



Social Epistemology

A Journal of Knowledge, Culture and Policy

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/tsep20

No-Regret Learning Supports Voters' Competence

Petr Spelda, Vit Stritecky & John Symons

To cite this article: Petr Spelda, Vit Stritecky & John Symons (11 Sep 2023): No-Regret Learning Supports Voters' Competence, Social Epistemology, DOI: [10.1080/02691728.2023.2252763](https://doi.org/10.1080/02691728.2023.2252763)

To link to this article: <https://doi.org/10.1080/02691728.2023.2252763>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 11 Sep 2023.



Submit your article to this journal [↗](#)



Article views: 800



View related articles [↗](#)



View Crossmark data [↗](#)

No-Regret Learning Supports Voters' Competence

Petr Spelda ^a, Vit Stritecky ^a and John Symons ^b

^aDepartment of Security Studies, Institute of Political Studies, Faculty of Social Sciences, Charles University, Prague, Czech Republic; ^bDepartment of Philosophy, The University of Kansas, Lawrence, KS, USA

ABSTRACT

Procedural justifications of democracy emphasize inclusiveness and respect and by doing so come into conflict with instrumental justifications that depend on voters' competence. This conflict raises questions about jury theorems and makes their standing in democratic theory contested. We show that a type of no-regret learning called meta-induction can help to satisfy the competence assumption without excluding voters or diverse opinion leaders on an a priori basis. Meta-induction assigns weights to opinion leaders based on their past predictive performance to determine the level of their inclusion in recommendations for voters. The weighting minimizes the difference between the performance of meta-induction and the best opinion leader in hindsight. The difference represents the regret of meta-induction whose minimization ensures that the recommendations are optimal in supporting voters' competence. Meta-induction has optimal truth-tracking properties that support voters' competence even if it is targeted by mis/disinformation and should be considered a tool for supporting democracy in hyper-plurality.

ARTICLE HISTORY

Received 18 December 2022

Accepted 18 August 2023


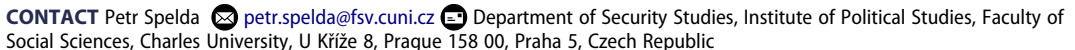
KEYWORDS

Jury theorems; meta-induction; no-regret learning; epistemic democracy

1. Introduction

The epistemic turn in the study of democracy assumes that the acceptability of a political order should be judged, at least in part, on its record of and prospects for epistemic success. The assumption has been that any justification of democracy would necessarily involve epistemic elements (see, Anderson 2006; Landmore 2012; Martí 2006). Epistemic democrats like Héléne Landmore (2017) defend the truth-tracking properties of democratic procedures. While opponents argue that epistocracy (Estlund 2003) is preferable to democracy because the vast majority of individuals in democratic societies are unable to form correct judgments concerning socio-economic matters (Brennan 2016).

In this paper, we provide an answer to why the wisdom of the crowds, the supposed ability of large groups to vote wisely, is at odds with election results. This general principle figures in many epistemic justifications of democracy (Schwartzberg 2015), but it seems belied by evidence from actual democratic societies. The formalization of the wisdom of the many, the Condorcet Jury Theorem (Condorcet 1976/1785), provides a mathematical guarantee that the many will be correct if among them there is a majority of correct votes (Goodin and Spiekermann 2018). Despite that formal guarantee, populists, revisionists and other malicious or incompetent political actors are frequently voted into office, negatively influencing policy decisions on almost every issue. Clearly, some assumption(s) required for the mathematical (formal) guarantee to hold in the real world has

CONTACT Petr Spelda  petr.spelda@fsv.cuni.cz 

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

failed. This should be of concern to epistemic democrats given that their formal models apparently fail to correspond to reality.

One might observe, for example, voter competence decreasing over time under pressure from online mis/disinformation and mistrust of science. Since a level of voter competence is required for the formal guarantee of the Condorcet Jury Theorem to work, decreased levels of competence make it more likely that the majority among the many does not consist of correct votes. In this paper, we focus on the divergence between the formally required level of voter competence – > 0.5 , i.e. voters are more likely than not to vote correctly – and its real-world value on different issues. We design a new approach to the enhancement of voter competence that is intended to help re-establish jury theorems as a means of justifying epistemic democracy. This is of interest to epistemic democrats because the enhancement can bridge traditional formal arguments and real-world voter behaviour.

We define voter competence as the ability to track the political actor most capable at predicting successfully on a given issue. When the issue becomes subject of a majority vote, our voter enhancement determines who is, according to the past predictive successes and failures, the most successful actor and suggests the actor to voters to follow. In Section Two, we explain why this does not break the voter independence assumption. Voters following the advice maximize their competence with respect to the evidence on the actors' predictive track record. The core of the competence enhancement is no-regret learning from the actors' predictive track record called multiple-favourite meta-induction developed by Schurz (2008, 2019). The learning method cannot be manipulated by populists, revisionists¹ and other actors who are harmful overall despite their potential limited predictive successes. The no-regret voter enhancement does not require that these harmful actors are a priori removed from consideration. Learning from their predictive track record has the same effect as exposing their falsehoods² with one all important difference. Instead of trying to remove the harmful actors from the political competition, they remain involved. Therefore, there is an appropriate level of respect and inclusion as required by epistemic democracy. However, the method producing what we call *no-regret voter enhancements* learns to decrease the importance of harmful actors as their predictive track record becomes increasingly unsuccessful over time. We describe this process in detail in Section Three.

The takeaway of the paper is twofold. First, we identify the failed assumption that separates the formal justification of epistemic democracy (Condorcet Jury Theorem) from real-world voter behaviour. Second, we define a voter enhancement that can help to satisfy the voter competence assumption and by that bring the formal justification of epistemic democracy and the real-world voter behaviour closer together. What we mean by enhancement here is fully theoretically defined, including provable guarantees on its effects along with implementation details for empirical studies.

Different models of preference aggregation exist, and epistemic democrats have closely attended to the study of their properties (Goodin and Spiekermann 2018; Schwartzberg 2015). The foundational framework, the Condorcet Jury Theorem (Condorcet 1976[1976/1785]), concerns a group of n voters deciding between two options by means of majority rule (Goodin and Spiekermann 2018, 17). If the voters satisfy the competence, statistical independence, and sincerity assumptions (more on their modern relaxation by Goodin and Spiekermann in Section Two), then the theorem shows when the majority of the group of n voters decides correctly (Goodin and Spiekermann 2018, 19–20). Crucially, to satisfy the competence assumption, each voter's probability of being correct must be $p_c > 0.5$. In its simplest form, the theorem describes a binary prediction game, where 'for every n and a constant p_c such that $0.5 < p_c < 1 : P_{n+2} > P_n$ ' (Goodin and Spiekermann 2018, 19–20). That is, the majority of the larger group, P_{n+2} , is more likely to be correct than the majority of the smaller group, P_n , with the asymptotic result given as $\lim_{n \rightarrow \infty} P_n = 1$, i.e. as the number of voters approaches infinity, the probability of the majority being correct converges to 1 (ibid. $n + 2$ in the non-asymptotic result to avoid ties, as n is usually assumed to be odd, Goodin and Spiekermann 2018, 17; 19).¹⁰

Thus, the voter competence assumption $p_c > 0.5$ is a prediction about the robustness of voters' predictive models, which allows for a particular aggregative mechanism with provable guarantees (see Section Two for a revision to the Condorcet Jury Theorem's assumption of independent voters, more realistically reflecting the epistemic situation of voters compared to Condorcet's original definition). The prediction about voter competence is, however, fallible if subjected to the pressures of deceptive actors who, by spreading mis/disinformation, aim to decrease voter competence and undermine the epistemic virtues of the democratic system (Symons and Elmer 2022).

We will focus on the situation when one of the assumptions, voter competence, necessary for the formal (mathematical) guarantee of Condorcet's framework and similar ones (Dietrich and Spiekermann 2022) becomes an unreliable prediction. This allows us to show that epistemic democracy justified in both the classical way and in the modern sense, where the original assumptions are relaxed (Goodin and Spiekermann 2018), is vulnerable to a version of Hume's problem of induction (1739[1978]). Our voter enhancement offers a way to reduce the unreliability of the voter competence prediction by using an optimal learning method that Gerhard Schurz (2019) used to obtain a solution to Hume's problem. The vulnerability is a pressing issue for epistemic democrats. Its exploitation by mis/disinformation produced by malicious political actors decreases voter competence and makes the mathematical justification of epistemic democracy unusable in the real world. Rising numbers of populists and revisionists in the ranks of political actors, exploiting online platforms to mislead voters (Bennett and Livingston 2020), decreases the likelihood of elections leading to beneficial outcomes and motivates the development of new voter enhancement techniques. We offer one such new enhancement that can help in the tug-of-war over voter competence playing out between epistemically successful political actors and populists, revisionists and alike.

1.1 Epistemic Democracy and Prediction Games

We now move to present some key concepts and definitions in order to explain the connection between epistemic democracy, no-regret learning and meta-induction. It is important to note from the outset that not all policy questions can be treated as (binary) prediction games. Furthermore, as cautioned by Schwartzberg (2015, 196), democratic decision-making cannot be modelled entirely in terms of prediction games, since the measurability of a policy's impact, e.g. the influence of a piece of welfare legislation on citizens' poverty, is not straightforward. On the other hand, by relaxing the assumptions of the Condorcet Jury Theorem and extending it to many-options cases, Goodin and Spiekermann (2018, 8) argue for treating voting as the epistemically foundational practice that researchers should focus on in investigations of the democratic form of government as a whole. We introduce some simplifying assumptions that limit our attention to problems which (i) can be cast as repeated binary prediction problems, admitting real-valued predictions whose successes are measured by some loss function (Schurz 2019, Section 5.5), and which (ii) involve extremely polarising issues, well-suited for revisionist attacks against democratic institutions. In such cases, the Condorcet Jury Theorem and similar frameworks (Dietrich and Spiekermann 2022) become vulnerable to deceptive actors (e.g. political revisionists) who attempt to lower voter competence below the critical threshold, $p_c < 0.5$. In this situation, as we shall see, adding more predictors (voters) to the pool decreases the probability of the majority reaching correct predictions (decisions).

An alternative is to introduce a regret-based weighting of opinion leaders, that is, of political actors influencing voters, using the actors' past predictive performance (see also an extensive form of prediction games described below, using fact-checking as a backend, applicable to a wider variety of policy issues), which allows meta-induction to become a type of no-regret learning supporting voters' competence.

Before moving to show an example prediction game, which has been heavily targeted by adversaries and showcases a possible transition from *plurality* to what we are call *hyper-plurality*, a summary explanation of some additional key concepts is provided.

Table 1. A typology of political actors.

Type	Actor Characterization
1a	Best Responder Actor – an actor capable of correctly interpreting scientific findings on a given issue and translating them into beneficial policy interventions, i.e. achieves optimal utilization of the available evidence.
1b	Arbitrarily Unreliable Actor – an actor failing to grasp complex and quickly shifting empirical realities, might be prone to failures even in situations where full, detailed actionable evidence is available.
2a	(Systematic) Deceiver-Revisionist Actor – an actor oscillating between two epistemic postures – (i) before gaining voters' trust, the actor signals a Type 1a or 1b affiliation, (ii) after gaining voters' trust, the actor's policy negates the available evidence, e.g. refusing to act on vital issues.
2b	Mis/Disinforming Revisionist Actor – an actor seeking to gain voters' trust in a success-independent manner, i.e. without producing successful or partially successful policies and covering their poor epistemic performance with further mis/disinformation.

First, apart from its formal definition in Proposition 1 below, hyper-plurality can be characterized as a change in the make-up of the pool of political actors, impacting voters' competence and threatening to turn majority vote against the best interests of democratic societies. Table 1 provides a simple typology of political actors for this purpose. A growing presence of Type 2a and 2b actors, undermining the consensus on a given issue based on scientific findings, together with an increasing failure rate of Type 1b actors and diminished influence of Type 1a actors signals a transition from plurality to hyper-plurality.

The three latter categories of actors have a negative impact on voter competence. Type 2a actors include also political populists who consider expert policies antagonistic to people's interests and using sovereignty of the people argue that experts should not have a decisive say in policy making (Stanley 2008). Populists can cause systematic deception as well. By acting according to the supposed interests of the people and, thus, considering crowds to be wise, populists make a strong assumption about voters' competence on all issues. Although this can work in trivial situations, on issues where voters lack competence (are not truth-biased), populists fail by acting according to instrumentally defined 'people's interests' that oppose experts. Even if populists can score some initial success, they are bound to fail as systematic deceivers because of disregarding expert advice in order to appeal as anti-elitist before their electorate.³ Reducing competence below the critical threshold transforms majority vote into a tool for adversaries. An intuitive countermove amounts to excluding as many of Type 1b, 2a, and 2b actors as possible from the pool of political actors to ensure that voter behavior can be influenced only by Type 1a actors, that is, by the Best Responders. This is an epistocratic solution that could be facilitated, for example, by persistent and wide-scale debunking efforts or by censorship of epistemically vicious actors. If successful, this would prevent the transition from plurality to hyper-plurality. Whether this can be achieved today is doubtful.

An alternative approach can use some actions of actors (see below for an explanation of how a certain type of fact-checking can serve as a backend for repeated prediction games) to recover repeated prediction games. Instead of excluding them, the point is to treat all actors as predictors and maintain respect and inclusivity based on their epistemic performance. Then, an evaluation method is epistemically democratic and blind to whether it operates in plurality or hyper-plurality if at any point in the game it does not depend on any judgment (prediction) with respect to which category from Table 1 the actors occupy and to the hardness of the game.

These conditions can be satisfied by no-regret learning (Cesa-Bianchi and Lugosi 2006) and by the meta-inductive methods in particular (Schurz 2008, 2019). The goal of any meta-inductive method is to minimize its regret, that is, the difference between its predictive performance and that of the best actor in hindsight (Schurz 2008, 2019). The requirement of no predictions about the actors' categories and the game's hardness is satisfied if a regret bound is proved such that it holds for an arbitrarily hard game and an arbitrarily unreliable set of actors whose predictions are available to the meta-inductive method. Crucially, any meta-inductive evaluator whose regret can grow linearly with the number of rounds in the prediction game (for linear losses) is unhelpful, because it can be

deceived and its performance possibly reduced to zero (cf. Schurz 2019, Sect. 6.3). For example, if it is impossible to ensure that Type 2a actors are not present in the pool, then the intuitive meta-inductive method, selecting in each round the so-far best-performing actor (to be recommended to voters to follow), can lower voter competence below the critical threshold (see Table 2). In hyper-plurality, Type 2a actors can cause any such one-favorite method to switch its selected actor in each round of the game, thereby reducing the performance of the method to (or near to) zero (see Table 1, Type 2a actors decrease their performance after being selected and since the selection comes always first, because meta-induction is a dependent learning method, one-favorite meta-induction has no remedy).

This weakness is overcome by weighted average forecasters (Schurz 2019, Sect. 6.6). These methods are safeguarded against deception by having ‘multiple-favorites’. That is, instead of assigning a weight equal to one to the selected actor and zero weights to the rest as in the one-favorite case, weighted average forecasters assign to each actor a weight based on the actor’s performance until the current round while predicting a weighted average of the actors’ predictions for the next. The best regret bound proved so far relies on weights defined as exponential functions of the actors’ predictive performance until the current round of the game (Schurz 2019, Sect. 6.6.2). As a result, multiple-favorite meta-induction approximates the maximal available predictive performance of any pool of actors without requiring information about actors’ categories and the hardness of the game. This makes the multiple-favorite meta-induction with exponential weights a suitable epistemically democratic evaluation method supporting voters’ competence. Section Three discusses no-regret voter enhancements in detail.

Therefore, in each round of the game, predictions of political actors need to be scored and the actors weighted accordingly (this can also be done for epistemic practices that they are using and that can be tracked). A round of the game can be alternatively formed by collecting actors’ statements and actions for a period of time and scoring their epistemic stances at the end of the period. This would give fact-checkers the ability to score political actors at their own accord by specifying what entails a round in the game. In Section 1.2, we discuss how this can be done practically in the context of fact-checking.⁴ The results should be then distributed as widely as possible in a bid to balance the attempts to decrease voter competence below the critical threshold. The argument for using meta-induction to maintain voter competence $p_c > 0.5$ necessary for majority rule to reach correct outcomes stems also from the fact that some of the existing methods, such as debunking of misinformation or disinformation, might fail to produce the intended results. Mosleh et al. (2021) found that when publicly corrected by a human fact-checker, users on online social media platforms⁵ are *not* less likely to share false political news in the future. In fact, debunking can provoke a deterioration in quality of the subsequently shared content and instead of fostering accuracy, such interventions negatively impact the users’ downstream sharing behaviours (Mosleh et al. 2021). It is possible that debunking might produce embarrassment or reduced social standing, and as a result it can be perceived as stigmatising by those who become its subjects (Mosleh et al. 2021). Even if fact-checkers provide corrections, the effect such interventions have on voters might become a further alienation from the status quo. The other important issue here is the threat of attitude polarization as discussed in the social psychology literature (Mackie and Cooper 1984). The examination of optimal debunking methods is ongoing and features a wide range of open questions and problems (Lewandowsky et al. 2020).

Consider the critical, extremely polarising political issue of anthropogenic climate change. Misinformation and disinformation about climate change have the effect of lowering voter competence below the threshold so that majority vote is not directly responsive to scientific evidence or understanding. For our purposes, climate change represents an instructive issue for three reasons.

First, it shows that frameworks of epistemic democracy such as the Condorcet Jury Theorem are not just about aggregating preferences or values (Lane 2016, 116–17). Facts established by the scientific method in disciplines that investigate climate change and misinformation or disinformation, contesting those facts by disputing the scientific method (or scientists), are a part of the picture

as well (Biddle and Leuschner 2015; Lewandowsky, Cook, and Lloyd 2018). Robustness of the voter competence inference $p_c > 0.5$ depends on voters' decisions who to follow, or rather imitate, while making their predictions expressed as votes, and this is no longer just scientists or domain experts in general.

Second, climate science showcases how uneasy, and sometimes even adversarial, the relation between scientists and other voters (citizens) can become (ibid., also cf. Landemore 2020, 191–92), causing some voters to follow advice of political actors who reject or undermine the consensus view established by the methods of the scientific community. The tension reflects difficulties of balancing the competing demands for equality among voters and prioritization of superior expert knowledge (Chambers 2017; Fuerstein 2008; Holst and Molander 2017, 2018, 2019; Schwartzberg 2015, 194–95). The expanding role of experts in policymaking seems to have encouraged the growth of political alternatives that diverge from findings provided by the scientific method. This is partly due to changes to public deliberation caused by social media, which offer a platform to such deceptive alternatives and allow an unprecedented targeting of voter competence (Bennett and Livingston 2020; O'Connor and Weatherall 2019; Persily and Tucker 2020; Zimdars and McLeod 2020). In sum, voters are able to freely select from a rich offering of alternatives, often antithetical to scientific findings, and can be unapologetic about following the alternative of their choosing. That is, letting it inform their voting, which reduces their competence and as per jury theorems also the probability that majority rule chooses the correct option. This is how we define 'interpretative individualism'. Epistemic democrats are then confronted with attempted misuses of majority rule that aim to hinder democracy.

Among the main responses to climate change mis/disinformation is debunking whose use and effects establish an active area of study (Lewandowsky et al. 2020). Other notable initiatives include educational efforts using massive open online courses (Cook 2016), gamification (Cook 2019), creation of platforms for citizen scientists to rebut climate change mis/disinformation (Winkler and Cook 2020) or popularization of basic arguments against climate change denialism (Cook 2020). All these efforts can be considered tools for increasing voters' competence. No-regret learning in prediction games among political actors is an addition to the toolbox, adhering strictly to the principles of epistemic democracy.

Third, changing values of climate change indicators can be used as an example for how to frame policy issues at least partially in terms of prediction games. Consider the 'Vital Signs of the Planet', which include periodic measurements of carbon dioxide, global temperature, arctic sea ice minimum, ice sheets, sea level and ocean heat content (NASA 2021). If, for example, a revisionist political actor claims that the sea level rise is slower or not occurring at all, then, apart from disinforming or misinforming the voters (Douven and Hegselmann 2021), they also predict that the value of the sea-level indicator evolves contrary to the predictions derived from the scientific method. Since the measurements are taken repeatedly, the revisionist actor can be understood to enter a repeated prediction game alongside other actors, such as climate scientists, democratic actors or revisionists with different goals, who all influence the voter competence p_c by persuading the voters to imitate their predictions. For the n -th round of the game, each (active) actor A issues a $pred_n(A)$, and after the measurement is taken, a value from the loss function $loss(pred_n(A), e_n)$ is returned, determining A 's success in the n -th round. In this toy example, we consider the natural loss function, $e_n \in \{0, 1\}$, where $0 = \textit{sea level stable or dropping}$ and $1 = \textit{sea level rising}$, a measurement threshold distinguishing the two events, and real-valued predictions $pred_n(A)$.⁶ There are two motivations for considering epistemic democracy, at least partially, in terms of prediction games. First, leaderboards of epistemic performance can be maintained, holding cumulative losses of individual actors until the game's present round n . Expanding on our toy example, each 'vital sign' would have its own leaderboard, tracking the epistemic performance of actors issuing predictions on the development of the sign's value. This is how we define 'issue-centric leaderboards', each holding a measurable quantity which can be used to create a binary prediction game. Second, if we focus solely on the issues which can be recast as prediction games, then we can reduce truth-tracking, the key concern

of epistemic democracy, to epistemic performance. This dissolves legitimate arguments that the usage of truth-tracking in the era of post-truth (Iyengar and Massey 2019; McDermott 2019; McIntyre 2020) creates more problems than it solves.

1.2 Prediction Games and Fact-Checking

Calculating the value of the loss incurred by each actor participating in a prediction game (recorded in an issue-centric leaderboard) until the present round n can equal to fact-checking when two conditions are jointly met. First, statements considered as submissions to each round of the game must be labelled by fact-checkers as TRUE/FALSE, i.e. $\{1, 0\}$, or with intermediate labels (e.g. mostly/half TRUE/FALSE equal to values from $[0, 1]$) in case the game admits real-valued predictions. This way even statements that do not carry explicit predictive content can be treated as submissions to the game and their losses calculated by comparing the assigned label with the ground-truth. That is, with a factual statement, here becoming $e_n = 1$ (always equal to 1 due to its factuality). Although this represents an extensive form of the prediction game, the form is justified by the fact that such statements depend on inductive inferences (predictions or generalizations) about the world. Even if they remain implicit, actors using failing or failed inferences can be scored appropriately by fact-checkers, as compared to epistemically successful actors, producing success records recorded in leaderboards.

Second, for each round of the game actors' statements that are scored should be comparable so that comparable implicit inductive inferences of about the world used by a stable actor pool (see Section Five on how to deal with evolving actor pools) are scored and the game remains consistent. Using Poynter's Environment Scoreboard⁷ as an example, it is clear that the second condition is not met. Therefore, this fact-checking format cannot support the meta-inductive evaluation of political actors out of the box because it is impossible to recover a consistent game, providing the desired leaderboard for each round of the game. Leaderboards for meta-inductive evaluations are distinct from many fact-checking formats. However, if the latter jointly satisfies the two conditions, then it can serve as a backend for leaderboards, thus supporting meta-inductive evaluations. This would also open up a possibility of using meta-inductive evaluations in policy areas where conventional prediction games are inapplicable. The number of rounds in a game per unit of time is limited by the fact-checkers' ability to extract actors' statements and by the availability of facts in case the issue at stake evolves rapidly. The fact-checkers' bandwidth can be increased by machine means (e.g. Bekoulis, Papagiannopoulou, and Deligiannis 2023). For the benefit of Type 1a actors, it is also vital to disseminate facts quickly, using automation as well (Moravec et al. 2020).

Compared to traditional prediction games, extensive form games, relying on suitable fact-checking backends, have higher potential to become politicized. Here, the connection of ground-truths, i.e. factual statements, to observable phenomena can be indirect and multifaceted, raising the likelihood that counter-fact-checking backends emerge and with them also counter-leaderboards. This does not mean that the meta-inductive evaluation is a doomed endeavour from the very beginning. Building on Goodin and Spiekermann's (2018) version of the Condorcet Jury Theorem, concerned with the average voter and their ability to track and imitate the best responder (predictor) political actor, it is unnecessary that the competence of each voter is increased in order to maintain the average voter's competence above random (Goodin and Spiekermann 2018, 92–94). The same obviously applies to political revisionists and their counter efforts, seeking to lower the average voter's competence below the 0.5 threshold. However, since the transition from plurality to hyper-plurality has only just begun, defenders are in a better position than attackers, trying to hurt democratic societies via turning majority vote against them by lowering the average voter tracking competence below random.

Although defenders seem to have an initial advantage, the situation might change in the future. In case the average voter competence approaches the 0.5 threshold due to mis/disinformation,

voters could have difficulties in identifying which actor scoring, and thus weighting of the actors, to believe. Apart from the legitimate one based on established facts, scorings could be based on information derived using pseudoscientific epistemic practices. Voters would be then facing a dilemma, which scoring to believe and, as a result, which weights allowing meta-induction to let inform their voting. Here, we see that our voter enhancement based on no-regret learning in prediction games is only one tool from a large toolbox. Among others we count education. Some enhancements require time to produce the desired results (e.g. education); others, such as no-regret voter enhancements discussed below, can take effect immediately. The richness of the voter enhancement toolbox and the assumption that individual tools can be combined effectively gives hope that voter competence, and with it the epistemic justification of democracy, can be secured in the future despite the revisionist, populist and other types of epistemic attacks on democracy.

Therefore, it is opportune to use every option that can safeguard majority vote from revisionist political actors. Compared to debunking, rooted in an epistocratic view on democracy, ours is a proposal inspired by epistemic democracy. This does not necessarily mean that their relation is mutually exclusive. Rather, the question is how best to adapt democratic societies for the arrival of hyper-plurality. Debunking attempts to postpone or avoid hyper-plurality altogether, whereas the meta-inductive evaluation embraces it. The interplay between debunking and meta-induction in democratic societies is a subject for further research.

2. Epistemic Democracy in Hyper-Plurality

We are now ready to connect the Condorcet Jury Theorem (and jury theorems in general) with our definition of individualistic interpretations to draw a distinction between political plurality and hyper-plurality.

Proposition 1: Until the rise of adversaries who exploit social networking platforms, the inductive inference on the mean voter competence was robust. That is, the average voter will more likely than not imitate the best predictor,⁸ i.e. $\bar{p}_{BP} > 0.5$, on all possible prediction problems (cf. Goodin and Spiekermann 2018, 76–79). This made majority rule in plurality dependable.⁹ After the rise of individualistic interpretations, such an inference ceased to be robust, because the revisionists seek to create hyper-plurality, i.e. $\bar{p}_{BP} < 0.5$, where majority rule decreases the probability of predicting correctly and hinders democracy.¹⁰ The situation is further worsened by unreliable legitimate political actors, who fail to grapple with increasingly complex and dynamically evolving (social) realities. If mishandled, voter enhancements such as debunking might generally deepen the problem even further.

Before Proposition 1 can be unpacked, the notion of individualistic interpretations needs to be explained more fully. Individualistic interpretations are encouraged by mis/disinformation and held by voters rejecting the epistemic successes of scientific inquiry. Instead, the voter is the supreme epistemic authority over what and whom to believe, free also to reject scientific findings in the characteristic populist or revisionist manner. To become this kind of epistemic authority is to ‘secede’ from majorities that can be found among voters thanks to the scientific consensus on facts. Epistemic secessions by voters from majorities secured by scientific findings make individualistic interpretations vehicles of hyper-plurality where majority rule fails to select the correct option because epistemic successes of scientific inquiry are simply ignored or denied.

Let us proceed by recapping the epistemic situation captured by Proposition 1. Political actors refer to predictors in the game sense described above. First, we include what Goodin and Spiekermann (2018, 76–79) call the best responder actors (Type 1a, Table 1) whose inferences reflect the best of what can be achieved, considering the available evidence (such actors correctly interpret scientific findings, which, for example, in the case of climate change brings about a formidable challenge, cf. Winsberg 2018). If the evidence about the state of the world is unreliable or misleading, then the best responders fail (Winsberg 2018). Second, there are unreliable legitimate political actors (Type 1b, Table 1), failing to adapt to rapidly changing (social) realities whose inferences might not be robust even in epistemically favourable conditions. Third, we include revisionist adversaries (Type

2a, Table 1), deceivers in Schurz's original sense motivating the multiple-favourite meta-induction (Schurz 2008, 285–295), who begin to purposefully fail once voters decide to imitate them. The deception could be subtle, amounting to not acting on critical issues without any prior indication, e.g. knowing that the sea level indicator is rising, yet once assuming power acting as if the predictions to the contrary were successful in the past. Finally, revisionists using misinformation and/or disinformation to deceive voters into following them, while denying or covering their predictive failures by the identically deceptive means (Type 2b, Table 1). The last form of deception is perhaps most disturbing, suggesting that a non-negligible number of voters is persuaded by and learns using criteria other than predictive success.

As a result, hyper-plurality contains several types of revisionist actors who (i) base their inferences on information in conflict with findings provided by the scientific method and (ii) are highly successful at attracting voters due to the radically open nature of public deliberation on online social networking platforms. Using the 'Best Responder Corollary' by Goodin and Spiekermann (2018, Sec. 5.3), we can establish that as the transition from plurality to hyper-plurality proceeds, a large volume of voters will no longer be described as best-responder-trackers. This means that instead of being at least slightly better than random at imitating the best responder (predictor), i.e. $\bar{p}_{BP} > 0.5$, the average voter will not be the best responder themselves but, more profoundly, will no longer be the best-responder-tracker, i.e. $\bar{p}_{BP} < 0.5$. As per the Condorcet Jury Theorem, the probability that majority rule leads to correct decisions then decreases and with it also the propensity of democracy to deliver beneficial solutions.

For this mechanism to approximate real-world voting, Goodin and Spiekermann (2018, Chap. 5) proposed an important revision to the independence assumption of the original Condorcet Jury Theorem. Independence is considered conditional on all common causes, including the state of the world, the evidence reflecting it and opinion leaders, heuristics, etc., processing the evidence. Independence conditional on all common causes still satisfies voter independence because voters are not encouraged to copy actions of their peers but instead are encouraged to become best responder (predictor) trackers, voting according to the choices of the best responders (Goodin and Spiekermann 2018, 76–78). The voter independence conditional on all common causes is more realistic than the original Condorcet Jury Theorem independence because it allows for upstream influences on voters and uses the influence to redefine competence as the voters' ability to track actors who are best at utilizing all the common cause (the upstream influence).

A simple solution to the voters' inability to robustly track best responders would be to try to remove failing legitimate and revisionist actors from the pool of actors so that they cannot be imitated by voters. Only the best responders (predictors), basing their inferences on correctly interpreted findings produced by the scientific method, would remain in the pool. Considering the meta-inductive research programme, such an intervention corresponds to allowing only the class of prediction games G for which the Imitate-the-Best (ITB) meta-induction remains optimal in the sense of regret minimization (Schurz 2019, Theorem 6.2). G is such that the prediction games include a unique best responder who begins to predict successfully sufficiently early and becomes the unchanging ITB's favourite (Schurz 2019, Theorem 6.1). G cannot include untrustworthy actors, i.e. Type 1b, 2a, and 2b, who increase or maximize the ITB's regret by failing (possibly on purpose) each time ITB selects them for their performance in the past rounds of the game. In G -class prediction games, voters following the ITB's favourite increase their competence by minimizing their regret with respect to alternatives, which in turn increases the probability that $\bar{p}_{BP} > 0.5$.

The removal of Type 1b, 2a, and 2b actors from the pool, ensuring that ITB is a regret minimization strategy, might be practically impossible and possibly undesirable. For example, in the context of social networking platforms, it would have to depend on large-scale debunking (or even prebunking) efforts combined with educational initiatives. There is ongoing research into optimal methods that can be used to present facts (Lewandowsky et al. 2020) to avoid recently observed negative effects of debunking (Mosleh et al. 2021). Increasing the probability that $\bar{p}_{BP} > 0.5$ by allowing only G -

class prediction games where ITB remains optimal violates epistemic democracy in the sense of a priori respect and a priori inclusion towards all actors regardless of their nature.

Here, we arrive at the central difference between one-favourite meta-inductive strategies, such as ITB and its variants, and their multiple-favourite alternatives, such as the weighted-average meta-inductivists (Schurz 2019, 138–40), which are discussed in the following sections with focus on the exponential strategy (Schurz 2019, 144–45). Under the weighted-average meta-induction, instead of removing Type 1b, 2a, and 2b actors from the pool, every actor is continually weighted according to their predictive performance. As per Schurz (2019, 6.8[i]), each actor's weight can be defined as the exponential function of their absolute success until now and used to determine the actor's impact on the meta-inductivist's prediction, which for each round corresponds to a weighted average of the predictions of all actors in the pool. A priori respect and inclusion towards all actors cannot be maintained if epistemic democracy becomes limited to *G*-class prediction games. This tension can be resolved as follows:

Proposition 2: Epistemic democracy in hyper-plurality described in Proposition 1 requires multiple-favourite meta-inductive strategies, which are universally optimal and, thus, not limited to a particular class of prediction games (Schurz 2019, Definition 5.4 [2]).

Table 2 compares one-favourite (OFMI) and multiple-favourite meta-inductive (MFMI) strategies in the political environment defined in Proposition 1. The difference lies in the OFMI's optimality limited to *G*-class prediction games compared to the MFMI's universal optimality.

Even if unsuccessful actors receive exponentially smaller weights than those close to the MFMI (optimal responder, considering the exponential version of multiple-favourite meta-induction [eMI], Schurz 2019, Sec. 6.6.2), and their impact thus converges to zero, epistemic democracy in the sense of no presuppositions about the actors holds. Should their predictive performance improve later, the impact of the unsuccessful actors on the weighted average will change as well. As a result, the

Table 2. OFMI v. MFMI (cf. also Schurz 2019, 229, Table 8.1).

Stage	OFMI	MFMI
Starting Point	Attempting to preserve plurality in situations in which individualistic interpretations are no longer negligible	Embracing hyper-plurality, i.e. neither expecting that plurality returns, nor expecting hyper-plurality to possess self-organising features
Goal	Consensus on policy issues by majority rule (either via direct or representative means)	Consensus on policy issues by majority rule (either via direct or representative means)
Intervention	A priori (prebunking) and/or a posteriori (debunking) efforts to exclude all but best responder (predictor) actors from the pool of candidates from which voters can choose who to follow based on the actors' successes in the past rounds of the game held in the respective leaderboard	Optimal utilisation of all of actors' predictions via weighted average (Equation (1)), where the weight $w_{i,n-1}$ of the prediction submitted by the actor A_i for the next round n is derived from an exponential function combining some learning rate $\eta > 0$ and the so far success $G_{i,n-1}$ of the actor A_i (Schurz 2019, Sec. 6.6.2); formally, $w_{i,n-1} = e^{\eta G_{i,n-1}}$ calculated for each round of a given game using the respective leaderboard
Outcome	Expected $\bar{p}_{BP} > 0.5$; voters' competence increased by the <i>G</i> -class optimality of OFMI	Expected $\bar{p}_{BP} > 0.5$; voters' competence increased by the universal optimality of MFMI
Impact on Voters and Governance	Epistocracy controlled by the following principle – growing exclusivity of the prebunking and/or debunking policies to increase the likelihood that $\bar{p}_{BP} > 0.5$	Epistemic democracy in the sense of no presuppositions about the actors – everybody is let in and according to their predictive performance allowed to influence the weighted average, i.e. predictive performance-based respect and inclusivity to increase the likelihood that $\bar{p}_{BP} > 0.5$
Long-Term Viability	Improbable; the exclusionary policies become a likely cause behind the accelerating transition to hyper-plurality, a paradoxical end-result	Possible; not everybody needs to be persuaded by the performance-based MFMI evaluations to avoid $\bar{p}_{BP} < 0.5$ (cf. Goodin and Spiekermann 2018, 92–94); the delivery mechanism of MFMI predictions/evaluations seems to be as important as the evaluation itself

MFMI's optimality is not limited to G -class prediction games. During each round of the game, MFMI can utilize the available predictive capability of the whole pool of actors instead of a single actor, as in the case of OFMI, which safeguards MFMI against Type 2a actors.

3. No-Regret Voter Enhancements

No-regret learning methods such as MFMI can support voters' competence in several ways. First, voters can become optimal responder followers by simply voting according to the eMI's predictions (a weighted average of the actors' predictions for the round n ; formally,

$$pred_n(eMI) = \frac{\sum_{i=1}^N w_{i,n-1} pred_n(A_i)}{\sum_{j=1}^N w_{j,n-1}} \quad (1)$$

where N is the number of actors, see Schurz (2019), Def. 6.4). The eMI's worst-case regret bounds¹¹ hold for all prediction games, including those where no actor predicts correctly. The ability to minimize voters' regret with respect to the best responder in hindsight in an optimal way in each round makes eMI a no-regret learning method that maximizes voters' competence in each round of the game and thus increases the likelihood that $\bar{p}_{BP} > 0.5$.

Second, in case voters are hesitant to follow an 'artificial' optimal responder, the eMI's weights can be used to determine the probability that an actor A_i is the best responder for the round n . Formally, the probability of following the actor A_i in the round n is given as

$$v_{i,n} = \frac{w_{i,n-1}}{\sum_{j=1}^N w_{j,n-1}} \quad (2)$$

The probability vector $\mathbf{v}_n = (v_{i,n})_{1 \leq i \leq N}$ can be used to determine the eMI's success in each round $\hat{g}_n = \mathbf{v}_n^T (\mathbf{g}_{i,n})_{1 \leq i \leq N}$ and minimizes voters' cumulative regret with respect to the best responder in hindsight, thus increasing the likelihood that $\bar{p}_{BP} > 0.5$ by maximizing voters' competence.

Third, if epistemic practices used by the actors are scored as well, Equation (2) can be applied in an analogical way to determine the probability that an epistemic practice EP_i is optimal for the round n . Voters can then decide among the actors who use it. This way other than epistemic factors could be considered, possibly increasing the political acceptability of MFMI. However, the probability that $\bar{p}_{BP} > 0.5$ could fail to increase because some voters might choose to follow actors misusing the optimal epistemic practice identified by MFMI.

Building on results by Mosleh et al. (2021) outlined in Section 1.1, the way in which no-regret voter enhancements, $pred_n(eMI)$ or v_n , are delivered and presented to voters is likely to determine the success of MFMI as a no-regret learning method that supports voters' competence. Diversity in the presentation is important because voters do not form a homogeneous group and adversaries (e.g. political revisionists) can target them in a highly adaptive manner. Online social media platforms are perfect for targeting voters as their recommender systems (cf. Eirinaki et al. 2018), making sure that the voters leaning towards unreliable or revisionist actors receive their message in full. There are generally two ways which could increase the success rate of no-regret voter enhancements in supporting voters' competence and thus decreasing the likelihood that $\bar{p}_{BP} < 0.5$.

First, no-regret voter enhancements should be delivered via private channels and should not bear signs of being personalized to individual voters. Public distribution is likely to increase the probability of reputational costs and in turn decrease voters' social capital and the willingness to reconsider their votes. Personalized enhancements could produce surveillance-related concerns. For diehard followers of Type 2b actors, no-regret voter enhancements are unlikely to make impact because these voters do not consider epistemic success but rather chose to follow authority. However, for voters on the fence who can be swayed, an unintrusive delivery and form of the enhancement, perhaps along the lines of public service announcements, could succeed.

Second, as per Goodin and Spiekermann (2018, 85–86), instead of everybody's tracking performance, only the average competence has to be maintained above random. Small improvements of a few percent points to the mean tracking performance have a significant positive impact on the probability that majority rule leads to the correct outcome even for small groups (Goodin and Spiekermann 2018, 85–86), provided that the best (optimal) responder predicts correctly with high probability (cf. Goodin and Spiekermann 2018, Sec. 5.3). Not only it would be naïve to think that everybody could be persuaded by no-regret voter enhancements, but it is also unnecessary. No-regret voter enhancements never increase the likelihood that $\bar{p}_{BP} < 0.5$ because the algorithm assigning weights to actors is provably optimal even in games with little to no achievable predictive success.

From this perspective, exclusionary voter enhancements like prebunking and debunking, limiting epistemic democracy to a particular class of prediction games, come out as epistemically undemocratic. Exclusionary voter enhancements could be viewed as dogmatic by voters who choose by following authority instead of predictive success. Mistrust could then motivate shifts from evidence produced by the scientific method to pseudoscientific epistemic practices and illiberal forms of government. Depending strictly on the actors' past predictive successes, no-regret voter enhancements are neither dogmatic in the sense of permanent exclusion nor circular in the sense of presuppositions about the actors whose reliability cannot be checked before each round of the game.

The unrestricted, epistemic performance-based inclusivity and respect alone might not be sufficient to protect democracy in hyper-plurality. Rejections or ignorance of scientific findings can no longer be explained away as public misinterpretation of or inattentiveness to facts established by the scientific method. Examining vaccine hesitancy, Maya Goldenberg (2016, 2021) found that mistrust of scientific experts and institutions, becoming seats of power and privilege, explains the increased receptivity to anti-vaccination messaging. In societies that tend towards hyper-plurality, this epistemic mistrust is likely to only deepen. Epistemically democratic voter enhancements (e.g. no-regret ones) should fare better than 'epistocratic' ones (e.g. prebunking/debunking), which contributed to the growth of mistrust even in plurality. Although the impact of no-regret voter enhancements on rebuilding the epistemic trust and democratic consensus remains an open question, compared to no-regret voter enhancements, exclusionary ones come with political costs.

Apart from mis/disinformation, mistrust of science can also result from undeclared or incompletely identified social and ethical value judgments involved in scientific modelling. Harvard et al. (2021) show in the case of a COVID-19 vaccination model that value judgements in scientific modelling are unavoidable and involving the public is necessary. Without it, expert value judgements inform modelling that's results can be mistaken for or even presented as facts obtained using supposedly value-free scientific methods. What to represent in a model and how to interpret the results of modelling are never value-free actions but rather social and ethical value judgements that are selective (Harvard et al. 2021). If experts perform the selection and interpret the results as facts, the epistocratic outcome will not increase trust in science by the public. Understanding science as a value-free enterprise suggests that scientific judgements can be separated from value judgements, which is an ideal hardly ever achievable or desirable (Douglas 2009). Trust in science will not increase by concealing value judgements behind the value-free ideal. Instead, trust in science and scientists involved in policymaking can increase if voters are involved in value judgements interfaced with scientific judgements in a way that is interpretable by the public.

The exclusionary enhancements can cause epistemic stigmatisation of voters and sometimes converge to a complete lack of empathy, indiscriminately accusing voters of participation in information or even hybrid warfare which aims to destroy liberal democracy (cf. Mälksoo 2018). Although understandable, such a response produces distrust and political 'scripturalism' rather than democratic respect and inclusion. Political revisionism is a method that seeks to create hyper-plurality by matching facts with mis/disinformation tailor-made to fit voters' grievances with the status quo. Attempts to debunk such realities and delegitimize

or possibly stigmatise their followers fail to engage the method itself. Even the most exhaustive debunking effort cannot prevent production of new pieces of mis/disinformation that branch out the hyper-plurality. Debunking can, however, transform a democratic plurality into an increasingly narrow ‘scripturalist lane’, desperately seeking to maintain a coherent and reflective enough notion of what the status quo represents vis-à-vis the alternative realities of interpretative individualists. If added to the toolbox of democratic theory, no-regret voter enhancements could put the political costs of exclusionary enhancements into perspective.

No-regret enhancements provide voters from societies with a high degree of epistemic specialization information on the predictive performance of political actors across the whole range of tracked issues (maintaining voters’ competence in networked societies is an issue in itself, see Hahn, Ulrik Hansen, and Olsson 2020). Centrally maintained voter enhancements solve the issue of the individual voters’ inability to track the prediction game and correctly score the actors. Leaderboards holding actors’ predictive performance on a single issue allow for context-dependent scoring rules (Douven 2020) that help with attuning the scoring of actors’ performance on different issues. In order to participate in the prediction game, every actor predicting on an issue is required to submit their prediction in each round of the game so that MFMI avoids intermittent submissions (Schurz 2019, Subsection 7.2.2).

It is sensible to expect that in hyper-plurality the number of actors will grow unboundedly even if some actors become inactive over time (cf. Schurz 2019, 180–81; also; Schurz 2021). The evaluation of predictive performance of actors entering the prediction game after it has started can be solved by self-completion (Schurz 2019, 180–83). The missing predictive performance of latecomers is filled with the predictive performance of MFMI (Schurz 2019, 181–83). Self-completion gives new actors a fair chance to quickly demonstrate their predictive performance while not endangering eMI’s universal optimality (Schurz 2019, 181–83). However, if the number of actors does not grow slower than exponential with the number of rounds in the prediction game, eMI’s asymptotic regret bound, and thus its universal optimality, does not hold (Schurz 2019, Theorem 7.3 [2a], 183). Such an increase in the number of actors is unlikely, as it would cause an epistemic anarchy that is hostile even for Type 2a and 2b actors and their revisionist aims.

No-regret voter enhancements require maintaining leaderboards of actors’ predictive performance for each of the tracked issues. In case the number of rounds in a game increase rapidly and the pool of actors is large, this can be a difficult and costly organizational task (more so if fact-checking is involved). In such situations, local meta-induction proposed by Schurz (2012) could be considered. Here, it is assumed that voters are willing to influence their views by epistemically successful actors with whom they are in contact and in exchange become sources of positive epistemic influence in their neighbourhoods (Schurz 2012, Theorem 3 for a model defining conditions under which voters converge to the success rate of epistemically successful actors). No-regret enhancements using local meta-induction are still compatible with voter independence. Voters are not encouraged to copy each other. Instead, upstream nodes in the communication network belong among common causes (see Section 2) for each epistemic voter neighbourhood. By way of simulation, Schurz’s (2012, Figure 4) results also show that if voter competence drops below a threshold (voters lack a truth-bias, 14–15), because no-regret voter enhancements are disregarded, i.e. local meta-induction is not practiced, positive epistemic impacts of successful actors do not spread through the network. Then, if majority rule is used to aggregate votes and voters are more likely to err than to vote correctly, increasing the number of voters decreases the likelihood that the election leads to the correct outcome (conditions for the convergence guaranteed by the Law of Large Numbers are not met, [ibid., also Dietrich and Kai 2022, Section 5.3]).

4. Conclusion

We showed that a type of no-regret learning called meta-induction can help to satisfy the competence assumption of jury theorems by no-regret voter enhancements. Their contribution is twofold. The theoretic one shows that no-regret enhancements can make the competence assumption less controversial and no-regret learning can, therefore, contribute to the epistemic analysis of democracy. The empirical one shows that the epistemic crisis of democracy caused by mis/disinformation does not necessitate a priori exclusions of supposedly unreliable or malicious political actors. Rather, the level of inclusion of the actors in no-regret enhancements is determined by their past predictive performance and contributes to maintaining an epistemic respect. The respect derives from avoiding any presuppositions about actors, possibly leading to their a priori exclusions. As a result, all actors can demonstrate their epistemic capabilities without any bias. As the difference between no-regret enhancements and the best actor in hindsight is unimprovable for certain no-regret algorithms, the former has optimal truth-tracking properties regardless of the epistemic conditions. The no-regret support of voters' competence makes it possible to consider jury theorems a useful tool for the epistemic analysis of democracy not only in plural but also hyper-plural political environments.

Notes

1. Among revisionists we count, for example, political actors subscribed to a cynical epistemic view that 'nothing is true and everything is possible' (Pomerantsev 2014) championed by the Russian state and its proxies today.
2. If falsehoods are exposed simultaneously with their circulation (or before), we speak about prebunking. In case the falsehoods are exposed after they reached their target audience, we speak about debunking.
3. We are grateful to a referee for suggesting populism as a cause of systematic deception.
4. We are thankful to a referee for the journal suggesting that such a relaxation should be emphasized.
5. in this case Twitter.
6. In case real-valued predictions are interpreted as actors' epistemic probabilities, prediction games as they are used here result in a meta-inductive probability aggregation of Bayesian predictors, including proper scoring rules (Schurz 2019, Sect. 7.1).
7. <https://www.politifact.com/environment/>
8. Relying on findings produced by the scientific method whenever possible.
9. The mean voter competence to follow the best predictor \bar{p}_{BP} is a relaxation of Condorcet's original condition, requiring each voter to have an equal competence (cf. Goodin and Spiekermann 2018, Sec. 3.1.1). For both non-asymptotic and asymptotic results to hold for heterogeneous voter competences, it is necessary that the individual voter's competences p_{BP} are distributed symmetrically around the mean \bar{p}_{BP} (Goodin and Spiekermann 2018, Sec. 3.1.1).
10. The results of prediction games can support voters' competence also in cases where only a part of the issue can be treated as a prediction game in terms explained earlier.
11. See Schurz (2019), Theorem 6.9 for the upper bound on the eMI's regret in case the number of rounds is known beforehand [i], for $n \geq 1$ [ii], and [iii] for the upper bound in the asymptotic case.

Acknowledgments

We would like to thank two anonymous reviewers for the journal for contributing detailed feedback allowing us to improve the original version of the manuscript.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This output was supported by the NPO 'Systemic Risk Institute' no. [LX22NPO5101] funded by European Union – Next Generation EU (Ministry of Education, Youth and Sports, NPO: EXCELES).

Notes on contributors

Petr Spelda is an Assistant Professor at the Department of Security Studies, Charles University. His research focuses on safe machine learning and building-blocks of optimal inductive inference for trustworthy artificial intelligence. He publishes in several fields, including computer and political science or philosophy. More about his interdisciplinary research can be found at <https://research.spelda.cz/>.

Vit Stritecky is an Associate Professor in International Security and Head of Department of Security Studies, Charles University. His research focuses on policy and regulatory issues connected to machine learning and on the role of standardization in addressing safety and security concerns about technology. He is author and editor of several books and journal special issues on international security and nearly 40 papers and book chapters whose topics range from strategic studies to sociotechnical understanding of artificial intelligence.

John Symons is Professor of Philosophy at The University of Kansas and Director of the Center for Cyber Social Research. His research areas include philosophy of technology and general philosophy of science. He is author of 11 books or edited volumes and over 60 articles and book chapters.

ORCID

Petr Spelda  <http://orcid.org/0000-0003-4199-645X>

Vit Stritecky  <http://orcid.org/0000-0003-1778-3657>

John Symons  <http://orcid.org/0000-0001-8173-2948>

References

- Anderson, Elizabeth. 2006. "The Epistemology of Democracy." *Episteme* 3 (1–2): 8–22. <https://doi.org/10.3366/epi.2006.3.1-2.8>.
- Bekoulis, Giannis, Christina Papagiannopoulou, and Nikos Deligiannis. 2023. "A Review on Fact Extraction and Verification." *ACM Computing Surveys* 55 (1): 1–35. <https://doi.org/10.1145/3485127>.
- Bennett, Lance, and Steven Livingston. 2020. *The Disinformation Age: Politics, Technology, and Disruptive Communication in the United States*. Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/9781108914628>.
- Biddle, Justin, and Anna Leuschner. 2015. Climate Skepticism and the Manufacture of Doubt: Can Dissent in Science Be Epistemically Detrimental? *European Journal for Philosophy of Science* 5 (3): 261–278. <https://doi.org/10.1007/s13194-014-0101-x>.
- Brennan, Jason. 2016. *Against Democracy*. Princeton: Princeton University Press.
- Cesa-Bianchi, Nicolò, and Gábor Lugosi. 2006. *Prediction, Learning, and Games*. Cambridge, UK: Cambridge University Press.
- Chambers, Simone. 2017. "Balancing Epistemic Quality and Equal Participation in a System Approach to Deliberative Democracy." *Social Epistemology* 31 (3): 266–276. <https://doi.org/10.1080/02691728.2017.1317867>.
- Condorcet, Marquis de. 1785/1785. "Essay on the Application of Mathematics to the Theory of Decision-Making." In *Selected Writings*, edited and translated by Keith Michael Baker, 33–70. Indianapolis, IN: Bobbs-Merrill.
- Cook, Jason. 2016. "Using an Interdisciplinary MOOC to Teach Climate Science and Science Communication to a Global Classroom." American Geophysical Union, Fall Meeting 2016. <https://ui.adsabs.harvard.edu/abs/2016AGUFMED13A0923C>.
- Cook, Jason. 2019. "Using Mobile Gaming to Improve Resilience Against Climate Misinformation." *American Geophysical Union, Fall Meeting 2019*. <https://ui.adsabs.harvard.edu/abs/2019AGUFMPA13A.10C>.
- Cook, Jason. 2020. *Cranky Uncle Vs. Climate Change: How to Understand and Respond to Climate Science Deniers*. New York: Citadel Press.
- Dietrich, Franz, and Kai Spiekermann. 2022. "Jury Theorems." In *The Stanford Encyclopedia of Philosophy (Summer 2022 Edition)*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/sum2022/entries/jury-theorems/>.
- Douglas, Heather E. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press.
- Douven, Igor. 2020. "Scoring in Context." *Synthese* 197 (4): 1565–1580. <https://doi.org/10.1007/s11229-018-1867-8>.
- Douven, Igor, and Rainer Hegselmann. 2021. "Mis- and Disinformation in a Bounded Confidence Model." *Artificial Intelligence* 291 (103415): 103415. <https://doi.org/10.1016/j.artint.2020.103415>.

- Eirinaki, Magdalini, Jerry Gao, Iraklis Varlamis, and Konstantinos Tserpes. 2018. "Recommender Systems for Large-Scale Social Networks: A Review of Challenges and Solutions." *Future Generation Computer Systems* 78:413–418. <https://doi.org/10.1016/j.future.2017.09.015>.
- Estlund, David. 2003. "Why Not Epistocracy?" In *Desire, Identity and Existence: Essays in Honor of T.M. Penner*, edited by Reshotko Naomi, 53–69. Berrima: Academic Printing and Publishing. <https://doi.org/10.2307/j.ctv10kmfns.8>.
- Fuerstein, Michael. 2008. "Epistemic Democracy and the Social Character of Knowledge." *Episteme* 5 (1): 74–93. <https://doi.org/10.3366/E1742360008000245>.
- Goldenberg, Maya J. 2016. "Public Misunderstanding of Science? Reframing the Problem of Vaccine Hesitancy." *Perspectives on Science* 24 (5): 552–581. https://doi.org/10.1162/POSC_a_00223.
- Goldenberg, Maya J. 2021. *Vaccine Hesitancy: Public Trust, Expertise, and the War on Science*. Pittsburgh, PA: University of Pittsburgh Press.
- Goodin, Robert E, and Kai Spiekermann. 2018. *An Epistemic Theory of Democracy*. Oxford: Oxford University Press.
- Hahn, Ulrike, Jens Ulrik Hansen, and Erik J Olsson. 2020. "Truth Tracking Performance of Social Networks: How Connectivity and Clustering Can Make Groups Less Competent." *Synthese* 197 (4): 1511–1541. <https://doi.org/10.1007/s11229-018-01936-6>.
- Harvard, Stephanie, Eric Winsberg, John Symons, and Amin Adibi. 2021. "Value Judgments in a COVID-19 Vaccination Model: A Case Study in the Need for Public Involvement in Health-Oriented Modelling." *Social Science & Medicine* 286:114323. <https://doi.org/10.1016/j.socscimed.2021.114323>.
- Holst, Cathrine, and Anders Molander. 2017. "Public Deliberation and the Fact of Expertise: Making Experts Accountable." *Social Epistemology* 31 (3): 235–250. <https://doi.org/10.1080/02691728.2017.1317865>.
- Holst, Cathrine, and Anders Molander. 2018. "Asymmetry, Disagreement and Biases: Epistemic Worries About Expertise." *Social Epistemology* 32 (6): 358–371. <https://doi.org/10.1080/02691728.2018.1546348>.
- Holst, Cathrine, and Anders Molander. 2019. "Epistemic Democracy and the Role of Experts." *Contemporary Political Theory* 18 (4): 541–561. <https://doi.org/10.1057/s41296-018-00299-4>.
- Hume, David. 1739 (1978). *A Treatise on Human Nature. Book I: On Human Understanding*. <https://doi.org/10.1093/oseo/instance.00046221>.
- Iyengar, Shanto, and Douglas S. Massey. 2019. "Scientific Communication in a Post-Truth Society." *The Proceedings of the National Academy of Sciences of the United States of America* 116 (16): 7656–7661. <https://doi.org/10.1073/pnas.1805868115>.
- Landemore, Hélène. 2012. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, NJ: Princeton University Press.
- Landemore, Hélène. 2017. "Beyond the Fact of Disagreement? The Epistemic Turn in Deliberative Democracy." *Social Epistemology* 31 (3): 277–295. <https://doi.org/10.1080/02691728.2017.1317868>.
- Landemore, Hélène. 2020. *Open Democracy: Reinventing Popular Rule for the Twenty-First Century*. Princeton, NJ: Princeton University Press.
- Lane, Melissa. 2016. "Political Theory on Climate Change." *Annual Review of Political Science* 19 (1): 107–123. <https://doi.org/10.1146/annurev-polisci-042114-015427>.
- Lewandowsky, Stephan, John Cook, Ullrich Ecker, Dolores Albarracín, Michelle A Amazeen, Panayiota Kendou, Doug Lombardi, et al. 2020. *The Debunking Handbook 2020*. <https://doi.org/10.17910/b7.1182>.
- Lewandowsky, Stephan, John Cook, and Elisabeth Lloyd. 2018. "The 'Alice in Wonderland' Mechanics of the Rejection of (Climate) Science: Simulating Coherence by Conspiracism." *Synthese* 195 (1): 175–196. <https://doi.org/10.1007/s11229-016-1198-6>.
- Mackie, Diane, and Joel Cooper. 1984. "Attitude Polarization: Effects of Group Membership." *Journal of Personality and Social Psychology* 46 (3): 575–585. <https://doi.org/10.1037/0022-3514.46.3.575>.
- Mälksoo, Maria. 2018. "Countering Hybrid Warfare as Ontological Security Management: The Emerging Practices of the EU and NATO." *European Security* 27 (3): 374–392. <https://doi.org/10.1080/09662839.2018.1497984>.
- Martí, Jose L. 2006. "The Epistemic Conception of Deliberative Democracy Defended Reasons, Rightness and Equal Political Autonomy." In *Deliberative Democracy and Its Discontents*, edited by Besson Samantha and Jose L. Martí, 27–56. Aldershot, UK: Ashgate.
- McDermott, Rose. 2019. "Psychological Underpinnings of Post-Truth in Political Beliefs." *PS: Political Science & Politics* 52 (2): 218–222. <https://doi.org/10.1017/S104909651800207X>.
- McIntyre, Lee. 2020. *Post-Truth*. Cambridge, MA: MIT Press.
- Moravec, Václav, Veronika Macková, Jakub Sido, and Kamil Ekštejn. 2020. "The Robotic Reporter in the Czech News Agency: Automated Journalism and Augmentation in the Newsroom." *Communication Today* 11 (1): 36–52.
- Mosleh, Mohsen, Cameron Martel, Dean Eckles, and David Rand. 2021. "Perverse Downstream Consequences of Debunking: Being Corrected by Another User for Posting False Political News Increases Subsequent Sharing of Low Quality, Partisan, and Toxic Content in a Twitter Field Experiment." *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, Yokohama, Japan, 1–13.
- NASA. 2021. "Vital Signs of the Planet, Global Climate Change." <https://climate.nasa.gov/vital-signs/carbon-dioxide/>.
- O'Connor, Cailin, and James O. Weatherall. 2019. *The Misinformation Age: How False Beliefs Spread*. New Haven, CT: Yale University Press.

- Persily, Nathaniel, and Joshua A. Tucker. 2020. *Social Media and Democracy: The State of the Field, Prospects for Reform*. Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/9781108890960>.
- Pomerantsev, Peter. 2014. *Nothing is True and Everything is Possible: The Surreal Heart of the New Russia*. New York: Public Affairs.
- Schurz, Gerhard. 2008. "The Meta-Inductivist's Winning Strategy in the Prediction Game: A New Approach to Hume's Problem." *Philosophy of Science* 75 (3): 278–305. <https://doi.org/10.1086/592550>.
- Schurz, Gerhard. 2012. "Meta-Induction in Epistemic Networks and the Social Spread of Knowledge." *Episteme* 9 (2): 151–170. <https://doi.org/10.1017/epi.2012.6>.
- Schurz, Gerhard. 2019. *Hume's Problem Solved: The Optimality of Meta-Induction*. Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/11964.001.0001>.
- Schurz, Gerhard. 2021. "Meta-Induction Over Unboundedly Many Prediction Methods: A Reply to Arnold and Sterkenburg." *Philosophy of Science* 88 (2): 320–340. <https://doi.org/10.1086/711587>.
- Schwartzberg, Melissa. 2015. "Epistemic Democracy and Its Challenges." *Annual Review of Political Science* 18 (1): 187–203. <https://doi.org/10.1146/annurev-polisci-110113-121908>.
- Stanley, Ben. 2008. "The Thin Ideology of Populism." *Journal of Political Ideologies* 13 (1): 95–110. <https://doi.org/10.1080/13569310701822289>.
- Symons, John, and Stacy Elmer. 2022. "Resilient Institutions and Social Norms: Some Notes on Ongoing Theoretical and Empirical Research." *Merrill Series on the Research Mission of Public Universities*, MASC Report No. 125, : 95–112. <https://doi.org/10.17161/merrill.2022.19584>
- Winkler, Bärbel, and John Cook. 2020. "The Story of Skeptical Science: How Citizen Science Helped to Turn a Website into a Go-To Resource for Climate Science." *EGU General Assembly 2020*. <https://doi.org/10.5194/egusphere-egu2020-562>.
- Winsberg, Eric. 2018. *Philosophy and Climate Science*. Cambridge: Cambridge University Press.
- Zimdars, Melissa, and Kembrew McLeod. 2020. *Fake News: Understanding Media and Misinformation in the Digital Age*. Cambridge, MA: The MIT Press. <https://doi.org/10.7551/mitpress/11807.001.0001>.