THE ECONOMICS OF BELIEFS UNDER FUNDAMENTAL UNCERTAINTY

Johannes Binswanger* Manuel Oechslin†

December 11, 2020

Abstract

Introducing the concept of “belief entrepreneur”, this paper offers a novel theory on endogenous belief formation under fundamental uncertainty in the sense of Knight (1921). We consider a generic setup in which individuals must choose between a tested approach (supplied by a “defender”) and a competing innovative approach (supplied by an “innovator”). While the innovation is promising, its true merits are uncertain (e.g., financial engineering in the 1990s). Facing an ambiguous choice, individuals are susceptible to narratives. The innovator and defender thus act as competing belief entrepreneurs who engage in a narrative contest whose outcome shapes individual prior beliefs. We clarify the conditions under which the contest outcome predominantly reflects information on the merits of the innovation—and when other factors, such as the entrepreneurs’ payoffs, dominate. Our analysis may be helpful to regulators that have to grapple with innovations whose fundamentals they do not know any better than the public.

JEL classification: D83, D91

Keywords: Uncertainty, belief entrepreneurs, endogenous beliefs, heuristics, narratives

*University of St. Gallen, Department of Economics, Switzerland. Email: johannes.binswanger@unisg.ch.

†University of Lucerne, Department of Economics, Switzerland. Email: manuel.oechslin@unilu.ch.

1 Introduction

Societies committed to liberal democracy and capitalism are often viewed as unrivaled “discovery engines”. Popper (1945) describes liberal democracy as an incubator for innovative policy ideas while, according to Schumpeter (1942), capitalism is a force that

(...) incessantly revolutionizes the economic structure from within, incessantly destroying the old one, incessantly creating a new one. (p. 83).

Because of these innovation forces, societies committed to liberal democracy and capitalism often encounter novel situations in which data and experiences from the past are of little help in forming beliefs about the consequences of different courses of action. Yet, in practice, decisions have to be taken—and deliberate decisions require some belief. How do people form beliefs in novel situations? It is unlikely that beliefs simply fall “out of the sky”, in a completely exogenous fashion. We argue that beliefs emerge endogenously, influenced by familiar economic and psychological forces. In this paper, we show how these forces interact to shape subjective beliefs in novel situations that arise as a result of innovations. To the best of our knowledge, we are the first to address belief formation from this angle.

Many discoveries—technological, policy-related, or institutional—are associated with a move into “uncharted territory”: when an innovation emerges, there is often a scarcity of relevant data, making it hard to form an “objective” belief about the innovation’s up- and downsides. Innovations thus often come with what we call *fundamental uncertainty*, i.e. with uncertainty (in the sense of Knight 1921) related to their economic fundamentals. As an example, consider the share of jobs that over the next two decades will be lost to automation and artificial intelligence. Existing suggestions for the distribution of possible outcomes vary widely (Segal 2018)—and it is all but impossible to discriminate between them based on existing data. In a similar vein, in the 1990s and early 2000s, there was fundamental uncertainty about the consequences of securitization of US residential mortgage loans. Obviously, there is a long list of fitting examples. At the macro level, it includes the consequences of large-scale quantitative easing by central banks in the aftermaths of the Great Recession and the Corona Crisis. At the micro level, ventures into uncharted territory include novel products, production methods, and business models in general—as well as their possible regulation (e.g., “big tech”).¹ Section 3 offers a detailed discussion of the securitization example.

A cornerstone of our analysis is that, in the context of any innovation, there are actors

¹The examples mentioned here are all among the “many phenomena of interest, [for which] probabilities cannot be claimed to be scientifically measurable quantities.” (Gilboa et al. 2014, p. 1409).

with a clear incentive to influence the formation of subjective beliefs. For instance, investment bankers may see their incomes rise if the demand for a novel product in the realm of financial engineering is high because market participants have adopted a generally positive belief about the innovation. Similarly, a generally positive belief about a proposal for a novel policy means a higher chance that the proposal will pass the responsible committee or legislature—and so advances the interests of the proposal’s sponsor. But often, when there are winners, there are also losers. And those who expect disadvantages from an innovation have an incentive to influence subjective beliefs in a negative way. Thus, the pattern that emerges is one of a competitive game between different *belief entrepreneurs* that try to instill the desired subjective belief into the relevant public. Section 3 exemplifies this pattern with the help of a particular innovation, the securitization of US residential mortgages.

But why do clearly self-interested actors obtain a chance to influence beliefs? Innovations involve fundamental uncertainty. It is often impossible to obtain an objective estimate $\hat{p}$ for a probability p that at time t innovation I produces outcome O —with the result that I cannot be valued objectively. Similarly, individuals are hardly in a position to immediately come up with a firm *subjective belief* about $\hat{p}$; they need to make up their minds first. So, when an innovation emerges, individuals a priori see themselves confronted with an ambiguous valuation and decision problem.² Yet, as we know from a large literature (Ellsberg 1961, Machina and Siniscalchi 2014), individuals dislike ambiguity and show a tendency to avoid ambiguous options. But do people generally sacrifice innovation just because of ambiguity? We think not. The human mind has an alternative approach: the literature on heuristics (Kahneman and Tversky 1973; Tversky and Kahneman 1974; Slovic et al. 2007; Kahneman 2011) shows that decision makers are prone to substituting difficult problems with seemingly analogous, yet simpler ones. We suggest that an ambiguous valuation and decision problem relating to an innovation is perceived as difficult—and therefore is replaced by what appears to be an analogous, but non-ambiguous problem. This mechanism of *heuristic substitution* offers a measure of influence to belief entrepreneurs: they are competing suppliers of narratives that carry analogies which aim to promote a desired belief.

The mechanism of heuristic substitution is consistent with recent research in neuroscience, in particular with predictive coding theory (Friston 2010; Clark 2013, 2016; Barrett 2017b) which understands the brain as a “prediction machinery” (in the sense of statistical learning) that uses past experiences as “training data”. Importantly, this data need not coincide with what

²Fundamental uncertainty or uncertainty in the sense of Knight (1921) means that event probabilities cannot be objectively established. Our use of “ambiguous” is based on Ellsberg (1961): a decision problem is ambiguous if the agent has little confidence in her estimates of probabilities (also see Machina and Siniscalchi 2014).

would count as relevant data under standard economic definitions of rationality. According to the theory, when encountering an unfamiliar instance (in our case, an innovation), the brain proactively employs analogies, associations, and mental simulations to generate predictions about the valuation of that instance based on its internal models (Bar 2007; Barsalou 2009). Narratives can influence valuations, and hence choices, by actively assisting the construction of analogies that are then used as “data” in the brain’s prediction process.

Existing theoretical work in economics offers little help as to how people form subjective (prior) beliefs in situations where a lack of experience and data means that a novel phenomenon cannot be assessed objectively (e.g., Brunnermeier et al. 2014).³ Yet, as we hope to demonstrate, it is precisely in those situations where economics can be of great help in understanding belief formation. The contribution of this paper is to offer a novel theory on how under fundamental uncertainty beliefs may emerge endogenously, shaped by such familiar economic concepts as supply, demand, and competition. Psychological concepts (ambiguity aversion, receptiveness for narratives) play a key role on the demand side. Using the existing microeconomic toolbox, we propose a framework that identifies key factors in the formation of beliefs under fundamental uncertainty and shows how these factors interact.

Our model considers a generic business or political domain in which there is an incumbent *defender* of the status quo. The defender is experienced in the use of a tested *approach* for producing an outcome. An approach can mean different things: a production method or business model, an economic policy, or an institution. We consider the domain in a particular moment: an *innovator* has emerged and now proposes an innovative approach for producing the outcome in question. Being novel, the innovative approach is subject to fundamental uncertainty: while it is clear that the innovation is superior in one state of the world and inferior in the other, the probabilities of the states are unknown. Yet objective information on the state need not be entirely absent: the defender and the innovator receive a signal whose informativeness, a key parameter of the analysis, can range from zero to perfect. The signal is private and cannot be credibly conveyed to third parties. A continued use of the tested approach, *irrespective* of the state of the world, implies an immediate gain for the defender (e.g., in the form of short-run profits in the case of a business domain), while a broad adoption of the innovative approach leads to an immediate gain for the innovator.

Besides the defender and the innovator, there is a *target audience*, a group of individuals who decide which one of the two approaches to use. The audience’s objective is to choose

³ “As economics does not offer much guidance on how individuals form their prior beliefs, economists tend to agree that prior beliefs probably depend on an individual’s background and experience.” (Brunnermeier et al. 2014, p. 1760). There will also be a role for individual backgrounds and experiences in our setup.

the superior approach, either individually or collectively. Audience members are averse to ambiguity (caused by fundamental uncertainty) and thus receptive for ambiguity-suppressing subjective beliefs. The latter are formed in response to the outcome of a competitive game between opposing narratives supplied by the innovator and the defender, a competition we capture by means of a contest game. The coveted prize in this *belief contest* is the broad adoption of the desired belief (positive or negative towards the innovation). The belief entrepreneurs spend resources—creativity, time, money—on their respective narratives. Besides these short-run expenses, acting as a belief entrepreneur may entail a “long-run” cost: promoting an approach that later is revealed as inferior causes reputational damage (while there is a long-run reputational benefit if the approach proves to be a success).⁴

According to our analysis, three factors influence the set of subjective beliefs the members of the target audience eventually adopt: the information contained in the signal to the belief entrepreneurs (“information”); the magnitudes of the gains and losses (both immediate and long-run) that may come with the participation in the belief contest (“payoffs”); and accidental forces, such as a narrative suddenly going viral (“luck”). Yet only the first factor is related to the actual state of the world, i.e., to the sole “fundamental” determining whether the innovative approach is superior or inferior. Our analysis clarifies how much influence each of these three factors commands under different circumstances. A key question concerns the extent to which the information contained in the signal to the belief entrepreneurs will eventually be reflected in the subjective beliefs adopted by members of the target audience. We show that the role of the signal depends on the potential gains and losses the belief entrepreneurs stand to incur: for a given level of informativeness, the signal’s influence is stronger if the weight of potential immediate gains relative to that of long-run losses is smaller. We further show that, as the signal’s informativeness falls, the reins are increasingly taken by the factors unrelated to the fundamental state, including accidental forces. Exploring the role of narratives, Shiller (2017) suggests that sometimes a narrative develops into a force of its own that, by influencing the decisions of many, even has the power to affect the aggregate economy. Our model characterizes circumstances that are particularly conducive to such a course of events.

In practice, as the securitization example in Section 3 suggests, innovations that come with fundamental uncertainty can have powerful implications at the macroeconomic level. Against this background, we see two applications for our model. First, it may be used as a diagnostic framework by policy makers that have to grapple with innovations whose fundamentals they

⁴To keep the exposition concise, the main part of the paper focuses on a variant of the model in which the members of the target audience decide individually on which approach to use. A sketch of a model variant with collective choice (committee decisions or decisions by the public at large) can be found in Appendix A.

do not know any better than the public at large. In particular, the model may help an impartial policy maker calibrate its response towards a fast-spreading innovation on the basis of observable variables and available knowledge. For instance, if the innovator stands to gain a large immediate profit from the broad adoption of the innovation and conclusive evidence is unlikely to emerge quickly, the positive subjective belief that propels the spread of the innovation cannot be considered informative about the fundamental (while it could be if the situation were reversed); so the policy maker has a particularly strong reason to prepare for the eventuality that the innovation will prove harmful later on.

Second, while our model is static and focused on partial equilibrium, it may serve as a building block that can be introduced into a suitable macroeconomic framework. In the model, there can be a severe disconnect between “the fundamentals” of the economy and what people believe them to be. Accordingly, introducing (a variant of) our model into a macro framework sets the stage for rich dynamics at the aggregate level.⁵ Consider, for instance, the following possibility. Thanks to a highly contagious narrative, many investors initially hold an exuberant belief about an innovation—and realize only after a lengthy time that, in fact, the innovation is a complete failure and investments must be written off. At that point, there could be a sharp fall in aggregate consumption, pushing the economy into a deep slump. So, as compared to more standard macro frameworks, a setup including fundamental uncertainty and belief contests along the lines of the model would show more pronounced fluctuations, including sporadic severe downturns that may be called “crisis episodes”.

The rest of this paper is organized as follows. The next section discusses the related literature. Section 3 provides a concrete example of the basic constellation that is captured by our model of belief formation. Section 4 presents the model. In Section 5, we derive the outcome of the belief formation process in a simplified setting in which the campaigning efforts chosen by the belief entrepreneurs are binary variables; the section also includes a discussion of a number of key insights. The analysis of the more general setup, along with a discussion of additional insights, is contained in Section 6. Finally, Section 7 concludes.

2 Related Literature

In our setting, the principal actors are belief entrepreneurs who compete in shaping beliefs about an innovation by crafting respective narratives. Mokyr (2013, 2017), analyzing the coevolution of culture and capitalism, develops the closely related concept of “cultural en-

⁵Clearly, consistent with the dynamic nature of macroeconomic frameworks, such a model variant would have to be more specific on the nature of the updating of beliefs by members of the target audience.

trepreneur”. Mokyr defines culture as a “system of beliefs, values and preferences that shape [...] institutions.” Cultural entrepreneurs are understood as “individuals who refuse to take the existing technology or market structure as given and try to change it.” (Mokyr 2013, p. 2). Their essential role is to provide a perspective on new social orders and on how nature can be analyzed and harnessed for the development of new technologies. While the concept of belief entrepreneur links our analysis with Mokyr’s work, there are several important differences. First, in our setup, belief entrepreneurs may be driven by material motives; rather than considering intellectuals, we focus on business people, politicians, and technocrats (i.e., on actors that more commonly populate economic models). Second, we investigate a contest between belief entrepreneurs with opposing interests; our hypothesis is that innovations create losers who have strong incentives to fight the broad adoption of a particular innovation. Third, the analysis here relies on formal modeling.

Bénabou et al. (2020) put forward the idea of a “narrative entrepreneur” in their analysis of narratives and moral behavior. They investigate the conditions under which individuals use narratives as exculpatory or responsabilizing arguments to either justify selfish behavior or encourage pro-social behavior that is individually costly. In an extension, Bénabou et al. (2020) allow agents to actively search for such narratives and thus become “narrative entrepreneurs”. In Bénabou et al. (2020), agents assume one of two possible unobserved types that assign either a low or high value to a moral action; both types share an equal concern for being perceived as morally responsible. Searching and sharing of narratives thus serves for strategic signaling about their types. In the setting of Bénabou et al. (2020), there is no competition between different narrative entrepreneurs. In our analysis, belief entrepreneurs aim to influence an audience’s belief about the merits of an innovation that is beneficial to the respective belief entrepreneur’s business (or otherwise). We do not distinguish between different unobserved types of entrepreneurs; rather, our focus is on how their competition affects prior beliefs about the innovation as an equilibrium outcome.

This paper rests on the idea that individuals deal with a difficult problem—the assessment of an innovation under fundamental uncertainty—by substituting it with a simpler one that features an analogy to the original problem. We suggest that narratives are carriers of such analogies. Shiller (2017, 2019) provides an extensive treatment of narratives and offers many examples of narratives that influence beliefs about (new) ideas by politicians, business people, and other types. Shiller argues that narratives have the potential to cause changes in economic behavior and even in aggregate economic dynamics. This paper analyzes one possible channel through which narratives may have a causal effect on economic outcomes. We further com-

plement Shiller’s analysis by providing a formal model of the incentives to invest in narrative campaigns. We also take into account the existence of competitors with an incentive to launch counter-narratives, giving rise to what we call belief contests.

Reasons for the prevalence and effectiveness of narratives have recently been considered from a neuroscience and evolutionary perspective (Armstrong 2020). Humans are a uniquely cooperative species (Fehr and Gächter 2002). An important question concerns how such cooperation is achieved. Key answers are altruistic punishment (Fehr and Gächter 2002; Fehr and Fischbacher 2003) and also that trust (towards members of an in-group) tends to “feel good” (Fehr et al. 2005). However, besides altruistic punishment and trust, narratives, too, are thought to have the potential to support (within-group) cooperation. They may do so by fostering “shared intentionality” (Tomasello and Rakoczy 2003; Armstrong 2020). This view is supported by the fact that narratives have a unique potential to turn viral—a major theme in Shiller (2019)—and thus can be highly effective in promoting a shared goal at the group level. This, in turn, may explain why humans are predisposed to “listen” to narratives. In our analysis, belief entrepreneurs make use of this predisposition.

Ambiguity-suppressing narratives can also be seen as belonging to the realm of motivated reasoning (Bénabou and Tirole 2002, 2016; Bénabou 2013). Agents may exhibit demand for an ambiguity-reduction device to avoid being dragged by the confusion and emotional burden of uncertainty-related ambiguity. Narratives may just serve that purpose. Their adoption could be understood as the result of a mental cost-benefit mechanism akin to “mindful economics”. The benefit of adopting a “plain” subjective prior belief is the suppression of ambiguity and the associated gain in motivation and readiness to make decisions; it may also be associated with a certain anticipatory utility. The cost is that ignoring ambiguity may lead to overconfidence of sorts. While it would be interesting to model this mechanism explicitly, we do not follow this route here but focus on our main point: the competition among belief entrepreneurs and how it shapes equilibrium prior beliefs about an innovation.

Eliaz and Spiegler (2020) offer an intriguing formalization of narratives as stories that arrange correlated variables into a specific causality chain. A narrative then generates a subjective belief about how policy actions map into consequences. While the correlation structure is consistent with available data, the particular causality chain may be spurious and lead to distorted beliefs. A representative agent adopts the narrative that maximize their anticipatory utility by creating hope. The authors consider stationary equilibria. While status quo policies—for which consequences can be unambiguously identified from the (sufficiently) recent past—cannot be associated with distorted hopeful beliefs, this possibility arises for counter-

factual policies. Narratives promoting a counterfactual policy may thus compete with the status quo and appear more attractive as they may be more “hopeful”. There are several important differences to our approach. First, in Eliaz and Spiegler (2020), there is no explicit modeling of narrators with personal or economic interests; such actors—in the form of belief entrepreneurs—are key in our model. Second, Eliaz and Spiegler (2020) consider a stationary setting where data from the past are relevant (and sufficient) for assessing future prospects; there is also a correct way for doing such an assessment (which, however, is either too difficult or otherwise unattractive for voters). We consider a setting of fundamental uncertainty with the defining property that data from the past may not be relevant for assessing future prospects. So narratives have many more degrees of freedom. Moreover, in our case, narratives need not be restricted to data-related entities but can be abstract and (partly) immeasurable concepts such as the “invisible hand”. Finally, we do not model narratives directly but only via the effort that belief entrepreneurs invest in crafting them.

We finally refer to some of our previous work. In Binswanger and Oechslin (2015), we consider a setting with fundamental uncertainty in which different societal groups hold dissenting subjective beliefs about the success probability of an economic reform. We show that fundamental uncertainty can lead to political gridlock and stagnation. But in this previous analysis, the subjective beliefs are taken as exogenously given.

3 Beneficial Innovation or “Hydrogen Bomb”?

This section discusses a concrete example of the basic constellation captured by our model. We consider securitization, a past innovation in the financial industry. The term securitization refers to a business model that in the late 1980s started to emerge in the domain of US commercial banking. Securitization promised a novel approach to dealing with the credit risks inextricably linked to the lending business—an approach that *potentially* would lead to a more efficient outcome in terms of the economy-wide risk allocation and so improve financial stability. For a bank, the tested approach was to manage a diversified portfolio of loans that it retained on its balance sheet. The innovative approach, securitization, consisted in distributing those loans, including large parts of the associated credit risk, to third-party investors. Specifically, a bank would pool a certain number of loans and then re-finance them through issuing securities—collateralized debt obligations—with different levels of “riskiness” (from super-senior to junior); the bank would then sell those securities to third-party investors with different levels of risk capacity. The latter, in turn, could redistribute the risks once again using an accompanying

innovation called credit derivatives (such as credit default swaps).

While securitization and credit derivatives were successfully applied in some segments of the credit market, it was fundamentally uncertain whether these innovations would improve the risk allocation in other segments, too. For instance, when it came to the securitization of residential mortgage loans, a scarcity of data bedeviled the estimation of the probabilities with which the resulting securities would default. A particular problem concerned the correlations of defaults among pooled mortgages. Home prices had been growing more or less continuously since World War II. So there was uncertainty regarding the extent to which the historical data would be meaningful in the event of a large common shock to the housing market (Duney et al. 2013). Industry insiders spotted signs of the severity of the problem early on. Evidence of this is that J.P. Morgan, a pioneer of securitization, abandoned the business of packaging mortgages as early as in the late 1990s (Flood 2009).⁶ But many others remained unperturbed and the market—driven by a large demand for “safe” assets—virtually exploded in the following years.

Referring to securitization in general, the role of the incumbent defender was assumed by the “community” of senior commercial bankers with a large stake (e.g., in terms of compensation and career perspectives) in the continuing dominance of traditional balance sheet lending. The part of the innovator, on the other hand, was played by the “community” of members of banks’ securitization and credit derivatives teams. Vying for the very same compensations and career perspectives, the latter wanted to scale up the innovative approach, i.e., turn it from an obscure method used in some segments of the credit market into the standard way of dealing with credit risk. The target audience consisted of people inside and outside commercial banking, including bank CEOs and board members (who would determine the scope for securitization within their institutions) as well as regulators, rating agencies, and professional investors (whose decisions would critically affect the profitability of the innovative approach).

Both communities were engaged in a contest whose prize was the wide adoption of the desired ambiguity-suppressing subjective belief about securitization by the target audience. The community of innovators supplied supporting narratives, while the community of defenders pushed opposing ones. For concreteness, consider narratives (primarily) aimed at the regulators. Supporting narratives were designed to nurture optimistic beliefs by promoting a particular substitution: when evaluating securitization and credit derivatives, regulators should substitute the complex and imperfect real-world financial system with something close to the well-understood and orderly perfect-markets benchmark. In that benchmark, market

⁶Later on, such data uncertainties started to attract broader attention. For instance, in 2005, the Bank for International Settlements’ Committee on the Global Financial System (BIS 2005, p. 24) warned: “As a result [of uncertain default correlations], model-based risk assessments can be a long way from ‘true’ values (...).”

discipline—with its self-regulating forces—would limit the demand for (and sale of) securities that are ill-priced due to misjudged or unquantifiable risks. And in doing so, the market would minimize any threat to financial stability. Tett (2009) quotes an influential member of the community of innovators (Mark Brickell, a J.P. Morgan banker and key contributor to a 1993 report on credit derivatives by the Group of Thirty) as remarking:⁷

Markets can correct excesses far better than any government. Market discipline is the best form of discipline there is. (p. 36).

From today’s perspective, the boldness of this statement may surprise. Yet it comes from the early 1990s, a time when governments were often the cause of economic illnesses, while “the market” was considered a miracle cure. Regulators should transfer this logic to the new field of financial engineering. And they did. Brickell’s opinion was echoed by Alan Greenspan, then the chairman of the Federal Reserve. In a speech on May 25, 1994, the Fed chairman (according to a [transcript](#) dated July 1) testified to selected members of the House of Representatives that:

Legislation directed at [credit] derivatives (...), absent broader reform, could actually increase risks in the US financial system by creating a regulatory regime that is itself ineffective and that diminishes the effectiveness of market discipline.

Opposing narratives, on the other hand, were designed to nurture pessimistic beliefs by promoting a different kind of substitution. According to those narratives, the real-world financial system would have little in common with the perfect-markets benchmark and thus would not curb the sale of ill-priced securities and credit derivatives. Over time, ill-priced securities and derivatives would do to the financial system what weapons of mass destruction do to the enemy in a war. According to Tett (2009), Felix Rohatyn—a “legendary Wall Street figure” and advocate of the tested approach to dealing with credit risk—called

(...) [credit] derivatives financial hydrogen bombs, built on personal computers by twenty-six-year olds with MBAs. (p. 36).

In a 2003 letter to shareholders (dated February 21), Warren Buffett—Chairman of Berkshire Hathaway, a company that was among other things a large provider of more traditional financial services—took a similar position by saying that credit derivatives were

(...) financial weapons of mass destruction, carrying dangers that, while now latent, are potentially lethal. (p. 15).

⁷The G30 report was a compendium written by banking industry representatives. It contained, among other things, advice and “best practice” norms and was an attempt to forestall regulation by the government.

In retrospect, it is fair to say that the belief contest was won by the innovators. Apparently, their narratives struck a chord with regulators and other influential members of the target audience. Regulation remained light-touch and the securitization of, most notably, residential mortgage loans became a thriving business. But when in 2007 US house prices started to slide across the board, this mortgage-related business was put to a serious test—and it failed. The early doubts about the default correlations were revealed to be well founded. Instead of improving the risk allocation, the novel approach had mostly masked the credit risks. As a result, the involved banks suffered a heavy reputational loss and have paid hundreds of billions of dollars in fines since the end of the financial crisis (e.g., BCG 2017).

4 The Model

4.1 Basic Setup

There are two possible approaches $a \in \{0, 1\}$ for producing an outcome: a tested approach ($a = 0$) and an innovative approach ($a = 1$). The tested approach leads to a known outcome $\rho > 0$. The innovative approach produces an unknown outcome θ . There are two states of the world, $F \in \{H, L\}$. In state H , the outcome of the innovative approach is $\bar{\theta} > \rho$, while in state L the outcome is 0. In the beginning, state F is not observed and subject to fundamental uncertainty in the sense of Knight (1921). In particular, there is a lack of data that would suggest any specific objectively-valid belief about $\Pr[F = H]$. However, if the demand for the innovative approach turns out to be positive, state F will be revealed eventually. In what follows, we will often refer to F as the innovation’s “fundamental”.

Decisions on the use of the two approaches are taken by the members of a “target audience” (which is of unit size). Yet not only is there no basis for an objective belief about $\Pr[F = H]$; in addition, audience members are not in a position to immediately come up with a subjective belief. Rather, they find themselves confronted with an ambiguous valuation and decision problem. Subsection 4.2 addresses how audience members cope with ambiguity. We consider two alternative types of decision making, *individual* choice and *collective* choice. In the main part of the paper, the focus is on individual choice: each audience member decides for herself which one of the two approaches to use. Appendix A assumes collective choice.

The two approaches are offered by competing suppliers, the innovator (I) and the defender (D). Obviously, the former provides the innovative approach, while the defender is the established supplier of the tested approach—and potentially faces displacement by the innovator. As we discuss in Subsection 4.3, each has an incentive to influence the process of belief formation

by the audience members. To this end, they devise and disseminate appropriately designed narratives. Therefore, we refer to them collectively as the belief entrepreneurs.

4.2 Belief Formation and Belief Contest

Members of the target audience need to choose between a tested approach with a known outcome and an innovative approach with an ambiguous outcome. So, with respect to the innovation, members of the audience a priori see themselves confronted with an ambiguous valuation and decision problem—or with an instance of “experiential blindness” (Barrett 2017a). What is a realistic behavioral assumption for these circumstances? We find it natural to draw on the psychological literature on heuristics (Kahneman and Tversky 1973; Tversky and Kahneman 1974; Slovic et al. 2007; Kahneman 2011). The robust pattern found in this literature is that, when decision makers face a difficult problem, they heuristically substitute it with a seemingly analogous, yet simpler, one—and then adopt the solution to the latter. Our premise is that an ambiguous valuation and decision problem relating to an innovation is perceived as difficult; accordingly, it will be replaced by what appears to be an analogous, but non-ambiguous, valuation and decision problem. This type of heuristic substitution is supported by recent research in neuroscience (see Bar 2007, Barsalou 2009). Narratives are instrumental in the substitution process as they are the carriers of analogies. Returning to the securitization example of Section 3, an analogy put forth by the innovators was a theoretical benchmark: the audience should evaluate securitization through the lens of the well-understood and orderly perfect-markets benchmark rather than against the backdrop of imperfect and complex real-world financial markets; if this substitution were embraced, securitization would be seen as strengthening Adam Smith’s invisible hand.

Shiller (2019) offers many more examples of how narratives can “help” in the process of substitution. Moreover, he points out that narratives often include an emotional component that makes them contagious. However, potency and contagiousness of a particular narrative are subject to substantial randomness and cannot be perfectly controlled. There are many examples of how a seemingly minor detail can decide whether a narrative “goes viral” or not. As a result, a competing belief entrepreneur can hardly ensure that her own narrative is superior to those pushed by her opponents—even if the latter are vastly outdone in terms of resources that go into the design and dissemination of a narrative.⁸

Why do people use substitution and listen to narratives? As is well known, ambiguity

⁸Comparing narratives to movies, Shiller (2019, p. 41) quotes a former president of the Motion Picture Association of America: “With all the experience, with all the creative instincts of the wisest people in our business, *no one, absolutely no one* can tell you what a movie is going to do in the marketplace...”.

tends to have a negative effect on valuations, inducing an a priori tendency to avoid ambiguous options (Ellsberg 1961, Machina and Siniscalchi 2014). However, from both an evolutionary and a historical perspective, such behavior—if it were to be universal—would be extremely costly in terms of foregone opportunities. Almost all important innovations come with a high degree of ambiguity. If ambiguous options were generally avoided, there would be little progress, economic and other. Ambiguity suppression by way of substitution is a key mechanism to allow for progress by weakening the tendency to avoid ambiguous options.

Formally, we capture narratives in an indirect way by relating their potency to the effort (in terms of time and/or money) that goes into the respective campaign. For instance, it takes time to devise a strong narrative and then to try to convince opinion leaders, journalists, or celebrities. The innovator’s and the defender’s efforts in this regard are denoted by e_I and e_D , respectively. Moreover, we call ζe_I the innovator’s *effective* effort, where $\zeta \geq 0$ is a random variable with density function $\psi(\zeta)$. The introduction of ζ , whose distribution is known to the belief entrepreneurs, captures that the potency of a narrative depends on accidental forces that are beyond the innovator’s control. Obviously, the same holds for the defender’s narratives relating to the innovation. But introducing a random variable that multiplies e_D would serve no purpose: it would always only be the ratio of the two random variables that matters. So we assume that only one effort level is affected by a random force.

Following the choice of e_I and e_D , all audience members adopt a subjective belief about the probability that the fundamental of the innovation is sound ($F = H$). We denote by $\Pr_i[F = H]$ the subjective belief adopted by member i and assume

$$\Pr_i[F = H] = \alpha_i \pi, \quad (1)$$

where $\alpha_i \in [0, \bar{\alpha}]$ is an individual parameter that is distributed according to a distribution function $H(\alpha)$. Moreover, suppose that π is determined by a standard contest success function:

$$\pi = \begin{cases} 0 & \text{if } e_I = e_D = 0 \\ \frac{1}{\bar{\alpha}} \frac{\zeta e_I^\eta}{\zeta e_I^\eta + e_D^\eta} & \text{otherwise} \end{cases}, \quad (2)$$

where $0 < \eta < 1$. A contest game is a widely used modeling device to capture situations in which two or more parties compete for the same prize.⁹ In our case, the prize an entrepreneur

⁹Standard applications include competitions for access to a natural resource and competitions for large vote or market shares (see Konrad 2009 and Corchón and Marini 2018). Falkinger (2007) develops an alternative concept of a “contest for attention” that is inspired by empirical findings on the psychology of attention. His framework is an interesting avenue for further research on “narrative competition”.

can win is friendly beliefs about their own approach. Since audience members select just one of the two competing approaches, a generally friendly belief about the innovation (π close to 1) means a win for the innovator and a loss for the defender (and the other way round if π is close to 0). Equation (2) implies that π takes a large value if the innovator's effective effort is high *relative* to the defender's. The parameter η governs the “decisiveness” of the belief contest, i.e., how fast differences in effort levels translate into outcome changes.¹⁰ Dividing by $\bar{\alpha}$ ensures that all subjective beliefs fall in the range $[0, 1]$.

Equation (1) captures the idea that the belief contest provides a common anchor for individual subjective beliefs about $\Pr[F = H]$. If π takes a high value, the perception of the innovation is generally positive; if π takes a low value, it is generally negative. In what follows, we call π a belief anchor. Actual subjective beliefs, $\alpha_i\pi$, deviate from the anchor in both directions. If $\alpha_i > 1$, audience member i holds a more positive belief about the innovation (compared to the anchor); if $\alpha_i < 1$, the belief is more negative. The α_i s reflect individual personality traits. For instance, a high α_i describes an individual that is “open to new experiences” or that shows little “risk aversion” (in a broad sense of the term). These characteristics can be innate or acquired through upbringing and personal experiences (see, e.g., McAdams and Pals 2006 for an illuminating discussion on personality characteristics).

Our assumptions so far imply that audience member i , when considering the two available approaches, (subjectively) expects the following outcome:

$$E_i\{U(a_i)\} = \begin{cases} \rho & \text{if } a_i = 0 \\ \alpha_i\pi \bar{\theta} & \text{if } a_i = 1 \end{cases}. \quad (3)$$

The second line in equation (3) reflects that $\theta = 0$ if $F = L$. In what follows, we refer to equation (3) as audience member i 's objective function.

4.3 Belief Entrepreneurs' objective functions

By choosing e_I and e_D , the belief entrepreneurs influence the distribution of beliefs in the target audience. In doing so, they affect the levels of *demand* for the innovative approach, D_I , and the tested approach, D_D . As the size of the audience is normalized to 1, we have $D_D = 1 - D_I$, where D_I is the measure of members with $a_i = 1$. Clearly, $D_I(\pi)$ is an increasing function of π , while $D_D(\pi)$ is a decreasing function of π . The total payoff incurred by a belief

¹⁰This is seen most easily if ζ is set to 1 and we start with $e_I = e_D > 0$. In this case, $\pi = 1/2$; η determines to what degree a change in one of the effort levels makes π move away from 1/2.

entrepreneur consists of several elements. First consider the innovator. We use $v_I \geq 0$ to denote the net value of supplying one unit of the innovative approach that materializes *before* state F is revealed. With a slight abuse of language, v_I is called the innovator's *before*-payoff. The parameter w_I^H denotes the additional value that materializes *after* F has been revealed to be high, while $w_I^L \geq 0$ is a cost that is incurred *after* F has been revealed to be low. We refer to these parameters as *after*-payoffs.¹¹ The innovator holds a subjective belief about $\Pr[F = H]$, called ϕ_I . Given this, we state the innovator's objective function as

$$\begin{aligned} V_I &= E\{D_I(\pi)\} v_I + \delta_I E\{D_I(\pi)\} [\phi_I w_I^H - (1 - \phi_I) w_I^L] - c_I e_I \\ &\equiv E\{D_I(\pi)\} \tilde{v}_I - c_I e_I, \end{aligned} \quad (4)$$

where $\tilde{v}_I \equiv v_I + \delta_I [\phi_I w_I^H - (1 - \phi_I) w_I^L]$. In equation (4), $E\{\cdot\}$ denotes the expectations operator, reflecting that π is influenced by ζ , a random variable whose distribution the belief entrepreneurs know. The term containing the after-payoffs is weighted by a discount factor, δ_I , as uncovering the fundamental takes time (and there may be turnover in position of decision maker). We assume that the after-payoffs are proportional to the demand for the innovative approach: having campaigned for an innovation that works (does not work) leads to a larger reputational gain (damage) if the innovation has been used more widely. Finally, campaigning costs are captured by $c_I e_I$, where $c_I > 0$ is the cost per unit of effort.

Now consider the defender, whose objective function precisely mirrors that of the innovator:

$$\begin{aligned} V_D &= E\{D_D(\pi)\} v_D + \delta_D E\{D_D(\pi)\} [(1 - \phi_D) w_D^L - \phi_D w_D^H] - c_D e_D \\ &\equiv E\{D_D(\pi)\} \tilde{v}_D - c_D e_D, \end{aligned} \quad (5)$$

with $\tilde{v}_D \equiv v_D + \delta_D [(1 - \phi_D) w_D^L - \phi_D w_D^H]$. Note that for the defender the state being revealed as low means a gain (since the innovation failed). Therefore, in equation (5), $w_D^H \geq 0$ enters with a negative sign, while $w_D^L \geq 0$ enters with a positive one.

Before the game starts, the belief entrepreneurs hold their own prior beliefs about $\Pr[F = H]$. We denote those priors, which are purely subjective, by ϕ_k^0 , $k = D, I$.¹² At the start of the game, the belief entrepreneurs receive an identical signal $S \in \{H, L\}$ about the realization of F . The quality of the signal is given by $\sigma \equiv \Pr[F = f | S = f]$, $f = H, L$, where σ ranges from a minimum of $1/2$ (no information about fundamental) to a maximum of 1 (no funda-

¹¹We do not use the terms “short-run” or “long-run” payoffs. As it potentially takes a long time before the fundamental is revealed, the term “short-run” could actually be misleading.

¹²Arguably, without an optimistic prior belief, a potential innovator would hardly put effort into the development of a novel approach. So, in practice, ϕ_I^0 is likely to take a high value. See Herz et al. (2014) for experimental evidence on the relationship between entrepreneurial optimism and innovation.

mental uncertainty). The signal cannot be credibly communicated to the target audience since its interpretation requires expertise that the audience lacks. Assuming an identical signal is not crucial: we could assume the belief entrepreneurs to receive individual signals that are merely correlated (and not necessarily identical) without changing the behavioral implications of the model. However, working with an identical signal allows us to parameterize fundamental uncertainty in a parsimonious way (with $\sigma = 1/2$ and $\sigma = 1$ delimiting the entire spectrum). Having observed S , the two entrepreneurs update their subjective priors using Bayes' rule. The resulting posterior (subjective) beliefs are given by ϕ_I and ϕ_D (as they appear in equations 4 and 5). In several parts of the analysis, we will directly start with the posteriors, while in other parts we will investigate the role of the signal's quality.

4.4 Timeline

The timing of the game is as follows:

1. Nature determines $F \in \{L, H\}$, the fundamental of the innovation; the two belief entrepreneurs draw their prior beliefs, ϕ_I^0 and ϕ_D^0 .
2. Nature sends a signal $S \in \{L, H\}$ about F to the belief entrepreneurs; observing S , the belief entrepreneurs form their posterior beliefs, ϕ_I and ϕ_D .
3. The belief entrepreneurs simultaneously choose the levels of their campaign efforts, e_I and e_D ; nature draws ζ ; the belief anchor, π , forms.
4. Each member i of the target audience adopts a subjective belief about $\Pr[F = H]$ and then decides on the approach, $a_i \in \{0, 1\}$; the before-payoffs materialize.
5. If the innovative approach finds users, nature reveals the fundamental of the innovation; the after-payoffs materialize.

The essential stages are 3 and 4 (there are no strategic elements in 1, 2, and 5). Moreover, the analysis of stage 4 is straightforward and can be integrated into that of 3 by means of a simple application of backward induction. De facto, solving the model means solving a static contest game with two contestants, the belief entrepreneurs.

5 Belief Formation in a Simplified Setup

The purpose of this section is to provide a simple and transparent illustration of the principal forces and factors shaping subjective beliefs under fundamental uncertainty. All of the key

insights developed here will also hold in the general setup that is studied in Section 6. The general setup, however, will provide some additional insights.

This section makes two simplifying assumptions. First, the belief entrepreneurs must pick one of just two levels of campaigning effort: $e_k \in \{0, 1\}$, $k = I, D$. Second, in contest success function (2), the random factor takes one of only two possible values: $\zeta \in \{0, \bar{\zeta}\}$, where $\bar{\zeta} > 0$; the probability of $\zeta = \bar{\zeta}$ is strictly positive and denoted by ψ .

5.1 Analysis

Given these simplifying assumptions, the outcome of the belief contest is described as follows:

$$\pi(e_I, e_D, \zeta) = \begin{cases} 0 & \text{if } e_I = 0; \text{ or if } \zeta = 0 \\ \frac{1}{\alpha} & \text{if } e_I = 1, e_D = 0, \zeta = \bar{\zeta} \\ \frac{1}{\alpha} \frac{\bar{\zeta}}{1+\bar{\zeta}} & \text{if } e_I = 1, e_D = 1, \zeta = \bar{\zeta} \end{cases} \quad (6)$$

The assumptions regarding ρ , $\bar{\theta}$, and $H(\alpha)$ are implicit: we assume a constellation that leads to the following “demand function” for the innovative approach:¹³

$$D_I(\pi) = \begin{cases} 0 & \text{if } \pi = 0 \\ \lambda & \text{if } \pi = \frac{1}{\alpha} \frac{\bar{\zeta}}{1+\bar{\zeta}}, \\ 1 & \text{if } \pi = \frac{1}{\alpha} \end{cases} \quad (7)$$

where $0 < \lambda < 1$. So, if the belief anchor takes one of the polar values, the winner serves the entire audience; if π takes the interior value, the audience is split.

Together, equations (4), (5), (6) and (7) imply the payoffs shown in Table 1. To understand the entries, note that the innovator’s choice of campaigning effort, e_I , affects the belief anchor only if $\zeta = \bar{\zeta}$ (which happens with probability ψ). As a result, even if the innovator campaigns, the defender still serves the entire audience if $\zeta = 0$. Moreover, observe that, if the outcome is such that the innovative approach is not used at all, F is never revealed and reputational effects are absent. Thus, in those cases, the after-payoffs are identical to zero, implying $\tilde{v}_k = v_k$. This is why some entries in Table 1 show a v_D instead of a $\tilde{v}_D$.

We are interested in what type of Nash equilibrium arises under which parameter constellation. For this discussion, we denote strategy profiles as (e_I, e_D) . For instance, $(1, 0)$ means

¹³The parameters ρ and $\bar{\theta}$ appear in objective function (3), on which individual decisions on the use of the approach are based. The distribution function $H(\alpha)$ determines how individual decisions translate into total demand. “Demand function” (7) requires that there be no mass point at $\alpha = 0$.

		Defender	
		$e_D = 0$	$e_D = 1$
Innovator	$e_I = 0$	0 v_D	0 $v_D - c_D$
	$e_I = 1$	$\psi \tilde{v}_I - c_I$ $(1 - \psi)v_D$	$\psi \lambda \tilde{v}_I - c_I$ $(1 - \psi)v_D + \psi(1 - \lambda)\tilde{v}_D - c_D$

Table 1: Game matrix for belief contest

that the innovator invests, while the defender does not. It is obvious that profile $(0, 1)$ cannot be a Nash equilibrium. If the innovator does not campaign, the defender does not want to do so either: as the incumbent, the defender anyway serves the entire audience if the innovator is inactive. However, all other strategy profiles can be equilibria under some parameter constellations. Concerning $\tilde{v}_k$ and c_k we find that it is always only the ratio of the two that matters for the equilibrium outcome. For brevity, we refer to $\tilde{v}_k/c_k$, $k = I, D$, as the *normalized payoff*. Note that $\tilde{v}_k$ combines before-payoffs and after-payoffs.¹⁴

Table 1 implies that profile $(0, 0)$ is an equilibrium if and only if $\tilde{v}_I/c_I \leq 1/\psi$, i.e., if and only if the innovator's normalized payoff is small (also see Figure 1). Profile $(1, 0)$ is an equilibrium if and only if $\tilde{v}_I/c_I \geq 1/\psi$ and $\tilde{v}_D/c_D \leq 1/(\psi(1 - \lambda))$, i.e., in a situation in which the innovator's normalized payoff is large and the defender's is relatively small. The remaining pure-strategy profile $(1, 1)$ is an equilibrium if and only if $\tilde{v}_I/c_I \geq 1/(\psi\lambda)$ and $\tilde{v}_D/c_D \geq 1/(\psi(1 - \lambda))$, i.e., if both entrepreneurs stand to win large normalized payoffs.

Under the constellation $1/\psi < \tilde{v}_I/c_I < 1/(\psi\lambda)$ and $\tilde{v}_D/c_D > 1/(\psi(1 - \lambda))$, a Nash equilibrium in pure strategies does not exist. The innovator's normalized payoff is too large for profile $(0, 0)$ to be an equilibrium. Profile $(1, 0)$ cannot be an equilibrium because the defender's normalized payoff is too large for inaction if the innovator campaigns. But what about $(1, 1)$? This profile cannot be an equilibrium either because the innovator's payoff is just not large enough to justify a fight against the defender. Since profile $(0, 1)$ can never be an equilibrium, it follows that there is no pure-strategy equilibrium at all. But because every game with a finite number of players and profiles has at least one Nash equilibrium, there must be an equilibrium in mixed strategies. Denote by p the probability of the innovator choosing $e_I = 1$ and by q the

¹⁴More precisely, $\tilde{v}_k/c_k$ is the sum of (discounted) payoffs, as expected by belief entrepreneur k , that come with supplying one unit of k 's approach, normalized by k 's per-unit cost of effort.

Figure 1: Equilibrium strategy profiles

probability of the defender choosing $e_D = 1$. Then:

$$p = \frac{1}{\psi(1-\lambda) \tilde{v}_D/c_D} \quad \text{and} \quad q = \frac{\psi - 1/(\tilde{v}_I/c_I)}{(1-\lambda)\psi}. \quad (8)$$

Again, it is normalized payoffs that matter. All else equal, a stronger defender (i.e., one with a larger normalized payoff) makes the innovator *less* likely to enter a belief campaign, while a stronger innovator means that the defender is *more* likely to campaign. Figure 1 summarizes what type of equilibrium arises under which parameter constellation. The x -axis and the y -axis measure the normalized payoffs that, respectively, the innovator and the defender stand to win.

5.2 Discussion

While simple, the toy model conveys many of the insights of the more general setup in Section 6. It assembles a set of factors influencing subjective beliefs under fundamental uncertainty and clarifies how they interact. The toy model offers guidance even in the limiting case of a complete absence of objective information on $\Pr[F = H]$. It demonstrates that economics can provide insights into belief formation even if there is little or no valuable information on the state of nature. In treating the distribution of subjective beliefs as an equilibrium outcome, our analysis deviates sharply from existing economic analyses of fundamental uncertainty which almost invariably treat subjective (prior) beliefs as primitives.

We now discuss how the different factors assembled by the model determine the campaigning efforts. First consider a scenario with full discounting of after-payoffs (i.e., $\delta_D = \delta_I = 0$, implying $\tilde{v}_k = v_k$ for $k = I, D$). In this case, the fundamental has “no voice” in the behavior of any of the players. In particular, the belief entrepreneurs’ incentives to campaign are driven by the normalized before-(resolution-of-uncertainty) payoffs, v_k/c_k , which in this scenario are measured on the two axes of Figure 1. The figure can be used to infer how v_I/c_I and v_D/c_D determine the type of equilibrium for given values of λ and ψ —and so influence the set of subjective beliefs adopted by the audience and the resulting demand for the innovation. Things are simple if the innovator’s normalized before-payoff is low. Then, there is no campaigning and all $\alpha_i\pi$ s, i.e., all subjective beliefs about $\Pr[F = H]$ held by the members of the audience, are identical to zero (equation 6). As a consequence, no one opts for the innovative approach (equation 7). If the innovator’s normalized before-payoff is higher, the defender’s payoff matters, too. First assume that v_D/c_D is low. Then, only the innovator campaigns. If $\zeta = \bar{\zeta}$, the contest produces $\pi = 1/\bar{\alpha}$, inducing all audience members to adopt a subjective belief that prompts the adoption of the innovative approach (equation 7); if $\zeta = 0$, the beliefs are again identical to zero and there is no demand for the innovative approach. Assume now that v_D/c_D takes a higher value. Then, both belief entrepreneurs have a strictly positive probability of campaigning—with the result that it is less likely that the beliefs about $\Pr[F = H]$ take a high value. Note, however, that in none of the cases discussed so far, there is any correlation between the fundamental F and the equilibrium beliefs about F .

This is different in the alternative scenario in which after-payoffs carry some weight (i.e., $\delta_k > 0$ for at least one $k = I, D$). Then, information about the fundamental influences the set of beliefs adopted by the audience. It does so by altering the normalized payoffs $\tilde{v}_k/c_k$, which in turn influence the type of equilibrium. To see how exactly information finds its way into subjective beliefs, consider Figure 2. In the figure, the quality of the signal to the belief entrepreneurs is varied from completely uninformative ($\sigma = 1/2$) to perfectly informative ($\sigma = 1$). With $\sigma = 1/2$, the prevailing $(\tilde{v}_I/c_I, \tilde{v}_D/c_D)$ -combination is assumed to be represented by point A . Note that the realization of the signal $S \in \{H, L\}$ does not affect the position of A since, with a completely uninformative signal, the belief entrepreneurs will not update their priors ϕ_I^0 and ϕ_D^0 (and so $\tilde{v}_I$ and $\tilde{v}_D$ are unaffected, too).

Yet, if $\sigma > 1/2$, the realization of S does affect the $(\tilde{v}_I/c_I, \tilde{v}_D/c_D)$ -combination. Suppose first that $S = H$. Then, for both belief entrepreneurs, updating means an increase in the subjective belief about the probability of $F = H$ (i.e., $\phi_I > \phi_I^0$ and $\phi_D > \phi_D^0$). So, as compared to the situation with a completely uninformative signal, $\tilde{v}_I$ takes a higher value and

$\tilde{v}_D$ takes a lower value. As result, $(\tilde{v}_I/c_I, \tilde{v}_D/c_D)$ must lie to the south-east of point A —by how much depends on the quality of the signal, as illustrated in the figure. A similar logic implies that $(\tilde{v}_I/c_I, \tilde{v}_D/c_D)$ must lie to the north-west of A if $S = L$. The conclusion to be drawn from Figure 2 is that the signal “pushes” the game towards an equilibrium that tends to produce subjective beliefs that are in line with the signal’s realization: positive towards the innovation if $S = H$ and negative towards the innovation if $S = L$. Clearly, the size of the push also depends on variables other than signal quality: all else equal, it is larger if the after-payoffs for adverse outcomes are more damaging (i.e., if w_I^L, w_D^H are larger) and if there is less discounting of after-payoffs (i.e., if δ_I, δ_D are larger).

5.3 Policy Implications

Unfortunately, there are good reasons to expect that for many innovations the discounting of after-payoffs is rather substantial. Conclusive evidence on the success or failure of an innovative approach often accumulates only slowly (e.g., as in the case of quantitative easing) and sometimes the merits of an innovation are not revealed at all before a rare “stress situation” emerges (e.g., as in the case of securitization). Moreover, the tenure of decision makers can be short, a factor that contributes to discounting. So, in practice, it is certainly not uncommon that before-payoffs, which show little relationship with the fundamentals of an innovation, dominate the incentives to campaign and the formation of subjective beliefs. As a consequence, the subjective beliefs that the members of a target audience adopt from a belief contest may often not relate in any substantial way to the fundamentals.

For a policy maker, it may be helpful to have an idea about the extent to which one can trust emerging positive (or even exuberant) subjective beliefs to reflect the fundamentals of an innovation. If the extent is small, the policy maker has to reckon with the possibility that at some point the belief will collide with the facts—and so produce a fallout for which preparations are required. Our theory offers policy makers a simple diagnostic framework for assessing such collision potentials. Importantly, the framework does not demand that policy makers have access to superior information. Our theory suggests that a policy maker consider information that helps in estimating the size of before-payoffs (i.e., v_I, v_D) and the extent of discounting (i.e., δ_I, δ_D). If the policy maker finds the discount factors δ_k to be small and/or the payoffs v_k to be large, there is no reason to expect that the adoption of positive beliefs by the target audience is informative of the true probability that the state of the world is favorable to the innovation. So fallout preparations are more urgent. In the opposite constellation, there may be a positive correlation between the subjective beliefs and the probability that the innovation

Figure 2: Information and equilibrium strategy profiles

is a success. Worries about a possible fallout may thus appear less pressing.

In essence, our theory suggests that—if objective information on the merits of an innovation is scarce or lacking—a sensible approach is to consider the *incentives* to influence beliefs. How big are the benefits that potentially await the contestants before the resolution of uncertainty? How far away is the resolution of uncertainty? For how long are decision makers typically held to account? It may not be trivial to answer these questions, not least because belief entrepreneurs may try to obscure the facts. But it is far from impossible to estimate these factors with some degree of precision.¹⁵ In any case, obtaining such estimates seems more feasible than detecting the fundamentals of an innovation at an early stage.

6 Belief Formation in the General Setup

We now relax the two simplifying assumptions stated at the beginning of Section 5. The campaigning efforts are now continuous choice variables ($e_k \in \mathbb{R}_{\geq 0}$, $k = I, D$) and the random variable ζ follows a continuous distribution. For tractability, we assume that ζ is uniformly distributed on $\zeta \in [0, u]$, with $u > 0$. For the very same reason, we assume a uniform distribution for the individual components of subjective beliefs, $\alpha_i \in [0, \bar{\alpha}]$.

¹⁵Non-standard approaches may be helpful in this regard. For instance, natural language processing (NLP) may help in the analysis of the language used in belief contests, with a view on identifying “hypes”.

6.1 Analysis

Individual decisions on the use of the approach follow immediately from equation (3). The equation implies that there exists a threshold $\tilde{\alpha} \equiv \rho/(\bar{\theta}\pi)$ such that $a_i = 1$ if and only if $\alpha_i > \tilde{\alpha}$. Using contest success function (2), and assuming $e_I > 0$, we find

$$\tilde{\alpha}(e_I, e_D, \zeta) = \frac{\bar{\alpha}\rho}{\theta} \left[1 + \frac{e_D^\eta}{e_I^\eta \zeta} \right]. \quad (9)$$

Threshold $\tilde{\alpha}$ determines how the two belief entrepreneurs split the target audience between them. Since the α_i s follows a uniform distribution on $[0, \bar{\alpha}]$, it follows that

$$D_I = \max\{1 - \tilde{\alpha}/\bar{\alpha}, 0\}. \quad (10)$$

Intuitively speaking, it is the position of $\tilde{\alpha}$ that is contested. In addition to the strategy variables e_I and e_D , $\tilde{\alpha}$ is affected by ζ , a random variable. Finding an optimal response to a given strategy by the opponent is thus an optimization problem under risk. As is illustrated in Figure 3, the innovator would like to push $\tilde{\alpha}$ to the left, while the defender wants to push it to the right. Yet both have only limited control over their “market share” because of random variable ζ . The minimum of $\tilde{\alpha}$ is $\bar{\alpha}\rho/\bar{\theta}$, a level also specified in the figure.

It is noteworthy that $\Pr[\tilde{\alpha} > \bar{\alpha}] = 1$ if e_I takes a small value *relative to* e_D (equation 9). This is because both α_i and ζ have strictly finite upper bounds. Therefore, if e_I is relatively small, the marginal benefit (in terms of a lowering of $\tilde{\alpha}$) associated with raising e_I is 0. When deriving best responses and equilibria, this feature must be taken care of. Moreover, we note that contest success function (2) is not continuous at $(0, 0)$, something that requires additional attention. As a result, objective functions (4) and (5) are not strictly concave, implying that local maxima (as identified by the first- and second-order conditions) do not necessarily mean positive expected payoffs. As the derivation of the equilibrium is tedious, we relegate it to the Appendix. But the characterization of the equilibrium is straightforward:

PROPOSITION 1 *Under individual choice, there exists a unique Nash equilibrium in pure strategies if and only if $\tilde{v}_I/c_I > \nu$ for some critical level $\nu > 0$. In this equilibrium, the ratio of campaigning efforts equals the ratio of normalized payoffs:*

$$\frac{e_I}{e_D} = \frac{\tilde{v}_I/c_I}{\tilde{v}_D/c_D}. \quad (11)$$

Figure 3: The battle for “market share”

Therefore:

$$\pi = \frac{1}{\bar{\alpha}} \left[1 + \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right)^\eta \frac{1}{\zeta} \right]^{-1} \quad \text{and} \quad \tilde{\alpha} = \frac{\bar{\alpha}\rho}{\theta} \left[1 + \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right)^\eta \frac{1}{\zeta} \right]. \quad (12)$$

Proof. See Appendix B. ■

The expressions in (12) follow immediately from equations (11), (2), and (9). The condition on the normalized payoff stated in the proposition rules out that an interior maximum of V_I is associated with a strictly negative payoff.¹⁶ If that condition is violated, an equilibrium in pure strategies does not exist. Since e_I and e_D are continuous choice variables, mixed equilibria are not straightforward to analyze. In particular, Nash’s existence theorem does not hold. One way of dealing with this constellation is to assume that nature enters as a player and discretizes the action space. The toy model of Section 5 is an extreme example of such a discretization—and its analysis provides insights into why an equilibrium in pure strategies may not exist: if the innovator’s normalized payoff is sufficiently small, they want to invest in a belief campaign if and only if the defender does not want to. Under pure strategies, this leads to mutually incompatible incentives because the defender invests in a belief campaign if and only if the innovator invests, too. Note that any use of mixed strategies strengthens the influence of randomness in the process of subjective belief formation.

¹⁶We could not find an easily interpretable algebraic expression for ν that is both necessary and sufficient. Conditions that are just sufficient are not particularly insightful.

6.2 Discussion

The expressions in (12) show how the three factors “information”, “payoffs” (both through $\tilde{v}_k$, $k = I, D$), and “luck” (through ζ) affect the outcome of a belief contest. Overall, the results of the general setup resemble those of the simplified one. But they are more nuanced due to the former setup’s continuous nature. In the general setup, the belief anchor takes an intermediate value unless the normalized payoffs are very dissimilar or unless ζ takes an extreme value. In the simplified setup, this feature is masked by the discretization (which pushes the belief anchor towards one of the polar values). An intermediate belief anchor means a balanced outcome in the use of the two approaches (intermediate value of D_I). The tendency towards the middle reduces the likelihood of situations in which there is either great stability at the expense of innovation (D_I close to 0) or a broad embrace of innovation at the expense of stability (D_I close to 1).¹⁷ From a policy perspective, this property could be desirable.

Going beyond the present framework leads to an additional reason for a possible desirability of moderate beliefs. Consider a dynamic setting in which beliefs are gradually updated, as a sequence of signals step by step unveils the state of the world. If a subjective prior belief is moderate, a string of unfavorable signals makes a user of an innovation change course in due time. However, if the user starts with a subjective prior that is highly favorable towards the innovation, even a string of clearly negative signals may not alter the user’s positive attitude. As a result, the user may hold on to the innovation for a long time—with corresponding consequences for the losses that eventually may fall due. If the group of users is large, this may even affect aggregate outcomes (to the point of triggering a recession).

7 Conclusion

This paper offers a new perspective on belief formation under fundamental uncertainty. We understand beliefs about novel phenomena as an equilibrium outcome, shaped in a contest between self-interested belief entrepreneurs. The major determinants of equilibrium beliefs are: the belief entrepreneurs’ expected benefits (“payoffs”), the degree of fundamental uncertainty (“information”), and randomness relating to the success of narratives (“luck”). A key contribution of this paper is to show that the microeconomic toolbox can be useful for explaining how beliefs emerge endogenously under fundamental uncertainty.

We see several applications for our framework, three of them briefly discussed here. First,

¹⁷The variance of the aggregate outcome is given by $D_I(\bar{\alpha})^2\bar{\theta}^2\Pr[F = H](1 - \Pr[F = H])$. Thus, unless $\Pr[F = H]$ is either 0 or 1, the variance of the outcome increases with the square of D_I .

it is natural to ask how endogenous beliefs, as modeled in this paper, affect macroeconomic time series in economies that regularly experience innovation-driven fundamental uncertainty. Clearly, looking into this question requires the specification of how beliefs respond to the gradual resolution of uncertainty. In practice, evidence on innovations often emerges slowly, and so uncertainty resolution takes time. As a result, favorable views about an innovation—even if proven unfounded eventually—may persist for a longer time. But once reality sinks in, the reactions may be severe. In our view, it is interesting to explore to what extent this pattern may help better understand business cycle dynamics, bubble patterns, and the possibilities and limitations of conditionally predicting booms, recessions, and crises.

Second, our framework may be helpful to regulators that have to assess and potentially regulate innovations that are within their purview. The present analysis identifies some observable factors—such as the innovator’s payoff structure—that can be considered in order to gauge the relationship between an emerging subjective belief about an innovation and the innovation’s true, but yet unobserved, fundamentals. So far, however, our analysis is stylized. A more elaborate framework useful to regulators confronted with fundamental uncertainty would be intertemporal and allow objective information to emerge stepwise.

Third, “ambiguity suppression” is not only observed in connection with innovation-driven fundamental uncertainty but may arise in the context of any complex choice situation. In particular, it is often difficult for voters to “objectively” evaluate the relative costs and benefits of competing policy proposals. As in the case of innovations, this may lead to a demand for heuristic short-cuts—and thus provide political parties (or other political players) with an opportunity to act as belief entrepreneurs that offer ready-made evaluation frames. Depending on the details of the contest game, the resulting “narrative campaigns” may lead to a strong polarization of the electorate. We find it interesting to explore to what extent this mechanism can shed light on the recent rise in political polarization and populism.

References

- Armstrong, Paul B. (2020). *Stories and the Brain. The Neuroscience of Narrative*. John Hopkins University Press.
- Bar, Moshe (2007). “The Proactive Brain: Using Analogies and Associations to Generate Predictions”. *Trends in Cognitive Sciences* 11.7, pp. 280–289.
- Barrett, Lisa Feldman (2017a). *How Emotions are Made. The Secret Life of the Brain*. Houghton Mifflin Harcourt.

- Barrett, Lisa Feldman (2017b). “The Theory of Constructed Emotion: An Active Inference Account of Interoception and Categorization”. *Social Cognitive and Affective Neuroscience* 12.1, pp. 1–23.
- Barsalou, Lawrence W. (2009). “Simulation, Situated Conceptualization, and Prediction”. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364, pp. 1281–1289.
- BCG (2017). *Global Risk 2017. Staying the Course in Banking*. The Boston Consulting Group, Inc.
- Bénabou, Roland (2013). “Groupthink: Collective Delusions in Organizations and Markets”. *Review of Economic Studies* 80, pp. 429–462.
- Bénabou, Roland, Armin Falk, and Jean Tirole (2020). “Narratives, Imperatives, and Moral Persuasion”. Working Paper; previous version available as NBER Working Paper 24798.
- Bénabou, Roland and Jean Tirole (2002). “Self-Confidence and Personal Motivation”. *Quarterly Journal of Economics* 117.3, pp. 871–915.
- (2016). “Mindful Economics: The Production, Consumption, and Value of Beliefs”. *Quarterly Journal of Economics* 30.3, pp. 141–164.
- Binswanger, Johannes and Manuel Oechslin (2015). “Disagreement and Learning about Reforms”. *The Economic Journal* 125, pp. 853–886.
- BIS (2005). *The Role of Ratings in Structured Finance: Issues and Implications. Report submitted by a Working Group established by the Committee on the Global Financial System*. Bank for International Settlements.
- Brunnermeier, Markus K., Alp Simsek, and Wei Xiong (2014). “A Welfare Criterion for Models with Distorted Beliefs”. *Quarterly Journal of Economics* 129.4, pp. 1753–1797.
- Clark, Andy (2013). “Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science”. *Behavioral and Brain Sciences* 36, pp. 281–253.
- (2016). *Surfing Uncertainty. Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Corchòn, Luis C. and Marco A. Marini (2018). “Contest Theory”. In: *Handbook of Game Theory and Industrial Organization*. Vol. Volume II. Edward Elgar, pp. 125–146.
- Dungey, Mardi, Gerald P. Dwyer, and Thomas Flavin (2013). “Systematic and Liquidity Risk in Subprime-Mortgage Backed Securities”. *Open Economies Review* 24, pp. 5–32.
- Eliasz, Kfir and Ran Spiegler (2020). “A Model of Competing Narratives”. *American Economic Review* 110.12, pp. 3786–3816.
- Ellsberg, Daniel (1961). “Risk, Ambiguity, and the Savage Axioms”. *Quarterly Journal of Economics* 75, pp. 643–669.

- Falkinger, Josef (2007). “Attention Economies”. *Journal of Economic Theory* 133, pp. 266–294.
- Fehr, Ernst and Urs Fischbacher (2003). “The Nature of Human Altruism”. *Nature* 425, pp. 785–791.
- Fehr, Ernst, Urs Fischbacher, and Michael Kosfeld (2005). “Neuroeconomic Foundations of Trust and Social Preferences: Initial Evidence”. *American Economic Review: AEA Papers and Proceedings* 95.2, pp. 346–3511.
- Fehr, Ernst and Simon Gächter (2002). “Altruistic Punishment in Humans”. *Nature* 415, pp. 137–140.
- Flood, John (2009). “Book Review: Fool’s Gold—How Unrestrained Greed Corrupted a Dream, Shattered Global Markets and Unleashed a Catastrophe by Gillian Tett”. *Journal of Law and Society* 36.4, pp. 579–584.
- Friston, Karl (2010). “The Free-Energy Principle: A Unified Brain Theory?”. *Nature Reviews Neuroscience* 11, pp. 127–138.
- Gilboa, Itzhak, Larry Samuelson, and David Schmeidler (2014). “No-Betting-Pareto Dominance”. *Econometrica* 82.4, pp. 1405–1442.
- Herz, Holger, Daniel Schunk, and Christian Zehnder (2014). “How Do Judgmental Overconfidence and Overoptimism Shape Innovative Activity?”. *Games and Economic Behavior* 83, pp. 1–23.
- Kahneman, Daniel (2011). *Thinking, Fast and Slow*. Random House.
- Kahneman, Daniel and Amos Tversky (1973). “On the Psychology of Prediction”. *Psychological Review* 80.4, pp. 237–251.
- Knight, Frank (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin Company.
- Konrad, Kai (2009). *Strategy and Dynamics in Contests*. Oxford University Press.
- Machina, Mark J. and Marciano Siniscalchi (2014). “Ambiguity and Ambiguity Aversion”. In: *Handbook of the Economics of Risk and Uncertainty*. Vol. Volume 1. Elsevier B.V., pp. 729–807.
- McAdams, Dan P. and Jennifer L. Pals (2006). “A New Big Five: Fundamental Principles for an Integrative Science of Personality”. *American Psychologist* 61.3, pp. 204–207.
- Mokyr, Joel (2013). “Cultural entrepreneurs and the origins of modern economic growth”. *Scandinavian Economic History Review* 61.1, pp. 1–33.
- (2017). *A Culture of Growth. The Origins of the Modern Economy*. Princeton University Press.
- Popper, Karl (1945). *The Open Society and Its Enemies*. Routledge.
- Schumpeter, Joseph (1942). *Capitalism, Socialism and Democracy*. Harper & Brothers.

- Segal, Michael (2018). “Automatic Pilots”. *Nature* 563.
- Shiller, Robert J. (2017). “Narrative Economics”. *American Economic Review* 107.4, pp. 967–1004.
- (2019). *Narrative Economics. How Stories Go Viral and Drive Major Economic Events*. Princeton University Press.
- Slovic, Paul et al. (2007). “The Affect Heuristic”. *European Journal of Operational Research* 177, pp. 1333–1352.
- Tett, Gillian (2009). *Fool’s Gold: How Unrestrained Greed Corrupted a Dream, Shattered Global Markets and Unleashed a Catastrophe*. Little, Brown.
- Tomasello, Michael and Hannes Rakoczy (2003). “What Makes Human Cognition Unique? From Individual to Shared to Collective Intentionality”. *Mind and Language* 18.2, pp. 121–147.
- Tversky, Amos and Daniel Kahneman (1974). “Judgment under Uncertainty: Heuristics and Biases”. *Science* 185.4157, pp. 1124–1131.

Appendix for

THE ECONOMICS OF BELIEFS UNDER
FUNDAMENTAL UNCERTAINTY

December 11, 2020

A. Collective Choice

In practice, there are many instances in which a decision on an innovation has to be taken collectively (e.g., by a board of directors, by a legislature, or in a popular vote). Appendix A outlines a variant of the model that is based on collective, instead of individual, choice. As this modification leaves the insights of the preceding analysis virtually unchanged, we limit the exposition to the simplified setup presented in Section 5. The analysis of the general setup with collective choice is available from the authors upon request.

Assumptions

Consider the following modification. Under collective choice, each audience member votes for the approach they prefer individually. The audience as a whole then adopts the innovative approach ($a = 1$) if the vote share of the latter (strictly) exceeds a certain threshold, γ ; otherwise, the audience chooses the tested approach ($a = 0$). Individual decisions (here: votes) continue to be based on objective function (3), with $a \in \{0, 1\}$ replacing a_i . Specifically, member i votes for the innovative approach if and only if $a = 1$ maximizes $E_i\{U(a)\}$. We denote by $X_I(\pi)$ the share of votes cast in favor of the innovative approach. The functional form of $X_I(\pi)$ is identical to that specified on the right-hand side of equation (7). It follows that the demand for the innovative approach, $D_I(X_I(\pi))$, is 1 if $D_I(X_I(\pi)) > \gamma$ and 0 otherwise. Below, we have to distinguish between two cases: $\lambda \leq \gamma$ and $\lambda > \gamma$.

		Defender	
		$e_D = 0$	$e_D = 1$
Innovator	$e_I = 0$	0 v_D	0 $v_D - c_D$
	$e_I = 1$	$\psi \tilde{v}_I - c_I$ $(1 - \psi)v_D$	$-c_I$ $v_D - c_D$

Table 2: Game matrix for belief contest with collective choice and $\lambda \leq \gamma$

		Defender	
		$e_D = 0$	$e_D = 1$
Innovator	$e_I = 0$	0 v_D	0 $v_D - c_D$
	$e_I = 1$	$\psi \tilde{v}_I - c_I$ $(1 - \psi)v_D$	$\psi \tilde{v}_I - c_I$ $(1 - \psi)v_D - c_D$

Table 3: Game matrix for belief contest with collective choice and $\lambda > \gamma$

Analysis

First assume that $\lambda \leq \gamma$. In this case, the audience chooses the innovative approach if and only if $\pi = 1/\bar{\alpha}$. The resulting payoffs are shown in Table 2. Compared to Table 1, only the payoffs for (1, 1) differ. Since with this strategy profile $X_I(\pi)$ necessarily falls short of γ , the innovator never wins and so has no chance of recovering any cost of campaigning. As a result, profile (1, 1) cannot be a Nash equilibrium. Profile (0, 1) cannot be an equilibrium either because the defender has an incentive to deviate. The two remaining profiles can be equilibria. Profile (0, 0) is an equilibrium if and only if $\tilde{v}_I/c_I \leq 1/\psi$. Profile (1, 0) is an equilibrium if and only if $\tilde{v}_I/c_I \geq 1/\psi$ and $v_D/c_D \leq 1/\psi$. If neither of the two latter profiles constitutes an equilibrium, there is an equilibrium in mixed strategies. The probabilities of campaigning are

$$p = \frac{1}{\psi(v_D/c_D)} \quad \text{and} \quad q = \frac{\psi - 1/(\tilde{v}_I/c_I)}{\psi}. \quad (\text{A1})$$

As in the case of individual choice, the probability that the innovator campaigns is smaller if the defender is stronger (i.e., has a larger normalized payoff), while the defender's probability of campaigning is larger if the innovator is stronger. Overall, we conclude that the pattern of equilibrium profiles is very similar to that under individual choice. In terms of Figure 1, the

most notable difference is that under collective choice the “mixed equilibrium area” extends to the upper-right corner (I and D both campaigning cannot be an equilibrium).

Now consider $\lambda > \gamma$. The payoffs for that case are shown in Table 3. Whenever $e_I = 1$, the innovator wins with probability ψ and the defender wins with the complementary probability. Clearly, for the defender, $e_D = 0$ dominates $e_D = 1$. As a consequence, the unique Nash equilibrium is given by profile $(1, 0)$ if $\tilde{v}_I/c_I > 1/\psi$ and by profile $(0, 0)$ otherwise.

B. Proof of Proposition 1

A profile (e_I, e_D) is an equilibrium if and only if e_k maximizes V_k , given e_{-k} , for both players $k = I, D$ —i.e., the profile consists of mutual best responses. In the proof below, we first state the objective functions V_k . We then characterize best responses. We show that any candidate equilibrium needs to fulfill standard first-order conditions, in spite of the fact that the objective functions are neither strictly concave nor everywhere differentiable. However, first-order conditions obviously are not sufficient for a best response. In the final part of the proof, we establish that a sufficiently large normalized payoff makes the first-order conditions also sufficient for a best response and thus an equilibrium.

Belief Entrepreneurs’ Objective Functions

We start with the innovator’s objective function. In stage 4 of the game, individuals decide on $a_i \in \{0, 1\}$. The resulting demand for the innovative approach, D_I , is given by equation (10). From equation (10), it follows that $D_I = 0$ if $\tilde{\alpha} \geq \bar{\alpha}$, where $\tilde{\alpha}$ is given by equation (9). (For simplicity, we assume that, in case of indifference, a member of the audience chooses $a_i = 0$). Now consider the formation of expectations about D_I in stage 3. From equations (9) and (10), and the fact that the maximum of ζ is u , it follows that

$$E[D_I] > 0 \text{ if and only if } e_I > \underline{e}_I(e_D), \quad (\text{A2})$$

where

$$\underline{e}_I(e_D) \equiv \left[\frac{\rho}{u(\bar{\theta} - \rho)} \right]^{1/\eta} e_D. \quad (\text{A3})$$

Clearly, $\underline{e}_I(e_D) > 0$ whenever $e_D > 0$. If $0 < e_I \leq \underline{e}_I(e_D)$, $\tilde{\alpha} < \bar{\alpha}$ would require random variable ζ to exceed its maximum of u . Conversely, if $e_I > \underline{e}_I(e_D)$, there is always a strictly positive probability that ζ takes a sufficiently large value such that $\tilde{\alpha} < \bar{\alpha}$ and therefore $D_I > 0$.

Provided that $e_I > \underline{e}_I(e_D)$, straightforward calculations show that

$$D_I > 0 \text{ if and only if } \zeta > \zeta_0, \quad (\text{A4})$$

where

$$\zeta_0 \equiv \frac{\rho}{\bar{\theta} - \rho} \left(\frac{e_D}{e_I} \right)^\eta. \quad (\text{A5})$$

Equation (10), together with (A2) to (A5), implies that—when stage 4 is anticipated—the innovator’s objective function (4) can be rewritten as follows:

$$V_I = \begin{cases} \int_{\zeta_0(e_I, e_D)}^u \left[1 - \frac{\tilde{\alpha}(e_I, e_D; \zeta)}{\bar{\alpha}} \right] \frac{1}{u} d\zeta \tilde{v}_I - c_I e_I & \text{if } e_I > \underline{e}_I(e_D), e_D > 0 \\ \left(1 - \frac{\rho}{\bar{\theta}} \right) \tilde{v}_I - c_I e_I, & \text{if } e_I > 0, e_D = 0 \\ -c_I e_I & \text{if } e_I \leq \underline{e}_I(e_D), e_D > 0 \\ 0 & \text{if } e_I = 0, e_D = 0 \end{cases}. \quad (\text{A6})$$

In equation (A6), the integral on the first line gives $E[D_I]$ for the specified (e_I, e_D) -combination. Because of (A4), ζ_0 marks the lower integration limit (which depends on the campaigning efforts, too). If $e_I > 0$ and $e_D = 0$, demand simplifies to $D_I = 1 - \rho/\bar{\theta}$, as shown on the second line. Because the minimum value of $\tilde{\alpha}$ is $\bar{\alpha}\rho/\bar{\theta} > 0$, the innovator does never server the entire target audience. The third line in equation (A6) refers to the situation in which e_I is not sufficiently large to generate positive demand for the innovative approach. To understand the final line, note that for this case equation (2) specifies $\pi = 0$ (and thus $\tilde{\alpha} \rightarrow \infty$).

Turning to the defender’s objective function, we note that in stage 4 she serves a fraction $D_D = \min\{\tilde{\alpha}/\bar{\alpha}, 1\}$ of the target audience. Because $\tilde{\alpha} \geq \bar{\alpha}\rho/\bar{\theta}$, there is always some demand for the tested approach. Moreover, $D_D = 1$ whenever $\zeta < \zeta_0$. Furthermore, if e_D is sufficiently large, the defender serves the entire audience even if $\zeta = u$:

$$\text{If } e_D \geq \bar{e}_D(e_I), \text{ then } D_D = 1, \quad (\text{A7})$$

where

$$\bar{e}_D(e_I) \equiv \left(\frac{u(\bar{\theta} - \rho)}{\rho} \right)^{\frac{1}{\eta}} e_I. \quad (\text{A8})$$

So, if $e_D \geq \bar{e}_D$, an increase in e_D is useless. All this implies that—when stage 4 is anticipated—

the defender's objective function (5) can be rewritten as follows:

$$V_D = \begin{cases} \frac{\zeta_0(e_I, e_D)}{u} \tilde{v}_D + \int_{\zeta_0(e_I, e_D)}^u \frac{\tilde{\alpha}(e_I, e_D; \zeta)}{\bar{\alpha}} \frac{1}{u} d\zeta \tilde{v}_D - c_D e_D & \text{if } e_D < \bar{e}_D(e_I), e_I > 0 \\ \tilde{v}_D - c_D e_D & \text{otherwise} \end{cases} \quad (\text{A9})$$

Note that for $e_D = 0$, we have $\zeta_0 = u$.

Best Responses

We start again with the innovator. Taking the partial derivative of V_I with respect to e_I yields:

$$\frac{\partial V_I}{\partial e_I} = \begin{cases} \int_{\zeta_0}^u \frac{\rho\eta}{\bar{\theta}u} \frac{e_D^\eta}{e_I^{1+\eta}} \frac{1}{\zeta} d\zeta \tilde{v}_I - c_I & \text{if } e_I \geq \underline{e}_I(e_D), e_D > 0 \\ -c_I & \text{if } 0 \leq e_I < \underline{e}_I(e_D), e_D > 0; \text{ or if } e_I > 0, e_D = 0 \\ \infty & \text{if } e_I = 0, e_D = 0 \end{cases} \quad (\text{A10})$$

To understand the first line of equation (A10), note that ζ_0 —defined in equation (A5)—is a function of e_I and marks the lower limit of the integral on the first line of equation (A6). In the partial derivative, this gives rise to a term $[1 - \tilde{\alpha}(e_I, e_D, \zeta_0)/\bar{\alpha}] u^{-1} (\partial\zeta_0/\partial e_I) \tilde{v}_I$. However, because $\tilde{\alpha}(e_I, e_D, \zeta_0) = \bar{\alpha}$ by definition of ζ_0 (see equations (9) and (A5)), this term vanishes. The third line of equation (A10) uses somewhat informal notation. It refers to the right-sided derivative (since $e_I \geq 0$) and expresses that V_I makes a discrete jump from 0 to a strictly positive value if e_I increases marginally from 0.

Consider the case $e_D = 0$. Then, the third line of the right-hand side of equation (A10) implies that $e_I = 0$ is not a best response, implying that the profile (0,0) cannot be an equilibrium. Furthermore, if $e_D = 0$, the second line implies that no $e_I > 0$ is a best response: the innovator could marginally decrease e_I to a level that is still strictly positive, with no reduction in market share. This follows from the discontinuity of the contest success function at zero. As a result, no profile with $e_D = 0$ can be an equilibrium.

For the case $e_D > 0$, equation (A10) implies that $0 < e_I < \underline{e}_I$ cannot be a best response because $\partial V_I/\partial e_I < 0$ and lowering e_I is feasible. Solving the integral on the first line yields

$$\frac{\partial V_I}{\partial e_I} = \frac{\rho\eta}{\bar{\theta}u} \frac{e_D^\eta}{e_I^{1+\eta}} \left[\ln \left(\frac{u(\bar{\theta} - \rho)}{\rho} \right) - \eta \ln \left(\frac{e_D}{e_I} \right) \right] \tilde{v}_I - c_I \quad \text{if } e_I \geq \underline{e}_I, e_D > 0. \quad (\text{A11})$$

The second partial derivative for this constellation is given by

$$\frac{\partial^2 V_I}{\partial e_I^2} = \frac{\rho\eta}{\bar{\theta}u} \frac{e_D^\eta}{e_I^{2+\eta}} \left[\eta - (1+\eta) \ln \left(\frac{u(\bar{\theta}-\rho)}{\rho} \frac{e_I^\eta}{e_D^\eta} \right) \right] \tilde{v}_I \quad \text{if } e_I > \underline{e}_I, e_D > 0. \quad (\text{A12})$$

Note that for $e_I = \underline{e}_I$ we have $u(\bar{\theta}-\rho)\rho^{-1}(e_I/e_D)^\eta = 1$. As a result, $\lim_{e_I \searrow \underline{e}_I} \partial V_I / \partial e_I = -c_I < 0$ and $\lim_{e_I \searrow \underline{e}_I} \partial^2 V_I / \partial e_I^2 > 0$. Moreover, the term in square brackets in equation (A12) is a strictly decreasing function of e_I and eventually turns negative. Finally, $\lim_{e_I \rightarrow \infty} \partial V_I / \partial e_I = -c_I$ since a power function grows sufficiently faster towards infinity than its logarithm. Thus, as e_I rises from $\underline{e}_I$ towards ∞ , $\partial V_I / \partial e_I$ (strictly) increases from $-c_I$, reaches a maximum, and then (strictly) falls towards the initial level. Depending on $\tilde{v}_I / c_I$, as well as on other parameters, $\partial V_I / \partial e_I = 0$ has either zero, one, or two solutions (the situation with only one solution is a limit case). If there are two solutions, $\partial V_I / \partial e_I$ crosses zero from below at the level of the smaller solution and from above at the level of the larger. So the second-order condition holds for the larger solution only. To summarize the results for the innovator's best response:

LEMMA 1 (i) If $e_D = 0$, no best response exists; as a result, no profile (e_I, e_D) with $e_D = 0$ can be an equilibrium. (ii) If $e_D > 0$, the innovator's best response is either $e_I = 0$ or given by the larger of the two solutions to $\partial V_I / \partial e_I = 0$, where $\partial V_I / \partial e_I$ is given by equation (A11).

We now turn to the defender. Taking the partial derivative of V_D with respect to e_D yields:

$$\frac{\partial V_D}{\partial e_D} = \begin{cases} \int_{\zeta_0}^u \frac{\rho\eta}{\bar{\theta}u} \frac{1}{e_I^\eta e_D^{1-\eta}} \frac{1}{\zeta} d\zeta \tilde{v}_D - c_D & \text{if } e_D < \bar{e}_D(e_I), e_I > 0 \\ -c_D & \text{otherwise} \end{cases} \quad (\text{A13})$$

Again note that e_D appears in the lower limit of the integral on the first line of equation (A9). The corresponding derivative term cancels with the one obtained for the first term of the same line of equation (A9).

Suppose that $e_I = 0$ (such that $\bar{e}_D(e_I) = 0$). Then, the second line of equation (A13) implies that no $e_D > 0$ can be a best response. Therefore, the best response is $e_D = 0$. Consider now $e_I > 0$. The second line of equation (A13) implies that $e_D > \bar{e}_D$ is never a best response in this case. Solving the integral in the first line yields

$$\frac{\partial V_D}{\partial e_D} = \frac{\rho\eta}{\bar{\theta}u} \frac{1}{e_I^\eta e_D^{1-\eta}} \left[\ln \frac{u(\bar{\theta}-\rho)}{\rho} - \eta \ln \left(\frac{e_D}{e_I} \right) \right] \tilde{v}_D - c_D \quad \text{if } e_D < \bar{e}_D, e_I > 0. \quad (\text{A14})$$

Clearly, $\lim_{e_D \rightarrow 0} \partial V_D / \partial e_D = \infty$ in this constellation (recall that $0 < \eta < 1$). It follows that

$e_D = 0$ is never a best response to any $e_I > 0$. As e_D approaches $\bar{e}_D$ from below, the expression in square brackets in equation (A14) approaches 0. Hence, we have $\lim_{e_D \nearrow \bar{e}_D} \partial V_D / \partial e_D = -c_D$ for the respective constellation. Moreover, it is straightforward to show that, for $e_D < \bar{e}_D$ and $e_I > 0$, we have $\partial^2 V_D / \partial e_D^2 < 0$. To summarize the results for the defender's best response:

LEMMA 2 (i) If $e_I = 0$, the defender's best response is $e_D = 0$. (ii) If $e_I > 0$, the first-order condition $\partial V_D / \partial e_D = 0$, where $\partial V_D / \partial e_D$ is given by equation (A14), has a unique solution, $e_D^* \in (0, \bar{e}_D)$. This solution is the unique best response and the second-order optimality condition for an interior optimum holds. Moreover, $e_D = 0$ is never a best response to any $e_I > 0$.

Equilibrium

Lemmas 1(i) and 2(i) imply that no equilibrium involves $e_I = 0$ or $e_D = 0$. Moreover, the lemmas' second parts imply that, if an equilibrium exists, the corresponding best responses are given by unique interior solutions characterized by standard first-order conditions:

LEMMA 3 If an equilibrium exists, it is unique and determined by the first-order conditions $\partial V_I / \partial e_I = 0$ and $\partial V_D / \partial e_D = 0$, where the two derivatives are given by equations (A11) and (A14), respectively.

Setting the expressions in equations (A11) and (A14) to zero, we obtain equation (11), the key equation of Proposition 1. The expressions in (12), which are also part of the proposition, follow immediately from equations (11), (2), and (9).

However, since the objective functions (A6) and (A9) are neither strictly concave nor everywhere differentiable, the first- and second-order conditions may not be sufficient for a best response (there may be corner solutions). To complete the proof, we need to verify three more properties of the candidate equilibrium choices as given by the first-order conditions: (i) $e_I > \underline{e}_I > 0$ and $0 < e_D < \bar{e}_D$ (as a consistency check); (ii) the second-order conditions (again as a consistency check); (iii) both players are better off by choosing $e_k > 0$ rather than $e_k = 0$, $k = I, D$. In other words, appropriate versions of "participation constraints" for the belief contest must hold. Specifically, we must show that these conditions hold if and only if $\tilde{v}_I / c_I$ is sufficiently large.

By setting the expressions in equations (A11) and (A14) to zero, and then solving the resulting system of equations for e_I and e_D , we obtain

$$e_I = \frac{\rho \eta \tilde{v}_I / c_I}{\bar{\theta} u} \left(\frac{\tilde{v}_D / c_D}{\tilde{v}_I / c_I} \right)^\eta \left[\ln \left(\frac{u(\bar{\theta} - \rho)}{\rho} \right) - \eta \ln \left(\frac{\tilde{v}_D / c_D}{\tilde{v}_I / c_I} \right) \right] \quad (\text{A15})$$

and

$$e_D = \frac{\rho\eta \tilde{v}_D/c_D}{\bar{\theta}u} \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right)^\eta \left[\ln \left(\frac{u(\bar{\theta} - \rho)}{\rho} \right) - \eta \ln \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right) \right]. \quad (\text{A16})$$

Clearly, $e_I > 0$, $e_D > 0$ if and only if in equations (A15) and (A16) the expression in square brackets is strictly positive. This is the case if and only if

$$\frac{\tilde{v}_I/c_I}{\tilde{v}_D/c_D} > \left(\frac{\rho}{u(\bar{\theta} - \rho)} \right)^{\frac{1}{\eta}}. \quad (\text{A17})$$

Moreover, using equations (A3), (A8), and (11), it follows that $e_I > \underline{e}_I$ and $e_D < \bar{e}_D$ if and only if equation (A17) holds. Since a sufficiently high value of $\tilde{v}_I/c_I$ implies condition (A17), it also implies $e_I > \underline{e}_I > 0$ and $0 < e_D < \bar{e}_D$. So we have checked condition (i) from above.

Next consider the second-order conditions. For e_D , the second-order condition obviously holds because of Lemma 2(ii). Regarding e_I , it is clear from equation (A12) that the sign of $\partial^2 V_I / \partial e_I^2$ exclusively depends on the sign of the term in square brackets. By using equation (11) in equation (A12), and then rearranging terms, we obtain that the term in square brackets is strictly negative if and only if

$$\frac{\tilde{v}_I/c_I}{\tilde{v}_D/c_D} > \left(\frac{\rho}{u(\bar{\theta} - \rho)} \right)^{\frac{1}{\eta}} e^{\frac{1}{1+\eta}}, \quad (\text{A18})$$

where e denotes Euler's number. Since $e^{1/(1+\eta)} > e^{1/2} > 1$, condition (A18) is stricter than condition (A17), implying that we can ignore the latter. This verifies condition (ii) from above.

Finally, we need to check whether the belief entrepreneurs' "participation constraints" hold, i.e., that they are not better off by choosing a corner solution $e_k = 0$, $k = I, D$. Using equations (A5), (9) (11), and (A15) in the first line of equation (A6) results in

$$V_I = \frac{\tilde{v}_I}{\bar{\theta}u} \left\{ (\bar{\theta} - \rho)u - \rho \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right)^\eta \left[1 + (1 + \eta) \left(\ln \frac{u(\bar{\theta} - \rho)}{\rho} - \eta \ln \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right) \right) \right] \right\}. \quad (\text{A19})$$

With $\tilde{v}_I/c_I$ increasing towards infinity, V_I approaches $\tilde{v}_I / (\bar{\theta}u) (\bar{\theta} - \rho)u > 0$. This follows from the fact that a power function grows sufficiently faster towards infinity than its logarithm. Therefore, $V_I > 0$ holds if $\tilde{v}_I/c_I$ is sufficiently large.

Note that V_D is always positive since the defender serves at least a fraction $\bar{\alpha}\rho/\bar{\theta}$ of the target audience (even without contest participation). Thus, the defender's "participation constraint" concerning the contest is $V_D > \rho/\bar{\theta}\tilde{v}_D$. However, since $\lim_{e_D \rightarrow 0} \partial V_D / \partial e_D = \infty$ for any $e_I > 0$, this participation constraint is implied by the first-order condition. Even so, for completeness,

we state the equilibrium value of V_D :

$$V_D = \frac{\rho \tilde{v}_D/c_D}{u} \left[\frac{1}{\bar{\theta} - \rho} \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right)^\eta + \frac{1}{\bar{\theta}} \left(u - \frac{\rho}{\bar{\theta} - \rho} \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right)^\eta \right) + \frac{1 - \eta}{\bar{\theta}} \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right)^\eta \left(\ln \frac{u(\bar{\theta} - \rho)}{\rho} - \eta \ln \left(\frac{\tilde{v}_D/c_D}{\tilde{v}_I/c_I} \right) \right) \right]. \quad (\text{A20})$$

We conclude that a sufficiently large value of $\tilde{v}_I/c_I$ also implies that condition (iii) holds.

In sum, if $\tilde{v}_I/c_I$ is sufficiently large, the profile specified by equations (A15) and (A16)—which is derived from first-order conditions—indeed consists of mutual best responses and hence constitutes an equilibrium in pure strategies. In this equilibrium, (11) and (12) hold.

To conclude the proof, we need to consider the converse: if $\tilde{v}_I/c_I$ is not sufficiently large, no equilibrium in pure strategies exists. Denote by ν' a critical level for $\tilde{v}_I/c_I$ such that, if $\tilde{v}_I/c_I > \nu'$, conditions (i) – (iii) as stated above hold. Define ν as the smallest such critical level. Then, whenever $\tilde{v}_I/c_I < \nu$, at least one of the conditions (i) – (iii) is violated and hence no (pure-strategy) equilibrium exists.