answers in the CRT; R: preference for the Republican party; W: white ethnicity; F: female; A: age. 95% confidence intervals plotted.

Moreover, when comparing workers before and during quarantine, we find no meaningful differences for the unrestricted ones. Yet, returning baseline workers are younger, less reflective, and less likely to be white (Table S3). Thus, we find converging evidence that samples during the pandemic tend to be somewhat more diverse but also less reflective, and that most of the differences observed are driven more by the influx of post-quarantine accounts than by changes in the returning ones.

Given the tendency for non-whites to disfavor the Republican party, the fact that so many new accounts identify as non-white Republicans raises the question of whether the apparent increase in diversity is simply the result of random responding by new workers (8). Thus, we also explore how the data gathered from each type of participant fares in five quality checks during quarantine, and test if the differences seen in Fig. 1b. hold when excluding any potential outliers from the sample (for details see Table S4).

First, one might wonder if post-quarantine accounts are more likely to provide bot-like responses. We find pre- and post-quarantine accounts to be more likely to connect from a suspicious internet service provider than baseline, and no difference between pre- and post-quarantine accounts (3% base; 14% pre; 10% post; $p < 0.001$ and $p = 0.068$; as classified by ref 9, which is updated periodically). Thus, we find some evidence that the quality produced by unrestricted accounts might be lower than baseline. When excluding suspicious accounts, all demographic differences reported in Fig. 1b remain significant though.

Second, we investigate whether new accounts rush through the tasks assigned by looking at the time taken to complete them (completion times are log-transformed due to right skew; study dummies are included). Post-quarantine participants actually take longer to complete studies relative to baseline but not to pre-quarantine ($p < 0.001$ and $p = 0.502$), suggesting that they are not rushing through questionnaires providing nonsense random answers. This is further confirmed when looking at deciles 1 and 10 of time taken per study, where the share of post-quarantine workers is lower and higher, respectively, than baseline workers ($p < 0.001$) but not pre-quarantine ones ($p > 0.510$). Thus, baseline differs from the other two samples in that they are faster

completing studies overall. Removing outliers (deciles 1 and 10) from the analysis does not change the demographic differences depicted in Fig. 1b, apart from post-quarantine samples being relatively less white than the pre-quarantine ones.

Third, we look at the fraction of accounts completing studies from mobile devices, as it is possible that attention, and data quality overall, is influenced by the nature and functions of such devices. We find baseline working on mobile devices relatively more often than unrestricted samples (4% base; 2% pre and post), but the shares are so small that the general results in Fig. 1b barely change when excluding them.

Fourth, we evaluate subjects' consistency by comparing their responses in a follow-up study surveying age, gender, ethnicity, and political preference, but not CRT (n=744; see Appendix C for details). In all, responses were fairly consistent (94%). Yet, we find baseline participants to be more likely to be consistent than post-quarantine workers (97% base; 91% pre; 88% post; $p < 0.001$), and no differences between pre- and post- or pre- and baseline accounts ($p > 0.067$). Restricting to only consistent records, we find that the differences in Fig. 1b for political party and age remain significant for baseline, but not the ones reported for gender and ethnicity – although note that relative sample sizes for this particular check are rather small.

Figure 2. a) Distribution of responses in the CRT; b) Distribution of responses in the CRT that are neither reflective nor intuitive; c) Distribution of correct answers to two attention checks administered in four studies (10).

Finally, we evaluate attentiveness as a fifth quality check by looking at the CRT responses that were neither intuitive nor reflective, and at participants' performance in attention checks. We differentiate between CRT responses that were wrong because they gave the intuitive response versus other unexpected wrong answers, because people providing incorrect, non-intuitive answers may have failed to read the question or responded randomly – rather than merely failing to reflect upon their intuitive response. We find that post-quarantine workers give unexpected answers more often than the rest ($p < 0.001$; Fig. 2a), but that the number of unexpected answers by new workers is rarely greater than 1 – that is, there is essentially no evidence of post-quarantine workers answering all items randomly (Fig. 2b).

With regards to performance on two attention checks (taken from ref 10) administered in four of the studies run during quarantine (n=2,407), unrestricted workers perform substantially worse on the attention checks overall but very few of them get both questions wrong (Fig. 2c). This again suggests inattention, but not completely random responding. Moreover, when comparing performance on a trivial attention check (“dog is to puppy as cat is to ___”) collected as part of the follow-up survey reported earlier, we find post-quarantine accounts rarely fail to provide “kitten” as an answer, but that they nevertheless do so less often than baseline and at a similar rate than pre-quarantine accounts (100% base; 92% pre; 87% post; $p < 0.001$ and $p = 0.468$). Of note, removing the outliers found in this section does not change any of the differences depicted in Fig. 1b. Moreover, when excluding all but the outliers reported in the fourth check, because of the small sample, we

only find two significance changes: relative to post-quarantine samples, pre-quarantine accounts are no longer whiter and become more favorable toward the Democratic party.

Discussion

Here we find that the sample composition of MTurk studies differs before versus during the COVID-19-induced national quarantine. New participants who entered during the quarantine are less reflective, and somewhat more likely to be non-white and Republican, even relative to other participants first recruited just a few weeks before. Indeed, these new participants are actually closer to the national average than baseline and pre-quarantine workers (most notably, scholars studying COVID-19 and related political issues using MTurk should note that, during quarantine, the traditional liberal bias of the MTurk subject pool can be practically eliminated, especially if conventional recruitment restrictions are removed). Thus, contrary to concerns that have been raised (2), quarantine has made MTurk more representative in numerous respects.

At the same time, there is some potential trade-off between representativeness and overall data quality, as post-quarantine workers are also more likely to be inconsistent, fail attention checks, and give responses on the CRT that are neither intuitive nor correct. To put the attentiveness of these new workers in context, it is useful to compare them to workers from other crowdsourcing platforms. To do so, we analyze large datasets collected recently using Prolific (11) and Lucid (12). Considering the one attention check question that was common to the studies run on all three platforms, we find a failure rate of 3% among baseline workers, 23% among pre- and post-quarantine workers, 21% among Prolific workers, and 75% among Lucid workers. The fraction of participants giving at least one unexpected CRT response is 15% among baseline workers, 29% among pre-quarantine workers, 38% among post-quarantine workers, 19% among Prolific workers, and 34% among Lucid workers. Furthermore, the fraction of MTurk workers with at least one unexpected response in a 2019 study with no restrictions other than being located in the US was of 25% (5). Thus, new MTurk workers do not seem appreciably worse than participants from Lucid but they do seem to perform worse on the CRT than Prolific users and their MTurk counterparts a year ago.

One might wonder if our findings on attentiveness were influenced by overall lack of experience on the platform (13), as new MTurk workers were not only less used to the demands of completing studies than more experienced ones but could also be less able to manage the competing demands on their time that quarantine brought with it. Although data on participants' total experience on the platform is not available, we can use our account's study count as a rough proxy and assess evolving performance in attention checks. Indeed, we note mild improvements with experience, but also that baseline workers perform better than unrestricted ones at any comparable level (Fig. S2). Furthermore, we find that including our proxy for experience as an additional control in two supplementary analyses (during pandemic and throughout 2020; Table S5), leaves our general findings consistently and qualitatively the same. Hence, although we cannot rule out a distinctive effect of (lack of) experience on new participants' performance, we find little support for it with our proxy.

Moreover, given that our recruitment has been unrestricted since February 24, it is possible that some workers first participating since then met the quality and quantity restrictions set in baseline, so the differences here reported between baseline and unrestricted (pre- and post-quarantine) samples might be underestimated.

Our results are seemingly in contrast with the one prior attempt we are aware of to investigate the COVID-19 pandemic's impact on the MTurk subject pool (3). This prior work used CloudResearch's publicly available Metrics tool to examine the demographics of nearly all MTurk subjects available on the platform, and found no differences in demographics due to the pandemic. Their work, however, did not specifically examine users

participating in social science studies, which represent a notable minority of their participants (14). Moreover, the modal social science researcher using CloudResearch is likely to be applying the conventional quality restrictions we used in our baseline cohort. Thus, as the changes we observe are mainly driven by new workers whereas our baseline workers display relatively more stable patterns over time, our findings are not necessarily in contrast with theirs.

In sum, our results suggest that there has actually been a meaningful change in social science participants, mainly driven by a steady influx of new participants. We hope these observations will be of use to researchers using MTurk, helping to allay concerns about demographic shifts while potentially raising attentiveness concerns. It seems likely that at least some of these changes will persist as the quarantine eases and many people are able to return to work.

Acknowledgments: We are grateful to Aaron Moss, Leib Litman, and the CloudResearch team for insightful conversations when comparing their findings with ours. We are also grateful to anonymous reviewers, Jillian J. Jordan, Michael Stagnaro, Gordon Pennycook, Gordon Kraft-Todd, Adam Bear, and Arber Tasimi for helpful comments.

References

1. Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental economics*, 14(3), 399–425.
2. Lourenco, S. F., & Tasimi, A. (2020). No participant left behind: conducting science during COVID-19. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2020.05.003>
3. Moss, A. J., Rosenzweig, C., Robinson, J., & Litman, L. (2020). Demographic Stability on Mechanical Turk Despite COVID-19. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2020.05.014>
4. Peer, E., Vosgerau, J., & Acquisti, A. (2014). Reputation as a sufficient condition for data quality on Amazon Mechanical Turk. *Behavior Research Methods*, 46(4), 1023–1031.
5. Robinson, J., Rosenzweig, C., Moss, A. J., & Litman, L. (2019). Tapped out or barely tapped? Recommendations for how to harness the vast and largely unused potential of the Mechanical Turk participant pool. *PloS One*, 14(12), e0226394. <https://doi.org/10.1371/journal.pone.0226394>
6. Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., et al. (2014). Social heuristics shape intuitive cooperation. *Nature Communications*, 5(1), 1–12.
7. Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25–42.
8. Chandler, J., Sisso, I., & Shapiro, D. (2020). Participant carelessness and fraud: Consequences for clinical research and potential solutions. *Journal of Abnormal Psychology*, 129(1), 49.
9. Prims, J. P., Sisso, I., & Bai, H. (2018). Suspicious IP Online Flagging Tool. Retrieved from <https://itaysisso.shinyapps.io/Bots>
10. Berinsky, A. J., Margolis, M. F., & Sances, M. W. (2014). Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys. *American Journal of Political Science*, 58(3), 739–753.
11. Pennycook, G., McPhetres, J., Bago, B., & Rand, D. G. (2020, June 3). Predictors of attitudes and misperceptions about COVID-19 in Canada, the U.K., and the U.S.A. Retrieved from osf.io/3a497
12. Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (In press). Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy nudge intervention. *Psychological Science*.
13. Chandler, J., Mueller, P., & Paolacci, G. (2014). Nonnaïveté among Amazon Mechanical Turk workers: Consequences and solutions for behavioral researchers. *Behavior Research Methods*, 46(1), 112–130.
14. Litman, L., Robinson, J., Rosen, Z., Rosenzweig, C., Waxman, J., & Bates, L. M. (2020). The persistence of pay inequality: The gender pay gap in an anonymous online labor market. *PloS One*, 15(2), e0229383.

Appendix A: Study Details

SN	Launched	Finished	N	Reward	% New	% Female	Age	% White	% Rep	CRT right	% Mobile
1	2/25/20	2/25/20	103	\$1.50	0.32	0.39	35.31	.	.	.	0.01
2	3/10/20	3/12/20	1050	\$2.25	0.19	0.45	37.07	0.77	0.45	1.59	0.02
3	3/14/20	3/14/20	998	\$2.00	0.28	0.43	38.53	.	0.46	1.55	0.04
4	3/16/20	3/17/20	845	\$1.50	0.38	0.38	36.39	0.73	0.44	1.57	0.04
5	3/16/20	3/16/20	1207	\$1.05	0.31	0.40	36.54	.	0.46	1.49	0.03
6	3/20/20	3/21/20	105	\$1.50	0.97	0.33	35.25	.	.	.	0.01
7	3/24/20	3/24/20	197	\$0.75	0.33	0.36	34.14	0.74	.	.	0.01
8	4/16/20	4/16/20	203	\$1.35	1.04	0.35	35.93	0.71	0.53	1.11	0.00
9	4/17/20	4/17/20	754	\$1.05	0.92	0.37	35.83	0.67	0.49	1.34	0.04
10	4/22/20	4/22/20	753	\$1.05	1.16	0.36	37.46	0.66	0.54	1.05	0.02
11	4/23/20	4/23/20	758	\$1.05	0.85	0.45	37.64	0.73	0.49	.	0.07
12	4/30/20	4/30/20	1793	\$0.60	0.93	0.52	38.62	0.71	0.45	.	0.05
13	5/2/20	5/2/20	307	\$0.45	0.63	0.38	37.98	0.66	.	1.49	0.05
14	5/8/20	5/8/20	107	\$2.25	0.76	.	.	0.68	.	.	0.01
15	5/8/20	5/14/20	1072	\$2.40	1.44	.	35.76	0.69	0.46	0.99	0.02
16	6/9/20	6/9/20	97	\$1.05	1.54	0.36	39.69	.	.	.	0.02
17	6/15/20	6/15/20	996	\$0.20	0.91	.	.	.	.	.	0.05
18	6/22/20	6/22/20	496	\$0.30	1.63	0.35	36.23	0.69	0.57	.	0.00
19	6/22/20	6/26/20	986	\$2.40	1.76	.	36.88	0.61	0.51	0.67	0.01
20	7/17/20	7/17/20	327	\$2.25	1.44	0.24	36.72	0.53	0.65	0.91	0.00
21	7/18/20	7/19/20	1507	\$2.25	1.24	0.40	38.13	0.68	0.51	1.18	0.00
22	7/28/20	7/29/20	285	\$0.90	1.04	.	.	0.82	.	.	0.01
23	7/31/20	7/31/20	593	\$1.50	1.43	.	.	.	.	.	0.01

Table S1. Study details by: date when the studies were launched and finished; reward paid; sample size; fraction of participants without prior experience before 2/24/2020; fraction of female participants; mean age; fraction of participants self-identified as white; fraction of participants leaning towards the Republican party; number of correct answers in the Cognitive Reflective Test (out of three); fraction of participants using a mobile device to respond to the survey.

Figure S1. Distribution of variables by type (gray=baseline; blue=unrestricted pre-quarantine; red=unrestricted post-quarantine) a) over time and b) overall.

Appendix B: Robustness checks

	CRT	R	W	F	A
baseline	0.913*** (0.026)	-0.429*** (0.023)	0.278*** (0.024)	0.155*** (0.025)	0.339*** (0.025)
pre-	0.242*** (0.069)	-0.104 (0.065)	0.193* (0.074)	-0.076 (0.063)	0.026 (0.061)
time trend	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.002*** (0.000)
Constant	-0.453*** (0.027)	0.175*** (0.027)	-0.133*** (0.029)	-0.074** (0.028)	-0.323*** (0.027)
R ²	0.195	0.047	0.018	0.007	0.026
N	7,643	10,440	10,085	9,146	10,454
p (B =U)	<0.001	<0.001	0.250	<0.001	<0.001

Table S2. Standardized variables. Include a linear time trend as control. <0.05 *;<0.01 **;<0.001 ***

	<i>Baseline</i>					<i>Pre-Quarantine</i>				
	CRT	R	W	F	A	CRT	R	W	F	A
3/16	-0.087** (0.032)	0.015 (0.032)	-0.116** (0.039)	-0.027 (0.031)	-0.073* (0.031)	0.134 (0.086)	-0.079 (0.080)	0.007 (0.106)	-0.138 (0.075)	-0.071 (0.073)
time trend	0.001** (0.000)	-0.001* (0.000)	0.001* (0.000)	0.000 (0.000)	0.002*** (0.000)	0.000 (0.002)	-0.001 (0.002)	-0.001 (0.002)	0.000 (0.001)	0.001 (0.001)
Constant	0.013 (0.029)	0.030 (0.030)	0.047 (0.035)	-0.003 (0.029)	-0.063* (0.028)	-0.056 (0.056)	0.072 (0.056)	0.024 (0.076)	0.060 (0.058)	-0.013 (0.057)
R ²	0.002	0.001	0.002	<0.001	0.007	0.005	0.004	0.001	0.005	0.001
N	4,977	6,310	5,162	6,388	6,602	790	827	506	876	892

Table S3. Standardized variables. Include a linear time trend as control. <0.05 *;<0.01 **;<0.001 ***

	All	Excl. 1	Excl. 2	Excl. 3	Excl. 4	Excl. 5	Excl. 1-3, 5
<i>CRT</i>							
base	0.913***	0.918***	0.929***	0.925***	-	0.837***	0.860***
pre-	0.242***	0.220**	0.277***	0.243***	-	0.209**	0.249**
N	7,643	7,191	6,208	7,495	-	5,564	4,220
<i>R</i>							
base	-0.429***	-0.442***	-0.418***	-0.441***	-0.472***	-0.446***	-0.470***
pre-	-0.104	-0.107	-0.106	-0.116	-0.265	-0.129	-0.222*
N	10,440	9,874	8,480	10,155	753	5,411	4,124
<i>W</i>							
base	0.278***	0.294***	0.293***	0.284***	0.175	0.280***	0.296***
pre-	0.193*	0.203*	0.152	0.186*	0.125	0.212*	0.226
N	10,085	9,534	8,209	9,822	892	4,574	3,480
<i>F</i>							
base	0.155***	0.134***	0.168***	0.160***	0.178	0.199***	0.165***
pre-	-0.086	-0.108	-0.075	-0.064	-0.064	0.052	-0.016
N	9,146	8,713	7,400	8,866	820	4,262	3,262
<i>A</i>							
base	0.339***	0.338***	0.415***	0.339***	0.258**	0.317***	0.374***
pre-	0.026	0.056	0.111	0.017	0.288	0.083	0.210
N	10,454	9,890	8,456	10,163	783	5,069	3,842

Table S4. Standardized variables. Include a linear time trend as control. <0.05 *;<0.01 **;<0.001 ***

Figure S2. Distribution of experience (i.e. study number in our account as proxy) by participant's type during the pandemic, and the corresponding mean values of the key variables reported, by type (gray=baseline; blue= pre-; red= post-) and experience.

	Excluding data before 3/16					Including data before 3/16				
	CRT	R	W	F	A	CRT	R	W	F	A
baseline	0.523*** (0.052)	-0.368*** (0.048)	0.267*** (0.050)	0.368*** (0.050)	0.294*** (0.046)	0.498*** (0.048)	-0.323*** (0.046)	0.271*** (0.048)	0.333*** (0.047)	0.278*** (0.043)
pre-	0.211** (0.074)	-0.136 (0.072)	0.209* (0.081)	0.065 (0.071)	0.088 (0.067)	0.139** (0.046)	-0.039 (0.050)	0.205** (0.062)	0.065 (0.050)	0.095 (0.049)
time trend	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.002*** (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.002*** (0.000)
first task	0.312*** (0.068)	-0.234*** (0.065)	0.138* (0.066)	0.141* (0.067)	0.392*** (0.063)	0.267*** (0.063)	-0.212** (0.062)	0.120 (0.063)	0.096 (0.064)	0.360*** (0.060)
second task	0.333*** (0.069)	-0.172** (0.064)	0.171** (0.064)	0.011 (0.066)	0.352*** (0.062)	0.325*** (0.064)	-0.178** (0.061)	0.164** (0.061)	-0.043 (0.063)	0.335*** (0.060)
third task	0.314*** (0.071)	-0.204** (0.064)	0.123 (0.065)	-0.071 (0.068)	0.347*** (0.064)	0.289*** (0.065)	-0.219*** (0.060)	0.097 (0.062)	-0.112 (0.064)	0.297*** (0.061)
log(task>3)	0.637*** (0.056)	-0.257*** (0.049)	0.137** (0.050)	-0.105* (0.052)	0.387*** (0.051)	0.628*** (0.052)	-0.283*** (0.047)	0.126** (0.048)	-0.107* (0.051)	0.384*** (0.050)
Constant	-0.741*** (0.073)	0.395*** (0.071)	-0.277*** (0.073)	-0.191** (0.073)	-0.709*** (0.069)	-0.758*** (0.069)	0.408*** (0.068)	-0.265*** (0.070)	-0.151* (0.070)	-0.674*** (0.066)
R ²	0.221	0.051	0.020	0.010	0.036	0.224	0.052	0.021	0.009	0.034
N	7,396	10,179	9,812	8,894	10,177	9,279	12,070	10,758	10,862	12,158
p(B = U)	<0.001	0.001	0.464	<0.001	0.002	<0.001	<0.001	0.319	<0.001	0.001

Table S5. Standardized variables. Include a linear time trend as control. <0.05 *;<0.01 **;<0.001 ***

Appendix C: Follow-up study

Methods

The aim of this follow-up was to test the consistency of participants' responses over time in a succinct way. Thus, we first granted an MTurk qualification to all workers in the aggregated dataset with records for the Cognitive Reflective Test (CRT). Between 6/15/2020 and 6/21/2020, we used that qualification to screen and recruit participants for a 1-minute study. A total of 744 participants took part. As per the classifications reported in the main text of this document, 469 participants were baseline, 33 pre- and 242 post-quarantine.

We surveyed the following six items, all on one screen: age, gender, identity with a political party, ethnicity, an estimate for the overall number of studies taken on MTurk, and an attention check ("dog is to puppy as cat is to ____"). For the question on identity, participants could see either a 6-point Likert scale with values starting with "Strongly Democrat" (1) and finishing with "Strongly Republican" (6), or the same scale but in reverse order. The order in which all questions appeared was randomized, as well as the options available for gender (male, female, or other) and ethnicity (White / Caucasian, Black, American Indian or Alaska Native, Asian, Native Hawaiian or Pacific Islander, Hispanic / Latino, Multicultural; allowing for more than one choice). To retrieve information on the device participants used, we relied on Qualtrics metadata.

To measure the level of discrepancy between values reported in the aggregated dataset and in this follow-up survey, we subtracted the difference between them. If the difference was equal to zero, then the participant's response was qualified as consistent, and assigned a value of 1 in an auxiliary dummy variable (with 0 otherwise), with an exemption for age where a difference of 1 was treated as a missing value as it is plausible that some of the participants' birthdays occurred between studies. For political preference, participants were treated as either Democrat (if responses were "3" at most) or Republican (if responses were at least "4"). The aggregated dataset had some studies on a 7-point Likert scale. We excluded participants who selected the mid-point in such cases. Finally, the self-reported number of studies taken on MTurk was larger for baseline workers than for unrestricted workers ($p < 0.001$), which further suggests that our methodology for classifying participants based on experience within our account is appropriate.