Institute for Research in Economic and Fiscal issues

IREF Working Paper Series

Media and the Selection of Experts

Marek Hudík

Petr Špecián

IREF Working Paper No. 202207

September 2022

In English: en.irefeurope.org
In French: fr.irefeurope.org

Institute for Research in Economic and Fiscal issues

Media and the Selection of Experts

Marek Hudík

Faculty of Business Administration, University of Economics, Prague

IREF – Institute for Research in Economic and Fiscal Issues

Petr Špecián

Faculty of Humanities, Charles University, Prague

IREF – Institute for Research in Economic and Fiscal Issues

Abstract: The paper addresses the media filtering of experts. We consider a media firm that asks experts to assess if a given problem is major or minor so it can report the type of the problem. The media firm always first asks a generalist who can be contacted with zero cost but has limited accuracy. After observing the generalist's report, the media firm decides whether to contact a specialist. The specialist can identify the type of the problem with certainty; however, finding one is costly. We analyze how equilibria depend on the media search costs and the probability that the true state later reveals itself to the public. We demonstrate that the probability that the true state will ultimately be revealed in a way obvious to the lay public is the critical determinant of the accuracy of the reported expert testimony. If the revelation probability is low, strategizing by the experts and the media bias may prevent accurate expertise from being broadcasted. The paper suggests implementing techniques that improve the falsifiability of expert forecasts to mitigate the problem.

Keywords: markets in expertise, media bias, expert accuracy, selection of experts, division of knowledge

JEL: D83, C72, L82

Introduction

In any modern society, there is a high degree of division of knowledge which forces the individuals to rely extensively on the testimonies of others when forming their judgments and making their choices. Most pronounced is the knowledge asymmetry between laypeople and experts. As with any asymmetric relationship, trust-building is fraught with difficulties, and—especially in emergencies—the transmission of knowledge from experts to laypeople may break down. Such an "epistemic crisis" (Špecián 2022b) forces the individuals to choose while bereft of expert guidance. It paralyzes collective choice as the popular opinion fails to converge towards a shared understanding of the situation at hand.

Our paper addresses the internal dynamics of the markets in expertise with a particular emphasis on the role of media in filtering of experts. We highlight that laypeople do not engage in the selection of experts unaided. The legacy news brands—such as the TV, radio, or newspapers—and the contemporary news media based on the internet play an indispensable role in pre-selecting the experts whose testimony is to be broadcasted to the audience. Thus, the media supply the public with expertise in expert selection (Collins and Evans 2017). To analyze the challenges inherent in such a task, we model a situation where a media firm can select experts. We define an expert as a person, who "knows sufficiently more" than an "average" person to be able to help them solve problems and accomplish various tasks (cf. Goldman 2018, 2). Under this definition, almost everyone is expert in something (Koppl 2018). We assume that experts are either generalists or specialists. Generalists are inaccurate in their assessments but can be approached without cost;

identification of the highly accurate specialists, in turn, is costly. Moreover, both the generalists and the specialists may strategically manipulate their testimony since they compete for the media spotlight.

Several research questions drive our modeling exercise: Why do the media often fail to broadcast testimonies of the leading experts? What incentives do media face and how do these incentives influence the reliability of media's filtering of experts? Does competition among the experts always increase the accuracy of their testimony?

We demonstrate that the critical determinant of the degree of media and expert honesty is the probability that the truth will ultimately be revealed in a public fashion. The search costs, in contrast, only play a secondary role. With a low probability that the true nature of the problem is revealed, reducing the cost of finding a specialist does not necessarily improve the accuracy of media reports. There are two reasons for this. First, it becomes safe enough for the media firm to reap the fruits of exaggeration without being caught red-handed. Second, the specialist becomes motivated to strategically distort his testimony because contradicting the generalist may increase his visibility. Thus, the probability of truth-revelation manifests itself as the critical determinant of the accuracy of the media broadcasted expertise.

The rest of the paper is structured as follows. *Section 1* explains our approach to modeling the markets for expertise against the background of the existing scholarly literature. *Section 2* presents our model and its main results. *Section 3* discusses the model's implications and suggests techniques that may mitigate media bias in expert selection. *Section 4* addresses the limitation and possible future extensions of our model. *Section 5* concludes.

1. Filtering the Experts

Tetlock (2006) has famously demonstrated an abysmal forecasting performance of various pundits. His research is often used to support the wholesale rejection of the validity of any expertise (Goldman 2018). Nonetheless, while much can be claimed about its death (Nichols 2017), expertise appears more abundant today than ever. There are indications that even during the recent COVID-19 pandemic, which revealed new depths of the partisan division and led to unprecedented misinformation spread (Uscinski et al. 2020), the degree of trust in science and scientists has increased on the global scale (Wellcome Trust and Gallup 2020). At the same time, trust in scientists has been found as the most significant determinant driving the support and compliance with the pandemic restrictions (Algan et al. 2021). Therefore, an across-the-board skepticism regarding the value of expertise or the willingness of the public to follow expert advice appears inadvisable.

However, it must still be considered how expertise can be identified by those who lack the specialized epistemic tools to directly evaluate the reliability of expert claims, that is, by a vast majority of the democratic public (Goldman 2001). In the modern networked world, trusting experts is essentially a question of how expertise gets mediated to reach the mass audience. From the chaos of the often-contradictory expert voices, it is the media who select either the fruits of expert knowledge (such as specific findings and facts) to be highlighted or "boost the signal" of the experts who are allowed to communicate their opinions directly.

Earlier literature has focused on the fact-filtering (e.g., Gentzkow and Shapiro 2006; Shapiro 2016). What has not been previously examined is the dynamics among the experts, the media filtering of experts, or the probability that the truth will ultimately be revealed for everyone to see with their own eyes.

Addressing these important yet neglected themes, we take Tetlock's (2006) conclusions to imply that the incentives present in the market for expertise often do not stimulate media to treat accuracy as the key consideration in their expert selection. Media pundits—be they the grand narrative "hedgehogs" or the versatile "foxes"—lean towards generalism: the scope of the topics they address in their analyses is overly broad for them to claim thorough expertise in their particulars. This is in line with Sowell's (2011) complaint that the most famous pundits are quite different from the most prominent academic experts or with Duflo's and Banerjee's (2019) claim that the lack of trust in economics stems from the fact that the economists in the media are not actual experts but rather the voice of organized interest groups. Survey data also support such assessment. The public often thinks that economists "express views based on personal and political opinion than on verifiable data and analysis" (ING-Economics Network Survey, 2019). In this particular survey, only 16% of the respondents could name an economist. Where names were given, policymakers or personal financial advisors have been mentioned predominantly, with less than 1% of the respondents naming a researcher.

In the light of these considerations, we distinguish two types of experts: the generalists and the specialists. Generalists have lower overall accuracy since they hold views on a broad range of topics. It is essentially costless for a media organization to contact them since they are already at any journalist's speed dial. Specialists, in contrast, concentrate narrowly on a specific theme. This means high accuracy but low visibility for the media. Therefore, identifying an appropriate specialist is costly for the media firm: there are search costs as well as costs of introducing a new "expert brand" into the public arena. Generalists—once well-known among the public—command their own audience and offer a guarantee of a certain degree of attention (cf. Maćkowiak, Matějka, and Wiederholt 2021). Specialists possess no such attention capital. Also, their communication skills in addressing the mass public are untested and may prove mediocre. In short, for the media firm, working with a specialist tends to be significantly more troublesome than working with a generalist.

Concurrently, the experts' intrinsic accuracy is distinct from the honesty of their public testimony. Their epistemic advantage over laypeople allows them to pursue ends that are not necessarily compatible with those of their clients (Akerlof 1970). In short, experts—no matter how accurate—may be incentivized to distort their situation assessment. This makes expert-client trust-building a challenging task. If trust cannot be established, or if the public ends up trusting low-quality or dishonest experts, the resultant welfare losses can be significant. For instance, one may end up rejecting vaccination in the face of a global pandemic. Therefore, it is essential to examine whether the experts' reputation concerns suffice to enforce their honesty and if the media meta-expertise facilitates the identification of not only accurate but honest experts for the public.

As far as reputation is concerned, faulty assessments and forecasts can be costly for an expert, but the situation is far from straightforward. After all, even in Tetlock's (2006) classic study, poor forecasting performance scarcely undermined the pundits' professional success. Multiple explanations have been offered. Some point to the bounds of human rationality, such as the hindsight bias, that hinder the public from adequately evaluating the experts' track record, especially in an uncertain, probabilistic environment (Kahneman 2011). Others hint at the importance of politically motivated reasoning that leads the public to evaluate experts not based on their forecasting performance but their compatibility with the cherished signature beliefs of the competing social groups (Kahan 2015). Our model emphasizes the mixed incentives experts face when it comes to truthfulness. Firstly, experts compete for the media spotlight. Secondly, the public—and the media in its stead—is not concerned solely with accuracy.

Consider the spotlight. While not necessarily an intrinsic good, it improves one's opportunities and can be translated into career success or increased income with relative ease. Even in academia, "impact" and "outreach" are important indicators of one's performance. Both can be substantially enhanced if one gets access to a broad audience that the media have to offer. We thus postulate that both the generalist and the specialist experts benefit from the media attention and—since it is scarce—are set to compete for it. To some extent, our approach is similar to that of a recent study by Kurvers et al. (2021), who model competition for social influence between two rival influencers. However, we posit an asymmetric position between a generalist and a specialist expert in our model, emphasizing the sequential aspects of their competition. The intuition here is that generalists who regularly appear in the media had once started as specialists who built on the success of their initial appearance and broadened their thematic scope sacrificing some of their accuracy in the process. In the marketplace for expertise, the older generation of generalists becomes discredited or retires. Their successors are then recruited from the specialists who achieved success in interacting with the media.

The competition for the spotlight introduces mixed incentives for a specialist. When consulted by the media, the specialist can benefit from an increased chance of replacing the generalist as the media's "go-to-person." However, since merely confirming the generalist's assessment cements their standing without improving the specialist's odds of replacing them, the specialist's position is conflicted. Two strategies appear viable. Either the specialist acts as a watchdog that calls out errors in the generalist's assessment or as a contrarian who always promotes an opinion different from that of a generalist, independent of the accuracy of the generalist's initial assessment.

Generalists who defend their acquired position with the media need to be wary of the specialists' possible strategic behavior. At the same time, however, they must remain sensitive to the fact that accuracy is not the media's sole interest. Most obviously, the media are profit-maximizing firms, and the benefits of the improved accuracy may not always be worth the costs of searching for the best expert. Also, as examined by Shapiro (2016), the media can be captured by special interests or driven to misrepresent the actual distribution of expert assessments by the overemphasis on "balanced treatment" of issues. Finally, we turn our attention toward the fact that media biases tend to mirror the biases of their audience. If news consumers judge reports that conform to their expectations as of higher quality (Gentzkow and Shapiro 2006), anything that skews their expectations will also push the media to skew their selection of experts in the corresponding direction. In this vein, ideological bias, wishful thinking, or pessimism bias are worth consideration.

The ideological bias in the demand for expertise—be it triggered by rational pursuit of non-epistemic goals (Caplan 2008; Kahan 2015) or by rationality's bounds, such as the confirmation bias (Sunstein 2017)—pushes the experts to adjust their reports to align with the signature beliefs of the consumer groups the media firm serves. In the politically charged issues, generalists are incentivized to report dishonestly, i.e., the opposite of their actual assessment, to advertise their loyalty to a partisan cause (cf. Špecián 2022a, chap. 3). Wishful thinking, in its turn, creates a different media bias, namely a demand for expertise that underestimates the seriousness of social issues (Caplin and Leahy 2019). In particular, a fear of large-scale risks that cannot be reduced by private action or whose reduction would require drastic immediate sacrifices for a merely probabilistic and possibly quite distant future gain, such as climate change mitigation or nuclear disarmament, may trigger a preference for soothing reports. If so, generalists are pressured to underplay their risk assessment to pander to the wishful thinking of the masses. Finally, the opposite bias, namely pessimism bias, triggers demand for "doom and gloom" forecasts. In this vein, McCloskey (2017, 64) refers to a question asked by a Whig historian T. B. Macaulay in 1830: "On what principle is it, that when we see nothing but betterment behind us, we are to expect nothing but deterioration before us?" As the question implies, prophets of doom are often in high demand.

To explore these intricacies, we model a situation where there is a social problem with an uncertain degree of severity. For instance, consider the current debate regarding the threat of artificial intelligence misalignment (Russell 2019). The problem can be minor, such as if artificial general intelligence (AGI) cannot be constructed at all, or it cannot be constructed without having a “hardwired” appreciation of human values. Or the problem can be major if the AGI is not only technically feasible but also exceptionally difficult to align with human goals (Bostrom 2014).

For a generalist, who possesses a better epistemic position than the media or the public but remains uncertain about the correct answer, there exist four possible strategies in the face of a problem that could be either minor or major. At the same time, only one of these strategies is such that the generalist announces an honest assessment of the nature of the problem. The other strategies include a degree of dissimulation in line with the media biases suggested above. The generalist may report the opposite of her true assessment, i.e., a minor issue when she perceives a major problem and vice versa. But the distortion in the report can also be more one-sided. Here, we take guidance from Mann (2018) and distinguish between *prophets*, who take a dismal, Malthusian perspective on the prospects of humankind, and *wizards*, for whom cornucopia is just behind the corner and technical fixes are all that is needed to achieve it. Accordingly, a prophet strategy means that the generalist always reports a major problem; a wizard strategy means always reporting a minor problem.

Expert strategizing and media biases notwithstanding, mistakes can come with a cost. We posit that the public's biases are what skews the media expert selection and incentivizes the experts to distort their assessment. Nevertheless, this does not mean that evident mistakes will go unpunished. The public may well appreciate the media pandering to its wishful thinking or pessimism bias, but merely up to the moment when it becomes apparent that it has been misinformed. If a problem that had been underplayed blows up in the public's face, or if a problem against which costly precautions had been implemented ultimately reveals itself as inconsequential, there will be a reputational price to be paid by the media and the experts. In this vein, we follow up on Goldman's (2001) observation that expert statements, which are "esoteric," that is, unverifiable for the laypeople, occasionally turn "exoteric" as time goes by, and anyone can see if the expert assessment had been correct. Consider a moon eclipse: many a sage has lost a livelihood by miscalculating its occurrence.

Therefore, the critical question is if the truth will ultimately be revealed. Clearly, the revelation probability differs with various social issues. Its determinants include complexity, where the causal chains behind various events remain obscure, or reflexivity, where predictions themselves affect the underlying reality and turn self-fulfilling or self-defeating. Even if the fact of climate change becomes evident, the same does not hold for its cause. Even where a social problem is claimed to have been successfully prevented, such as the Y2K millennial computer meltdown, doubts may linger if there ever was a problem in the first place. In brief, since the probability of truth-revelation represents a crucial determinant of the trustworthiness of the experts and a constraint on their strategizing (cf. Koppl 2010, 221), we will treat it as a central variable in our model.

2. The model

Assume that a problem is major, with probability p, or minor, with probability $1 - p$. A media firm can access the opinion of a generalist expert who can imperfectly identify the type of problem. In each state, the generalist expert receives a signal M or m, with $P_G(M|\text{major}) = \alpha$, $P_G(M|\text{minor}) = 1 - \alpha$ and $P_G(m|\text{minor}) = \beta$ and $P_G(m|\text{major}) = 1 - \beta$, where $\alpha > 1 - p$ and $\beta > p$. The generalist reports either a major or a minor problem, and the media firm decides whether to accept this report or ask a specialist. If the media firm accepts the generalist's report, the true state is revealed with probability t, and the game ends. Finding a specialist involves search cost, S. The specialist can recognize the type of the problem with certainty, i.e., $P_S(M|\text{major}) = P_S(m|\text{minor}) = 1$. After the specialist reports the type of the problem, the true state is revealed with probability t, and the game ends. The timing of decisions is shown in Figure 1. Let us now specify the payoffs and analyze each player's decision.

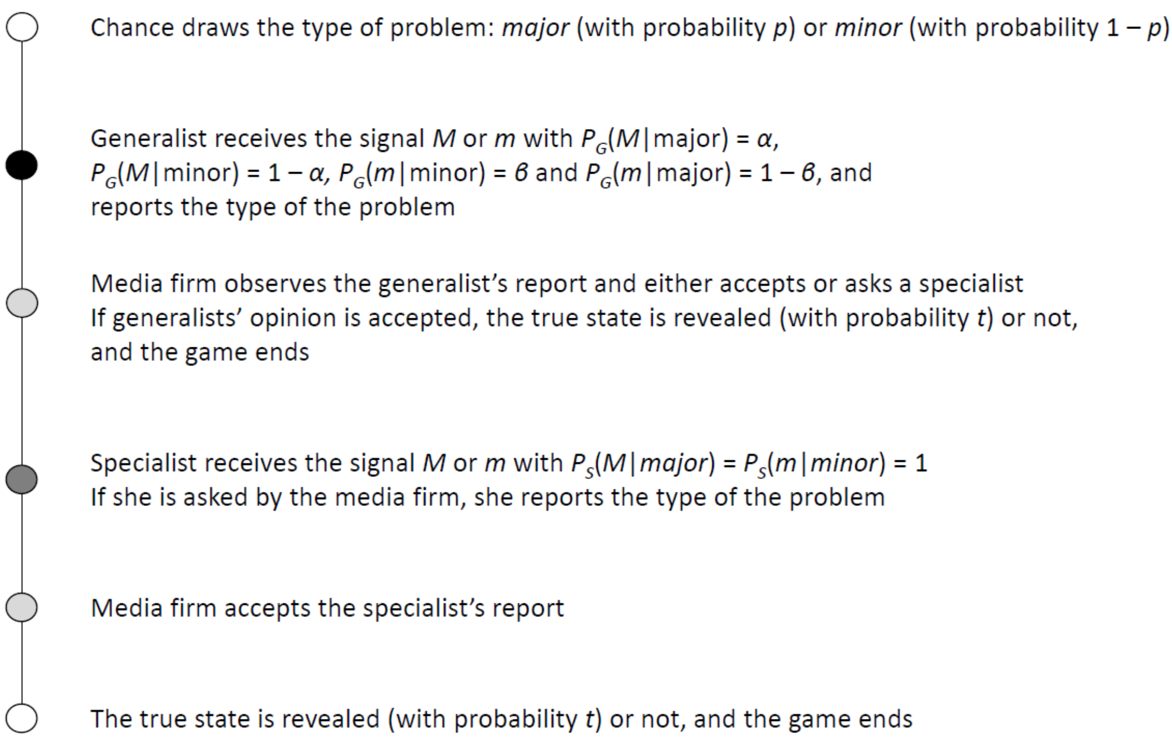

Figure 1: Timing of decisions

2.1. Specialist

The specialist's strategy specifies her decision in four situations: the generalist reports a major or minor problem, and the report is correct or incorrect. We assume that the specialist's payoffs are the same regardless of whether the generalist reports a major or minor problem. However, it makes a difference for a specialist if the generalist's report correctly identifies the type of the problem (Figure 2).

If the specialist confirms the generalist's correct report, she receives 0 since she neither benefits from the media spotlight nor risks her reputation. If the specialist opposes the generalist's correct opinion, she receives $v > 0$, given the true state is not revealed, or $-r < 0$, given the true state is revealed. The rationale behind v is that the media are potentially interested in the specialist's opinion only if her opinion differs from the opinion of the generalist. The positive value of v means that the specialist cares about the media interest. r represents the reputational costs of committing a public blunder.

If the specialist confirms the generalist's incorrect opinion, she receives 0, given the true state is not revealed, or $-r$, given the true state is revealed. If the specialist opposes the generalist's incorrect opinion, she receives v, if the true state is not revealed, or w, with $w > 0$, if the true state is revealed.

		Generalist's report	
		Correct	Incorrect
Specialist	Confirm	0	$-rt$
	Oppose	$v(1-t)-rt$	$v(1-t)+wt$

Figure 2: Specialist's payoffs

The specialist's best-response strategy is specified by the following proposition:

Proposition 1*. If $t \leq v/(v + r)$, the specialist's best response strategy is always to oppose the generalist's report. This strategy shall be called Contrarian. If $t \geq v/(v + r)$, the specialist's best response strategy is to confirm the generalist's correct report and oppose the generalist's incorrect report. This strategy shall be called Watchdog.*

Proof. See Figure 2.

2.2. Media firm

For conciseness, we limit our exploration to the pessimism bias. Accordingly, the media firm is motivated to report a major problem rather than a minor one in our model. The media firm thus receives $B > 0$ if it reports the major problem correctly or if it incorrectly identifies a minor problem as major, but the true state is not revealed. If the minor problem was incorrectly identified as major and the true state is revealed, the media firm receives $-c < 0$. The media firm receives 0 if it reports the minor problem correctly or if it incorrectly identifies a major problem as minor, but the true state is not revealed. If the major problem was incorrectly identified as minor and the true state is revealed, the media firm receives $-C < 0$. The rationale behind the assumption $B > 0$ is that the media firm will attract more consumers if it reports major problems than when it reports minor problems.

The media firm only has access to the public information on whether the problem is major (i.e., to probability p). However, in the perfect Bayesian equilibrium, the generalist's behavior can provide the firm with additional information about the type of the problem. Therefore, we consider four different situations based on the strategy of the generalist: he always reports a major problem (strategy *Prophet*); he always reports a minor problem (strategy *Wizard*); he reports a major problem if he receives the signal M and a minor problem if he receives the signal m (strategy *Honest*); he reports a minor problem if he receives M and a major problem if he receives m (strategy *Dishonest*).

2.2.1. Generalist chooses *Prophet*

The generalist's strategy *Prophet* does not provide the media firm with any new information about the true state. It believes that the problem is major with the probability p and minor with the probability $1 - p$. The payoffs are shown in Figure 3.

		Specialist	
		Contrarian	*Watchdog*
Media firm	*Accept*	$[1 - (1 - p)t]B - (1 - p)tc$	$[1 - (1 - p)t]B - (1 - p)tc$
	Ask	$- ptC - S$	$pB - S$

Figure 3: Media firm's payoffs if the generalist's strategy is *Prophet*

The media firm's best-response strategy is specified in the following proposition:

Proposition 2.

a) Assume that the specialist's strategy is Contrarian. If $(1 - p)tc \geq [1 - (1 - p)t]B + S + ptC$*, the media firm's best response strategy involves asking the specialist if the generalist reports that the problem is major. If* $(1 - p)tc \leq [1 - (1 - p)t]B + S + ptC$*, the media firm's best response strategy involves accepting the generalist's report that the problem is major.*

b) Assume that the specialist's strategy is Watchdog. If $(1 - p)tc \geq (1 - p)(1 - t)B + S$*, the media firm's best response strategy involves asking the specialist if the generalist reports that the problem is major. If* $(1 - p)tc \leq (1 - p)(1 - t)B + S$*, the media firm's best response strategy involves accepting the generalist's report that the problem is major.*

Proof. See Figure 3.

2.2.2. Generalist chooses strategy *Wizard*

The generalist's strategy *Wizard* again does not provide the media firm with any new information about the type of problem as the expert always reports minor problem. The media firm's payoffs are shown in Figure 4.

		Specialist: *Contrarian*	Specialist: *Watchdog*
Media firm	*Accept*	$-ptC$	$-ptC$
	Ask	$[1-(1-p)t]B-(1-p)tc-S$	$pB-S$

Figure 4: Media firm's payoffs if the generalist's strategy is *Wizard*

The media firm's best-response strategy is specified in the following proposition.

Proposition 3.

a) Assume that the specialist's strategy is Contrarian. If $ptC + [1-t(1-p)]B \geq (1-p)tc + S$*, the media firm's best response strategy involves asking the specialist if the generalist reports that the problem is minor. If* $ptC + [1-t(1-p)]B \leq (1-p)tc + S$*, the media firm's best response strategy involves accepting the generalist's report that the problem is minor.*

b) Assume that the specialist's strategy is Watchdog. If $ptC + pB \geq S$, the media firm's best response strategy involves asking the specialist if the generalist reports that the problem is minor. If $ptC + pB \leq S$, the media firm's best response strategy involves accepting the generalist's report that the problem is minor.

Proof. See Figure 4.

2.2.3. Generalist chooses strategy *Honest*

In the perfect Bayesian equilibrium, the generalist's strategy *Honest* provides the media firm with information about the true state. If the media firm observes that the generalist reports major, it believes that the problem is indeed major with the following probability: $Q \equiv p\alpha/[p\alpha + (1 - p)(1 - \alpha)]$. If the generalist reports a minor problem, the media firm believes that the problem is minor with the following probability: $q \equiv (1 - p)\beta/[(1 - p)\beta + p(1 - \beta)]$. Since we assume that $\alpha > 1 - p$ and $\beta > p$, then $Q > 1/2$ and $q > 1/2$. This means that if the generalist receives *M*, the posterior probability of a major problem is greater than the posterior probability of a minor problem. Accordingly, if the generalist receives *m*, the posterior probability of a minor problem is greater than the posterior probability of a major problem.

The media firm's payoffs are shown in Figure 6.

		Specialist	
		Contrarian	*Watchdog*
Media firm	*Accept*	$[1 - (1 - Q)t]B - (1 - Q)tc$	$[1 - (1 - Q)t]B - (1 - Q)tc$
	Ask	$- QtC - S$	$QB - S$

a) Generalist reports a major problem

		Specialist	
		Contrarian	*Watchdog*
Media firm	*Accept*	$- (1 - q)tC$	$- (1 - q)tC$
	Ask	$(1 - qt)B - qtc - S$	$- S$

b) Generalist reports a minor problem

Figure 6: Media firm's payoffs if the generalist's strategy is *Major if M*, *Minor if m*

The media firm's best-response strategy is specified in the following proposition.

***Proposition 4*.**

a) Assume that the specialist's strategy is Contrarian and the generalist honestly reports a major problem. If $(1 - Q)tc \geq [1 - (1 - Q)t]B + S + QtC$, *the media firm's best response strategy involves asking the specialist. If* $(1 - Q)tc \leq [1 - (1 - Q)t]B + S + QtC$, *the media firm's best response strategy involves accepting the generalist's report.*

b) Assume that the specialist's strategy is Watchdog and that the generalist honestly reports a major problem. If $(1 - Q)tc \geq (1 - Q)(1 - t)B + S$, *the media firm's best response strategy* *involves asking the specialist. If* $(1 - Q)tc \leq (1 - Q)(1 - t)B + S$*, the media firm's best response strategy involves accepting the generalist's report.*

c) Assume that the specialist's strategy is Contrarian and the generalist honestly reports a minor problem. If $(1 - q)tC + (1 - qt)B \geq qtc + S$*, the media firm's best response strategy involves asking the specialist. If* $(1 - q)tC + (1 - qt)B \leq qtc + S$*, the media firm's best response strategy involves accepting the generalist's report.*

d) Assume that the specialist's strategy is Watchdog and that the generalist honestly reports a minor problem. If $(1 - q)tC \geq S$*, the media firm's best response strategy involves asking the specialist. If* $(1 - q)tC \leq S$*, the media firm's best response strategy involves accepting the generalist's report.*

Proof. See Figure 6.

2.2.4. Generalist chooses strategy *Dishonest*

If the generalist's strategy is *Dishonest*, then in the perfect Bayesian equilibrium, the media firm believes that the problem is major with probability $1 - q$ if the generalist reports major. If the generalist reports a minor problem, the media firm believes that the problem is minor with the probability $1 - Q$. The media firm's payoffs are shown in Figure 7.

		Specialist	
		Contrarian	*Watchdog*
Media firm	*Accept*	$(1 - qt)B + qt(- c)$	$(1 - qt)B + qt(- c)$
	Ask	$- (1 - q)tC - S$	$(1 - q)B - S$

a) Generalist reports a major problem

		Specialist	
		Contrarian	*Watchdog*
Media firm	*Accept*	$-QtC$	$-QtC$
	Ask	$[1-(1-Q)t]B-Qtc-S$	$-S$

b) Generalist reports a minor problem

Figure 7: Media firm's payoffs if the generalist's strategy is *Minor, Major*

The media firm's best-response strategy is specified in the following proposition.

***Proposition 5*.**

a) Assume that the specialist's strategy is Contrarian and the generalist dishonestly reports a major problem. If $qtc \geq (1-qt)B+S+qtC$*, the media firm's best response strategy involves asking the specialist. If* $qtc \leq (1-qt)B+S+qtC$*, the media firm's best response strategy involves accepting the generalist's report.*

b) Assume that the specialist's strategy is Watchdog and that the generalist dishonestly reports a major problem. If $qtc \geq q(1-t)B+S$*, the media firm's best response strategy involves asking the specialist. If* $qtc \leq q(1-t)B+S$*, the media firm's best response strategy involves accepting the generalist's report.*

c) Assume that the specialist's strategy is Contrarian and the generalist dishonestly reports a minor problem. If $QtC+[1-(1-Q)t]B \geq (1-Q)tc+S$*, the media firm's best response strategy* *involves asking the specialist. If $QtC + [1 - (1 - Q)t]B \leq (1 - Q)tc + S$, the media firm's best response strategy involves accepting the generalist's report.*

d) Assume that the specialist's strategy is Watchdog and that the generalist dishonestly reports a minor problem. If $QtC \geq S$, the media firm's best response strategy involves asking the specialist. If $QtC \leq S$, the media firm's best response strategy involves accepting the generalist's report.

Proof. See Figure 7.

2.3. Generalist

The generalist receives 0 if he correctly identifies the problem or if he identifies the problem incorrectly, but the true state is not revealed. If he identifies the problem incorrectly and the true state is revealed, he receives $-R < 0$. If his report is opposed by the specialist and the true state is not revealed, he receives $-L < 0$.

2.3.1. Specialist chooses strategy *Contrarian*

We first consider the situation when the specialist's strategy is *Contrarian*. The situation is shown in Figure 8. Columns represent the media firm's strategies.

	Always accept	Ask if minor	Ask if major	Always ask
Major	$-(1-Q)tR$	$-(1-Q)tR$	$-(1-t)L-(1-Q)tR$	$-(1-t)L-(1-Q)tR$
Minor	$-QtR$	$-(1-t)L-QtR$	$-QtR$	$-(1-t)L-QtR$

a) Generalist receives M

	Always accept	Ask if minor	Ask if major	Always ask
Major	$-qtR$	$-qtR$	$-(1-t)L-qtR$	$-(1-t)L-qtR$
Minor	$-(1-q)tR$	$-(1-t)L-(1-q)tR$	$-(1-q)tR$	$-(1-t)L-(1-q)tR$

b) Generalist receives m

Figure 8: Generalist's payoffs if specialist's strategy *Contrarian*

The following proposition specifies the generalist's best-response strategy.

Proposition 6. *Assume that the specialist's strategy is Contrarian.*

a) Assume that the media firm's strategy is Always accept. Then the generalist's best-response strategy involves reporting a major problem if $QtR \geq (1-Q)tR$ *and a minor problem if* $QtR \leq (1-Q)tR$ *after he receives M. If he receives m, the best-response strategy involves reporting a major problem if* $qtR \leq (1-q)tR$*, and a minor problem if* $qtR \geq (1-q)tR$*.*

b) Assume that the media firm's strategy is Ask if minor. Then the generalist's best-response strategy involves reporting a major problem after he receives M. If he receives m, the best-response strategy involves reporting a major problem if $(1-t)L \geq (2q-1)tR$*, and a minor problem if* $(1-t)L \leq (2q-1)tR$*.*

c) Assume that the media firm's strategy is Ask if major. Then the generalist's best-response strategy involves reporting a major problem if $(2Q-1)tR \geq (1-t)L$ *and a minor problem if* $(2Q-1)tR \leq (1-t)L$ *after he receives M. If he receives m, the best-response strategy involves reporting a minor problem.*

d) Assume that the media firm's strategy is Always ask. Then the generalist's best-response strategy involves reporting a major problem if $QtR \geq (1 - Q)tR$ and a minor problem if $QtR \leq (1 - Q)tR$ after he receives M. If he receives m, the best-response strategy involves reporting a major problem if $qtR \leq (1 - q)tR$, and a minor problem if $qtR \geq (1 - q)tR$.

Proof. See Figure 8.

2.3.2. Specialist chooses strategy *Watchdog*

Figure 9 then shows the situations where the specialist's strategy is *Watchdog*.

	Always accept	Ask if minor	Ask if major	Always ask
Major	$-(1-Q)tR$	$-(1-Q)tR$	$-(1-Q)[(1-t)L + tR]$	$-(1-Q)[(1-t)L + tR]$
Minor	$-QtR$	$-Q[(1-t)L + tR]$	$-QtR$	$-Q[(1-t)L + tR]$

a) Generalist receives *M*

	Always accept	Ask if minor	Ask if major	Always ask
Major	$-qtR$	$-qtR$	$-q[(1-t)L + tR]$	$-q[(1-t)L + tR]$
Minor	$-(1-q)tR$	$-(1-q)[(1-t)L + tR]$	$-(1-q)tR$	$-(1-q)[(1-t)L + tR]$

b) Generalist receives *m*

Figure 9: Generalist's payoffs if specialist's strategy is *Watchdog*

The following proposition specifies the generalist's best-response strategy.

Proposition 7. *Assume that the specialist's strategy is Watchdog.*

a) Assume that the media firm's strategy is Always accept. Then the generalist's best-response strategy involves reporting a major problem if $QtR \geq (1 - Q)tR$ *and a minor problem if* $QtR \leq (1 - Q)tR$ *after he receives M. If he receives m, the best-response strategy involves reporting a major problem if* $qtR \leq (1 - q)tR$*, and a minor problem if* $qtR \geq (1 - q)tR$*.*

b) Assume that the media firm's strategy is Ask if minor. Then the generalist's best-response strategy involves reporting a major problem after he receives M. If he receives m, the best-response strategy involves reporting a major problem if $(1 - q)(1 - t)L \geq (2q - 1)tR$*, and a minor problem if* $(1 - q)(1 - t)L \leq (2q - 1)tR$*.*

c) Assume that the media firm's strategy is Ask if major. Then the generalist's best-response strategy involves reporting a major problem if $(2Q - 1)tR \geq (1 - Q)(1 - t)L$ *and a minor problem if* $(2Q - 1)tR \leq (1 - Q)(1 - t)L$ *after he receives M. If he receives m, the best-response strategy involves reporting a minor problem.*

d) Assume that the media firm's strategy is Always ask. Then the generalist's best-response strategy is Honest.

Proof. See Figure 9.

2.4. Equilibria

There are sixteen outcomes that can constitute a perfect Bayesian equilibrium depending on the values of parameters (see Table 1 and Appendix). There are five equilibria in which media report honestly, either because they accept an honest report by the generalist, or the specialist, or both (III, XI, XIII, XV, and XVI). In all these equilibria, the generalist chooses *Honest*, even in the equilibria where the media firm asks a specialist.

Specialist	**Media**	**Generalist**	**Equilibria**	***S* = 0**	***t* = 1**	***t* = 0**
Contrarian	Always accept	Wizard	I-equilibrium [NS]			Y
		Prophet	II-equilibrium [NS]			Y
		Honest	III-equilibrium	Y		Y
		Dishonest	-			
	Ask if minor	Wizard	-			
		Prophet	IV-equilibrium	Y		Y
		Honest	V-equilibrium	Y		
		Dishonest	VI-equilibrium	Y		
	Ask if major	Wizard	VII-equilibrium	Y		
		Prophet	-			
		Honest	VIII-equilibrium	Y		
		Dishonest	-			
	Always ask	Wizard	-			
		Prophet	-			
		Honest	IX-equilibrium	Y		
		Dishonest	X-equilibrium [NS]	Y		Y
Watchdog	Always accept	Wizard	-			
		Prophet	-			
		Honest	XI-equilibrium		Y	
		Dishonest	-			
	Ask if minor	Wizard	-			
		Prophet	XII-equilibrium	Y	Y	
		Honest	XIII-equilibrium	Y	Y	
		Dishonest	-			
	Ask if major	Wizard	XIV-equilibrium		Y	
		Prophet	-			
		Honest	XV-equilibrium		Y	
		Dishonest	-			
	Always ask	Wizard	-			
		Prophet	-			
		Honest	XVI-equilibrium	Y	Y	
		Dishonest	-			

NS refers to equilibria that are never strict.
Equilibria, in which media report honestly (but perhaps inaccurately), are highlighted in grey.

Table 1: Perfect Bayesian equilibria

XVI-equilibrium stands out because, in this equilibrium, the media firm reports the specialist's accurate and honest information. This equilibrium is attained if t is high enough for the specialist to choose *Watchdog* [$t \geq v/(v + r)$] while the generalist chooses *Honest* [$(1 - Q)(1 - t)L \leq (2Q - 1)tR$]. At the same time, the expected cost of misreporting for the media firm exceeds the expected benefits from exaggerating and its search cost [$(1 - Q)tc \geq (1 - Q)(1 - t)B + S$ and $(1 - q)tC \geq S$].

If the specialist chooses *Contrarian* and $S \geq 0$ and $t > 0$, then the media firm mediates dishonest reports in all the equilibria except III. In this equilibrium, the media firm always accepts the generalist's honest report because the expected benefits from exaggerating the problem and the expected costs of underestimating the problem are lower than the search costs and the expected costs from exaggerating the problem [$(1 - q)tC + (1 - qt)B \leq qtc + S$].

Notably, zero search cost ($S = 0$) does not imply that the media will always ask the specialist to obtain a more accurate report. The reason is that the media firm prefers to report a major rather than a minor problem ($B > 0$).

Proposition 8. *Assume $S = 0$, and $t > v/(v + r)$. Then there is no perfect Bayesian equilibrium in which the media firm accepts the generalist's report that the problem is minor.*

Proof. See the Table 1 and Appendix.

Note that the validity of Proposition 8 is conditional on the specialist choosing *Watchdog* [i.e., if $t > v/(v + r)$]. Therefore, the specialist will oppose the generalist's report that the problem is minor whenever this report is incorrect, which benefits the media firm. However, when the specialist chooses *Contrarian* [$t < v/(v + r)$], she will oppose all of the generalist's reports. If so, the media may or may not benefit from the report. The outcome depends on the relevant probabilities, the cost of reporting the wrong type of problem, and the benefit of reporting a major problem. More specifically, if $S = 0$, all equilibria with a contrarian specialist are feasible, except for the two non-strict equilibria I and II that require that the true state is never revealed ($t = 0$), but the media firm always accepts the generalist's report, nevertheless.

We now examine how the probability that the true state is revealed influences the feasibility of the equilibria.

Proposition 9. *Assume that the true state is always revealed ($t = 1$).*

a) In all feasible equilibria, the specialist chooses the strategy Watchdog.

b) For any media firm strategy, there is a feasible equilibrium in which the generalist chooses the strategy Honest.

Proof. See the Table 1 and Appendix.

If the true state is revealed with certainty, the generalist, the specialist, and the media firm are incentivized to report honestly to avoid a reputation loss. However, for certain parametric values, the media firm also accepts *Prophet* (equilibrium XII) or *Wizard* (equilibrium XIV) generalist's report. The intuition behind these equilibria is that if the generalist does not exaggerate (*Prophet*) or underestimate (*Wizard*) the problem, the media firm will ask the specialist since the expected cost of incorrectly identifying major (*Ask if major*) or minor (*Ask if minor*) problem is greater than the search cost. Note that the feasibility of these equilibria is independent of the assumption that media tend to exaggerate the problem ($B > 0$). Here, parameter B only affects the threshold when the media firm asks the specialist.

As our analysis suggests, if the aim is to pressure the media firm to report honestly, increasing the probability that the true state is revealed bears a greater promise than lowering the cost of finding the specialist (S) or increasing the reputational cost for the specialist (r). Low S enables the media to use the presence of the specialist opportunistically to increase the pressure on the generalist to provide a report to the media's liking. High r only disciplines the specialist, but not the other agents. High t, on the other hand, motivates all the agents—i.e., the specialist, the generalist, and the media firm—towards honesty. Still, neither $t = 1$ nor $S = 0$ provide its full guarantee, as *Table 1* demonstrates.

In contrast, if the true state is never revealed ($t = 0$), there exist merely five feasible equilibria (I, II, III, IV, and X). The following proposition establishes their properties.

Proposition 10. *Assume $t = 0$.*

a) In all feasible equilibria, the specialist chooses the strategy Contrarian.

b) There is a unique feasible strict equilibrium, (Prophet; Ask if minor; Contrarian).

Proof. See Table 1 and Appendix.

The intuition behind the IV-equilibrium *(Prophet; Ask if minor; Contrarian)* is simple: since $t = 0$, players are not concerned with reputational cost. Therefore, the specialist chooses the *Contrarian* strategy since it improves her chance of displacing the generalist. The media firm's preference to report a major problem also becomes unhindered. If the search cost is sufficiently low ($B \geq S$), the media firm chooses the strategy *Ask if minor*. Since the generalist suffers a loss if his report is contradicted by the specialist ($L > 0$), he always reports a major problem.

3. Discussion

As suggested by our results, the most important determinant of the experts' and media behavior is the probability that the truth will be publicly revealed. The specialist, who is tempted toward opportunistic contrarianism by her competition for the spotlight with the generalist, only resorts to dishonest reporting—that is, to the contrarian strategy—if the probability of the truth being revealed (t) is sufficiently low. Specifically, she weights the t relative to the reputational concerns and the benefits of the media attention as described by the inequality $t \lesseqgtr v/(v + r)$. With high enough t, she chooses the watchdog strategy, which in turn stimulates the generalist towards honesty across a broader range of parameter values.

In contrast, the extreme case where the truth is never revealed not only motivates the specialist towards choosing the contrarian strategy in her reporting but opens more possibilities for the media to shape the expert testimony to conform to its bias toward overblowing the social problems. Given that the search costs of identifying the specialist are lower than the benefit of reporting a major problem—which does not appear to be an exceedingly strong assumption—only the IV-equilibrium remains feasible. Here, the media firm chooses to *Ask if minor*, effectively threatening the generalist that his report will be contradicted by a contrarian specialist if he were to report a minor one. The generalist then chooses the *Prophet* strategy.

At the same time, low search costs (S) do not fundamentally alter the game towards honest behavior by its participants: they neither imply that the media will always contact a specialist nor that it will always broadcast an honest report. Notably, $S = 0$ disproportionally eliminates the equilibria that include the specialist's watchdog strategy. It also makes it simpler for the media to opportunistically doubt the generalist's report if it inconveniently states that the problem is minor. The only constraint upon such opportunism is provided by the reputational costs the media would face if the truth were revealed and their misreporting became obvious.

What are the implications of these results? We take our discussion as a contribution to the debate on the welfare implications of the competition among experts. In many contexts, such as in the paradigmatic case of perfect competition, competition is welfare enhancing. It disciplines the market agents and prevents them from exploiting their position for a selfish gain, i.e., from resorting to a strategy that creates social waste. In the marketplace of expertise, the situation is more nuanced, however. What determines the outcomes in terms of the welfare of the consumers of expert advice is not so much the degree of competition among the experts as the ability of the public to independently judge the accuracy of the expert testimony, if only in hindsight. Where the judgment is impossible for a non-expert (t = 0), even the generalists' strong incentives towards honesty can be easily undermined by the threat of rival specialists' strategizing. Also, the media's power to shape the expert opinion is bolstered significantly.

As we show, the t-parameter is thus far more important that the search costs. This may explain why the decrease of the search costs introduced by the digital technologies that make any information—including the information regarding the availability of the specialists in virtually any field—more accessible, did not lead to the improvement of the reporting quality. Quite the opposite, the digital era may have led to a more polarizing and ideological behavior by the media firms. This is consistent with our findings, especially if the t was not fortuitously affected by the digitalization. In this context, it appears that the jury is still out: on the one hand, witness reports are more easily collected, transmitted, and supported by audio-visual evidence (Tufekci 2017; Gurri 2018); on the other hand, it has also become simpler to cherry-pick the testimonies and manipulate the evidence in sophisticated ways (Schick 2020).

Be that as it may, the determinants of t are crucially important. In our model, t is exogenous but not necessarily the same for various issues. Problems differ in their inherent propensity to reveal their nature in plain sight. For instance, the existence of a pandemic does become evident at some point, while the severity of its long-term health outcomes—perhaps with the grim exception of deaths—may not, especially if the frequency of their occurrence remains below a certain threshold. When it comes to the concerns of misinformation by the media, we should thus be especially cautious with the kinds of problems where the assessment of their significance requires attention to subtle hints and a substantial amount of cognitive effort. For such issues, the public discourse is more likely to lose touch with reality and fail to identify the relevant expertise as well as to discipline the media filtering of experts.

At the same time, t is exogenous by assumption only. In fact, it is codetermined by how expert testimonies are framed (cf. Kelly 2012). As Tetlock and Gardner (Tetlock and Gardner 2015) have shown, one of the main problems with expert testimonies is that they are often deliberately vague and equivocal. They tend to lack specificity regarding effect magnitudes, probabilities, or time horizons. As ancient prophecies, they are prone to ambiguity while leaving themselves open to reinterpretation. One may claim, for instance, that his forecasts were meant as conditional upon an unspecified set of preconditions. In short, the experts often intentionally strive to make their predictions unfalsifiable. Our model also suggests that the media cannot be expected to push for corrections in this regard. To the contrary, they are incentivized to become complicit since pushing t to zero increases their ability to deliver such expert testimony that best panders to their biases and the biases of their audience.

This is quite unfortunate, since the media could do much to increase the *t* of various social issues. If successful, they could significantly bolster the public's benefits from the existing expertise. The problem is that the bias towards exaggeration stands in their way and undermines their credibility in the filtering of experts. That is, they are compromised by their mixed incentives. At the same time, fixing the incentives is challenging since the bias is not necessarily generated endogenously—from the inside of the media organization—but characterizes the approach of the media audience, making its attenuation a strategy that would conflict with profit maximization by the media firm. Fixing the media is not possible without fixing the public.

In this context, we see one possibility worth exploring: one could try to "change the game" of the mediated expertise. We suspect that the current biases projected by the public into the media filtering of the experts have to do with the "infotaining" aspects of the media experience. People watch the news not just to become informed but also to receive emotional satisfaction from confirming their preexisting beliefs, to soothe their worries, or to experience a pleasant thrill. Various techniques of improving the accuracy of forecasts while gamifying the forecasting experience are currently being employed with some success by prediction markets. In line with their approach, the media could generate the non-epistemic benefits from reporting expertise by transforming it into an accuracy competition. Using the modern methods of reputation-tracking, such as the Grier Score, in combination with clear-cut questions, unconditional probabilities, and objective criteria for deciding which forecasts were valid, the media could edge closer to the prediction markets.

Another road toward mitigating the welfare costs of expert competition for the spotlight would be to target the incentives that trigger contrarianism. Increasing the reputation costs for a dishonest specialist and introducing benefits from collaboration with a generalist may represent an attractive solution. As far as the reputation costs are concerned, the problem could be alleviated by an appropriate update to the ethical codes at the prominent scholarly institutions since the specialists' usual background is in academia. Accordingly, rules of proper conduct could require the specialist testimony to be provided while adhering to stricter standards of falsifiability.

Also, the possibility of an adversarial collaboration between the generalist and the specialist could limit the welfare costs of expert competition for the spotlight. After all, forecasting tends to be most accurate when it becomes a team effort. At least if the team possesses a certain degree of internal diversity (Tetlock and Gardner 2015; cf. Landemore 2012). While hiding in the shadow of generalists is of no interest for a specialist given our assumptions regarding their incentives, benefiting from their attention capital could be. Thus, if the specialist could be brought to the spotlight due to their collaboration with a generalist, the motivation towards dishonesty would be blunted. It may also be less costly for a generalist to locate a suitable specialist than it is for the media firm, providing it with an incentive to allow the collaboration.

4. Limitations and Extensions

While our approach can account for some of the observed phenomena, further extensions are possible.

First, generalists may have preferences for reporting a specific type of problem. For instance, reporting honestly may clash with their desire to promote a specific ideology or serve some special interest (Oreskes and Conway 2010). Our model can easily incorporate such preference by associating a specific strategy (i.e., reporting major or minor problem) with a payoff. A more general model would include various types of experts that vary in their rates of substitution between ideology and honesty.

Second, our model also assumes that media tend to exaggerate problems. However, we consider the determinants of neither the intensity nor the direction of this bias. In some cases, the media can be motivated to downplay the importance of an issue. Within our model, this issue can be dealt with by simply relabeling major and minor problems. Parameter B can be interpreted not only as a pessimism (or optimism) bias, but also as an ideological bias.

Third, in our model, if media is unhappy with the generalist's testimony, it can ask a specialist. Alternatively, a media firm can directly incentivize the generalists to provide their support for the media promoted view. If it prefers to report a major problem, it can "pay" the generalist to report accordingly. This may be cheaper for the media than searching for another expert. For instance, if media's preferred view is that the problem is major, then it can transfer a part of B to the generalist to provide them with a higher payoff from reporting accordingly.

Fourth, we assume that specialists report more accurately than generalists. However, this may not be universally plausible, such as with complex problems that transcend traditional disciplinary boundaries. This problem may be addressed by interpreting the specialist more broadly as someone who has issue-specific but not necessarily specialized human capital. For instance, health-economists combine expertise from various disciplines, but their human capital is issue-specific.

Fifth, we did not model an expert's choice whether to become a generalist or a specialist. We assumed that the pathway of an expert's career development always leads from specialism to generalism, conditional on the success of her media appearance. However, this pathway—while clearly important—is not the only feasible one. How early in their career may the experts choose the generalist path, depends on the specifics of the market for expertise, such as its degree of credentialism. If prior excellence in a particular field is not required for the sake of the initial media appearance, a pundit may channel her resources into increasing her attention capital instead of her accuracy in a specific field.

Sixth, our model assumes a monopoly media firm. A more realistic model would include competition among the media. Its effects would depend on whether media firms compete on the speed or quality of reporting. If the speed of reporting matters, the effect of competition would be captured through the search cost, *s*. The more competitive market, the costlier it is to search for an expert because media searching for a specialist risk that their competitors will report before them. Therefore, in the competition for the speed of reporting means that media are less likely to ask a specialist. Competition for the quality would be reflected through higher cost of misreporting, *c* and *C*, and hence, more accurate reporting. At the same time, competition may also trigger product differentiation where quality is less of a concern for some producers that others.

Finally, we have assumed rational Bayesian players. Our results may change if some form of bounded rationality is assumed.

5. Conclusion

Our model addresses the situations where the media prefer to exaggerate problems but require expert testimony to support their claims. However, the established experts (generalists) do not gain from reporting dishonestly. We show that even in such a setting, the generalists can be pushed towards distorting their testimonies due to the competition for the media spotlight between them and the less known, but more specialized experts eager to take their place. These specialists possess superior abilities when it comes to identifying if a given social issue is of minor or major significance, but their honesty can be undermined by their strategic behavior. To displace the current generalists, specialists may dissimulate their assessment to gain increased media attention.

Our chief finding is that neither the search costs of the media nor the reputation price to be paid for a mistaken forecast suffice to enforce honesty. The single most important determinant of the outcomes on the meta-expertise market we model is the probability that the truth will be revealed to the lay observers. When it comes to the prospects of the reform of the meta-expertise markets—which have not served many societies particularly well during the recent pandemic, leading to a lingering discord—we thus stress the importance of framing the expert testimonies in ways that maximize their falsifiability.

References

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *The Quarterly Journal of Economics* 84 (3): 488–500. https://doi.org/10.2307/1879431.

Algan, Yann, Daniel Cohen, Eva Davoine, Martial Foucault, and Stefanie Stantcheva. 2021. "Trust in Scientists in Times of Pandemic: Panel Evidence from 12 Countries." *Proceedings of the National Academy of Sciences* 118 (40): e2108576118. https://doi.org/10.1073/pnas.2108576118.

Banerjee, Abhijit V., and Esther Duflo. 2019. *Good Economics for Hard Times*. First edition. New York: PublicAffairs.

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. 1 edition. Oxford: Oxford University Press.

Caplan, Bryan. 2008. *The Myth of the Rational Voter: Why Democracies Choose Bad Policies*. New edition. Princeton, N.J.; Woodstock: Princeton University Press.

Caplin, Andrew, and John V. Leahy. 2019. "Wishful Thinking." Working Paper 25707. Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w25707.

Collins, Harry, and Robert Evans. 2017. *Why Democracies Need Science*. Cambridge, UK ; Malden, MA: Polity Press.

Gentzkow, Matthew, and Jesse M. Shapiro. 2006. “Media Bias and Reputation.” *Journal of Political Economy* 114 (2): 280–316. https://doi.org/10.1086/499414.

Goldman, Alvin I. 2001. “Experts: Which Ones Should You Trust?” *Philosophy and Phenomenological Research* 63 (1): 85–110. https://doi.org/10.2307/3071090.

———. 2018. “Expertise.” *Topoi* 37 (1): 3–10. https://doi.org/10.1007/s11245-016-9410-3.

Gurri, Martin. 2018. *The Revolt of the Public and the Crisis of Authority in the New Millennium*. San Francisco: Stripe Press.

Kahan, Dan M. 2015. “The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It.” In *Emerging Trends in the Social and Behavioral Sciences*, edited by Robert A Scott and Stephan M Kosslyn, 1–16. Hoboken, NJ, USA: John Wiley & Sons, Inc. https://doi.org/10.1002/9781118900772.etrds0417.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. London: Allen Lane.

Kelly, Jamie Terence. 2012. *Framing Democracy: A Behavioral Approach to Democratic Theory*. Princeton: Princeton University Press.

Koppl, Roger. 2010. “The Social Construction of Expertise.” *Society* 47 (3): 220–26. https://doi.org/10.1007/s12115-010-9313-7.

———. 2018. *Expert Failure*. 1 Edition. Cambridge Studies in Economics, Choice, and Society. New York: Cambridge University Press.

Kurvers, Ralf H.J.M., Uri Hertz, Jurgis Karpus, Marta P. Balode, Bertrand Jayles, Ken Binmore, and Bahador Bahrami. 2021. “Strategic Disinformation Outperforms Honesty in Competition for Social Influence.” *IScience* 24 (12): 103505.

https://doi.org/10.1016/j.isci.2021.103505.

Landemore, Hélène. 2012. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton University Press.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2021. “Rational Inattention: A Review.” *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3871534.

Mann, Charles C. 2018. *The Wizard and the Prophet: Two Remarkable Scientists and Their Dueling Visions to Shape Tomorrow’s World*. New York: Alfred A. Knopf.

McCloskey, Deirdre N. 2017. *Bourgeois Equality: How Ideas, Not Capital or Institutions, Enriched the World*. Paperback edition. Chicago London: The University of Chicago Press.

Nichols, Thomas M. 2017. *The Death of Expertise: The Campaign Against Established Knowledge and Why It Matters*. 1 edition. New York, NY: Oxford University Press.

Oreskes, Naomi, and Erik M. Conway. 2010. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. 1st edition. New York: Bloomsbury Press.

Russell, Stuart J. 2019. *Human Compatible: Artificial Intelligence and the Problem of Control*. New York: Viking.

Schick, Nina. 2020. *Deep Fakes and the Infocalypse*. London: Monoray.

Shapiro, Jesse M. 2016. “Special Interests and the Media: Theory and an Application to Climate Change.” *Journal of Public Economics* 144 (December): 91–108. https://doi.org/10.1016/j.jpubeco.2016.10.004.

Sowell, Thomas. 2011. *Intellectuals and Society*. Rev. and enl. Ed. New York, NY: Basic Books, A Member of the Perseus Books Group.

Špecián, Petr. 2022a. *Behavioral Political Economy and Democratic Theory: Fortifying Democracy for the Digital Age*. 1 Edition. Routledge Frontiers of Political Economy. New York, NY: Routledge.

———. 2022b. "Epistemology and the Pandemic: Lessons from an Epistemic Crisis." *Social Epistemology* 36 (2): 167–79. https://doi.org/10.1080/02691728.2021.2009931.

Sunstein, Cass R. 2017. *#Republic: Divided Democracy in the Age of Social Media*. Princeton ; Oxford: Princeton University Press.

Tetlock, Philip E. 2006. *Expert Political Judgment: How Good Is It? How Can We Know?* New Ed edition. Princeton, N.J.: Princeton University Press.

Tetlock, Philip E., and Dan Gardner. 2015. *Superforecasting: The Art and Science of Prediction*. F First Edition edition. New York: Crown.

Tufekci, Zeynep. 2017. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. New Haven ; London: Yale University Press.

Uscinski, Joseph E., Adam M. Enders, Casey Klofstad, Michelle Seelig, John Funchion, Caleb Everett, Stefan Wuchty, Kamal Premaratne, and Manohar Murthi. 2020. "Why Do People Believe COVID-19 Conspiracy Theories?" *Harvard Kennedy School Misinformation Review* 1 (3). https://doi.org/10.37016/mr-2020-015.

Wellcome Trust, and Gallup. 2020. "Wellcome Global Monitor: How Covid-19 Affected People's Lives and Their Views about Science." *Wellcome*.

https://wellcome.org/reports/wellcome-global-monitor-covid-19/2020.

Appendix 1. Perfect Bayesian Equilibria

A.1 Equilibria with a contrarian specialist

A.1.1 Media firm chooses *Always accept*

i) Assume that $t \leq v/(v + r)$, $t = 0$, and $pB \leq S$; *then there is a perfect Bayesian equilibrium (Wizard; Always accept; Contrarian). Denote this equilibrium I-equilibrium.*

ii) Assume that $t = 0$ *and* $B \leq S$; *then there is a perfect Bayesian equilibrium (Prophet; Always accept; Contrarian). Denote this equilibrium II-equilibrium.*

iii) Assume that $t \leq v/(v + r)$, $(1 - Q)tc \leq [1 - (1 - Q)t]B + S + QtC$, *and* $(1 - q)tC + (1 - qt)B \leq qtc + S$; *then there is a perfect Bayesian equilibrium (Honest; Always accept; Contrarian). Denote this equilibrium III-equilibrium.*

A.1.2 Media firm chooses Ask if minor

iv) Assume that $t \leq v/(v + r)$, $[1 - (1 - p)t]B + S + ptC \geq (1 - p)tc$, $(1 - t)L \geq (1 - 2Q)tR$, *and* $(1 - t)L \geq (2q - 1)tR$; *then there is a perfect Bayesian equilibrium (Prophet; Ask if minor; Contrarian). Denote this equilibrium IV-equilibrium.*

v) Assume that $t \leq v/(v + r)$, $(1 - qt)B + (1 - q)tC \geq qtc + S$, $(1 - Q)tc \leq [1 - (1 - Q)t]B + S + QtC$, $(1 - t)L \geq (1 - 2Q)tR$, *and* $(1 - t)L \leq (2q - 1)tR$; *then there is a perfect Bayesian equilibrium (Honest; Ask if minor; Contrarian). Denote this equilibrium V-equilibrium.*

vi) Assume that $t \leq v/(v + r)$, $t = 0$, *and* $B \geq S$; *then there is a perfect Bayesian equilibrium (Dishonest; Ask if minor; Contrarian). Denote this equilibrium VI-equilibrium.*

A.1.3 Media firm chooses *Ask if major*

vii) Assume that $t \le v/(v + r)$, $(1 - p)tc + S \ge [1 - (1 - p)t]B + ptC$, $(1 - t)L \ge (2Q - 1)tR$, $(1 - t)L \ge (1 - 2q)tR$,*; then there is a perfect Bayesian equilibrium (Wizard; Ask if major; Contrarian). Denote this equilibrium VII-equilibrium.*

viii) Assume that $t \le v/(v + r)$, $(1 - Q)tc \ge [1 - (1 - Q)t]B + S + QtC$, $(1 - q)tC + (1 - qt)B \le qtc + S$, $(1 - t)L \le (2Q - 1)tR$, *and* $(1 - t)L \ge (1 - 2q)tR$*; then there is a perfect Bayesian equilibrium (Honest; Ask if major; Contrarian). Denote this equilibrium VIII-equilibrium.*

A.1.4 Media firm chooses Always ask

ix) Assume that $t \le v/(v + r)$, $(1 - Q)tc \ge [1 - (1 - Q)t]B + S + QtC$, *and* $(1 - q)tC + (1 - qt)B \le qtc + S$*; then there is a perfect Bayesian equilibrium (Honest; Always ask; Contrarian). Denote this equilibrium IX-equilibrium.*

x) Assume that $t \le v/(v + r)$, $t = 0$, and $B \ge S$ *; then there is a perfect Bayesian equilibrium (Dishonest; Always ask; Contrarian). Denote this equilibrium X-equilibrium.*

A.2 Equilibria with a watchdog specialist

A.2.1 Media firm chooses *Always accept*

xi) Assume that $t \ge v/(v + r)$, $(1 - Q)tc \le (1 - Q)(1 - t)B + S$, *and* $(1 - q)tC \le S$*; then there is a perfect Bayesian equilibrium (Honest; Always accept; Watchdog). Denote this equilibrium XI-equilibrium.*

A.2.2 Media firm chooses *Ask if minor*

xii) Assume that $t \geq v/(v + r)$, $(1 - p)tc \leq (1 - p)(1 - t)B + S$, $p(tC + B) \geq S$, *and* $(1 - q)(1 - t)L \geq (2q - 1)tR$; *then there is a perfect Bayesian equilibrium (Prophet; Ask if minor; Watchdog). Denote this equilibrium XII-equilibrium.*

xiii) Assume that $t \geq v/(v + r)$, $(1 - Q)tc \leq (1 - Q)(1 - t)B + S$, $(1 - q)tC \geq S$, *and* $(1 - q)(1 - t)L \leq (2q - 1)tR$; *then there is a perfect Bayesian equilibrium (Honest; Ask if minor; Watchdog). Denote this equilibrium XIII-equilibrium.*

A.2.3 Media firm chooses *Ask if major*

xiv) Assume that $t \geq v/(v + r)$, $(1 - p)tc \geq (1 - p)(1 - t)B + S$, $p(tC + B) \leq S$, *and* $(1 - Q)(1 - t)L \geq (2Q - 1)tR$; *then there is a perfect Bayesian equilibrium (Wizard; Ask if major; Watchdog). Denote this equilibrium XIV-equilibrium.*

xv) Assume that $t \geq v/(v + r)$, $(1 - Q)tc \geq (1 - Q)(1 - t)B + S$, $(1 - q)tC \leq S$, *and* $(1 - Q)(1 - t)L \leq (2Q - 1)tR$; *then there is a perfect Bayesian equilibrium (Honest; Ask if major; Watchdog). Denote this equilibrium XV-equilibrium.*

A.2.4 Media firm chooses *Always ask*

xvi) Assume that $t \geq v/(v + r)$, $(1 - Q)tc \geq (1 - Q)(1 - t)B + S$, $(1 - q)tC \geq S$, *and* $(1 - Q)(1 - t)L \leq (2Q - 1)tR$; *then there is a perfect Bayesian equilibrium (Honest; Always ask; Watchdog). Denote this equilibrium XVI-equilibrium.*