

# How Knowledge Brokers Shape the Evidence-Based Policy Landscape\*

Hannah Rubin      Remco Heesen      Mike D. Schneider  
Alejandro Bortolus      Emelda E. Chukwu      Chad L. Hewitt  
Ricardo Kaufer      Hanna Metzen      Veli Mitova  
Anne Schwenkenbecher      Evangelina Schwindt  
Temitope O. Sogbanmu      Helena Slanickova  
Katie Woolaston      Li-an Yu

May 29, 2026

## Abstract

Knowledge brokers, usually conceptualized as passive intermediaries between scientists and policymakers in evidence-based policymaking, are understudied in philosophy of science. Here, we challenge that usual conceptualization. As agents in their own right, knowledge brokers have their own goals and incentives, which complicate the effects of their presence at the science-policy interface. We illustrate this in an agent-based model and suggest several avenues for further exploration of the role of knowledge brokers in evidence-based policy.

## 1 Introduction

The accumulated products of science have a dominant effect on our wider social world of believers, learners, and the like. Science, more than anything else, shapes what *we* know. This observation is not new. It is, for instance, the linchpin of Karl Popper’s massively influential program in methodology, as he clarifies in the Preface to the 1959 English edition of *The Logic of Discovery*: “The central problem of epistemology has always been and still is the problem of the growth of knowledge. *And the growth of knowledge can be studied best by studying the growth of scientific knowledge.*” [Popper, 2002, p. xix, original emphasis].

---

\*A shortened version of this paper is forthcoming in *Philosophy of Science*. Please cite from that version if possible. This work was supported by the Zentrum für interdisziplinäre Forschung (Centre for Interdisciplinary Research, ZiF) at Bielefeld University (Grant RG2023/1) and by the Volkswagen Foundation. The authors would like to thank them as well as audiences at the Philosophy of Science Association biennial conference in New Orleans, at the PPE conference in London, and at Kansas State University and the University of Missouri.

What does it mean to consider scientific knowledge as the center point of the question of the overall growth of our knowledge? Perhaps that the ends of scientific knowledge are ultimately found not in ongoing science, but in the context of that science’s embedding in society. Indeed, billions of dollars of economic activity are dedicated to the notion that the connection between scientific knowledge and the education and betterment of society needs maintenance. The maintenance crews are impossible to miss: governmental expert groups, international scientific committees, lobbyist think tanks, research institutes, and so on. Whether local, regional, national, international, or global, they go by the label *Knowledge Brokers* (KBs). The KB label derives from a nominal purpose: to broker knowledge between those at its source (typically, scientific research) and societal users (often policymakers).

Despite the strategic position of KBs at the interface between scientific research and society, philosophers of science have generally neglected how KBs get involved in brokering the knowledge that shapes what we know.<sup>1</sup> Not so in adjacent fields: in the science and technology studies, political science, and knowledge management literatures that research *knowledge brokering*, there is a long history of discussion about the ‘linear model’ of the science-policy interface, which posits that science feeds into policymaking activities through KBs acting as intermediaries (though see [Edgerton, 2004] for historical context). That KBs perform in the role of intermediaries according to the linear model allegedly gives historical context for the naming convention that emphasizes a brokering arrangement between science and society. But in these literatures, it is generally recognized that the linear model suppresses crucial complexities: at the science-policy interface, all relevant parties, KBs among them, *co-produce* policy. Hence, one commonly finds endorsements of a rival ‘co-production model’ [Jasanoff, 2004, Bandola-Gill et al., 2023]. Examining the roles of KBs within the context of either the linear or co-production model is an active area of ongoing research.

To catch up with adjacent literatures, and to properly study the growth of scientific knowledge in society, philosophers of science require an epistemology of KBs. In this article, we get started by asking: what are the general effects of KBs on how policymakers might learn and decide how to act, given ongoing scientific research? We argue that we must go beyond a passive characterization of KBs given in terms of their ideal roles (‘relays’ in the linear model, or ‘facilitators’ in the co-production model). Instead, KBs are strategic agents, with distinctive

---

<sup>1</sup>Some philosophers of science have, however, considered the epistemic consequences of particular knowledge brokering arrangements, e.g., philosophers of climate science have considered the IPCC (an example of a KB), as in [Katzav, 2014] and references therein. Others have focused on epistemic consequences of science’s embedding in a larger social context, e.g., work on science communication [Dethier, 2025] or misinformation [O’Connor and Weatherall, 2018, Friedman and Šešelja, 2023] or science funding [Sanders and Robison, 1992, Pinto, 2021, Shaw, 2022]. And in the epistemology of data, Sabina Leonelli has argued that data curation has epistemic significance, in that it generates theoretical knowledge [Leonelli, 2009, 2015, Leonelli and Tempini, 2020]. The model we develop below treats KBs, in effect, as data curators with a similar moral: they generate scientific knowledge found at the outcome of processes that amount to packaging science for societal uptake.

incentives surrounding their participation in an overall knowledge ecosystem (§2). On this basis, and given our interest in studying the effects of KBs, we develop a model of information flow through that knowledge ecosystem, which lets us track the effects of the decisions motivated by the diversity of incentives at play (§3). We show that KBs have effects on the overall epistemic health of the society in which they operate that are not captured by the existing literatures (§4). Finally, we discuss the future research that is needed to more fully study the impacts of KBs (§5).

## 2 Incentives Surrounding Knowledge Brokering

There are three types of agents operating at the science-policy interface: scientists, or more generally, knowledge producers; policymakers, or more generally, knowledge users; and KBs. In a particular situation, each of these three types is instantiated by some number of agents, where agents can be either individuals or organizations.<sup>2</sup> These agents each have their own goals and implicit or explicit beliefs about how best to achieve these goals. They are therefore *incentivized* to act in some ways as opposed to others.

The epistemic consequences of the goals and incentives of scientists are quite well studied [Strevens, 2003, Heesen, 2017, Zollman, 2018, O’Connor and Bruner, 2019, Rubin and Schneider, 2021]. For instance, it can matter in questions of pursuit that an important goal for most scientists is to build a reputation for important discoveries, as this is key to a successful academic career [Merton, 1973]. Policymakers have their own goals and incentives, quite different in nature from those of scientists [Choi et al., 2005, Uneke et al., 2011, Beesley et al., 2022]. It is worth noting here that the phrase ‘policymaker’ is used in different ways. Here we are primarily interested in the people who organize relevant knowledge and write and implement detailed policy proposals (think: civil servants or bureaucrats) rather than politicians who set broad policy directions.

The linear model portrays KBs as (merely) transmitting information from scientists to policy makers, perhaps reducing uncertainty or improving clarity while doing so. But KBs, being agents in their own rights, have their own goals and incentives. These will shape what information they are willing and able to receive, and what information they are interested in passing on. Even before knowing any specifics about their goals and incentives (which we get into below), we should expect the presence of KBs to affect the science-policy interface in more complicated and interesting ways than if they passively transmitted information. This point appears to be underappreciated in the literatures on KBs [see, e.g., MacKillop et al., 2020, for an overview].

Our aim then is to study the effects of KBs with different sorts of incentives

---

<sup>2</sup>Two caveats. First, the number of individuals or organizations instantiating a type in a particular situation could be zero, though our focus is on situations where all three types are represented. Second, we understand organizations broadly to include any group of individuals acting sufficiently coherently to appear as a single agent to others, regardless of whether this group is formally incorporated.

on the science-policy interface, specifically in terms of what science gets done and what policies get implemented. In a companion paper with a more empirical focus, we distinguish a wide range of different types of KBs, the goals or incentives for people to get involved with these KBs, and the goals or incentives of the KBs themselves [Bortolus et al., 2025]. Types of KBs include non-profit organizations that support the professional development of their members (professional societies, academies), expert groups with governmental mandate (standing committees, task forces), non-profit organizations that perform research and advocacy (Indigenous organizations, NGOs, think tanks), industrial organizations that do similar (lobbyists, consultants), and more.

As discussed in Bortolus et al. [2025], each of these KB types have incentives for why they want to interact with policymakers, and policymakers have different incentives for wanting to interact with them. To briefly illustrate: policymakers may be looking for clarity on their questions or translation of research into language they understand. But policymakers may also expect other benefits from interacting with a particular type of KB: perhaps social benefits from even being seen to interact with the KB (e.g., interacting with academies may help create an impression of rigor in the policymaking process, and interacting with Indigenous organizations may be viewed positively by those supporting their causes). Equally, consider the incentives that different types of KBs have to interact with policymakers: in some cases, financial incentives matter; in others, influence.

Moreover, consider the incentives for different types of KBs to interact with researchers, and for researchers to interact with them. KBs may get free labor or new knowledge from researchers, or they may derive prestige from having fancy scientists as members. Conversely, researchers may look forward to prestige to be gained, e.g., by being associated with a highly-regarded academy or a government-mandated expert group that receives a lot of media attention. Alternatively, interacting with KBs may yield resources for researchers (KBs may pay members or provide access to research equipment or conference space) or offer paths to greater influence and impact for their research.

The previous two paragraphs hint at the complexity that is ultimately of concern in developing an epistemology of KBs. Below, we construct a first model to furnish insight into the epistemic consequences of the presence of KBs. For this purpose, we simplify the broad diversity of KB types and incentives while aiming to retain some of the major distinctions just outlined. Specifically, while the KB motivations we discuss are not mutually exclusive or exhaustive, we will distinguish three broad types of KBs, each characterized by a dominant motivation. The first type are primarily motivated by good science, whatever they take that to be, and they want the ‘best’ scientists as members. We will refer to these as *productivity-driven* KBs. The second type are primarily motivated by having some policy impact. Here we are thinking for example of standing committees who really only have a reason to exist if the policymaker is listening to what they have to say. Such KBs, we assume, care about being listened to by policymakers more than the actual policy directions chosen. We call these *impact-driven* KBs. The third type, in contrast, are primarily

motivated by their policy agenda or mission. Think of an industrial lobbying organization whose goal is to achieve looser regulations. These will be referred to as *agenda-driven* KBs.

We assume, moreover, that individual researchers also vary in their motivations, in a way that matches the KB types. Thus, productivity-driven researchers care primarily about producing the most high-quality science. Impact-driven researchers care primarily about influencing policy regardless of direction, perhaps because they want the prestige associated with policy impact. And agenda-driven researchers care primarily about steering policy outcomes in a particular direction, e.g., a researcher who chooses their career path because they are motivated to preserve ecosystems.

### 3 Model

The basic setup of the model is a community of scientists performing research, which KBs then process and communicate to the policymaker. As these interactions take place, the incentives of both scientists and KBs affect which scientists associate with which KBs and the communication channels through which the policymaker receives evidence relevant to policy decisions. Given the diversity of incentives among researchers and KBs, using agent-based modeling can help us track the ultimate effects arising out of the decisions made.

Two things to note about this setup: First, the KBs do not simply tell the policymakers the contents of all the research they have collected. They process this information and communicate what they take to be important to tell the policymaker. This is a necessary simplification of what the research says; if they did not simplify in any way, they would not be doing their job. We refer to their communication with the policymaker as a ‘signal’, to capture that there is information being transmitted along a channel.

Second, policymakers do not always listen to every KB that is trying to communicate with them. They might see some of them as making claims that are totally implausible, outside their expectations or ability to act, or at odds with their own agenda. In what follows, we assume that policymakers do not have an agenda that causes them to be unresponsive to evidence, but we do assume they have some prior expectations about the kind of evidence they will receive, and they might ignore signals that are too far out of line with those expectations. Note that this is a minimal (and, hopefully, minimally controversial!) notion of policymaker commitment to practices ordinarily associated with ‘evidence-based policy’. And recall: policymakers are mostly government bureaucrats, not politicians.

Here is how the model works, also summarized in the schematic in Figure 1.<sup>3</sup> We start with a particular question scientists are trying to answer, like how much of an invasion threat some species is, and there is some fact of the matter that research aims to uncover, e.g., the chances the population will rise to a

---

<sup>3</sup>MATLAB code is available at: [https://osf.io/5cbnf/?view\\_only=46aa69f06d204af0adc318bc22aaf5e4](https://osf.io/5cbnf/?view_only=46aa69f06d204af0adc318bc22aaf5e4).

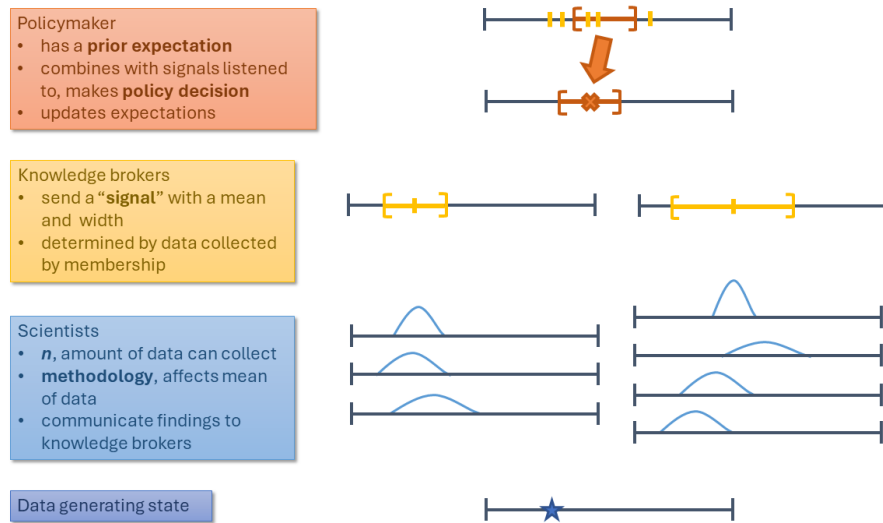


Figure 1: A schematic overview of the model. The process is depicted from the bottom up, starting with the data generating state (indigo star) and ending in policymaker learning and action. Scientists gather data (blue distributions), which the KBs consolidate into a signal (yellow interval and vertical line) to send to the policymaker. The policymaker incorporates those KB signals inside their expectations (orange interval) and makes a policy decision (orange X) based on those signals. The policymaker then updates their expectations (orange arrow) for the next round, with a new interval surrounding this decision.

certain level and scale of negative effects should it rise to that level. This fact of the matter is represented by a *data-generating state* in the model. As depicted in Figure 1, the kinds of questions we are considering have a range of possible answers along a spectrum. We assume the data-generating state is chosen uniformly at random between 0 and 1 (representing, e.g., no invasion threat versus the worst invasion threat readily imaginable).

Scientists' evidence is related to this data-generating state, but they don't all gather data in the same way. Each scientist is assigned one of the three drivers, and agenda-driven scientists additionally have an agenda, 0 or 1, representing, e.g., a position on environmental regulation or deregulation.<sup>4</sup> Each scientist has a different *productivity*, or amount of data they collect,  $n$ , relevant to the research question at hand.<sup>5</sup> They also each have a different *methodology* they employ, which is captured by slightly adjusting the mean of their data away from the data-generating state.<sup>6</sup> When scientists gather data, they draw  $n$  data points

<sup>4</sup>The distribution of scientist and KB types is drawn uniformly at random from within the simplex of possible type distributions given a fixed total number of scientists and KBs.

<sup>5</sup>Scientist productivity is drawn uniformly at random between 10 and 100.

<sup>6</sup>Scientist methodology was determined by a draw from a normal distribution with the

from a distribution with mean equal to their methodology and variance 0.1. This variation in methodology need not reflect poorly on any of the scientists or represent any kind of illegitimate science going on<sup>7</sup>, it is just a feature of diverse methods at work to tackle complex problems. Some methods might be better suited to investigating short term effects on fish populations, others on long term tree growth, and others on economic impact. The negative effects uncovered might be of different magnitudes, leading to different estimates of the severity of the invasion threat. The mean of all the data drawn (given the diversity of methodologies and productivity across scientists) is the best possible estimate of what is ‘really’ going on with the underlying question, so we call this ‘the voice of the science’.

Scientists do research, then communicate their findings to KBs who process it into a signal to send to the policymakers. The signal a KB sends to the policymaker has a mean — what they take to be the voice of the science, or what the evidence they have access to on average recommends — and a width — a measure of how uncertain they think that evidence is. Both of these are determined by what the data collected by their membership says: the KB simply reports the mean and 95% confidence interval of all the data brought to them by their members.

These signals are sent to the policymaker, who, again, has some prior expectation about what they will hear, also initially randomly chosen between 0 and 1. They will listen to those signals they take to be plausible, given their expectations, using that to make their policy decision. We varied the *policy width*, or the area around their expectation where they are still willing to listen to a KB’s signal. This captures the fact that policymakers might have more or less specific expectations regarding particular types of questions. The more specific their expectations are, i.e., the narrower their policy width, the less willing they are to accept different kinds of signals.

Policymakers then average the mean of each signal they listened to in order to form a belief about the voice of the science, which they subsequently act upon by taking the action that is most appropriate.<sup>8</sup> We assume that their belief translates directly into action. For instance, they respond with the level of intensity to a given invasion threat that they would regard as the right level if they knew exactly the voice of the science, e.g., on how much to restrict trade in a particular area or how strict regulations need to be on water purity, given an actual invasion threat. Finally, their expectations are responsive to the evidence, so they update their policy expectations for the next round based on what they have learned by setting next round’s expectation equal to this

---

mean equal to the data generating state and variance .1. Agenda-driven agents further had their methodology moved by .5 to capture a tendency to adopt methodologies which support their agenda. Either of these adjustments could lead to scientists’ methodologies falling outside the [0,1] interval.

<sup>7</sup>See also Holman and Bruner [2017] for a defense of assuming the existence of multiple legitimate methodologies.

<sup>8</sup>Note that though the KBs send information about uncertainty (through confidence intervals), this information is ignored. We think this is not an unreasonable assumption. Potential effects of reporting uncertainty are, in any case, reserved for future investigation.

round’s decision. To capture that the problem at hand may be different, the data-generating state and policymaker expectations shift before the start of the next round.<sup>9</sup>

Incentives come into play in determining which scientists are members of which KBs. Researchers are more likely to be assigned to KBs if their type matches (and for agenda-driven agents, if they share an agenda) and if they ‘have what the KB wants’, so to speak. Membership of each scientist is determined by a weighted draw among KBs and an additional ‘no KB’ option capturing failure of some scientists with interests in evidence-based policymaking to get involved in the process. Like the scientists, each KB is assigned one of the three drivers, and agenda-driven KBs additionally have an agenda, 0 or 1. Having a type match increases the weight on KBs of the same type. For agenda-driven scientists and KBs both the type and agenda must match for this weight to be increased.

The weight is also increased if the scientist under consideration produces research with features the KB finds desirable. KBs driven by good science want the ‘best’, most productive scientists, so will be more likely to have members who are more productive than average. KBs driven by their agenda will be more likely to have members whose methodology is closer to 0 or 1, depending on which agenda they have (where, for example, agenda 0 may represent wanting deregulation and agenda 1 wanting more environmental regulation and conservation). KBs driven by impact will be more likely to have members whose methodology is in line with what they think the policymaker will listen to, i.e., relatively close to the policymaker’s expectations about the voice of the science.

We implement this matching by having, for each scientist, a vector determining their matching probabilities. We start with a vector with entries of 1 for each KB and the additional ‘no KB’ option. Entries in the vector are multiplied by some amount  $m$  if their type (or for agenda-driven scientists, their agenda) matches the KB that entry designates. Then, for each scientist: their entries for productivity-driven KBs are increased by  $m$  times their relative productivity (their productivity divided by the average productivity); their entries for impact-driven KBs are increased by  $m$  times their relative policy closeness (their methodology’s distance to the policymaker’s expectations divided by the average distance to the policymaker’s expectations of all scientists’ methodologies); and their entries for agenda-driven KBs are increased by  $m$  times their

---

<sup>9</sup>The data-generating state and policymaker expectations were shifted independently by drawing from a double exponential distribution with scale parameter  $\theta$  and mean equal to the starting point (last round’s state or decision, respectively). This means small shifts are highly likely and large shifts happen only occasionally (if this approach generated a shift outside the  $[0,1]$  interval we discarded the result in favor of a new random draw). Results below are for  $\theta = 1/6$  but similar results were obtained for other values of  $\theta$ . Note that while the data-generating state and policymaker’s expectations are always between 0 and 1, the scientists’ data (and therefore KBs’ signals) are not constrained in the same way. That we, the modelers, stipulate the policymaker’s expectations and data-generating state are each within the unit interval can be understood as our setting an initial condition, which bounds how far apart initial expectations could be from the overall voice of the science. That scientists’ individually favored methods could yield findings for them that are systematically outside the unit interval reflects the fact that sometimes, science can surprise us all.

relative agenda closeness (their methodology’s distance to the KB’s agenda, either 0 or 1, divided by the average distance to the KB’s agenda of all scientists’ methodologies). Each scientist is then independently assigned to a broker with probability proportional to the vector entries. Note that brokers may end up with (substantially) different numbers of members.

This model simplifies in many ways. To name a few (more will be discussed in the conclusion): we focus on knowledge flow within an evidence-based policy ecosystem, and ignore how that knowledge was produced. Thus our model is consistent with both the simplified linear view of KBs, where KBs merely collect data and process it, and the co-production view, where KBs facilitate which data scientists will collect. The model only considers quantitative scientific evidence, which is not the only kind of evidence relevant to evidence-based policymaking. We also only include one policymaker, though in many cases there are likely multiple policymakers. Finally, we assume a fairly rosy picture of the policymaking process; there is no intentional misrepresentation of evidence by any party involved. Therefore, we are describing the influence KB incentives have on the science-policy interface while restricting their possible actions to accurately announcing an aggregation of their members’ signals. This is likely a lower bound on the amount of ‘distortion’ KBs create on the information flow from scientists to policy makers.

## 4 Results

Because of KBs, the total body of scientific knowledge learned by the policymaker is structured by features of the professional social structure of scientific research, in addition to the world itself.

Here, we present results from six different variations of the model. First, we present results from *Random*, a base case with no incentives in play, where matching of scientists and KBs is random. This allows us to evaluate how KBs affect policymaking just by the fact that they process and simplify in their communication of scientific evidence. We then present results from *Matched*, a variation where there is matching based on incentives, as described above. The differences in outcomes between this variation and *Random* show how KB incentives further distort the policymaking process (even though they continue to faithfully represent the evidence at hand).

We then present four variations where some KBs are more influential than others, which we captured by allowing their signals (if listened to) to be given more weight when the policymaker averages signals<sup>10</sup> to form a belief: *Dominant*

---

<sup>10</sup>Each broker is given the same base influence (10 for the results presented here), which is then multiplied by some value if their type is being given extra influence in the condition. How much more influential the broker is than the others being listened to determines the weight their signal is given. Let  $b_i$  be the influence of KB  $i$  minus the average KB influence in the population and  $w_i = \frac{1}{1 + \text{Exp}\{-b_i\}}$ . This equation for  $w_i$  ensures that the term ranges from 0 to 1, where a KB with average influence sits at .5. This term is then normalized by dividing by the sum of the  $w$  of other KBs whose signals are listened to so that the weight a

Variation	Membership	Extra influence
Random	random	n/a
Matched	matched	n/a
Dominant Agenda	matched	agenda-driven KBs w/ agenda 1
External Organizations	matched	agenda-driven KBs
Good Science	matched	productivity-driven KBs
Government Investment	matched	impact-driven KBs

Table 1: Summary of model variations

*Agenda* where all KBs of one agenda have extra influence, *External Organizations* where all agenda driven KBs have extra influence, *Good Science* where productivity driven KBs have extra influence, and *Government Investment* where impact driven KBs have extra influence.<sup>11</sup> These variations are summarized in Table 1. These variations let us explore the effects on policymaking if some KB incentives are allowed to dominate in the process.

In this model, we are interested in how KBs impact policy decisions. The main measure we use to evaluate the effects of KBs is the distance from the ‘truth’. For our purposes, we equate truth with the voice of the science: the mean of all the data actually gathered by the research community. Though this notion of truth is not completely accurate to some underlying state of the world (the data-generating state in the model), it is what one would learn if one enjoyed unfettered access to the science. That is, this notion of truth is the truth from the perspective of what society could in some abstract idealized sense be getting from science. We then think of distance from the truth as how far off the policy decision is from what it would be if the policymaker knew the truth. We look at how far off the policymakers are (on average, over many rounds of simulating this communication process) from what they would be doing if they knew what the science actually says.

Results are presented for simulations with ten KBs, 100 scientists, and 1000 rounds. For easy visualization we set the extra influence multiplier to 5 and the extra matching chance  $m$  to 2000. We choose these large values so that we can easily see the differences between variations as we vary policy width (since policy width has such a large impact on outcomes). Importantly, similar qualitative results occur for less extreme choices of these variables. Policy width was varied from 0 to 2, where 0 represents a case where the policymaker never listens to any evidence and 2 represents a case where the policymaker is essentially guaranteed to listen to every signal. Each parameter combination was run 1000 times in order to get an estimate of the expected distance from truth.

---

KB’s signal is given is  $\frac{w_i}{\sum_{j=1}^n w_j}$  where  $n$  is the total number of signals listened to.

<sup>11</sup>We also ran simulations for an alternative version of the model where influence instead impacts the likelihood that KBs are listened to by policymaker. In this version, KBs must first be listened to (given a foot in the door of the policymaker) and then be heard (have their signal be within the policymakers’ expectations) in order to impact policy. Results for this version of the model were similar.

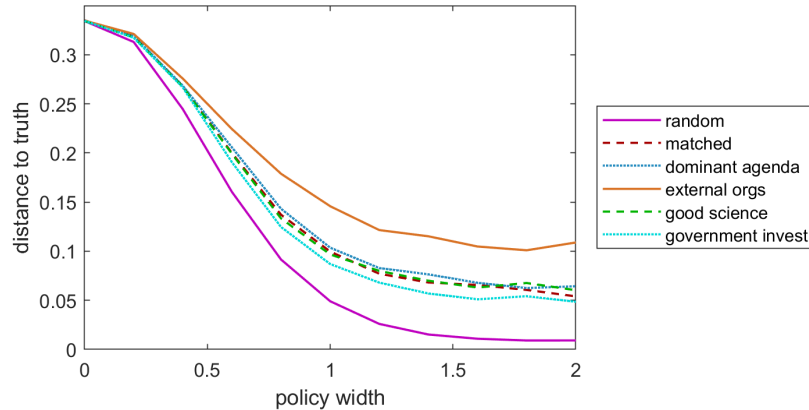


Figure 2: Distance from truth, over various policy widths for each of the six variations of the model

As Figure 2 shows, the distance from truth is affected by the width of the policymakers' expectations for all variations. This width represents how strong the policymakers' expectations are, or how strongly their expectations shape their view of the evidence they are presented, and it ultimately affects how many KB signals they listen to. When their expectations are narrow, they generally make decisions that are pretty far from what they would do with complete understanding of the scientific evidence. Larger policy widths lead to better decision making in this model.

Looking at *Random* (in pink), where there are no incentives or differences in influence, we can see that the policy width has a large impact on distance to truth. As the policy width increases, the policymaker gets very close, but not all the way, to perfect decision making. By definition, taking the average of all data gathered by the scientists would yield a distance of zero. The deviation from that we observe here is the result of KBs grouping the data (the policymaker averages KB signals which themselves average different quantities of data) and of the data lost because some scientists join no KB.

What happens when we add in incentives, where KBs tend to choose members that align with their aims? In *Matched* (in red), the distance from truth now increases. While similar when policymaker expectations are narrow, the impact increases further from the truth once the policymaker is listening to most signals.

So: KBs distort the evidence just by being part of the process and filtering the data. The *Random* variation shows that introducing this kind of filtration leads to some small distance from truth. The *Matched* variation shows that once that filtration is not completely random, this leads to an even greater distance to truth. So, the incentives of these agents matter. We think of *Matched* as a sort of minimal or baseline description of the impact of incentives.

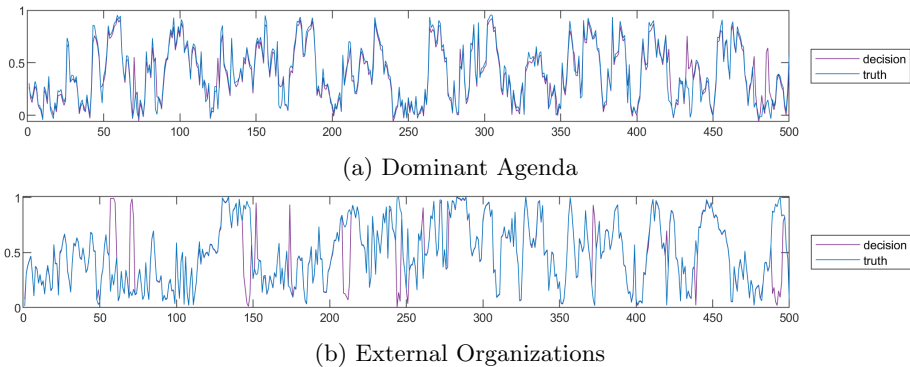


Figure 3: Values for truth and policymaker decision over time in an example simulation run for the (a) Dominant Agenda and (b) External Organizations variations of the model. Results from the first 500 rounds are shown in order to more clearly see the relations between the two lines.

Results for when certain types of KBs are more influential in that their signals are given more weight by the policymaker are also shown in Figure 2. Two interesting results: First, though it is unsurprising that the two variations where agenda-driven KBs are influential are the furthest from the truth, it was somewhat surprising that *External Organizations* (in orange) was worse than *Dominant Agenda* (in blue). That is, having influential people pushing both sides of the issue turned out to be worse than one side having all the influence. This is because, in our model, with only one side of the issue having extra influence, the policymaker can only be pushed into really bad decision making in one direction (e.g., they’re generally looser on regulations than they should be, but only occasionally drastically under-prepared to deal with an invasion threat) versus with both sides of the issue having extra influence, the policymaker can be pushed very far from the truth in both directions. Occasionally they drastically overreact and occasionally they drastically under-react and overall spend more time doing wildly the wrong thing.

Figure 3 shows example simulation runs from these variations of the model, where we can see this playing out. In Figure 3a, we can see how the policymaker’s decision fluctuates in comparison to the truth in *Dominant Agenda*. While the decision is generally close to the truth, large discrepancies in the policymakers’ decision are to one side of what they’d do if appropriately responding to the truth (e.g., they are too reactive to invasion threats). Compare this to figure 3b, which shows an example run for *External Organizations*. In this case, the policymaker’s decision is often too extreme, and the direction in which they are wrong (e.g., too over-reactive or too under-reactive to an invasion threat) alternates over the course of the simulation.

Second, *Government Investment* (in cyan) is better than *Good Science* (in green). This is a somewhat surprising result — why should an eye towards

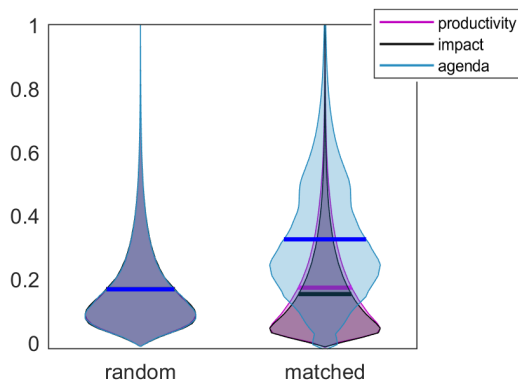


Figure 4: Distance to truth by KB type, for the first two model variations. In the Matched variation, both the mean distance from truth and the variance of the productivity-driven broker signals is higher than the impact-driven brokers.

telling the policymaker what they want to hear do better than picking based on the most productive scientists? You might think that the policymaker is closer to the truth because they are listening to more signals when impact-driven KBs are influential, but once the policy width is sufficiently wide, the policymaker is guaranteed to listen to everything, and government investment is still better. Instead, looking at the signals sent by each KB type can help us get a sense of why this might be the case.

Figure 4 shows the quality of signals sent by the different types of KBs, i.e., the distance from truth of the signals they sent, in the first two variations. The horizontal lines show the mean distance from the truth for each KB type, and the shaded areas show the variation. In the *Random* variation, there is no difference between KB types and their signals are all pretty close to the truth. This makes sense as there is really no difference between the KBs when everything is random and no incentives are interfering with the communication process.

Once KBs and researchers are matched according to incentives, we begin to see differences. Figure 4 shows results for *Matched*, but the other four conditions give the same results.<sup>12</sup> Agenda driven KBs (in blue) are sending consistently worse signals. This also makes sense — they have as members researchers whose methodology aligns with their agenda so are consistently sending signals that either over- or underestimate invasion threats. There is less difference between productivity and impact driven KBs. However, while there is substantial overlap in the kinds of signals sent, there is more variance in what the productivity-driven KBs send. Because they pick members without regard to methodology, as long as they are productive, there is greater variance in what their membership’s

<sup>12</sup>Because these results are only for signals sent, they are not affected by introducing influence of KBs into the model.

research says about the situation at hand.

So, we might say something like: we have to conditionalize our expectations on the brokering landscape, not just go with what we *a priori* think should be best. If there's a sufficient industry presence to skew the overall availability of methods in the rest of the population, government investment is the best. In that case, we might want more publicly funded research, with an eye toward impact.

## 5 Conclusion

KBs are major players in society's overall knowledge ecosystem. How we conceptualize them matters in the study of the growth of scientific knowledge in society. Yet, little is known about the epistemology of KBs: how these players' incentives impact knowledge flow from its origin in scientific research to its endpoint in use by society, especially in the decisions of policymakers. To develop an adequate epistemology of KBs, it is necessary to depart from simplistic narratives about KBs motivated by the linear and co-production models of the science-policy interface. Agent-based modeling is well suited to the task.

Of course, the model we have presented is still highly simplified. In reality, there may be multiple policymakers, with different incentives or agendas, and KBs may employ different strategies to try to be heard by some or all of them. Scientists may also have more than one KB membership, and how many memberships they have may depend on their productivity, methodology, or something else. There might also be multiple, competing incentives for a given KB or for a given scientist, which means they have to make trade-offs when they think about their membership. There might also be multiple questions to gather evidence about and co-production of knowledge in which KBs facilitate what research is produced by scientists. In this paper, we have assumed that all the science is in fact relevant to the policy, but this is not generally the case. KBs having influence over what science gets done could allow the different KB types to have a much greater influence on ultimate policy decisions.

We regard the initial results we have presented here as a launching off point. Considered as a base case, our model establishes how effects of KBs being part of a knowledge ecosystem are hidden when they are conceptualized only in terms of carefully delineated roles. In our model, all KBs have the same role. They all process and summarize scientific evidence in exactly the same way. Yet, the consequences for policymakers' decision-making is different in virtue of differences in dominant incentives and possible differences in influence. Moreover, our results indicate that facts about the overall composition of the brokering landscape non-trivially impact our strategizing for how to improve quality of knowledge available in policymaking.

There is still much to be done in developing an epistemology of KBs, but the model presented here can itself inform some next steps. Thus, we end by discussing two topics for future research, including suggestions for how to investigate them within the modeling framework inspired by this first model.

There is a trade-off for many researchers between what the globally-minded research community incentivizes and what policymakers demand. In our modeling framework, this might be captured by having scientists with multiple competing incentives, who are allowed to join multiple KBs. In general, productivity-driven KBs, incentivized by gaining well-regarded and productive members, are more attuned to credit incentives of the international scientific community, for instance valuing English publications on ‘global’ problems. Impact-driven KBs are more incentivized by the policymakers in that region, and so care more about locally relevant results, preferring those published in the local language. However, due to power differences, it ends up being that local or regional issues in the Global North are often taken to be what is of ‘global’ interest by prestigious journals [Bortolus, 2012]. This leads to more of an inherent trade-off for some researchers. Using our modeling framework, we can investigate how the different KB types would have different effects across global contexts, and how researchers in the Global South or where English is not the native language might be particularly affected.

This framework can also be used to better theorize about the impact of power dynamics that are wholly internal to descriptions of the processes at the science-policy interface — what has been here referred to in terms of researcher or KB ‘influence’. Namely: KBs might gain and lose influence based on how often they have been listened to or the quality of their membership. This might also impact their members, as being part of a more influential KB comes with benefits, like prestige by association, networking opportunities, and, often, funding [Bortolus et al., 2025]. So, the prestige of KBs and scientists feed into each other within the system, potentially amplifying the effects observed here and affecting the overall distribution of methods within the scientific community (similar to the effect of industry funding in Holman and Bruner [2017]).

## References

- Justyna Bandola-Gill, Megan Arthur, and Rhodri Ivor Leng. What is co-production? Conceptualising and understanding co-production of knowledge and policy across different theoretical perspectives. *Evidence & Policy*, 19(2): 275–298, 2023.
- Celeste Beesley, Darren Hawkins, and Nicholas Moffitt. What leads government officials to use impact evidence? *Journal of Public Policy*, 42(1):20–42, 2022. doi: 10.1017/S0143814X21000015. URL <https://doi.org/10.1017/S0143814X21000015>.
- Alejandro Bortolus. Running like Alice and losing good ideas: On the quasi-compulsive use of English by non-native English speaking scientists. *Ambio*, 41(7):769–772, 2012.
- Alejandro Bortolus, Chad L. Hewitt, Veli Mitova, Evangelina Schwindt, Temitope O. Sogbanmu, Emelda E. Chukwu, Remco Heesen, Ricardo Kaufer, Han-

- nah Rubin, Mike D. Schneider, Anne Schwenkenbecher, Helena Slanickova, Katie Woolaston, and Li-An Yu. Knowledge brokers at the science–policy interface: Insights from biosecurity and environmental management. *Ambio*, 54(12):2079–2091, 2025.
- Bernard C. K. Choi, Tikki Pang, Vivian Lin, Pekka Puska, Gregory Sherman, Michael Goddard, Michael J. Ackland, Peter Sainsbury, Sylvie Stachenko, Howard Morrison, and Clarence Clottey. Can scientists and policy makers work together? *Journal of Epidemiology & Community Health*, 59(8):632–637, 2005. ISSN 0143-005X. doi: 10.1136/jech.2004.031765. URL <https://doi.org/10.1136/jech.2004.031765>.
- Corey Dethier. How do you assert a graph? Towards an account of depictions in scientific testimony. *Noûs*, 59(3):678–699, 2025.
- David Edgerton. The linear model did not exist: Reflections on the history and historiography of science and research in industry in the twentieth century. *The science-industry nexus: History, policy, implications*, pages 31–57, 2004.
- Daniel C Friedman and Dunja Šešelja. Scientific disagreements, fast science and higher-order evidence. *Philosophy of Science*, 90(4):937–957, 2023.
- Remco Heesen. Communism and the incentive to share in science. *Philosophy of Science*, 84(4):698–716, 2017. ISSN 0031-8248. doi: 10.1086/693875. URL <https://doi.org/10.1086/693875>.
- Bennett Holman and Justin Bruner. Experimentation by industrial selection. *Philosophy of Science*, 84(5):1008–1019, 2017.
- Sheila Jasanoff. The idiom of co-production. In Sheila Jasanoff, editor, *States of Knowledge*, chapter 1, pages 1–12. Routledge, London, 2004.
- Joel Katzav. The epistemology of climate models and some of its implications for climate science and the philosophy of science. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 46:228–238, 2014.
- Sabina Leonelli. On the locality of data and claims about phenomena. *Philosophy of Science*, 76(5):737–749, 2009.
- Sabina Leonelli. What counts as scientific data? A relational framework. *Philosophy of Science*, 82(5):810–821, 2015.
- Sabina Leonelli and Niccolò Tempini. *Data journeys in the sciences*. Springer Nature, 2020.
- Eleanor MacKillop, Sarah Quarmby, and James Downe. Does knowledge brokering facilitate evidence-based policy? A review of existing knowledge and an agenda for future research. *Policy & Politics*, 48(2):335–353, 2020. ISSN 0305-5736, 1470-8442. doi: 10.1332/030557319X15740848311069. URL <https://doi.org/10.1332/030557319X15740848311069>.

- Robert K. Merton. *The Sociology of Science: Theoretical and Empirical Investigations*. The University of Chicago Press, Chicago, 1973. ISBN 0226520919.
- Cailin O'Connor and Justin Bruner. Dynamics and diversity in epistemic communities. *Erkenntnis*, 84(1):101–119, 2019. ISSN 1572-8420. doi: 10.1007/s10670-017-9950-y. URL <https://doi.org/10.1007/s10670-017-9950-y>.
- Cailin O'Connor and James Owen Weatherall. Scientific polarization. *European Journal for Philosophy of Science*, 8:855–875, 2018.
- Manuela Fernández Pinto. Science and industry funding. In *Global epistemologies and philosophies of science*, pages 164–173. Routledge, 2021.
- Karl R. Popper. *The Logic of Scientific Discovery*. Routledge, 2002.
- Hannah Rubin and Mike D Schneider. Priority and privilege in scientific discovery. *Studies in History and Philosophy of Science Part A*, 89:202–211, 2021.
- John T Sanders and Wade L Robison. Research funding and the value-dependence of science. *Business & Professional Ethics Journal*, pages 33–50, 1992.
- Jamie Shaw. Revisiting the basic/applied science distinction: The significance of urgent science for science funding policy. *Journal for general philosophy of science*, 53(4):477–499, 2022.
- Michael Strevens. The role of the priority rule in science. *The Journal of Philosophy*, 100(2):55–79, 2003. ISSN 0022-362X. doi: 10.5840/jphil2003100224. URL <https://doi.org/10.5840/jphil2003100224>.
- C. J. Uneke, A. E. Ezeoha, C. D. Ndukwe, P. G. Oyibo, F. Onwe, E. B. Igbinedion, and P. N. Chukwu. Individual and organisational capacity for evidence use in policy making in Nigeria: an exploratory study of the perceptions of Nigeria health policy makers. *Evidence & Policy*, 7(3):251–276, 2011. doi: 10.1332/174426411X591744. URL <https://doi.org/10.1332/174426411X591744>.
- Kevin J. S. Zollman. The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1):5–33, 2018. ISSN 0022-362X. doi: 10.5840/jphil201811511. URL <https://doi.org/10.5840/jphil201811511>.