

DISCUSSION PAPER SERIES

DP17710
(v. 2)

FACT-CHECKING POLITICIANS

Andrea Mattozzi, Samuel Nocito and Francesco
Sobbrio

POLITICAL ECONOMY

CEPR

FACT-CHECKING POLITICIANS

Andrea Mattozzi, Samuel Nocito and Francesco Sobbrío

Discussion Paper DP17710
First Published 28 November 2022
This Revision 27 February 2023

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Political Economy

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Andrea Mattozzi, Samuel Nocito and Francesco Sobbrío

FACT-CHECKING POLITICIANS

Abstract

We study the reaction of national politicians to a rigorous fact-checking of their public statements. Our research design relies on a novel randomized field experiment in collaboration with a leading fact-checking company. Our results show that politicians are responsive to negative fact-checking. We observe a reduction in the number of incorrect statements made by politicians after being treated in the order of one fourth of a standard deviation. This effect persists for at least two months. We also observe a reduction in the probability of politicians making verifiable statements, suggesting that fact-checking may also increase the ambiguity of politicians' statements.

JEL Classification: D72, D78, D8, D91

Keywords: Politicians, Accountability, Ambiguity

Andrea Mattozzi - andrea.mattozzi@unibo.it

Department of Economics, University of Bologna and CEPR

Samuel Nocito - samuel.nocito@uniroma1.it

Department of Social and Economic Sciences, Sapienza Università di Roma (Italy).

Francesco Sobbrío - francesco.sobbrío@uniroma2.it

Department of Economics and Finance - Tor Vergata University and CESifo

Acknowledgements

We thank Pietro Biroli, Nicola Mastrococco, Pietro Ortoleva, Franco Scoglio, Erik Snowberg, Francesco Trebbi and audience seminars at the 2022 Barcelona Summer Forum (PE), 2022 Perth Workshop, Poleconuk Webinar, Tilburg University, EIEF, Collegio Carlo Alberto, University of Padova, Universitat Complutense de Madrid, and Tor Vergata University of Rome for useful comments. We are grateful to Tommaso Canetta and Giovanni Zagni of Pagella Politica, Valerio Riavez of Electica, and Daria Moretti and Bartolomeo Sciacaluga. Simona Mandile and Giuseppe Spataro provided excellent research assistance. The study was approved by the Institutional Review Board of Luiss University (June 5, 2020) and it was partially financed thanks to a grant from the Research Council of the European University Institute. This experiment is registered in the American Economic Association Registry for randomized control trials; <https://www.socialscienceregistry.org/trials/6432>

Fact-checking Politicians*

Andrea Mattozzi[†] Samuel Nocito[‡] Francesco Sobbrío[§]

This version: February 2023

First version: October 2022

Abstract

We study the reaction of national politicians to a rigorous fact-checking of their public statements. Our research design relies on a novel randomized field experiment in collaboration with a leading fact-checking company. Our results show that politicians are responsive to negative fact-checking. We observe a reduction in the number of incorrect statements made by politicians after being treated in the order of one-fourth of a standard deviation. This effect persists for at least two months. We also observe a reduction in the probability of politicians making verifiable statements, suggesting that fact-checking may also increase the ambiguity of politicians' statements.

Keywords: Fact-Checking, Politicians, Accountability, Verifiability, Ambiguity, RCT

JEL Classification: D72; D78; D8; D91

*We thank Pietro Biroli, Nicola Mastrorocco, Pietro Ortoleva, Franco Scoglio, Erik Snowberg, Francesco Trebbi and audience seminars at the 2022 Barcelona Summer Forum (PE), 2022 Petralia Workshop, Poleconuk Webinar, Tilburg University, EIEF, Collegio Carlo Alberto, University of Padova, Universitat Complutense de Madrid, and Tor Vergata University of Rome for useful comments. We are grateful to Tommaso Canetta and Giovanni Zagni of *Pagella Politica*, Valerio Riavez of *Electica*, and Daria Moretti and Bartolomeo Sciacaluga. Simona Mandile and Giuseppe Spataro provided excellent research assistance. The study was approved by the Institutional Review Board of Luiss University (June 5, 2020) and it was partially financed thanks to a grant from the Research Council of the European University Institute. This experiment is registered in the American Economic Association Registry for randomized control trials; <https://www.socialscienceregistry.org/trials/6432>

[†]University of Bologna & CEPR, andrea.mattozzi@unibo.it

[‡]Sapienza University of Rome, samuel.nocito@uniroma1.it

[§]Tor Vergata University of Rome & CESifo, francesco.sobbrio@uniroma2.it

1 Introduction

The use of false or unsubstantiated claims by politicians is hardly a new or unexpected phenomenon. Otto Von Bismark contended that “people never lie so much as after a hunt, during a war or before an election”, and a century later Ronald Reagan ventured himself into claiming that “trees cause more pollution than automobiles do”. Only in the last decade, however, expressions such as *fake-news*, *alternative-facts* or *post-truth politics* started making regular appearances in the public discourse; a proliferation of incorrect or blatantly false public statements, which may “threaten to warp mass opinion, undermine democratic debate, and distort public policy”.¹

During the same period and arguably in response to such phenomena, a number of independent organizations committed to verify the factual accuracy of public statements appeared all over the world. Indeed, in 2016 there were more than 100 independent *fact-checking* organizations in more than 50 countries.² What appears to remain an unsettled issue is whether fact-checking is indeed an effective tool to curb fake news and misleading statements that may compromise the accuracy of individuals’ beliefs and ultimately the efficiency of democracy.³ The objective of this study is to investigate whether and how politicians react to a rigorous fact-checking of their political statements.

Assessing the impact of fact-checking on politicians poses two key empirical challenges. First, fact-checking organization endogenously choose which politician/statement to fact-check.⁴ This might lead to a bias in the estimates when simply taking the timing of fact-checking at face value. For example, assessing the impact of a fact-checking on a very newsworthy statement might lead to

¹Nyhan (2020), Fact and myths about misperception, *Journal of Economic Perspectives*, 34(3), p. 220.

²The 90% of these organizations have been established after 2010 and more than half are not affiliated with media companies (Graves and Cherubini, 2016). Fact-checkers are trained to acquire specific skills allowing them to judge the quality of information quickly and accurately (Wineburg and McGrew, 2017). Fact-checkers tend to outperforms crowd-sourced evaluations of news stories (Godel *et al.*, 2021).

³DellaVigna and Gentzkow (2010): “The efficiency of market economies and democratic political systems depends on the accuracy of individuals’ beliefs.” *Annual Review of Economics*, 2(1), p. 644.

⁴The *Washington Post* states that “We especially try to examine statements that are newsworthy or concern issues of importance. [...] We strive to be dispassionate and non-partisan, drawing attention to inaccurate statements on both left and right.” See, <https://www.washingtonpost.com/politics/2019/01/07/about-fact-checker/> Similar criteria are applied by the Italian fact-checking agency *Pagella Politica*, which selects politicians statements “on the basis of their relevance and resonance in the media and in the political debate, always trying to avoid an undue concentration of our articles on a single politician or political party”, translation from <https://pagellapolitica.it/progetto>.

an upward bias of the estimates. Likewise, fact-checking organizations might at time focus on less controversial and less newsworthy statements for fear of being perceived as politically biased and this might potentially lead to a downward bias of the estimated effect. Second, there is unobserved heterogeneity across politicians and over time in terms of political communication. For example, politicians might have different underlying likelihood of using misleading statements and there might be periods of more or less intense political communication. To address such empirical issues we combine two key elements:

i) A randomized “business as usual” field experiment in collaboration with the leading Italian fact-checking company *Pagella Politica* (“Political report card”), which provides us with the unique opportunity of addressing the issue of endogenous selection of statements by fact-checkers;

ii) A difference-in-differences analysis focusing on political statements of treated politicians before and after fact-checking, compared to not fact-checked politicians, which allows us to control for unobserved heterogeneity across politicians and over time.

Our study is based on a detailed dataset of political statements, in fact the universe of statements publicly released by a sample of 55 Italian MPs in a period of sixteen weeks (3 pre-intervention, 10 intervention and 3 post-intervention), starting from March 2021. In the core part of the study, that is during each intervention week, we first randomly select a politician among those in our sample who made at least one verifiable, i.e., fact-checkable, and incorrect statement during the previous week. Then, we randomly select one of these statements, which is rigorously fact-checked by our partner *Pagella Politica*. Following the usual practice of *Pagella Politica*, the verdict is published on the fact-checking company’s webpage and on its social media pages (Twitter, Facebook and Instagram). To maximize the effectiveness of the campaign and to make sure treated politicians are aware of the fact-checking, the Tweets advertising the fact-checking also mention the official Twitter account of the politician. Furthermore, a video, advertising the fact-checking, is disseminated on a number of popular websites and social media, geo-targeting the two zip-codes surrounding the Italian Parliament.

Our study focuses on mid-rank national politicians. The main reason follows from our purpose of interfering as little as possible with the standard *modus operandi* of a fact-checking organiza-

tion.⁵ There are other important reasons, however, to choose this specific focus. First, it enhances the external validity of our results, as mid-rank politicians are more similar both within Italy and across countries. Second, it minimizes the risk of spillovers on non fact-checked politicians, which we test anyway in our empirical analysis. Third, as *Pagella Politica* is de-facto the monopolist supplier of fact-checking on mid-rank politicians in Italy, we do not have to worry about potential general equilibrium effects of our experimental design. An additional advantage of our design, which only influences the selection process of statements usually handled by *Pagella Politica*, is that the concern for experimenter demand effects appears to be limited. As such our results are likely to preserve a high external validity in other contexts.

We estimate a difference-in-differences model to assess the causal effect of fact-checking on treated (fact-checked) politicians relative to the pre-treatment period and compared to our control group of non-fact checked politicians. Specifically, given the staggered nature of our treatment with ten different treated politicians fact-checked during ten different weeks, we follow Cengiz *et al.* (2019, 2022) and estimate a stacked difference-in-differences controlling for politician-event and time-to-event fixed effects.

Our results show that politicians respond to fact-checking. We observe a significant reduction in the number of incorrect statements made by a politician in the weeks after the fact-checking, in the order of more than a quarter of a standard deviation. This is the case both in terms of absolute number of incorrect statements and as a share of fact-checkable statements. Noticeably, the effects of the treatment are not short-lived as they last at least eight weeks. We show that the observed reduction in the number of incorrect statements is robust to alternative specifications such as implementing a two-way fixed effects model, using the Callaway and Sant’Anna (2020) doubly robust estimators and a Poisson fixed effect model. Furthermore, the results are robust to the use of different or more restrictive control samples such as, among others, using only never treated as controls, focusing on a homogenous time window across event datasets or excluding politicians with no incorrect statements over the sample period. The results are also robust to

⁵Otherwise, by randomizing fact-checking of incorrect statements made by, say, the prime minister, a party leader or politicians with key government positions, we could have run the risk of affecting the public perception of impartiality of the fact-checking organization.

controlling for the number of verifiable statements made by a politician in a week or when using the probability of making an incorrect statement as an alternative outcome.

Notice that the very purpose of our study, aiming to provide a policy evaluation on the effectiveness of fact-checking, somewhat “naturally” constraints the sample size of treated politicians. Indeed, a large set of treated politicians would require either a very long time span of the experiment or multiple treated politicians within a week. The latter may confound the estimates due to potential spillover effects across treated politicians. The former is simply unfeasible since it would jeopardize the “business as usual” perception of the fact-checking activities.⁶ To show, however, that our results are unlikely to be driven by the idiosyncratic characteristics of politicians in our sample, we provide two key pieces of additional evidence. First, we show the robustness and external validity of our results by using a different dataset obtained from a pilot experiment run in November 2020. Despite a different time period and a different set of politicians, the estimates that we obtain are remarkably similar to the ones of our baseline sample. Second, we validate the causal interpretation of our estimates by providing a random inference test in the spirit of Young (2018) and Dell and Olken (2020). Specifically, we exploit the existence of a randomization pool of politicians that could have been subject to fact-checking in a given week but that were not.

Why do politicians respond to fact-checking? After all, politicians belonging to the treated and control group in a given period face the same probability of being fact-checked in the future. Furthermore, we do not find evidence that the treatment might have spillover effects across politicians in the same party or in the same party-chamber, that is, politicians in the control group do not respond to fact-checking on a party peer. This suggests that fact-checking does not operate through a simple information channel (e.g., making politicians aware of being potentially fact-checked or increasing the salience of *Pagella Politica* per se). Rather, fact-checking seems to have a specific direct impact on the treated politician. We can think of three possible narratives, which rationalize the observed behavior. First, politicians might have convex costs from being repeatedly

⁶In order to provide causal evidence on the impact of fact-checking on politicians we need a commitment by the fact-checking company not to fact-check politicians in our sample other than the ones randomly selected in the experiment for the entire duration of the intervention. As such, a very large sample of politicians would come with a very large shift (and drop) in the usual activities of the fact-checkers. This would both deviate from the “business as usual” design and from what would be acceptable for any major fact-checking company.

exposed to negative fact-checking. This could be either due to self-image or career concerns, or it might be related to voters becoming progressively less forbearing with politicians repeatedly making false statements. This narrative is consistent with the fact that we observe more significant effects for politicians that have been already fact-checked in the past as compared with “rookies” in the fact-checking league. An alternative narrative could be that treated politicians revise upward their perceived probability of being fact-checked in the future. Unfortunately, we do not have direct evidence to test this hypothesis. Yet, one would expect that such updating might be stronger for “rookies”, which runs contrary to the evidence that we observe. Finally, fact-checking might have a behavioral impact on politicians by priming normative concerns against lying (Nyhan and Reifler, 2015; Cohn and Maréchal, 2016). This mechanism is consistent with the stronger response observed among politicians who received worse fact-checking scores.

A related question is whether the observed behavior is due to politicians’ electoral concerns or it is a direct response that does not transit through concern for voters’ reactions. To this end, a preliminary observation is in order. Independently on the fact that electoral concern is relevant or not for the underlying mechanism, the potential welfare consequences of fact-checking ultimately depend on how politicians react to it. Indeed, voters’ response to fact-checking may be neither a sufficient nor a necessary condition to trigger a politician’s reaction. On the one hand, politicians may react to fact-checking only if the effect on voters is sufficiently strong to induce an electoral concern and such response will likely be heterogeneous across politicians. On the other hand, even if electoral concerns and voters’ reactions are meagre, politicians may still be responsive to fact-checking for different reasons. For example, a negative fact-checking may harm the politician reputation and credibility since being negatively exposed may potentially lower the politician’s chances of entering party leadership, or seeking higher office, or simply hurt the politician due to self-image concerns (Bursztyn and Jensen, 2017). As such the risk of incurring in a negative fact-checking may increase the cost of lying in politics.

As for the specific issue of whether or not the observed behavior is due to politicians’ electoral concerns, we find mixed evidence. Politicians elected in single-member districts, where electoral accountability has more bite, are the ones most responsive to fact-checking in the order of almost

one standard deviation. Yet, we also observe a significant response by fact-checked politicians elected in multi-member districts (one fourth of a standard deviation). As the electoral fate of such politicians very much depends on their position on the electoral list, which is decided by the party, such evidence suggests that career concerns inside and/or outside the party or self-image concerns (Bursztyn and Jensen, 2017) could indeed play an important role.⁷

In terms of overall heterogeneity, we observe a significant response to fact-checking across politicians with different level of experience and ideology. Yet, we observe a stronger response in terms of reduction of false statements after being fact-checked by inexperienced (in their first term in parliament), by left-wing, and by non-populist politicians. As pointed out before, we observe a stronger response by politicians who obtained a worse fact-checking score, and by those elected in single-member districts. Instead, we do not observe much of a clear difference in terms of measures of diffusion of the fact-checking video.

Finally, turning our attention to other observed measures of politicians' behavior, we find evidence of treated politicians reducing the overall number of weekly statements and, importantly, also the probability of making any verifiable statement. These findings suggest that politicians, when exposed to negative fact-checking, also respond by resorting to non-factual claims or vacuous political rhetoric, that is they also increase the ambiguity of their statements.

The remainder of the paper is structured as follows. Section 2 discusses the contribution of the paper with respect to the existing literature. Section 3 presents background information on the fact-checking organization involved in the experiment. Section 4 provides detailed information on the design of our experiment (sample of politicians, selection procedure relative to politicians' statements, fact-checking campaign). Section 5 outlines the empirical strategy. Section 6 presents the main results, validation tests and the evidence on the mechanism. Section 7 concludes and discusses possible welfare implications.

⁷The relevance of such *direct* effects of fact-checking on politicians, i.e., not operating through voters, seems sensible given the somewhat limited share of voters following *Pagella Politica*.

2 Related Literature

Our paper relates to the recent literature assessing the impact of fact-checking. The existing economic literature has mainly focused on the impact of fact-checking on voters, providing mixed evidence of its effectiveness in influencing voters’ beliefs and attitudes towards “lying” politicians (Barrera *et al.*, 2020; Swire *et al.*, 2017; Nyhan *et al.*, 2020; Henry *et al.*, 2022). Conversely, there is very limited evidence on the impact of fact-checking on the supply of misinformation and in particular on politicians (Tucker *et al.*, 2018). A notable exception is Nyhan and Reifler (2015), who run a field experiment consisting in mailing messages to a random subset of state legislators in nine US states. The messages were meant to remind politicians about the political costs of having false claims identified by fact-checkers. Nyhan and Reifler (2015) show that politicians who received such messages had a lower ex-post probability of negative fact-checking. Our paper differs in a number of important ways. First, as mentioned before we only randomize the selection of statements and thus we do not affect, and certainly not directly, the politicians’ perceived probability of being exposed to fact-checking. Second, rather than looking at the ex-post probability of being fact-checked, we gather and analyze the universe of our politicians’ statements. As such, we are able to assess the impact of fact-checking both on the truthfulness of politicians’ statements and on their verifiability. Indeed, as pointed out above, while we find a (arguably welfare-improving) decrease in the number of incorrect statements we also provide evidence of politicians also reducing the verifiability of their statements which may instead be potentially detrimental from a welfare perspective. More broadly, the paper relates to the literature looking at misinformation in terms of its drivers, consequences and tools to contrast it (Nyhan, 2020; Bursztyn and Yang, 2022; Bursztyn *et al.*, 2022).

3 Background: *Pagella Politica*

Pagella Politica is the first and most important Italian fact-checking company. It has been online since 2012, and is the only Italian website entirely dedicated to political fact-checking. The mission of the company is “to monitor statements made by Italian politicians and verify their

truthfulness according to reliable data and sources.” Since 2017 Pagella Politica is an active member of the International Fact-Checking Network (IFCN) and one of the signatory of the related Code of Principles.

None of the founders of *Pagella Politica* or staffers are members of political parties and/or organizations and entities related to political parties, and the lack of political involvement is a key prerequisite in order to work or cooperate with them. *Pagella Politica* is financed mainly by selling content and services to third parties and by participating in international projects and calls.⁸ Their usual business activity consist of monitoring political statements from traditional and online media, social media, and news agencies. Clearly, they only focus on *verifiable* statements, that is propositions based on verifiable facts or numbers. Importantly, while a number of news media and websites (e.g., *lavoce.info*) sometimes provide fact-checking on key statetements made by political leaders or politicians in key government position, *Pagella Politica* is, *de facto*, the monopolist supplier of fact-checking on mid-rank politicians.

Pagella Politica selects political statements non-randomly and in particular they are primarily chosen “on the basis of their relevance and resonance in the media and in the political debate, always trying to avoid an undue concentration of our articles on a single politician or political party”. This drives the rationale behind our RCT intervention aiming at introducing an exogenous selection of the statements to be fact-checked. Each fact-checked statement gets a “truthfulness” score on a five-point scale similar to “Pinocchio Test” of the *Washington Post*: 1) Utterly False (*Panzana Pazzesca*); 2) Mostly false (*Pinocchio Andante*); 3) Half False (*Ni*); 4) Almost True (*C’eri quasi*); 5) Completely True (*Vero*).⁹

⁸According to *Pagella Politica*, in 2020, “the main sources of financing (those that accounted for more than 5% of revenues) were, in order of importance; Facebook (within the Third-Party Fact-checking Program); some calls promoted by the International Fact-checking Network; the AGI news agency (not renewed for 2021); the Italian public broadcaster RAI (not renewed for 2021); the funding obtained as part of the European projects SOMA (winner of a Horizon 2020 call from the European Commission, grant agreement no. 825469) and Fakespotting.” *Pagella Politica* does not receive any funding from political parties or entities related to political parties.

⁹*Pagella Politica* makes publicly available the data sources on which they base their assessments of the falsehood/veracity of the declarations. “If those data are updated later on, *Pagella Politica* will not change its assessment, because our assessments are made in the light of the data available to us (and therefore to the politician making the declaration) at the time of the declaration itself.” The score is attributed according to the following criteria. 1) Utterly False - The datum or fact is absolutely invented or reported in a totally distorted way, to support a substantially false thesis; 2) Mostly False -The statement starts with a fact or assertion that is not entirely unrealistic, but at least vague or overly general, and then draws incorrect conclusions; 3) Half False -

4 Experimental Design

The experiment lasted 16 weeks between April and July 2021 and comprises 3 pre-intervention weeks, 10 intervention weeks and 3 post-intervention weeks. For the entire length of the experiment (16 weeks) *Pagella Politica* committed not to publish any fact-checking involving any politician in our sample other than those that were part of the experiment.

4.1 Sample of Politicians

Our sample is composed of 55 Italian MPs. In what follows we first describe the different steps taken to construct this sample and then discuss the rationale behind each of them.

The first step consists in creating a list of Italian MPs present in the report on political pluralism published monthly by the Italian public communication authority (AGCOM, i.e., the Italian equivalent of the FCC in the US or OFCOM in the UK). The report lists the name and total time of monthly exposure of the 20 most present politicians in each news program of the Italian national TV channels (generalists or news channels).¹⁰ As a second step, we exclude politicians holding key institutional roles (i.e., President of the Republic, Speakers of the two Italian chambers, Prime Minister, and Ministers) and party leaders. Furthermore, we remove politicians in the upper and lower tail of the distribution of exposure time (i.e., above and below the 90 and 10 percentile, respectively, of total TV exposure in a month). As a third and final step, we exclude politicians without an active Twitter account and more specifically those without an account or with an inactive account during the month before the experiment.

The first step (i.e., our focus on AGCOM list) ensures to have a list of politicians that are likely to make public statements (and possibly verifiable statements) during our sample period, as many MPs seldom make public statements, let alone verifiable ones.¹¹ The second step (i.e., our focus on

Multiple facts or data are cited at the same time, some of which are not reported accurately: the statement is therefore only partially correct; 4) Almost True - The data or facts are reported slightly inaccurately, but close to the truth, or the data are nearly correct but the conclusions drawn distort their meaning; 5) Completely True - Data or facts are accurately reported, or rounded up correctly, and found in official documents or other reliable sources. From *Pagella Politica* website.

¹⁰This list covers 10 TV-Channels: Rai1, Rai2, Rai3, Rainews24, Rete4, Canale5, Italia1, TGCOM24, La7 and SkyNews. For each channel, it reports the list of the 20 most present politicians in the news segment of the channel and another list of those most present in all the other political news programs.

¹¹For example, even focusing on our sample, the 4% of MPs did not make any verifiable statement over the 16

mid-rank MPs) is instrumental to different purposes. First, it leads to a relatively homogeneous and comparable sample as politicians holding key institutional roles are a very selected group and this is also the case for party leaders or their main collaborators. This also enhances the external validity of our results as Italian mid-rank MPs are arguably more comparable to their counterparts in other democratic countries, with respect to party leaders or politicians in key government positions. Moreover, the focus on mid-rank MPs was part of the agreement with our fact-checking partner. Indeed, a key and rather unique feature of our RCT intervention is represented by the commitment by *Pagella Politica* not to publish any fact-checking on the politicians of our sample during the experiment period other than those that are part of the experiment. If such commitment would have involved party leaders and politicians with key government position, it could not have been acceptable on their side nor desirable from our perspective of maintaining a business as usual framework. That is, our focus on mid-ranked politicians avoids a potentially large change in the size of their usual fact-checking activity. Finally, by focusing on mid-rank politicians we minimize concerns regarding general equilibrium effects as *Pagella Politica* is de-facto the monopolist in fact-checking of such group of politicians and it also reduces potential spillover effects between the treated politicians and the control group. The third and final step (i.e., focusing only on politicians with an active Twitter account) ensures that all politicians are comparable in terms of social media engagement and are able to release public statements even in the absence of media exposure. Most importantly, this criterion allows making sure that treated politicians are aware of being fact-checked since, as we describe below, when publishing the fact-checking on Twitter, *Pagella Politica* mentions the fact-checked politician.

Given the above selection procedure, our final sample is composed of 55 mid-rank politicians, relatively balanced across the main six political parties present in the Italian parliament at the time of the experiment (see Panel A of Table 1).

weeks sample period, 25% of them made less than three verifiable statements and only the 51% of them made at least one verifiable and incorrect statement.

Table 1: Sample of Politicians

	Number	Percentage
Panel A: Whole Sample		
<i>Brothers of Italy</i>	7	12.72
<i>Northern League</i>	8	14.54
<i>Forza Italia</i>	8	14.54
<i>Italia Viva</i>	10	18.18
<i>Democratic Party</i>	11	20.00
<i>Five Stars Movement</i>	11	20.00
Total	55	100.00
Panel B: Fact-Checked Politicians		
<i>Brothers of Italy</i>	2	20
<i>Northern League</i>	2	20
<i>Forza Italia</i>	3	30
<i>Italia Viva</i>	0	0
<i>Democratic Party</i>	1	10
<i>Five Star Movement</i>	2	20
Total	10	100.00

4.2 Politicians' Statements

We collect the universe of politicians' statements provided by the main Italian news agencies (ANSA; AGI; Adnkronos; Askanews). The focus on news agencies allows us to avoid issues of biased coverage as the core business of such news agencies is exactly to monitor any possible public statements made by politicians on any type of news media or on the politicians' social media accounts.

During the intervention weeks, we pre-screen verifiable or fact-checkable (FC henceforth) statements using a 2-steps procedure. First we identify FC statements by using a machine-learning classifier. Specifically, we make use of a Support Vector Machine (SVM) trained on a dataset of all statements released by politicians in our sample in the third week of each month from January

to June 2020.¹² Second, *Pagella Politica* validates the FC statements identified by the classifier.¹³ Given this final set of FC statements, *Pagella Politica* identifies the incorrect statements (i.e., those with a truthfulness score below ≤ 3). Then, we randomly draw one politician from those who released at least one incorrect statement in the previous week. Finally, we randomly select one incorrect statement made by the selected politician.¹⁴ Once the statement is selected, *Pagella Politica* produces a fact-checking of the selected statement according to its standard rules and procedures.

Accordingly, our final sample is composed of ten treated politicians (staggered over time) and 45 never-treated politicians. Panel B of Table 1 presents the break-down of treated politicians by political party affiliation. Table 2 provides descriptive statistics on the sample of treated and control politicians also showing how the observable characteristics of the two groups are rather balanced. The only exceptions are right-wing and populist politicians which are more likely to be part of the treated group. This difference is simply due to the fact that, as described above, a necessary condition to be in the treated group is to make a verifiable and incorrect statement during the intervention weeks. It so happen that in such a period right-wing and populist politicians had a higher likelihood of making such statements with respect to left-wing or non-populist politicians.

¹²In this way we assemble a training sample of 12,070 statements and *Pagella Politica* classifies each statement as fact-checkable or not. To avoid issues related to class imbalances (Kubat *et al.*, 1997; Japkowicz and Stephen, 2002; Wasikowski and Chen, 2009), we adopt the undersampling procedure to lower the majority class (the unverifiable statements). We stop undersampling at the 65 percent level to minimize the drawbacks of discarding a potentially too large set of useful data (McCarthy *et al.*, 2005). We train the SVM by adopting different approaches (i.e., word-based and symbol-based) among which we chose the one with the highest accuracy and F_2 -score levels. Along this line, we use a word-based approach with uni-grams and bi-grams and a SVM with Radial Basis Function kernel with optimal parameters $\Gamma = 0.0001$ and $C = 1000$ obtained via a grid-search from a 5-folds cross-validation procedure. This approach provides an accuracy of the 82.5 percent with a *Recall* equal to 0.614 (F_2 -score = 0.5).

¹³ The following are examples of FC and non-FC statements made by politicians in our sample. Politician A) FC: “The Alitalia Brand is worth a billion Euro”; Non-FC: “most of the anti-covid measures presented in the last 12 months by the government have been ineffective from a health perspective and deadly for our economy”. Politician B) FC: “we have lost almost a million jobs in the last year and the employment rate in pre-covid Italy was still one of the lowest in Europe.”; Non-FC: “Setting to zero the contributions to be paid by companies for the new hiring of workers for two or more years could help, on the one hand, the growth of companies and, on the other hand, mitigate the employment crisis.”

¹⁴We adopt this two-step procedure to avoid oversampling politicians who make more verifiable but incorrect statements in a week.

Table 2: Balancing

	Control			Treatment			Diff
	n	mean	sd	n	mean	sd	
Woman	45	0.29	0.46	10	0.20	0.42	-0.089
Age	45	51.42	9.00	10	54.60	7.68	3.178
Education level	45	1.89	0.53	10	2.00	0.47	0.111
Ever fact-checked	45	0.71	0.46	10	0.70	0.48	-0.011
Lower-chamber MP	45	0.71	0.46	10	0.50	0.53	-0.211
1st parliamentary experience	45	0.40	0.50	10	0.40	0.52	0.000
Elected single-member district	45	0.31	0.47	10	0.30	0.48	-0.011
N. of parliamentary commissions	45	1.67	1.00	10	2.10	1.37	0.433
Right wing	45	0.36	0.48	10	0.70	0.48	0.344**
Populist	45	0.56	0.50	10	0.90	0.32	0.344***
Opposition	45	0.11	0.32	10	0.20	0.42	0.089

Notes: Test of differences in means across samples. Standard errors are clustered at the politician level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Figure 1: Fact-checking page



Notes: The fact-checking article (in Italian) is available at the following link: <https://pagellapolitica.it/fact-checking/rampelli-sbaglia-parecchio-sul-valore-miliardario-del-marchio-alitalia>. Since the time of the experiment, Pagella Politica has renewed its website with slight changes in the style with respect to what shown in this figure. However, we report the page style as appeared at the time of the experiment.

4.3 Fact-checking campaign

The fact-checking campaign has the following components. *Pagella Politica* publishes the fact-checking on its website (see Figure 1 for an example). Furthermore, to maximize the probability

Figure 2: Fact-checking tweets




Notes: The two tweets are available at: <https://twitter.com/PagellaPolitica/status/1386588433724973056?s=20> and <https://twitter.com/PagellaPolitica/status/1386621919731101696>.

that the politician is aware of being exposed to fact-checking, *Pagella Politica* simultaneously sends two Tweets from its official account—with a link to the fact-checking—mentioning the politician’s twitter account (see Figure 2 for an example). Finally, we advertise a video featuring the politician, its statement and the “truthfulness” score attributed by *Pagella Politica* on a number of popular websites and social media.¹⁵ In particular, we geo-targeted the two zip-codes (00186 and 00187) around the Italian Parliament to increase the effectiveness of the campaign (see Figure 3). This media-marketing campaign was run on our behalf by the UK based media company *Electica*. For each fact-checked politicians the campaign started right after the publication of the fact-checking on the Pagella Politica website (Monday afternoon) and it lasted five days (i.e., until the Friday evening of the week when the fact-checking was published).¹⁶

¹⁵An English version of the video can be found at the following link <https://tinyurl.com/283yut75>

¹⁶See also Galasso *et al.* (2022) for another scholarly work in collaboration with *Electica* also running a programmatic video advertisement.

Figure 3: Geo-targeted campaign



5 Empirical Strategy

As described in Section 4, the design of our experiment involves a staggered treatment of politicians: the first treated politician is fact-checked in intervention week one, the second in week two and so on. Given that the two-way fixed effects (TWFE) parameters are potentially biased due to the staggered nature of our treatment¹⁷ and that our design is essentially involving ten sequential “experiments”, we follow the stacked-diff-in-diffs approach proposed by Cengiz *et al.* (2019, 2022). Specifically, we create ten event-specific datasets (h) aligning events by event time (and not calendar time) and using only within event variation between the treated unit and clean control units. That is, each of the ten event-specific dataset is composed by one treated politician and her/his corresponding “clean” controls, that is, never treated and not-yet treated politicians. We then stack such ten event datasets. The alignment by event-time makes it equivalent to a setting where all events happened all at once and were not staggered (Cengiz *et al.*, 2022). Accordingly, we estimate:

$$Y_{h,i,t} = \beta D_{h,i,t} + \delta_{h,i} + \delta_{h,t} + \varepsilon_{h,i,t}$$

¹⁷See Goodman-Bacon (2018), Callaway and Sant’Anna (2020), De Chaisemartin and d’Haultfoeulle (2020), and Athey and Imbens (2021).

where $Y_{h,i,t}$ is the observed outcome of politician i (e.g., number of incorrect statements) at time t in the event-level h . Furthermore, $D_{h,i,t} = \mathbb{1}\{t \geq G_{h,i}\}$, where $G_{h,i}$ is time when i is fact-checked in event-level h . Finally $\delta_{h,i}$ and $\delta_{h,t}$ represent politician-event and time-event fixed effects, respectively. Standard errors are clustered at the politician-event level. As shown by Gardner (2022) such stacked difference-in-differences estimator is not affected by negative weighting bias in the presence of heterogeneous treatment effects.

6 Results

Table 3 presents our main results. Column 1 shows that fact-checking leads to fewer overall incorrect statements by politicians in the order of around a quarter of a standard deviation. Column 2 presents estimates when clustering standard errors at the politician level, rather than politician-event level as in the baseline specification. Column 3 provides estimates when using a simple two-way-fixed effects model (i.e., using a panel specification with time and politician fixed effects). Column 4 presents estimates when implementing the Callaway and Sant’Anna (2020) doubly robust estimator. Finally, in Column 5, as the number of incorrect statements is a count variable, we estimate a fixed effect Poisson model (Correia *et al.*, 2019) on the stacked-events sample.¹⁸

In Table 4, we assess both the robustness and external validity of our results by running the same analysis on a completely different sample. This sample is obtained from an analogous pilot experiment that we run between September and November 2020. Despite focusing on a different time period and a different set of control and treated politicians, the estimates are in line with the ones of our baseline analysis, lending support to the external validity and robustness of our results.¹⁹

Table 5 presents a series of additional exercises to further corroborate the robustness of our

¹⁸The number of observations is lower than that of Column 1 as the fixed effect Poisson model is estimated excluding observations that are either singletons or separated by fixed effects (Correia *et al.*, 2019). Also, notice that the point estimate of Colum 5 has to be interpreted analogously to that of a log-linear model; given the dichotomic nature of the treatment variable, it represents a semi-elasticity.

¹⁹The pilot experiment was conducted on a sample of 69 mid-rank politicians, the treatments occurred over 5 intervention weeks, and the experiment lasted 10 weeks in total.

Table 3: Negative Fact-Checking and Number of Incorrect Statements

	(1)	(2)	(3)	(4)	(5)
	CDLZ Baseline	CDLZ Cluster Pol	TWFE Model	CS DiD	Poisson Model
Fact-Checked	-0.378** (0.146)	-0.378** (0.146)	-0.370** (0.147)	-0.778*** (0.199)	-0.901*** (0.267)
Observations	8,035	8,035	880	880	3,715
Statistics on treated before treatment:					
Mean	0.67	0.67	0.67	0.67	0.67
SD	1.45	1.45	1.47	1.47	1.45
Politician-event FE	YES	YES	NO	NO	YES
Time-event FE	YES	YES	NO	NO	YES
Cluster SE at politician-event	YES	NO	NO	NO	YES
Politician FE	NO	NO	YES	YES	NO
Time FE	NO	NO	YES	YES	NO
Cluster SE at politician	NO	YES	YES	YES	NO

Notes: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Panel data at the politician-time-event level in columns 1, 2 and 5, and at the politician-time level in columns 3 and 4, over the period from March 22 to July 11, 2021 (16-weeks). *** p<0.01, ** p<0.05, * p<0.1

Table 4: Negative Fact-Checking and Number of Incorrect Statements - Pilot Dataset

	(1)	(2)	(3)	(4)	(5)
	CDLZ Baseline	CDLZ Cluster Pol	TWFE Model	CS DiD	Poisson Model
Fact-Checked	-0.400*** (0.101)	-0.400*** (0.108)	-0.399*** (0.107)	-0.655** (0.312)	-1.161** (0.525)
Observations	3,350	3,350	690	690	1,650
Statistics on treated before treatment:					
Mean	0.47	0.47	0.47	0.47	0.47
SD	0.72	0.72	0.74	0.74	0.72
Politician-event FE	YES	YES	NO	NO	YES
Time-event FE	YES	YES	NO	NO	YES
Cluster SE at politician-event	YES	NO	NO	NO	YES
Politician FE	NO	NO	YES	YES	NO
Time FE	NO	NO	YES	YES	NO
Cluster SE at politician	NO	YES	YES	YES	NO

Notes: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Panel data at the politician-time-event level in columns 1, 2 and 5, and at the politician-time level in columns 3 and 4, over the period from September 24 to November 24, 2020 (10-weeks). *** p<0.01, ** p<0.05, * p<0.1

baseline results. Column 1 shows that the number of incorrect statements decreases also when controlling for the number of verifiable statements made by the politician in a given week. Notice

that this specification includes an endogenous control as the number of verifiable statements is *per se* a decision variable that might be affected by the fact-checking (which we later analyze in Section 6.2). Yet, such a specification shows that the decrease in the number of incorrect statements is not simply a by-product of a possible decrease in the number of verifiable statements. That is, fact-checked politicians decrease the share of incorrect statements among the verifiable ones and thus increase the share of true or almost true statements. Column 2 presents estimates when adding an interaction terms between a dummy capturing whether a politician belongs to the randomization pool of a given event dataset and time-to-event fixed effects. Such specification is thus estimating the impact of fact-checking within the pool of randomizable politicians within each event-dataset and within each week to/from the fact-checking event. Columns 3 and 4 provides results when restricting the control group to never treated politicians only and when also restricting to an homogeneous time window across event-datasets ($-3/+3$ weeks from the fact-checking event), respectively. As Column 4 imposes a rather drastic reduction in the estimation window, as expected, it leads to less precise estimates. Column 5 excludes from the control group politicians who do not make any incorrect statements over the sample period, thus focusing on a more homogeneous sample. Finally, Column 6 shows that fact-checking has a similar impact on reducing falsehood when focusing on the probability of politicians making an incorrect statement.

Figure 4 presents the results of an event-study specification estimating leads and lags from/to the fact-checking event.²⁰ The graph shows the absence of significative pre-trends in the number of incorrect statements by fact-checked politicians. It also points out that the reduction in the number of incorrect statements induced by fact-checking is not short-lived.

Finally, we discuss the presence of potential spillovers of the fact-checking on control politicians since in principle politicians in the control group may also respond to the observed fact-checking on treated politicians. First, we point out that, if present, such spillovers are likely to “work against us” leading to a downward bias of our estimates on the impact of fact-checking on the number of incorrect statements of treated politicians. Second, we test whether there is any


²⁰Formally, we estimate the equation: $Y_{h,i,t} = \sum_{j=a}^b \beta_j D_{h,i,t}^j + \delta_{h,i} + \delta_{h,t} + \varepsilon_{h,i,t}$ where $D_{h,i,t}^j$ are leads and lags of the treatment variable defined in equation 5 with respect to the time of the fact-checking in the event-level h . Given the time structure of our experimental design, we fix $a = -6$ and $b = 8$. We follow McCrary (2008) and bind up end-points, results are similar for alternative windows and are available upon request.

Table 5: Negative Fact-Checking & Number of Incorrect Statements - Robustness

	(1)	(2)	(3)	(4)	(5)	(6)
	Control for Number Verifiable	Add int. time-event random. pol	Only Never treated	Hom. Time (-3/+3) (& only NT)	Exclude pol with no incorrect	Any Incorrect statement
Fact-checked	-0.320*** (0.110)	-0.268** (0.124)	-0.369** (0.145)	-0.287* (0.156)	-0.377** (0.148)	-0.228*** (0.063)
Observations	8,035	8,009	7,360	3,220	3,715	8,035
Statistics on treated before treatment:						
Mean	0.67	0.67	0.67	0.63	0.67	0.33
SD	1.45	1.45	1.47	1.29	1.45	0.47
Politician-event FE	YES	YES	YES	YES	YES	YES
Time-event FE	YES	YES	YES	YES	YES	YES

Notes: Column 1-4: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Column 4: the dependent variable is a dummy variable taking value one if the politician makes an incorrect statement in a week. Panel data at the politician-time-event level over the period from March 22 to July 11, 2021 (16 weeks). Standard errors in parentheses are clustered at the politician-event level. *** p<0.01, ** p<0.05, * p<0.1

Figure 4: Dynamic effect of fact-checking on the number of incorrect statements



Notes: Stacked panel data at the politician-week level over the period from March 22 to July 11 (16-weeks) per event level. Cluster standard errors at the politician-event level.

evidence of such potential spillovers. To do so, for each event-dataset we drop the fact-checked politician and attribute the treatment to a politician belonging to the same party of the fact-checked politicians or to the same party-chamber, including never-treated and not-yet treated

politicians in the control group. Table 6 presents the estimates of such exercise. The results show no evidence of potential spillover effects across party-peers or party-chamber peers.

Table 6: Negative Fact-Checking & Number of Incorrect Statements - Spillovers

	(1)	(2)	(3)	(4)
FC party peer	0.003 (0.022)	0.006 (0.020)		
FC party-chamber peer			0.005 (0.023)	0.009 (0.021)
Observations	4,416	4,416	5,600	5,600
Statistics on treated before treatment:				
Mean	0.08	0.08	0.05	0.05
SD	0.33	0.33	0.21	0.21
Control for N. FC	NO	YES	NO	YES
Politician-event FE	YES	YES	YES	YES
Time-event FE	YES	YES	YES	YES

Notes: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Panel data at the politician-time-event level over the period from March 22 to July 11, 2021 (16-weeks). Standard errors in parentheses are clustered at the politician-event level. *** p<0.01, ** p<0.05, * p<0.1


6.1 Placebo Test for Mean Reversion

One possible alternative interpretation of our results might be that rather than identifying the causal effect of fact-checking, we are just mechanically picking up a negative effect on false statements due to an underlying cycle in political communication (a reversion to the mean effect). If so, we should observe similar effect if we fictionally assign the fact-checking treatment to politicians after they made a false statement. To test this alternative hypothesis we perform a placebo test in the spirit of Random Inference Tests (Young, 2018; Dell and Olken, 2020). Specifically, we exploit the fact that in each intervention week there is a randomization pool of politicians that could have been subject to fact-checking in that week but they were not. Accordingly, we estimate a number of placebo counterfactuals by randomly and fictitiously attributing the treatment in week t to a politician in the randomization pool of that week excluding the actually treated politician.²¹ Figure 5 plots the distribution of estimated coefficients. The graph suggests that

²¹With respect to Dell and Olken (2020) we adopt a randomization without replacement since we are interested in estimating all possible combinations of placebo treatments across treatment weeks and politicians entering in

such reversion to the mean effect is unlikely to be the key driver of our observed effect. Indeed, none of the placebo coefficients is larger in magnitude than the actual estimates. Furthermore, the actual estimates are more than twice larger than the average placebo coefficient. The p-value is calculated as the fraction of the absolute value placebo coefficients greater than the absolute value of the actual estimated coefficient that is shown as a vertical line.²²

Figure 5: Placebo



6.2 Mechanism

We now turn to the discussion of potential mechanisms linking fact-checking to the observed reduction in the number of incorrect statements by fact-checked politicians. We proceed in steps trying to answer four separate yet related questions. First: Is the observed politicians' response driven by electoral concerns or is it due to a direct effect on politicians related to career or self-image concerns? Second: Why are fact-checked politicians behaving differently than control politicians after being exposed to the fact-checking, that is, in the ensuing continuation game?

the randomization pool.

²²The actual coefficient showed as a vertical line in figure 5 differs from what reported in column 1 of table 3 because for this analysis we do not take into account the last event level (i.e., the tenth intervention week). This choice is forced by the fact that in the last intervention week there were no politicians in the randomization pool but the actual treated. For this reason, we estimate both placebo coefficients and the actual benchmark by ruling out event level $h = 10$.

Third: Is the video-advertising campaign playing a key role in triggering the observed response by treated politicians? Fourth: Are fact-checked politicians changing other aspects of their political communication strategy, such as the number of statements and their verifiability?

To address the first question we exploit the fact that the Italian electoral system involves both single- and multi-member districts. As single-member districts are typically associated with strong electoral accountability, it is reasonable to expect politicians elected in these districts having a stronger electoral concern.²³ Finding a larger magnitude of the effects when looking at politicians elected in single-member districts would be evidence consistent with politicians responding to negative fact-checking because of their concerns about voters' reaction to such an event. Table 7 provides precisely this evidence.

Table 7: Heterogeneity (1): by politicians' characteristics

	(1) Single member district	(2) Multi members district	(3) Never Fact Checked	(4) Ever Fact Checked
Fact-checked	-0.277** (0.123)	-0.410** (0.200)	-0.134 (0.096)	-0.466** (0.185)
Observations	2,558	5,477	2,290	5,745
Statistics on treated before treatment:				
Mean	0.11	0.90	0.22	0.86
SD	0.32	1.66	0.42	1.67
Politician-event FE	YES	YES	YES	YES
Time-event FE	YES	YES	YES	YES

Notes: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Panel data at the politician-time-event level over the period from March 22 to July 11, 2021 (16-weeks). Standard errors in parentheses are clustered at the politician-event level. *** p<0.01, ** p<0.05, * p<0.1

In particular, Columns 1 and 2 present estimates when focusing on the subsample of politicians elected in single-member and multi-members districts, respectively. The results show that in the case of politicians elected in single-member districts, fact-checking reduces the number of incorrect statements in the order of 0.86 of a standard deviation. In the case of multi-member districts, such effect is in the order of 0.25 of a standard deviation. Accordingly, these results suggest

²³See, e.g., Buisseret and Prato (2017) and Raffler (2019).

that electoral concerns may indeed play a role in inducing politicians to reduce the number of incorrect statements. At the same time, electoral concerns are unlikely to be the only driver of such response for two main reasons. First, *Pagella Politica* does not have a huge audience (e.g., around 55000 followers on Instagram, 41000 on Facebook and 25000 on Twitter, respectively) and thus, while it might be used as an information source by major media outlets, it cannot play a pivotal role in influencing a large share of the electorate, even factoring in a reasonable multiplier effect. More importantly, as shown in Column 2 of Table 7, also politicians elected in multi-member districts respond to fact-checking, albeit to a lesser extent. As the electoral fate of such politicians largely depends on their positions in the ballot list, which is decided by the party, this latter evidence suggests that the observed effects are not only driven by electoral concerns but may be due to career or self-image concerns.

Next, we turn to the second question. We begin with a preliminary consideration. There are two conceivable counterfactuals when thinking about the potential impact of fact-checking. The first would require comparing a world where politicians face the risk of incurring in a negative fact-checking with one where they do not. Clearly assessing this type of counterfactual is unfeasible as all politicians face the risk of being fact-checked. The second counterfactual involves comparing the behavior of politicians after being fact-checked, with the behavior of those that were not fact-checked but could be fact-checked in the future. This counterfactual is the one that is observed in reality and that we focus on.²⁴ Yet one may wonder why treated politicians respond to fact-checking differently with respect to politicians in the control group as both groups face the same risk of being fact-checked in the future. That is, in principle, in the ensuing continuation game, a negative fact-checking at time t could well be irrelevant when comparing the future behavior of treated and control politicians. Furthermore, as shown in Table 6, there are no spillover effects across politicians. Accordingly, fact-checking does not seem to operate through a simple information channel (e.g., making politicians aware of being potentially fact-checked or increasing the salience of *Pagella Politica* per se). Rather, fact-checking has a specific impact

²⁴Notice that the observed response to this second counterfactual might be seen as a lower bound of the general response of politicians to the first counterfactual, i.e., to the effect on politicians behavior of the existence of fact-checking *per se*.

on the fact-checked politician. Here we discuss three possible narratives that may rationalize the observed differential response between treated and untreated politicians. First, politicians might have a convex cost from being repeatedly exposed to negative fact-checking (either directly due to convexity in self-image or career concerns, or indirectly due to voters becoming progressively and less tolerant with politicians reporting incorrect statements). This narrative is consistent with the fact that we observe more significant effects for politicians that have been already fact-checked in the past as compared to politicians never fact-checked. This difference is shown in Columns 3 and 4 of Table 7. The second possible narrative is that treated politicians revise upward their perceived probability of being fact-checked in the future. We do not have direct evidence to test this hypothesis. Yet, one would expect that such updating might be stronger for politicians who had never been fact-checked in the past, and this runs contrary to the evidence that we observe in Table 7. Finally, fact-checking might have a behavioral impact on politicians by priming normative concerns against “lying” (Nyhan and Reifler, 2015; Cohn and Maréchal, 2016). This mechanism is consistent with the stronger response observed among politicians who received “worse” fact-checking scores (see Table 8).²⁵

Table 8: Heterogeneity (2): by Fact-checking score

	(1) Half-false	(2) Mostly-false	(3) Utterly-false
Fact-checked	-0.159*** (0.039)	-0.500* (0.262)	-0.538*** (0.016)
Observations	3,209	4,022	804
Statistics on treated before treatment:			
Mean	0.67	0.67	0.67
SD	1.45	1.45	1.47
Politician-event FE	YES	YES	YES
Time-event FE	YES	YES	YES

Notes: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Panel data at the politician-time-event level over the period from March 22 to July 11, 2021 (16-weeks). Standard errors in parentheses are clustered at the politician-event level. *** p<0.01, ** p<0.05, * p<0.1

²⁵Appendix Table A.1 presents further heterogeneities results by political experience, political ideology and by whether politicians belong to a populists or non-populist party, following the classification by Rooduijn *et al.* 2019.

The third question concerns the role played by the video advertising campaign in triggering the observed response. This question is potentially relevant as the video advertising campaign is the only point of departure from how *Pagella Politica* normally advertises its contents. As such determining the relevance of this advertising “booster” is relevant for a proper policy evaluation of fact-checking. To this end, we exploit the heterogeneity present across treated politicians in terms of public exposure to the video ads. In particular, Table 9 illustrates how the response of politicians varies when looking at politicians above and below the median in terms of total video impressions, clicks and click-through-rate. The results do not suggest a clear pattern and, if anything, politicians in the lower end of the distribution in terms of video exposure seem to be the ones responding more. All in all, this suggests that video ads are not a key driver of the observed effects.

Table 9: Heterogeneity (3): Electica

	(1)	(2)	(3)	(4)	(5)	(6)
	Above median impressions	Below median impressions	Above median clicks	Below median clicks	Above median CTR	Below median CTR
Fact-checked	-0.449* (0.256)	-0.308** (0.135)	-0.137 (0.121)	-0.590*** (0.219)	-0.137 (0.121)	-0.590*** (0.219)
Observations	4,017	4,018	4,018	4,017	4,018	4,017
Statistics on treated before treatment:						
Mean	0.67	0.67	0.67	0.67	0.67	0.67
SD	1.45	1.45	1.45	1.45	1.45	1.45
Politician-event FE	YES	YES	YES	YES	YES	YES
Time-event FE	YES	YES	YES	YES	YES	YES

Notes: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Panel data at the politician-time-event level over the period from March 22 to July 11, 2021 (16-weeks). Standard errors in parentheses are clustered at the politician-event level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Finally, Table 10 provides evidence on other outcomes affected by fact-checking. Column 1 shows that politicians respond to fact-checking by also reducing their overall number of statements, in the order of one eighth of a standard deviation. Interestingly, Columns 4 and 5 provide evidence that politicians seem also to respond to fact-checking by reducing the likelihood of making any verifiable statement in a week. This is suggestive of a nuanced response by politicians

to fact-checking: a change in the type of political rhetoric adopted with a shift toward using non verifiable statements.

Table 10: Politicians’ Response to Negative Fact-Checking

	(1)	(2)	(3)	(4)	(5)
	Number of Statements	Number of Verifiable Statements	Number Verifiable Statements	Any of Verifiable Statements	Any Verifiable Statements
Fact-checked	-6.526** (2.857)	-0.417 (0.382)	-0.418 (0.381)	-0.153* (0.089)	-0.154* (0.088)
Observations	8,035	8,035	8,035	8,035	8,035
Statistics on treated before treatment:					
Mean	52.17	1.73	1.73	0.63	0.63
SD	49.54	2.44	2.44	0.48	0.48
Politician-event FE	YES	YES	YES	YES	YES
Time-event FE	YES	YES	YES	YES	YES
Control for Any ST	NO	NO	YES	NO	YES

Notes: Panel data at the politician-time-event level over the period from March 22 to July 11, 2021 (16-weeks). Standard errors in parentheses are clustered at the politician-event level. *** p<0.01, ** p<0.05, * p<0.1

7 Conclusion

We see our results as a first step in the direction of informing the academic and public debate on the impact of fact-checking on politicians’ behavior. We performed the first randomized field experiment on fact-checking in a business as usual environment as our intervention only changes the selection process of the statements to be fact-checked by the Italian leading fact-checking company. As such, our empirical design embeds two appealing features at the same time. First, the randomization allows us to overcome the potential biases arising from the endogenous selection of statements by fact-checkers. This, combined with a difference-in-differences analysis, allows to provide reliable causal estimates of the effect of fact-checking on politicians’ behavior. Second, as our design did not alter *per se* the politicians’ perception of the fact-checking process, the experimenter demand effect is not a concern in our setting and hence our results are likely to preserve a high external validity in other contexts.

Our results show that fact-checking discourages politicians from making factually incorrect

statements with effects lasting months. Since fact-checking is conducted on a regular basis, with a frequency typically increasing during electoral campaigns, our results provide the first evidence that fact-checking is indeed an effective tool in reducing the level of “factual misinformation” in politics. At the same time, we also document some unexpected drawbacks since fact-checked politicians are also less likely to make verifiable statements. This suggests that they deliberately increase the “ambiguity” of their language to escape the possibility of public scrutiny. The overall effect on welfare is therefore ambiguous. One of the pillars of representative democracy is a functioning mechanism of accountability and this mechanism also relies on the possibility of a constant and thorough scrutiny of politicians’ public statements. In this respect, constructing reliable individual level measures of ambiguity or non verifiability in political rhetoric seems a particularly interesting avenue for future research.

References

- ATHEY, S. and IMBENS, G. W. (2021). Design-based analysis in difference-in-differences settings with staggered adoption. *Journal of Econometrics*.
- BARRERA, O., GURIEV, S., HENRY, E. and ZHURAVSKAYA, E. (2020). Facts, alternative facts, and fact checking in times of post-truth politics. *Journal of Public Economics*, **182**, 104123.
- BUISSERET, P. and PRATO, C. (2017). Electoral accountability in multi-member districts. *Working Paper*.
- BURSZTYN, L. and JENSEN, R. (2017). Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure. *Annual Review of Economics*, **9**, 131–153.
- , RAO, A., ROTH, C. and YANAGIZAWA-DROTT, D. (2022). Opinions as facts. *Review of Economic Studies*, **forthcoming**.
- and YANG, D. Y. (2022). Misperceptions about others. *Annual Review of Economics*, **14**, 425–452.
- CALLAWAY, B. and SANT’ANNA, P. H. (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.
- CENGIZ, D., DUBE, A., LINDNER, A. and ZENTLER-MUNRO, D. (2022). Seeing beyond the trees: using machine learning to estimate the impact of minimum wages on labor market outcomes. *Journal of Labor Economics*, **40** (S1), S203–S247.
- , —, — and ZIPPERER, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, **134** (3), 1405–1454.
- COHN, A. and MARÉCHAL, M. A. (2016). Priming in economics. *Current Opinion in Psychology*, **12**, 17–21.
- CORREIA, S., GUIMARÃES, P. and ZYLKIN, T. (2019). Verifying the existence of maximum likelihood estimates for generalized linear models. *arXiv preprint arXiv:1903.01633*.
- DE CHAISEMARTIN, C. and D’HAULTFOEUILLE, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, **110** (9), 2964–96.
- DELL, M. and OLKEN, B. A. (2020). The development effects of the extractive colonial economy: The dutch cultivation system in java. *The Review of Economic Studies*, **87** (1), 164–203.
- DELLAVIGNA, S. and GENTZKOW, M. (2010). Persuasion: empirical evidence. *Annual Review of Economics*, **2** (1), 643–669.
- GALASSO, V., MORELLI, M., NANNICINI, T. and STANIG, P. (2022). Fighting populism on its own turf: Experimental evidence.

- GARDNER, J. (2022). Two-stage differences in differences. *arXiv preprint arXiv:2207.05943*.
- GODEL, W., SANDERSON, Z., ASLETT, K., NAGLER, J., BONNEAU, R., PERSILY, N. and TUCKER, J. A. (2021). Moderating with the mob: Evaluating the efficacy of real-time crowd-sourced fact-checking. *Journal of Online Trust and Safety*, **1** (1).
- GOODMAN-BACON, A. (2018). Difference-in-differences with variation in treatment timing. *NBER Working Paper*.
- GRAVES, L. and CHERUBINI, F. (2016). *The rise of fact-checking sites in Europe*. Reuters Institute for the Study of Journalism.
- HENRY, E., ZHURAVSKAYA, E. and GURIEV, S. (2022). Checking and sharing alt-facts. *American Economic Journal: Economic Policy*, **14** (3), 55–86.
- JAPKOWICZ, N. and STEPHEN, S. (2002). The class imbalance problem: A systematic study. *Intelligent data analysis*, **6** (5), 429–449.
- KUBAT, M., MATWIN, S. *et al.* (1997). Addressing the curse of imbalanced training sets: one-sided selection. In *Icml*, Citeseer, vol. 97, pp. 179–186.
- MCCARTHY, K., ZABAR, B. and WEISS, G. (2005). Does cost-sensitive learning beat sampling for classifying rare classes? In *Proceedings of the 1st international workshop on Utility-based data mining*, pp. 69–77.
- MCCRARY, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, **142** (2), 698–714.
- NYHAN, B. (2020). Facts and myths about misperceptions. *Journal of Economic Perspectives*, **34** (3), 220–36.
- , PORTER, E., REIFLER, J. and WOOD, T. J. (2020). Taking fact-checks literally but not seriously? the effects of journalistic fact-checking on factual beliefs and candidate favorability. *Political Behavior*, **42** (3), 939–960.
- and REIFLER, J. (2015). The effect of fact-checking on elites: A field experiment on us state legislators. *American Journal of Political Science*, **59** (3), 628–640.
- RAFFLER, P. (2019). Accountability and electoral systems. *Working Paper*.
- ROODUIJN, M., VAN KESSEL, S., FROIO, C., PIRRO, A., DE LANGE, S., HALIKIOPOULOU, D., LEWIS, P., MUDDE, C. and TAGGART, P. (2019). The populist: An overview of populist, far right, far left and eurosceptic parties in europe.
- SWIRE, B., BERINSKY, A. J., LEWANDOWSKY, S. and ECKER, U. K. (2017). Processing political misinformation: comprehending the trump phenomenon. *Royal Society open science*, **4** (3), 160802.

- TUCKER, J. A., GUESS, A., BARBERÁ, P., VACCARI, C., SIEGEL, A., SANOVICH, S., STUKAL, D. and NYHAN, B. (2018). Social media, political polarization, and political disinformation: A review of the scientific literature. *Political polarization, and political disinformation: a review of the scientific literature (March 19, 2018)*.
- WASIKOWSKI, M. and CHEN, X.-w. (2009). Combating the small sample class imbalance problem using feature selection. *IEEE Transactions on knowledge and data engineering*, **22** (10), 1388–1400.
- WINEBURG, S. and MCGREW, S. (2017). Lateral reading: Reading less and learning more when evaluating digital information.
- YOUNG, A. (2018). Channeling fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results. *The Quarterly Journal of Economics*, **134** (2), 557–598.

Appendix.

Table A.1: Heterogeneity: by experience, ideology & populism

	(1) First Legislature	(2) Experienced	(3) Right Wing	(4) Left Wing	(5) Populist	(6) Non Populist
Fact-checked	-0.351*** (0.123)	-0.388* (0.234)	-0.431** (0.188)	-0.186*** (0.051)	-0.410*** (0.157)	-0.071*** (0.024)
Observations	3,250	4,785	3,113	4,922	4,711	3,324
Statistics on treated before treatment:						
Mean	0.42	0.83	0.90	0.11	0.74	0.00
SD	0.64	1.78	1.66	0.32	1.51	0.00
Politician-event FE	YES	YES	YES	YES	YES	YES
Time-event FE	YES	YES	YES	YES	YES	YES

Notes: The dependent variable is the number of incorrect statements (grade 3 and below: half-false, mostly-false, utterly-false) made by a politician in a week. Panel data at the politician-time-event level over the period from March 22 to July 11, 2021 (16-weeks). Standard errors in parentheses are clustered at the politician-event level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$